INNOVATION SPILLOVERS, APPROPRIABILITY, AND ECONOMIC GROWTH

Douglas Hanley

A DISSERTATION

in

Economics

Presented to the Faculties of the University of Pennsylvania

in

Partial Fulfillment of the Requirements for the

Degree of Doctor of Philosophy

2014

Supervisor of Dissertation

____

Ufuk Akcigit, Assistant Professor of Economics

Graduate Group Chairperson

____

George J. Mailath, Professor of Economics, Walter H. Annenberg Professor in the Social Sciences

Dissertation Committee

Dirk Krueger, Professor of Economics

Iourii Manovskii, Associate Professor of Economics

To my friends and family.

Acknowledgements

I am deeply indebted to my advisor, Ufuk Akcigit, who has been a tremendous source of guidance, inspiration, and motivation throughout my graduate studies. Meanwhile, my dissertation committee members, Dirk Krueger and Iourii Manovskii, have been indispensable in providing honest and thoughtful advice on all aspects of the research process. I am truly grateful that they were willing to devote their time and energy to assisting me.

The Penn Economics faculty as a whole has supported me in ways too numerous and diverse to list here, but I would like to thank Hangming Fang, Andy Postlewaite, Hal Cole, and Ken Burdett in particular for their insights and commentary. Equally important have been my wonderful coauthors Daron Acemoglu, William Kerr, and Nicolas Serrano-Velarde. I feel extremely lucky to have had the chance to work with them, and I hope to continue to do so in the future.

I had the good fortune of having a group of smart and ambitious colleagues at Penn who aided me immensely during this process. Of particular note in this regard are Anton Badev, Mark Bognanni, Kurt Mitman, and David Weiss. To me, they exemplify the high ideals of the research enterprise of which we are a part. I would also like to thank Hongseok Choi for his clear-eyed and sage advice over the years.

Finally, I would like to thank my good friend and long-time partner Karam Kang, who gave me courage and hope.

ABSTRACT

INNOVATION SPILLOVERS, APPROPRIABILITY, AND ECONOMIC GROWTH

Douglas Hanley

Ufuk Akcigit

Innovation and technological change are important drivers of economic growth. There is strong evidence that various types of innovation, whether they differ by source, goal, or field, have differing implications for economic outcomes. These arise primarily because of differences in the level of associated externalities (spillovers) and in the ability of innovators to internalize the public benefits from these activities (appropriability). In my research, I focus on identifying the nature and magnitude of these spillovers. Additionally, building on recent advances in the structural modeling of firm incentives, I quantify the extent of appropriation by innovators, particularly as it varies across innovation types. This allows one to provide a detailed accounting of misallocation in the economy and consider policies which can alleviate this.

In the first chapter, entitled “Technological Interdependence”, I study theoretically and empirically how the level of interdependence between new and old technology affects firm dynamics and the incentives for innovation. In the second chapter, entitled “Back to Basics” (joint work with Ufuk Akcigit and Nicolas Serrano-Velarde), we propose and utilize a novel strategy for quantifying the spillovers associated with basic research as they differ from applied research. Finally, in the third chapter, entitled “Transition to Clean Technology” (joint work with Daron Acemoglu, Ufuk Akcigit, and William Kerr), we construct and estimate a joint model of the climate-economy system and investigate the effects of various carbon policies.

Contents

List of Tables

List of Figures

Chapter 1

Technological Interdependence

1.1 Introduction

The innovation process is known to be highly cumulative. New ideas are created because inventors "stand on the shoulders of giants" that preceded them. However, the extent to which new technology is dependent upon old technology varies substantially from field to field. In some areas, such as pharmaceuticals, new technologies often replace existing ones, rendering them obsolete. Here creative destruction is a natural byproduct of innovation. In other areas, such as computer software, new technologies complement existing ones and are integrated with them into a final product. In this setting, technologies are generated incrementally, potentially across multiple firms and over long periods of time, necessitating some form of technology transfer between firms.

Motivated by this last consideration, I proxy the level of technological interdependence in an industry by the rate of patent transfer between firms, which, following the literature, I refer to as sequentiality. Using this *index of sequentiality*, I find that more sequential industries have higher profitability, higher variance of firm growth, lower exit rates, and lower rates of patent expiry[1]. These trends may at first seem puzzling, but as I will show, they in fact arise naturally from a model of firm-driven technological progress featuring heterogeneity across industries in the level of sequentiality. By studying such a model and constraining it with data, we can address a number of important questions. For instance, how does the appropriability of the returns to innovation vary with sequentiality? Does cross-industry heterogeneity in sequentiality produce substantial research misallocation? And finally, what role can patent policy play in this setting?

Estimating this model using firm-level data on patenting and balance sheet information, each of the trends noted above is matched qualitatively, and a large fraction of the variation is accounted for quantitatively. I show theoretically that the larger the sequentiality in a particular industry, the more severe the monopoly distortions induced by a particular level of innovation. This leads to an overallocation of research inputs into more sequential industries. In line with this result, I find that implementing an optimal industry dependent patent policy, which features weaker patent protection in more sequential industries, can remedy a substantial fraction of this misallocation, over and above an optimal uniform patent policy.

This paper contributes to the existing body of literature along both empirical and theoretical dimensions. First, regarding theory, I construct a parsimonious, micro-founded model of sequential innovation and endogenous technological change that formalizes the process by which new ideas are generated, built upon, and subsequently transferred between firms or rendered obsolete. Sequential innovation has already been given treatment in the literature on innovation and endogenous growth, notably in Green and Scotchmer (1995) and Bessen and Maskin (2009), as well as Hopenhayn, Llobet, and Mitchell (2006), who analyze the inherent trade-off present between rewarding incumbents and subsequent innovators that will replace them. This model captures the same trade-off while incorporat-

[1]Patent holders must pay maintenance fees at 4, 8, and 12 years after granting or face permanent expiry.

ing features of more empirically focused models of firm dynamics such as those of Klette and Kortum (2004) and Lentz and Mortensen (2008).[2] I characterize the innovation decisions of firms in a manner that provides intuition for the various economic forces at play and solve for a variety of observable quantities.

In the model, new innovations have differing degrees of dependence on existing technology. High levels of dependence (sequential innovations) necessitate some form of patent sales agreement between the owners of existing technology and new innovators.[3] Conversely, low levels of dependence (independent innovations) necessitate little adjudication of rights between firms as the new innovation simply renders the old one obsolete, leading to the expiration of the original patent. Laitner and Stolyarov (2013) entertain a similar distinction in a model of exogenous innovation. In my model, innovation is endogenously determined and the frequency of sequential innovation varies from industry to industry.

As has been done in Akcigit and Kerr (2010) and Atkeson and Burstein (2010), firms can engage in two types of innovation: external, where they innovate on product lines owned by other firms, and internal, where they innovate on their own product lines. The effect of sequentiality on the rate of external innovation is *ex ante* ambiguous due to the presence of two countervailing forces. First, the value of owning an existing product line is larger in more sequential industries as they feature lower rates of creative destruction from competitors and larger streams of payments from subsequent sequential innovators who buy their patents. However, because of the increased probability of sequential innovation, which necessitates a payout to the existing incumbent, the net effect on innovators will be ambiguous. This stands in contrast to the model presented in Akcigit, Celik, and Greenwood (2013), which features the positive effect of revenues from patent sales, but not

[2]These in turn build upon foundational works such as Romer (1990), Aghion and Howitt (1992), and Grossman and Helpman (1991), as well as numerous other works produced in the interim. See Aghion, Akcigit, and Howitt (2014) for a very recent survey.

[3]I assume that firms always sell their patents rather than licensing them. In the model, this will always be the optimal type of agreement due to monopolistic distortions.

the inhibition of follow on innovation due to continuing patent protection. In the case of internal innovation, the picture becomes clearer, as only the positive effect described above remains.

Broadly speaking, the sequentiality dimension introduced here fills a gap between two classes of models commonly studied in the endogenous growth literature. That is, most models either feature firms that face no threat of replacement from other innovators at the product line level, as in the expanding variety model of Romer (1990), or just the opposite, that firms innovate solely on other firms' product lines and can take over production at will upon a successful innovation, as in Aghion and Howitt (1992), Grossman and Helpman (1991), or Klette and Kortum (2004). The model presented in this paper will act as a bridge between the two extremes presented above. In the extreme of full sequentiality, much of the gains from innovation will be internalized through repeated selling of patent rights down the quality ladder, though distortions from bargaining will complicate this process slightly. In contrast to the Romer model, however, this will come at the cost of a buildup of monopoly power. In the extreme of no sequentiality, we find ourselves with a standard model of creative destruction.

The second contribution of this paper is to enrich our understanding of the data on patenting and innovation by firms. To study cross-industry differences in the sequentiality of innovation, I propose a method of classifying technology classes–an index of sequentiality–based upon the fraction of patents that are transferred in their lifetime. Looking back to the introductory examples, patents in the major pharmaceutical patent classes are transferred 15% of the time, while the same figure for telecommunications is twice as large at 30%. Using this ordering, I document a variety of trends in both the patent data and in linked firm-level data. One would naturally expect the level of sequentiality in a particular industry to have an effect upon innovations dynamics in that industry. For instance, highly sequential industries should feature lower rates of patent obsolescence, as patents are more likely to be built upon and integrated into a larger portfolio rather than being replaced by a new type of technology. This buildup of larger patent portfolios should in turn cause profits to rise, as leading firms will have a larger technological lead over their nearest competitor.

In the data description section below, I document that these trends are in fact present in the data, and I enumerate other trends observed in the cross-industry data, namely that more sequential industries feature lower exit rates and higher variance in firm growth. The former is a natural implication of the reduction in the rate of creative destruction in more sequential industries. The latter effect arises from interaction with firm heterogeneity. Higher sequentiality and the concomitant patent transfers result in the agglomeration of more productive control into the hands of high quality firms. This magnifies the persistent growth differences between firms of differing quality, resulting in more volatile firm growth overall, particularly over longer time periods.

In addition to cross-industry statistics, there are notable trends occurring at the firm and patent level. There is a strong tendency for patents to flow from older and larger firms to younger and smaller firms, with the age dimension showing a distinctly stronger trend than the size dimension. This echoes the finding of Figueroa and Serrano (2013) that small firms receive a disproportionate amount of patent transfers.[4] These facts, in conjunction with the cross-industry trend in firm growth volatility lend support to the notion that patent transfers reflect an underlying process of reallocation amongst firms. This is particularly compelling given the strong evidence that younger, smaller firms excel in many measures of firm performance such as growth and profitability (both in the data presented here and in other works such as Akcigit and Kerr (2010) and Acemoglu, Akcigit, Bloom, and Kerr (2013)).

[4]They use firm size information (greater or less than 500 employees) contained in patent renewal applications.

The third contribution of this paper is to estimate the proposed model using data on public firms and patents in the US, provide a detailed quantitative analysis of the results, and study the implications for optimal patent policy. I use a Simulated Method of Moments (SMM) estimator to match various features of the US data on patent grants, transfers, and expiry and on firm level growth rates and profitability. The patent data comes from the USPTO/Google database and includes data on filings, grants, expiry, and transfers. Data on patent transfers in particular has not been utilized extensively in the literature, especially in a structural setting. The patent data is aggregated to the firm level and matched to Compustat balance sheet data using sophisticated name matching routines.

The estimated model is able to match the targeted moments quite well. In addition, the model can match various non-targeted features of the data, including some of the major trends noted above. The resulting eight-year patent expiration rate over all industries is 39%, compared to 34% in the data, while the standard deviation is 15% compared to 13% in the data. Thus the model captures the proportional variation in the data, while slightly overshooting the magnitude. Other cross-industry trends, such as the relationship between transfer rates and profitability, firm growth variance, and exit rate are qualitatively captured, and the model is able to quantitatively account for a large fraction of these tends.

As predicted by theory, the level of internal innovation rises with sequentiality, with a 52% increase moving from the least to most sequential industry. The level of external innovation, whose dependence on sequentiality was theoretically ambiguous, falls modestly along this dimension, largely due to bargaining distortions that limit the appropriability of back-loaded profit streams. To assess potential misallocation of production and research labor, both within and between industries, I consider a constrained social planner who can choose innovation rates but is still subject to monopoly distortions induced by patenting. I show theoretically that the larger the sequentiality in a particular industry, the more severe the monopoly distortions induced by an increase in innovation rates. Thus for either type of innovation, the planner optimally choses a profile that falls sharply with sequentiality–in the case of external innovation, much more so than in the equilibrium–so as to limit the monopoly distortions caused by the buildup of large, protected technological leads by firms. The equilibrium yields a consumption equivalent welfare 2.5% lower than that of the constrained social planner.

Finally, I investigate the implications of the model for patent policy. I consider both a uniform patent policy and one that depends upon the sequentiality of the industry in question. For the purposes of implementation, the sequentiality of a particular industry can be inferred using the monotonic relationship between sequentiality and the patent transfer rate in the equilibrium of the estimated model. The above discussion of the social planner's optimum leads one to suspect that the optimal patent policy would feature weaker protection in more sequential industries. Indeed, I find that for certain very low levels of sequentiality, an infinite patent is called for. The optimal patent length then decreases from infinity to a minimal value of 6 years in the most sequential industry. This policy results in welfare gains of 1.7% in consumption equivalent terms. For comparison, the optimal constant patent policy calls for a mean patent length of 12 years and delivers welfare gains of only 0.9%.

The remainder of the paper is laid out as follows: in Section 1.2, I describe the data set used and enumerate the notable trends in the data; in Section 1.3, I construct a model capable of matching these facts and describe its equilibrium properties; in Section 1.4, I describe the estimation procedure and results; in Section 1.5, I provide a detailed quantitative analysis of the estimated model with accompanying decompositions and policy experiments; and finally in Section 1.6, I conclude the analysis.

1.2 Empirical Findings

Data on patent grants, expirations, and transfers was acquired from the USPTO Bulk Download site (hosted by Google). Firm names are matched and aggregated into persistent entities based on a name matching algorithm described in Appendix A.1.[5]

For each patent transfer, the following information is provided: (1) the name of the origin firm and destination firm (assignor and assignee), (2) the date that the patent was legally transferred, (3) the date that the transfer was recorded by the patent office, and (4) the purpose of the transfer, amongst other things. In particular, the information on the purpose of the transfer (known as the conveyance text) is used to filter out mergers, licensing agreements, and collateralizations, leaving only simple patent sales, which account for about 85% of the original data points.

The names of the origin and destination firm were matched to the set of entities produced from the patent grant data using the same name matching algorithms. In order to focus on innovating firms and not firms that are simply acquiring patents for other reasons (such as resale), I keep only transfers to firms that have already acquired patents through filing and granting. This eliminates firms that act solely as patent brokers. Furthermore, to exclude instances where conglomerates are transferring patents amongst their constituent units, I eliminate transfers where the origin and destination firm names are sufficiently close, using a more aggressive version of the original name matching algorithm (this is also described in Appendix A.1).

To enrich the data on patent grants, I also use data on the payment of patent maintenance fees. Firms must pay fees to the US patent office after 4, 8, and 12 years from the time of granting. If these fees are not paid, the patent expires permanently. If a patent is maintained through the initial 12 year period, it remains valid until 20 years from its filing date.[6] This

[5] Python code to parse, match, and aggregate the USPTO patent data (along with Compustat data) can be found at https://github.com/iamlemec/patents.

[6] Traditionally, the maximum patent length was 17 years from the grant date. The 1994 Uruguay Round

data thus gives discretized information on the active lifespan of a patent. In total, 35% of patents expire after 8 years, while 48% make it to the natural expiry date. Using the data on patent expiration gives us fairly direct information on the rate of patent obsolescence and hence a window into the level of product market competition faced by firms. Using this in conjunction with the data on patent transfers helps us understand the importance of sequential innovation and its impact on firm dynamics and the incentives to innovate.

To register a patent reassignment with the USPTO, a firm must pay a one-time $\$40$ flat fee. The bulk of the cost is likely to be found in simply filling out the paperwork. Firms already incur legal fees to arrange the contracts for the transfer deals, so going the extra step to register with the patent office is probably not a huge effort. Patent maintenance fees are slightly higher but still not large compared to the common estimates of patent value in the literature. Some studies, such as Pakes (1986), Pakes and Schankerman (1984), and Bessen and Meurer (2008), have used patent renewal patterns to estimate the distribution of patent valuations. A survey by Griliches (1990) reports that various studies found a highly skewed distribution of patent valuations, with mean valuation estimates in the hundreds of thousands of current US dollars, and an obsolescence rate of between 10% and 20% per year. The fees required for renewal at 4, 8, and 12 years are $\$1600$, $\$3600$, and $\$7200$, respectively. These fees are cut in half for small entities (less than 500 employees), and halved again for "micro entities" (targeted towards individual inventors). This self-reported size information provides useful data on the actual size of particular patenting firms. Figueroa and Serrano (2013) utilize this to study the relationship between firm size and patent transferring activity.

An important consideration is the possibility that firms license the patents of other firms rather than buying them outright. Firms are required to register patent sales or transfers in

Agreements Act changed this to the above criterion. See Graham and Vishnubhakat (2013) for a review of the relevant statutes.

order to retain patent rights for a particular technology. However, there is no such requirement for patent licensing, which is regulated by state law in the US.[7] Approximately 1% of patent transfer entries list licensing as the documented activity, though this cannot be guaranteed to be a complete record. Looking across industries, there is no systematic variation in this fraction of reported licensing activity.

1.2.1 Major Trends

Testing various cross-industry predictions necessitates dividing the sample of patents and firms into particular industry level categories. For this exercise, I employ the level one technology class utilized by the US patent office. There are 714 such classes represented in the full patent grant dataset, with a median size of around one thousand patents. Using the first-level classification provides sufficient granularity to capture the specific features of various technological fields while being large enough to avoid excessive noise in aggregate statistics due to small within-industry samples.

It is also possible to extend patent classification information to the firm level. By assigning to a firm the modal patent class amongst its portfolio of patents, we can look at how various firm characteristics vary with technological field. Though most firms have patents in multiple patent classes, the modal patent class accounts for an average of 50% of a firm's patents. This extension to firm characteristics will be important for analyzing trends in balance sheet data from Compustat. For patent data, much of the analysis can be done purely on the patent level. However, I do analyze the patent data using firm-level class assignment for robustness and find similar results.

Industry level regressions are done using weighted least squares. The weighting used is simply the size of the particular technology class in terms of total patents granted. Data on patenting is available from 1976 to the present. For the facts below, I look at the five-year

[7] See Dykeman and Kopko (2004) for an overview of the relevant statutes and case law.

period from 1995 to 2000. This allows sufficient lead time to have realistic values for firm patent stocks, which is the count of patents that are unexpired at any given time. In a steady state world, a lead in time equal to or greater than the patent length suffices. Additionally, it allows sufficient lag time to analyze future transfer, maintenance, and citation activity. The correlations presented below are also weighted by patent class size. In each of the figures accompanying the following facts, the point size represents the total number of patents granted (the weight) and the color represents the the numerical patent class, which ranges from 1 (lightest) to 800 (darkest). Because patent classes have been added incrementally over time, the patent class (color) also provides a very good proxy for how recently the patent class was created.

Fact 1. *There is a negative correlation between patent transfer rates and patent expiry rates across industries.*

FIGURE 1.1: TRANSFER–EXPIRY RELATIONSHIP

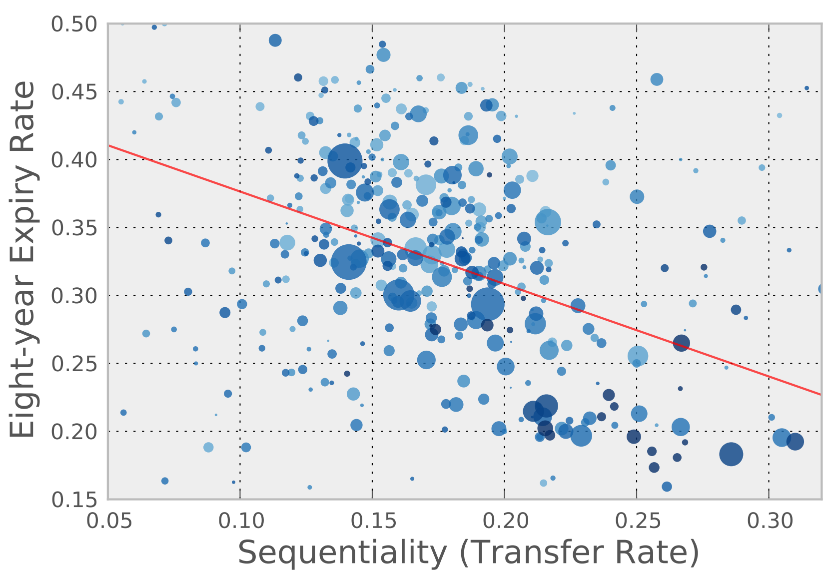

Relationship between patent transfer and expiry. Red line: WLS regression $\beta = -0.33$. Correlation is $\rho = -0.34$.

The transfer rate is the fraction of patents granted in the data window that are transferred in their lifetime, while the expiry rate is the fraction of patents granted in the data window that are not renewed after the first eight-year window and hence expire. The extent of the negative relationship is portrayed in Figure 1.1. This trend highlights a central feature of the model proposed herein, namely that industries with innovation that is more sequential in nature will see higher transfer rates due to higher levels of technological interdependence and lower levels of creative destruction for the same reason.

Fact 2. *There is a positive correlation between patent transfer rates and firm profitability across industries.*

FIGURE 1.2: TRANSFER–PROFITABILITY RELATIONSHIP

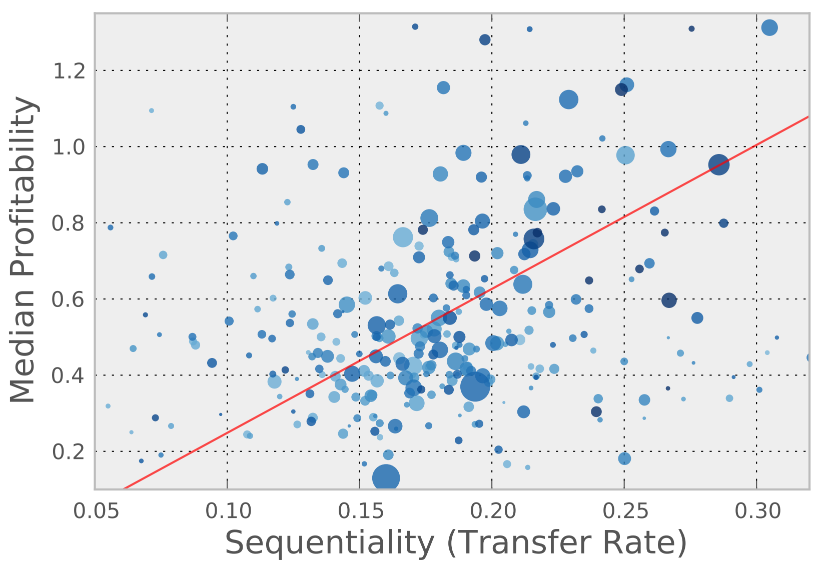

Relationship between patent transfer and median profitability. Red line: WLS regression $\beta = 3.85$. Correlation is $\rho = 0.40$.

Here I calculate the firm-level profitability as the ratio of revenue to the cost of goods sold, as given in the Compustat data. This excludes operating costs and allows us to look purely at variable cost relationships, which are a primary variable of interest in quality ladder models. I then look at the median value within each industry. Intuitively speaking, we would expect that industries with high transfer rates see more aggregation of monopoly power and less reduction through creative destruction (from Fact 1). This then leads to higher markups over cost being charged and higher profitability. This trend is portrayed in Figure 1.2. Looking at the relationship between log return on sales and transfer rate yields a similar trend, though with more noise on account of sales being in the denominator.

Fact 3. *The variance of firm growth rates is positively correlated with patent transfer rates across industries.*

FIGURE 1.3: TRANSFER–GROWTH VOLATILITY RELATIONSHIP

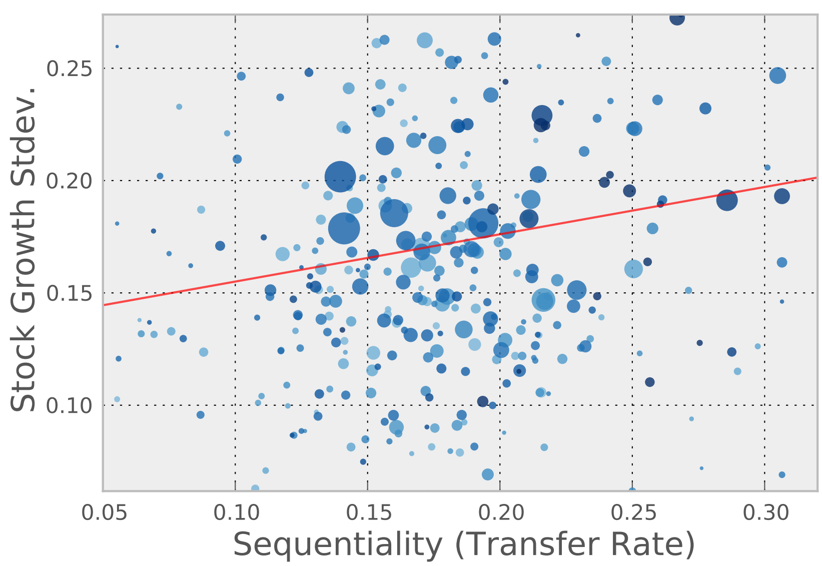

Relationship between patent transfer rate and growth volatility. Red line: WLS regression $\beta = 0.21$. Correlation is $\rho = 0.19$.

Here the variance of firm growth is calculated using a number of common firm size statistics, including patent stock, employment level, and earnings. Regardless of the metric used, industries with higher transfer rates display higher variance in firm growth rates. One mechanism that could generate this trend is that patent transfers allow higher quality firms to grow faster at the expense of lower quality firms. Furthermore, because of lower levels of creative destruction present in industries with higher transfer rates, larger firms are also more able to protect their market share, thus widening the gap further between the performance of high and low quality firms. The fact that this gap is larger in industries with higher transfer rates leads to a higher variance of firm growth in these industries. In order to account for this potentially being driven by level differences, due to certain sectors growing or shrinking in the aggregate, I look at the variance of log growth rates within industry. Using this measure, changes in industry size will come out as a common additive factor for each firm and thus will not affect the computed variance.

Fact 4. *There is a negative correlation between firm exit rates and patent transfer rates across industries.*

FIGURE 1.4: TRANSFER–EXIT RATE RELATIONSHIP

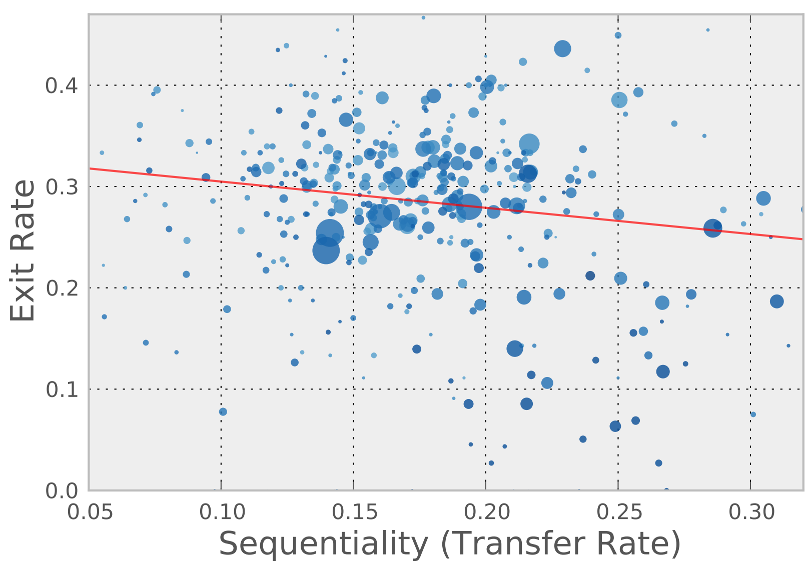

Relationship between patent transfer rate and five-year exit rate. Red line: WLS regression $\beta = -0.26$. Correlation is $\rho = -0.16$.

In line with Fact 1, interpreting lower expiry rates as indicating lower rates of creative destruction, one would also expect lower exit rates in industries with high transfer rates. This is indeed born out in the data, where exiting from the sample of patenting firms is taken to occur when a firm no longer displays any patenting activity.

Fact 5. *Patent transfers are directed primarily toward young and small firms. Firms aged less than* 10 *years account for only* 14% *of the patent stock, while receiving approximately* 57% *of patent transfers.*

FIGURE 1.5: AGE/SIZE CUMULATIVE FRACTIONS

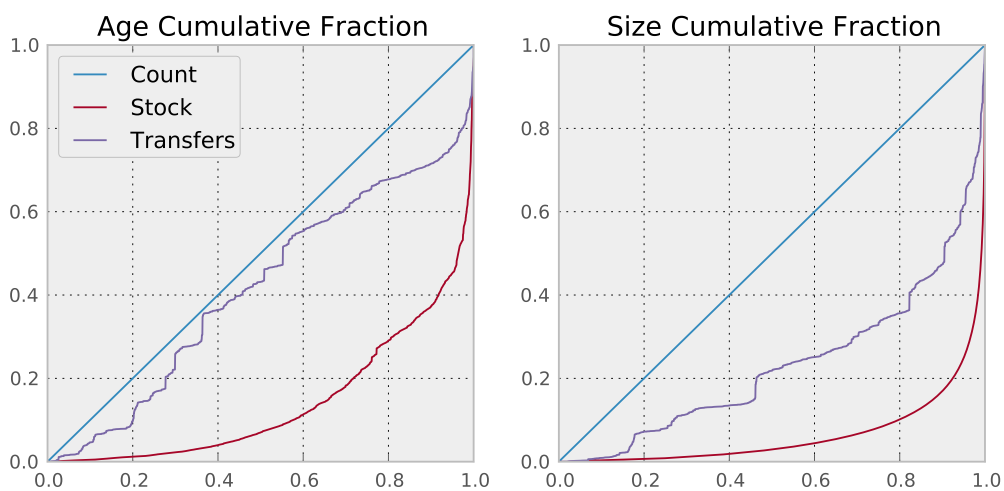

Cumulative fractions of patenting activity by firm age and size (patent stock).

Small and young firms account for a disproportionate share of patent transfer receipts relative to their size. The size distribution of firms is highly skewed. Therefore small and young firms will invariably constitute a relatively small fraction of the patent stock. The analogous numbers to those presented in the fact above for firms below the 80^{th} size percentile are that they account for 11% of the patent stock and 36% of patent receipts. This figure is still disproportionate in terms of firm size but not nearly as much as that for firm age. Looking at patent filings and patent transfer origination, we see similar but less extreme trends. Young firms account for 30% of filings and 33% of originations. The analogous figures for small firms are 19% of filings and 29% of originations.

Conventional wisdom dictates that small firms sell technologies to larger firms who are in a better position to bring products to market or integrate them into existing production processes (for instance, see Phillips and Zhdanov (2013) for a theoretical discussion of this dynamic). However, the data indicate a bulk flow towards small and, to a greater extent, young firms. This is consistent with a model where firms are imbued with persistent (though mutable) types and patent reassignment is a mechanism by which production control is transferred to higher quality firms.

1.2.2 Mechanism Evidence

In addition to the cross-industry trends presented above, I also provide more detailed evidence on the proposed mechanism. Of central importance is documenting that sequential innovation, in the sense of direct technological dependence, is the primary driving force behind the observed transfers of patent ownership. One might naturally expect citation patterns to shed light on this issue.

As described above, I classify firms into various technological categories by using the modal patent class in their portfolio. The average firm cites only 3% of the other firms in its patent class. Relating this to the data on patent transfers, 52% of firm pairs that have transfers between them also cite each other. Breaking this down by the direction of transfer, the destination firm cites the origin firm in 50% of cases, while the reverse happens in only 23% of cases. Thus it is quite rare that the origin firm cites the destination firm without the reverse also happening. Higher citation rates between firms that transfer patents, in of itself, may merely indicate that these firms are closer together in a technological sense and that such firms are more likely to cite one another. However, the asymmetry in citation rates between the different directions of transfer lends further support to the notion that transfer are acting as a mechanism for reallocation of production and research towards higher quality firms.

Fact 6. *The internal citation rate is uncorrelated with transfer rates, while the external citation rate is highly correlated with transfer rates across industries.*

FIGURE 1.6: TRANSFER–EXTERNAL CITATION RELATIONSHIP

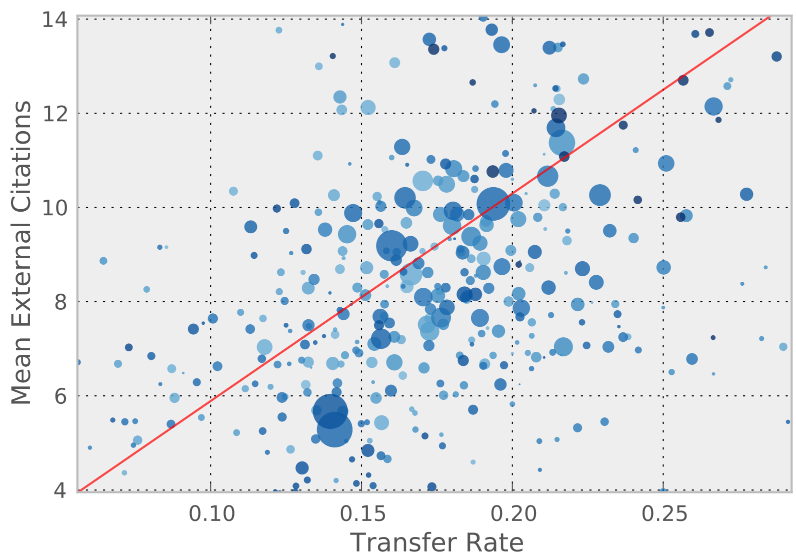

Relationship between patent transfer and external citations. Red line: WLS regression $\beta = 44.1$. Correlation is $\rho = 0.49$.

The data on citations show that the average number of citations per patent is higher in more sequential fields. Breaking these citations down into those that cite within firm and those that cite other firms, we see that internal citations are not related to industry sequentiality, while external citations are strongly related. This is consistent with the notion that technological dependency is the primary determinant of whether a new innovator must purchase the rights to existing ideas in the field. I use two measures of internal/external citations classification in this instance, both of which display the same trends across industries. First, I simply look at the aggregate number of internal and external citations per patent by industry. Second, I follow Akcigit and Kerr (2010) in classifying a patent as internally oriented if more than 50% of its citations are self-directed. Regardless of the measure used, I observe the trends noted above, as portrayed in Figure 1.6.

Fact 7. *There is a positive correlation across industries between the fraction of patents acquired by transfer and acquisition expenditures.*

The mechanism presented in this paper operates at the patent and product line level and firms are largely just collections of various product lines. However, it may be the case that resolution of patenting rights occurs not simply through the direct buying and selling of patents but at higher levels of aggregation such as the subsidiary of a conglomerate or entire firms of varying size. Particularly in the case of small firms or entrants, whose value is often encompassed in a single product line, this may be an important dynamic. And indeed, looking at the data on acquisition activity, we see a positive relationship between that and patent transfer rates. While it is certainly the case that there are other forces that can drive M&A activity, this trend indicates that the technological landscape plays an important role. Discussion of how this data may be mapped into the model is deferred until the section on estimation.

1.3 Model

In this section, I present a continuous-time model of firm dynamics and endogenous technological growth. After specifying the various elements of the model, I characterize the dynamic equilibrium. I then focus on the case of the steady state, with the objective of producing predictions that map into the trends described in the previous section.

1.3.1 Consumers

There is a unit mass of immortal consumer-workers in the economy. Each has one unit of labor that they supply inelastically. Their utility is a function on the infinite flow stream of consumption starting at time $t = 0$. In particular, they discount the future at rate δ and have an instantaneous utility function of $u(c)$ with constant relative risk aversion parameter σ. Thus their utility function can be expressed as

$$U(c) = \int_0^{\infty} \left[\frac{c(t)^{1-\sigma} - 1}{1 - \sigma} \right] \exp(-\delta t) dt$$

where c is a consumption profile that specifies the level of consumption at each point in time. All agents earn a wage w from employment. They also have access to a risk-free bond paying interest r and having zero net supply in the aggregate. Let their bond holding profile be the function a. Their budget equation is then given by

$$c + \dot{a} = w + ra$$

where time dependence is suppressed for notational convenience. There is a single final good Y for consumption, which is normalized to have a unit price at each point in time. Because all costs are purely in terms of labor, the final good resource constraint for the economy is simply $c = Y$. The associated Euler equation for this problem the delivers the result

$$g \equiv \frac{\dot{Y}}{Y} = \frac{\dot{c}}{c} = \frac{r - \delta}{\sigma}$$

Letting the common growth rate of Y and c be denoted by g, we arrive at $r = \delta + \sigma g$.

Finally, each worker can choose to be a production worker or a research worker. Let the respective masses of each type be L_P and L_R. The labor market clearing condition is then

$$L_P + L_R = 1$$

Given the costless choice between being employed as a production worker and a research worker, any equilibrium of the model will feature a common wage w for these two occupations.

1.3.2 Production

The final good is produced by combining a unit continuum of intermediate goods with the well-known Dixit-Stiglitz aggregator with unit elasticity

$$Y = \exp\left[\int_0^1 \ln(y_j)\, dj\right]$$

This technology is operated competitively by a continuum of firms. Each buys up certain quantities intermediate goods for respective prices p_j, combines them into a final good, and sells that for the normalized unit price. The objective of one such firm is then

$$\Pi = \max_{y_j}\left\{\exp\left[\int_0^1 \ln(y_j)\, dj\right] - \int_0^1 p_j y_j\, dj\right\}$$

Optimality dictates that $Y = p_j y_j$ for all j. Constant returns to scale ensure that these firms make zero profit in equilibrium.

Each intermediate good is in turn produced using a linear technology of the form $y_{jf} = q_{jf}\ell_{jf}$, where the f subscript allows for the fact that different firms have different know-how in producing each particular good. Now consider the firm with the most advanced production technology and simply let $q_j = \max_f \{q_{jf}\}$. Furthermore, let the next best producer be $q_{-j} = q_j/\lambda_j$, where $\lambda_j \geqslant 1$. The leading firm can then price the runner up out of the market by charging a price $p_j = w/q_{-j} = w\lambda_j/q_j$, thus selling $y_j = Yq_j/(w\lambda_j)$ and employing $\ell_j = Y/(w\lambda_j)$ labor. This leads to profits of

$$\pi_j = p_j y_j - w\ell_j = (1 - \lambda_j^{-1})Y$$

Thus the labor utilization and profit of each product line are purely a function of the technological lead λ_j and the not the absolute productivity value q_j. Using tilde to denote values normalized by Y, the labor utilization and profit for a product line with technological lead λ are given by

$$\ell(\lambda) = \frac{\lambda^{-1}}{\widetilde{w}} \qquad \text{and} \qquad \widetilde{\pi}(\lambda) = 1 - \lambda^{-1}$$

Having fully characterized the production decisions of intermediate goods producing firms, we can now use these production values to address the innovations decisions of firms.

1.3.3 Innovation

It was shown above that the only firm relevant variable for a particular product line is the technological lead λ_j. Each firm in the economy can thus be characterized simply as a portfolio of technological lead values for product lines in which it is the leading producer. For a firm with n product lines, denote such a vector by

$$\vec{\lambda} = (\lambda_1, \dots, \lambda_n)$$

Following the model presented in Klette and Kortum (2004), the external innovation production technology specified here uses only labor as an input and scales up linearly with firm size. In particular, a firm with n product lines can achieve a Poisson flow rate of innovation X by employing

$$C(n, X) = nc(X/n)$$

researchers. In other words, a firm must use $c(x)$ researchers per product line to achieve an innovation rate of x per product line, where $x = X/n$. Firms can also undertake internal innovation on one of their existing product lines. Here I allow the cost of internally oriented innovation to scale with a firm's technological lead (λ). This is motivated partly by tractability and partly through existing empirical evidence. Akcigit and Kerr (2010) find that the intensity of internally oriented innovation does not scale strongly with firm size. In order to generate such a result, we can use the form

$$d(\lambda, z) = \lambda^{-1}d(z)$$

where z is the flow rate of internal innovation. This ensures that the internal innovation rate will actually be constant across firms, regardless of their technological lead for a given product line.

The functional forms given allow one to treat each firm simply as a collection of research labs, each associated with a particular product line. Denote the value of a research lab with technological lead λ by $V(\lambda)$. A firm with portfolio $\vec{\lambda}$ will then have value

$$V(\vec{\lambda}) = \sum_{i=1}^{\infty} V(\lambda_i)$$

A research lab will accrue profits from production and generate innovations. Successful innovations will garner new research labs with their associated production and innovation capabilities. Now we can characterize all firm decisions by addressing the problem at a product line level.

When an external innovation occurs, the state-of-the-art productivity of a random product line is incremented by a random factor β. For internal innovations, the productivity in the target product line is incremented by a factor drawn from the same distribution. Measuring and constructing systematic data on innovation sizes is difficult. However, in a limited sample, Scherer (1965) finds evidence that a Pareto distribution is appropriate. Meanwhile, Pakes and Schankerman (1984) are able to fit data on patent expiry in multiple countries using a Pareto distributed innovation size distribution, while Kortum (1997) find a Pareto distribution to be consistent with aggregate trends in research, growth, and patenting. Thus I assume that, β is drawn from a Pareto distribution $F(\cdot)$ with tail index $1/\kappa$ and having cumulative density

$$F(\beta) = 1 - \beta^{-1/\kappa}$$

The inverted tail index is used as parameter so as to facilitate analogy to the step size parameter typically present in endogenous growth models. It will be of use later to know that the expected value of $\log(\beta)$ is simply κ, meaning a variable receiving such increments at Poisson rate x will have expected growth rate κx.

Upon the arrival of an internal innovation, with probability α the innovation is sequential and is dependent upon previous innovations. In this case, the innovating firm and the incumbent firm initiate a bargaining process by which either the existing patents of the incumbent are sold to the new innovator or the incumbent buys the new innovation and incorporates it into its portfolio. Conversely, with probability $1 - \alpha$, an innovation is independent. In this case, the new innovator assumes production responsibilities and the incumbent is summarily displaced.

Firms also face the rate of incoming external innovations by other firms. Let these events arrive at rate τ. Finally, all patents in a particular product line expire at rate b, meaning the technological lead goes to zero and production profits vanish. When this happens, the firm retains its research capacity in that product line but is displaced upon any subsequent innovation by another firm. Denote the present expected value of successful innovation by $\bar{V}$. Because both profits and labor costs scale up with output, I consider the output-normalized value of a patent protected product line with technological lead λ

$$\begin{aligned}\delta_F V(\lambda) - \dot{V}(\lambda) = {} & \widetilde{\pi}(\lambda) \\ & + \max_x \left\{-\widetilde{w}c(x) + x\bar{V}\right\} + \max_z \left\{-\widetilde{w}\lambda^{-1}d(z) + z(\mathbb{E}V(\beta\lambda) - V(\lambda))\right\} \\ & + \alpha\tau p(\mathbb{E}V(\beta\lambda) - V(\lambda)) + (1-\alpha)\tau(0 - V(\lambda)) + b(V_0 - V(\lambda))\end{aligned}$$

where $\delta_F = r - g$ is the effective discount rate used by the firm. The value of a product line without patent protection is simply

$$\delta_F V_0 - \dot{V}_0 = \max_x \left\{-\widetilde{w}c(x) + x\bar{V}\right\} + \max_z \left\{-\widetilde{w}\lambda^{-1}d(z) + z(\mathbb{E}V(\beta\lambda) - V_0)\right\} + \tau(0 - V_0)$$

Notice that $V(1) \neq V_0$, as expired product lines still retain their research capacity. To know the value of successful innovation, we must know the economy-wide distribution over λ. For now, denote the cumulative density for this variable by $\mu(\cdot)$. Furthermore, let μ_0 be the mass of products whose patent has expired (meaning $\lambda = 1$) and $\mu_+(\cdot)$ be the cumulative density over those products whose patents are not expired. The value of successful external innovation is then given by

$$\bar{V} = [(1-\alpha) + \alpha\mu_0)]\, \mathbb{E}V(\beta) + \alpha(1-p)\int_1^\infty (\mathbb{E}[V(\beta\lambda) - V(\lambda))d\mu_+(\lambda)$$

As discussed earlier, each product line has a production value and a research value. The production value and internal research value will be a function of the technological lead, while the external research value will be independent of that variable since future innovations are undertaken on random external product lines. Thus it is useful to define the option values

$$\Omega_x = \max_x \left\{-\widetilde{w}c(x) + x\bar{V}\right\}$$
$$\Omega_z(\lambda) = \max_z \left\{-\widetilde{w}\lambda^{-1}d(z) + z(\mathbb{E}V(\beta\lambda) - V(\lambda))\right\}$$
$$\Omega_0 = \max_{z_0} \{-\widetilde{w}d(z_0) + z_0(\mathbb{E}V(\beta\lambda) - V_0)\}$$

Notice that because successful internal innovation in an expired product line results in increased protection from external innovation, the incentive structure is slightly different. The product line value function expressions can then be simplified to

$$(\delta_F + (1-\alpha)\tau)V(\lambda) - \dot{V}(\lambda) = \pi(\lambda) + \Omega_x + \Omega_z(\lambda) + \alpha\tau p(\mathbb{E}V(\beta\lambda) - V(\lambda)) + b(V_0 - V(\lambda))$$
$$(\delta_F + \tau)V_0 - \dot{V}_0 = \Omega_x + \Omega_0$$

Having characterized the firm value functions and their dynamics, we must also address the evolution of the state space, which in this case consists of the technological lead distributions. First, the respective masses of expired and unexpired product lines will satisfy the flow equations

$$\dot{\mu}_0 = b\mu_+ - (\tau + z_0)\mu_0 \qquad \text{and} \qquad \dot{\mu}_+ = (\tau + z_0)\mu_0 - b\mu_+$$

Focusing on unexpired product lines (where $\lambda > 0$), the distribution will satisfy the flow equation

$$\dot{\mu}_+(\lambda) = (b + (1-\alpha)\tau)\left[F(\lambda) - \mu_+(\lambda)\right] - (\alpha\tau + z)\int_1^\infty \left[1 - F(\lambda/\lambda')\right] d\mu_+(\lambda') \tag{1.1}$$
$$+ \left(\frac{\dot{\mu}_+}{\mu_+}\right)\left[F(\lambda) - \mu_+(\lambda)\right] \tag{1.2}$$

The first two terms are what we would expect in the case without patent expiry. Independent innovations arrive at rate $(1-\alpha)\tau$ and reset the technological lead to some random value β, and similarly for patent expiry b. Meanwhile, sequential and internal innovations arrive at rate $\alpha\tau + y$ and increment the technological lead by some random value β. The last term simply deals with the fact that there are also product lines flowing into and out of expiry.

1.3.4 Equilibrium

Having described the optimization problems faced by consumers and firms, we can now move on to characterizing their optimal behavior and setting forth conditions for aggregate consistency given the equilibrium variables we have introduced. The aggregate information needed by the firm to make decisions includes the rate of creative destruction τ, the wage rate w, and the interest rate r. Finally, the firm needs to know the state space, namely the distribution over technological leads, which is fully described by the respective masses of expired and unexpired product lines μ_0 and μ_+ and the distribution of technological leads over unexpired product lines $\mu_+(\cdot)$. Eventually, it will be shown that the mean inverse over $\mu_+(\cdot)$

$$\Gamma_+ = \int_0^\infty \lambda^{-1} d\mu_+(\lambda)$$

will suffice for the purposes of the firm and for aggregate consistency. Now posit a linearly separable ansatz for the unexpired product line value function

$$V(\lambda) = A - B\lambda^{-1}$$

Recall that $\tilde{\pi}(\lambda) = 1 - \lambda^{-1}$. Inserting the above into the product line value function and equating coefficients on the constant term and the λ^{-1} terms yields the following characterization of the coefficients

$$(\delta_F + b + (1-\alpha)\tau)A - \dot{A} = 1 + \Omega_x + bV_0$$
$$(\delta_F + b + (1-\alpha)\tau)B - \dot{B} = 1 - \Omega_z - \alpha\tau p B/(1+\kappa^{-1})$$

Here Ω_x is the option value of external innovation. Because of the concavity of the profit function in the technological lead, the gross returns to internal innovation are decreasing. However, since the the cost also decreases by the same proportion, the net returns also scale down with the technological lead. Thus internal innovation shows up in the variable portion of the value function as

$$\Omega_z = \max_z \left\{ -\widetilde{w}d(z) + zB/(1+\kappa^{-1}) \right\}$$
$$\Omega_0 = \max_{z_0} \left\{ -\widetilde{w}d(z_0) + z_0(A - B/(1+\kappa) - V_0) \right\}$$

with $\Omega_z(\lambda) = \lambda^{-1}\Omega_z$. Using these expressions, the expected gain from innovation can be simplified to

$$\bar{V} = ((1-\alpha) + \alpha\mu_0)(A - B/(1+\kappa)) + \alpha(1-p)\mu_+\Gamma_+ B/(1+\kappa^{-1})$$

The labor market clearing condition will include contributions from production, external innovation (from incumbents and entrants) and internal innovation (on expired and unexpired product lines) as delineated below

$$1 = \frac{\Gamma}{\widetilde{w}} + (1+e)c(x) + \mu_0 d(z_0) + \mu_+\Gamma_+ d(z) \tag{1.3}$$

where Γ is the average inverse technological lead over all product lines and satisfies $\Gamma = \mu_0 + \mu_+\Gamma_+$. The flow equation for Γ_+ is described in the next section and depends only on

μ_0, μ_+, and Γ_+ itself. Therefore, with regards to solving the equilibrium by determining the evolution of the state space, Γ_+ is a sufficient statistic for $\mu_+(\cdot)$. In fact, just μ_0 and Γ would be a sufficient state space. However, the three variable specification proves to be notationally cleaner.

Aggregate consistency of the rate of external innovation requires that $\tau = (1+e)x$. Though it is not necessary for the equilibrium solution, the growth rate will naturally be of interest as an implication of this model. Each innovation, regardless of whether it is sequential or independent furthers the state of the art for a particular intermediate good by a random factor β drawn from F. Because of the log-log aggregation in producing the final good, output can be decomposed into

$$Y = QL_P/\Delta$$

where Q is the log aggregate productivity $\log(Q) = \int_0^1 \log(q_j)dj$ and Δ is a measure of labor misallocation given by

$$\begin{aligned}\log(\Delta) &= \log\left[\int_0^1 \ell_j dj\right] - \int_0^1 \log(\ell_j)dj \\ &= \log\left[\int_1^\infty \lambda^{-1} d\mu(\lambda)\right] - \int_1^\infty \log(\lambda^{-1})d\mu(\lambda) \geqslant 0\end{aligned}$$

where the inequality above follows from Jensen's inequality. It is straightforward to show that the growth rate of the aggregated productivity will simply be

$$g = \kappa(\tau + \bar{z}) \tag{1.4}$$

where $\bar{z} = \mu_0 z_0 + \mu_+ z$ is the aggregate rate of internal innovation. Outside of steady state the quantities L_P and Δ can of course also change. So the overall growth rate of output will be composed of contributions from these three factors. In steady state, however, the growth rate of Y will simply be g.

1.3.5 Steady State

The above section fully characterized the dynamic equilibrium of the model. In principle, this characterization could be used to describe the path of the economy starting from any given point in the state space. The usefulness of this capability is dampened by the inherent difficulty in simultaneously identifying the parameters of the model and the position in the state space. Therefore, I focus on the case of steady state.

In steady state, all normalized figures, such as those comprising the value function, will be constant. In addition, the position in the aggregate state space, as defined above, will be invariant. Proceeding from this basis, the firm value function coefficients simplify to

$$A = \frac{1 + \Omega_x + bV_0}{\delta_F + b + (1-\alpha)\tau} \qquad \text{and} \qquad B = \frac{1 - \Omega_z}{\delta_F + b + ((1-\alpha) + \alpha p \kappa/(1+\kappa))\tau} \tag{1.5}$$

The value of a product line where patent protection has expired becomes simply

$$V_0 = \frac{\Omega_x + \Omega_0}{\delta_F + \tau} \tag{1.6}$$

These quantities, in conjunction with the state space position will determine the expected present value from successful innovation, which will in turn determine innovation rates, the growth rate, wages, and other observables of interest.

We now move on to the task of characterizing the steady distribution of technological leads. Imposing steady state on the flow equation for $\mu_+(\cdot)$ object given in Equation 1.1, I find

$$(b + (1-\alpha)\tau)\left[F(\lambda) - \mu_+(\lambda)\right] = (\alpha\tau + z)\int_1^\infty \left[1 - F(\lambda/\lambda')\right] d\mu_+(\lambda')$$

For arbitrary F, the resulting distribution is intractable. However, given our assumption of Pareto distributed step sizes, one can show that the steady state distribution will in fact be Pareto as well.

Proposition 1. *The distribution of technological leads for patent protected product lines is Pareto with*

$$\mu_+(\lambda) = 1 - \lambda^{-1/m}$$

where the tail index parameter satisfies

$$m = \kappa \left[\frac{b + \tau + z}{b + (1-\alpha)\tau} \right] \tag{1.7}$$

Proof. Recall that the cumulative density for β is simply $F(\beta) = 1 - \beta^{-1/\kappa}$. Now posit a similar form for the technological lead distribution with shape parameter m. Plugging this into the flow equation, one can verify that this shape parameter is given the above expression. □

Thus the expected value of $\log(\lambda)$, conditional on being strictly positive is simply m. Here one can see that increasing either α and τ serves to attenuate the technological lead distribution while increasing the patent length b draws it closer to unity. Furthermore, the mean inverse technological lead for unexpired product lines can then be expressed as

$$\Gamma_+ = \frac{1}{1+m}$$

Note. *As an aside, I will note that the assumption of Pareto distributed step sizes is not critical to the equilibrium solution, but is needed to ensure tractability of the technological lead distribution, which simplifies notation in various places. For arbitrarily distributed* β*,* *one can write the flow equation for the quantity* Γ_+ *as*

$$\dot{\Gamma}_+ = (b + (1-\alpha)\tau)\left(\mathbb{E}[\beta^{-1}] - \Gamma_+\right) - (\alpha\tau + z)\left(1 - \mathbb{E}[\beta^{-1}]\right)\Gamma_+ + \left(\mathbb{E}[\beta^{-1}] - \Gamma_+\right)\left(\frac{\dot{\mu}_+}{\mu_+}\right)$$

Imposing $\dot{\Gamma}_+ = \dot{\mu}_+ = 0$ *and solving then yields the equations*

$$\Gamma_+ = \frac{1}{1+m} \qquad \textit{where} \qquad m = (1 - \mathbb{E}[\beta^{-1}])\left[\frac{b+\tau+z}{b+(1-\alpha)\tau}\right]$$

though m *can no longer be interpreted as the tail index of the distribution* μ_+.

The only remaining elements of the state space to be solved for are the aggregate shares of expired and unexpired product lines. Equating the flow equations for these quantities to zero yields the simple solution

$$\mu_0 = \frac{b}{b+\tau+z_0} \qquad \text{and} \qquad \mu_+ = \frac{\tau+z_0}{b+\tau+z_0}$$

Combining these with the results above, the inverse technological lead over all product lines, expired or unexpired, is then

$$\Gamma = \mu_0 + \mu_+\Gamma_+ = \frac{b + (\tau+z_0)/(1+m)}{b+\tau+z_0}$$

Existence A balanced growth path equilibrium of this model is characterized by a vector $(\widetilde{w}, g, A, B, V_0)$ consisting of the wage rate $\widetilde{w}$ satisfying Equation 1.3, the aggregate growth rate g satisfying Equation 1.4, the unexpired product line value coefficients A and B satisfying Equation 1.5, and the unexpired product line value V_0 satisfying Equation 1.6.

Proposition 2. *A balance growth path equilibrium for this economy exists.*

Proof. See Appendix. □

1.3.6 Welfare

As discussed in the previous section, aggregate output can be decomposed into contributions from three components

$$\begin{aligned}\log(Y) &= \int_0^1 \log(q_j \ell_j) dj = \log(Q) + \int_1^\infty \mu(\lambda) \log(\ell(\lambda)) d\lambda \\ &= \log(Q) + \log(L_P) - \log(\Delta)\end{aligned}$$

The term $\log(\Delta)$ is a measure of labor usage heterogeneity, which leads to productive misallocation. This implies that Q is the maximum possible output of the economy and QL_P is the maximal output given a certain amount of production labor L_P. In steady state, this takes on the value

$$\log(\Delta) = \log(\Gamma) + \mu_+ m = \log\left[\frac{b + (\tau + z_0)/(1+m)}{b + \tau + z_0}\right] + \left(\frac{\tau + z_0}{b + \tau + z_0}\right) m \tag{1.8}$$

Notice that, holding innovation rates constant, this is decreasing in the patent length b and increasing in sequentiality α, since m is also increasing in α. Welfare is given according to

$$W = \int_0^\infty \left[\frac{Y(t)^{1-\sigma} - 1}{1-\sigma}\right] \exp(-\delta t) dt$$

Without loss of generality, I can assume that $Q(0) = 1$. Furthermore, we know that $\dot{Q}/Q = g$. Plugging in for Y and evaluating the integral then yields

$$\delta W = \left(\frac{\delta}{\delta + (\sigma - 1)g}\right) \left[\frac{g}{\delta} + \frac{(L_P/\Delta)^{1-\sigma} - 1}{1-\sigma}\right] \tag{1.9}$$

Thus welfare can be easily expressed purely as a function of τ, z, and z_0. One of the main implications of this model is that the welfare effects of monopoly distortions are more severe in more sequential industries. To see this, consider the effect of varying the aggregate innovation rates on the labor misallocation factor Δ.

Proposition 3. *The effect of τ, z, and z_0 on monopoly distortions is greater for more sequential industries. That is, $\frac{\partial \Delta}{\partial \tau}$, $\frac{\partial \Delta}{\partial z}$, and $\frac{\partial \Delta}{\partial z_0}$ are increasing in α. Furthermore, the latter two derivatives are always positive, while the first is positive if $\alpha b > (1-\alpha)z$.*

Proof. See Appendix. □

This is one of the major implications of the model presented here. Not only does more innovation induce higher production labor misallocation in most cases, this effects is larger for more sequential industries. Thus in considering patent policy, where the fundamental trade-off is between incentivizing innovation at the cost of monopoly distortions, the benefit is the same while the cost is larger in more sequential industries.

1.3.7 Social Optima

Before considering various policy interventions or changes, it is important to study this model from a social planner's prospective in order to gain insight into the types and levels of inefficiency present in the decentralized equilibrium. Two types of social planners will be considered. The first is a partially constrained social planner who can control the innovation decisions of firms but is still subject to the outcome of the static product market equilibrium with its associated monopoly distortions. In this case, the patent length is assumed to be the same as in the decentralized equilibrium. Choosing external innovation rate τ, internal unexpired innovation rate z, and internal expired innovation rate z_0 allows one to determine the growth rate g, the production labor utilization L_P, and the labor misallocation factor Δ. Using these, one can compute steady state welfare using Equation 1.9.

The second type is an unconstrained planner who makes both innovation and production decisions for firms. This results in a simple closed form expressions for L_P and g as functions of τ, z, and z_0. The unconstrained planner will optimally choose labor utilization to be equal across product lines, meaning $\ell_j = L_P$ for all j and $\Delta = 1$.

1.3.8 Predictions

The only aggregate variables of interest in this setting are the wage $\widetilde{w}$ and the growth rate g. The following predictions will not be functions of these variables. They will depend only on the within industry variables, which are indexed by α. In general, both τ, z, and z_0 will be functions of α, however their dependence is suppressed here for the sake of brevity.

Transfer Rates The direction of transfer is indeterminate in this model. Let the probability that a transfer goes towards the innovator be q. Knowing this, what fraction of patents can one expect to be transferred in their lifetime? Patents are born at rate $\tau + \bar{z}$. They die at rate b and an rendered obsolete at rate $(1-\alpha)\tau$. Additionally, an external patent (fraction $\tau/(\tau+\bar{y})$) has a probability $\alpha(1-q)$ of being transferred immediately upon birth, implying

$$P(0) = \left(\frac{\tau}{\tau + \bar{z}}\right) \alpha(1-q)$$

Once born, patents of any type have a flow rate of transfer $\alpha q \tau$ for their entire lifetime. The probability of a particular patent surviving to age t and being transferred for the first time is then

$$P(t) = \left[1 - \left(\frac{\tau}{\tau + \bar{z}}\right) \alpha(1-q)\right] \alpha q \tau \exp\left(-\left(b + (1-\alpha)\tau + \alpha q \tau\right) t\right)$$

This exponential form is very close to what is seen in the data. A detailed evaluation of the match is given in the quantitative section. Finding the probability that a patent is never transferred is then a matter of evaluating the above expression in the limit as $t \to \infty$. This yields the expression

$$P = P(0) + (1 - P(0))\left(\frac{\alpha q \tau}{b + (1-\alpha)\tau + \alpha q \tau}\right) \in \left[\alpha\left(\frac{\tau}{\max\{\bar{z}, b\} + \tau}\right), \alpha\right]$$

Thus it is very closely related to the incidence of sequential innovation. Notice that the lower bound decreases with the fraction of innovations that are within-firm, since these innovations result in no transfers. Having only external innovation would result in implausibly high fractions of patents being transferred. Additionally, the lower bound decreases with the rate of patent expiry, since innovation on an expired product line induce no transfers as well.

Patent Expiry Now consider the process of patent obsolescence. This occurs due to patent expiry at rate b and due to technological replacement at the rate $(1-\alpha)\tau$. Therefore, the distribution over the productive lifetime of a patent, $E(\cdot)$, is given by

$$E(t) = (b + (1-\alpha)\tau)\exp(-(b + (1-\alpha)\tau)t)$$

which arises from the properties of continuous time Poisson processes. Therefore, when looking across industries or patent classes, one would expect to see a negative relationship between the fraction of patents that are not renewed after a given period of time and the fraction of patents that are transferred at least once.

Another figure of interest the fraction of patents that become obsolete through creative destruction, rather than patent expiry. This value can be shown to be

$$E = \frac{(1-\alpha)\tau}{b + (1-\alpha)\tau}$$

Interpreting the loss of a patent due to failure to pay maintenance fees as creative destruction, this figure is roughly $1/2$ in the data.

Markups and Profits We must also address the effect of α on markups over cost. Recall that for a given product line, total revenues are always Y, while production labor costs are to $\lambda^{-1}Y$. Thus the markup for any one product line is simply λ and average markup at the product line level is equal to

$$\bar{\Lambda} = \mu_0 + \mu_+ \left(\frac{1}{1-m}\right) = \frac{b + (\tau + z_0)/(1-m)}{b + \tau + z_0}$$

Unfortunately, this is not always well defined since we cannot guarantee that $m < 1$. We can however give a well-defined expression for the median markup at the product line level. The complete cumulative density for markups is given by $\mu(\lambda) = \mu_0 + \mu_+(1 - \lambda^{-1/m})$. Equating this to one half yields

$$\widehat{\Lambda} = \max\left\{1, \left[2\left(\frac{\tau + z_0}{b + \tau + z_0}\right)\right]^m\right\} \tag{1.10}$$

These cannot be directly observed in the Compustat data. However, one can look at total sales over variable cost at the firm level and target that using simulations. Meanwhile, the aggregate ratio of sales to cost over the entire economy (Λ) is simply the harmonic mean of individual product line markups

$$\Lambda = \left[\mu_0 + \mu_+ \left(\frac{1}{1+m}\right)\right]^{-1} = \frac{b + \tau + z_0}{b + (\tau + z_0)/(1+m)}$$

This measure can be calculated directly from the Compustat data set or taken from existing estimates. It can also be shown to be increasing in α. However, the endogenous response of τ, z, and z_0 will determine the net effect.

Patent Portfolios A closely related figure is the size of a firm's patent portfolio. In this case I will work with portfolios at the per-product line level, which can then be aggregated to the firm level. Let the mean portfolio size in a given industry be denoted by $\bar{n}$. The flow equation for this quantity is given by

$$\dot{\bar{n}} = \sum_{n=0}^{\infty} \dot{\mu}_n n = (1-\alpha)\tau + (\bar{n}+1)(\alpha\tau + \bar{z}) - (b + \tau + \bar{z})\bar{n}$$

which leads to the solution

$$\bar{n} = \frac{\tau + \bar{z}}{b + (1-\alpha)\tau}$$

This can be used to calculate a number of pertinent values. For instance, assuming independent innovations make no citations, the average number of external citations per patent will simply be $\left(\frac{\tau}{\tau+\bar{z}}\right)\alpha\bar{n}$.

1.3.9 Industry and Firm Heterogeneity

In order to match the cross-industry variation in quantities such as transfer rates, expiry rates, and profitability, I introduce the possibility of heterogeneity across industries in the sequentiality of innovation. Industries are segmented in terms of innovation but share a common pool of labor. Let there be M equal-mass segments in total. Each will have a particular value for α and associated innovation rates $\tau(\alpha)$, $z(\alpha)$, and $z_0(\alpha)$ and average inverse technological lead $\Gamma(\alpha)$. The initial analysis then carries through unchanged at the industry level and we need only be concerned with aggregate labor and bond market clearing. The labor market clearing condition is given by

$$1 = \mathbb{E}_\alpha \left[\frac{\Gamma(\alpha)}{\widetilde{w}} + (1+e)c(x(\alpha)) + \mu_0(\alpha)d(z_0(\alpha)) + \mu_+(\alpha)\Gamma_+(\alpha)d(z(\alpha)) \right]$$

In practice, one can deal in the limit where $M \to \infty$, using a continuous distribution over α. Each industry is then treated as an infinitesimally small segment of the overall spectrum of products.

One potentially unsatisfying implication of the above model is that the direction of transfer is indeterminate and is effectively decided by a weighted coin flip q between incumbent and challenger. By introducing firm level heterogeneity, either in production or research capability, we can not only break this indeterminacy, we can investigate the potential allocative implications of the transferring of patent rights. This modification of the model is motivated by the trends presented above as well. In particular, Fact 5 regarding the predominance of patent flows towards younger and smaller firms leads one to believe such a dynamic is an important determinant of trends in the transfer of patent rights.

A two-type extension of the model presented herein is exposited in detail in Appendix A.2. This partially resolves the indeterminacy of patent transfer direction in the sense that interactions between firms of unlike types will result in final ownership being vested in the high type firm. However, interactions of like type firms will still require the introduction of *ad hoc* randomness. Firms differ in their cost of external innovation, while the cost of internal innovation and production capacity remain identical across firms. As a result, there will be differential firm-level external innovation rates x^H and x^L and expired internal innovation rates z_0^H and z_0^L but a common unexpired internal innovation rate z. Entering firms are high type with probability ζ and decay to low type at the Poisson flow rate ν. The associated aggregate external innovation rate is $\tau = (\zeta e + \mu^H)x^H + ((1-\zeta)e + \mu^L)x^L$, where μ^i is the mass of products owned by a type i firm, while the aggregate internal innovation rate is $\bar{z} = \mu_0^H z_0^H + \mu_0^L z_0^L + \mu_+ z$.

1.4 Estimation

In this section I bring the proposed model to the firm-level data. First, I provide a summary of the challenges associated with identifying the key aspects of the model quantitatively. Then I present the results of the estimation, assess the quality of the fit, and give some interpretation to the resulting parameters values.

1.4.1 Identification

In the basic model, there are seven parameters that need to be identified. First, those associated with the innovation production process, namely the step size distribution tail index κ and the cost parameters c and η. Additionally, there is the mass of entrants e and the discount rate δ, which is set to 0.05. Finally, there are two parameters specific to the phenomenon studied in this paper, the probability of sequential innovation α and the bargaining power of the incumbent p. Moving to the setting with own-product innovation will add an additional cost parameter d, and the introducing industry level heterogeneity in α will add in two distributional parameters. The specific form used will be the two parameter Beta distribution, a flexible choice for random variables on the unit interval.

First, the the innovation production parameters can be jointly constrained using aggregate moments. Naturally, their exact values will dependent upon the estimated values of other parameters in the model. However, this partitioning is still useful at a conceptual level. The tail index parameter κ will be a strong determinant of profits in the economy. The expression given in Equation 1.10 is calculated at the product line level. Aggregating this to the firm level has no apparent analytical form, so simulation must be used. Meanwhile, the R&D production function parameters can be effectively constrained using the aggregate growth rate and the share of R&D spending in the economy. Both these moments are readily available from BEA data. For the sake of consistency, analogues at the firm level, namely the incumbent innovation rate and R&D intensity, can be used as well.

Introducing the possibility of own-product innovation adds one additional cost scaling parameter d, while the elasticity η is assumed to be common with the external innovation cost function. This can be constrained in a straightforward manner by using self-citation activity by firms. In particular, one can look at the fraction of citations that are internally directed versus externally directed. Alternatively, one can classify patents as primarily internal or external and look at the composition thereof. In the US patent data, approximately 23% of patents cite other patents from their filing firm. Akcigit and Kerr (2010) study these patterns extensively, finding that approximately 20% of patents are self-citing. They also have access to analogous information at an R&D expenditure level (there termed product versus process innovation) and find a similar fraction.

The mass of entrants parameter can readily be determined by looking at the fraction of the patent stock owned by entrants over a five year period, for instance. Alternatively, one can undertake a growth decomposition and target the standard number given by Davis, Haltiwanger, and Schuh (1998) of $1/3$. As innovation and growth are one-to-one in this framework, we would then simply set $e = 1/2$. However, due to selective exit of low type firms early in the life cycle, targeting simulated numbers over a number of years is more accurate.

Capturing information about α and the distribution over various industries proves not to be too difficult. As shown above, the fraction of patents that are transferred in their lifetime varies with α. This response is not provably one-to-one due to variability in the endogenous response of innovation rates, but in practice it proves to be robustly monotone. Thus we can effectively use the inverse of this function to relate data on patent transfer rates to sequentiality measures α and the industry-level distribution thereof.

The most difficult parameter to estimate is the bargaining power parameter p. One possible stance is that, there being no difference *ex ante* between firms in terms of bargaining position, symmetry dictates the value be set to $1/2$. However, if one imagines a model in which bargaining does not happen instantaneously, the incumbent might be more patient if it thinks it can produce a workaround innovation. Alternatively, there may be asymmetries in patent protection for incumbents and new innovators that lead to asymmetries that can be captured by the bargaining power parameter. In the baseline case, I simply set $p = 1/2$, however, I also investigate the effects of changing this parameter on innovation rates and aggregate outcomes.

To match the facts regarding patent flows to younger and smaller firms, I also introduce firm-level heterogeneity in the cost of external innovation. This introduces three new parameters. First, there are now two cost factors for external innovation (c^H and c^L) rather than one. Additionally, there is the initial fraction of entrants that are high type (ζ) and the rate at which high type firms decay into low type firms (ν). By looking at growth rate differentials between young and old firms, as well as the fraction of the patents stock owned by young firms, we can constrain these various parameter values.

1.4.2 Results

As noted at the outset, the previous discussion is merely a conceptual overview of identification. Each of the parameters will affect each potential moment in varying degrees, as determined by the model. In order to match all of these simultaneously and intelligently trade off prediction errors in the case of a less-than-perfect match, I employ a simulated method of moments (SMM) estimator. This has the additional advantage of allowing for bootstrapping to determine parameter standard errors. See Bloom (2009) for further details on the usage of SMM. The details of the equilibrium solver and simulation algorithm are described in Appendix A.3.

The SMM objective function is simply quadratic form of the differences between the data and model predictions for a certain vector of moments. The requirements for identification discussed above will largely determine which moments are used. However, there are still some specific implementation details that should be explained.

First, the notion of entry employed here counts any previously unobserved firm that files for a patent. As the primary focus is on the innovation process, this would seem to be the most relevant statistic for our purposes. Additionally, it does not suffer from the selection issues that using Compustat would entail. It is susceptible to misclassification of firms in the sense that any novel name matching failure would show up as an entrant rather than simply a filing from an existing firm. Including a measure of the patent stock that recent entrants comprise could ease these concerns and potentially provide valuable information about the type distribution of entrants (the parameter ζ).

To measure internal citations, one must take a stance on the exact mapping between the model and the data in this regard. Independent innovations of course cannot be entirely disparate from work that has come before. There is a sense in which general knowledge informs innovations within and across fields. The extent to which citations reflect general versus specific influence is not clear. Assuming that sequential and independent innovations have roughly similar amounts of external citations, the internal citation ratio would simply be $\bar{z}/(\tau + \bar{z})$. Note that here I assume that a new addition to a patent portfolio cites all of the patents below it, or at least a fixed fraction thereof.

The transfer rate mean and variance statistics are calculated as the probability that a patent in a particular industry filed for during the period in question is ever transferred (including beyond the end of the period). The three year transfer rate is simply the same probability but restricting to the transfer occurring within three years of filing. The analog in the model is taken to be the probability of immediate transfer. The delay is chosen to allow for the fact that this process will not be instantaneous in a real-world setting.

The remaining aggregate moments are fairly straightforward. The aggregate growth figure is taken from the FRED data on output per person. The median profit is computed

TABLE 1.1: MOMENT VALUES

Name	Data	Model
Aggregate Growth	0.034	0.034
Entrant Stock Frac	0.085	0.081
Transfer Prob Mean	0.180	0.187
Transfer Prob Std	0.105	0.104
Internal Cite Frac	0.229	0.268
Median Profit	0.164	0.175
3-year Transfer	0.365	0.356
Transfer To Younger	0.699	0.621
Transfer To Young	0.573	0.546
Young Stock Fraction	0.139	0.143
Young Filing Fraction	0.197	0.198
Average Filing Fraction	0.166	0.182
R&D Intensity	0.092	0.051

using the sample of Compustat firms. The moment values at the optimum are given in Table 1.1. It is evident that the fit of the model to the data is quite close in many dimensions, but misses the mark in some cases. The moments on aggregate growth, transfer statistics, citations, profits, and R&D levels are all very closely matched.

The parameter estimates themselves are summarized in the Table 1.2. Of prime importance are the sequentiality distribution parameters, which indicate that in the average industry, 35% of innovations are sequential as opposed to independent. Furthermore, looking across industries, this quantity has a standard deviation of 22%.

The R&D production function parameters are in line with those found in the existing literature. In particular, the curvature implies an elasticity of $53\%(= 1/1.880)$. This is similar to the value of 0.61 found in Pakes and Griliches (1980). As documented by Kortum (1993), estimates of this parameter generally lie between 0.1 and 0.6. External innovation is more than twice as costly for low type firms compared to high type firms. Meanwhile, the common cost of internal innovation is nearly twice as costly again as low type innovation. The entrant mass of 0.19 is lower than what one might guess from directly calibrating

TABLE 1.2: ESTIMATED PARAMETER VALUES

Name	**Symbol**	**Value**
Discount Rate	δ	0.050
Bargaining Power	p	0.500
CRRA Parameter	σ	2.819
Step Distribution	κ	0.339
External R&D (High)	c^H	6.206
External R&D (Low)	c^L	12.659
Internal R&D Cost	d	20.257
R&D Cost Curvature	η	1.880
Mean Sequentiality	Mean(α)	0.347
Std Sequentiality	Std(α)	0.216
Entrant Mass	e	0.161
Entry High Type	ζ	0.580
Type Decay Rate	ν	0.126
Transfer Direction	q	0.730

to a growth decomposition attributing $1/3$ of growth to entrants. However, as noted in Foster, Haltiwanger, and Krizan (2001), the time horizon over which we measure entry contributions will affect even annualized figures due to selective exit of recent entrants. The firm selection dynamic present in the model induces a similar effect in simulated results.

The firm type dynamics parameters are consistent with previous structural studies. In particular, Acemoglu, Akcigit, Bloom, and Kerr (2013), on whose model of reallocation I build, estimate that the proportional variance in the cost faced by entrants is 1.21.[8] The analogous number using my estimates is 1.06. In their model, ignoring selection due to exit, the average cost faced by incumbent firms rises by 5.5% over the course of one year, while I find that quantity to be 5.3%.

[8]The fraction of entrants that are high type is not directly comparable, as the respective R&D cost parameters also differ.

1.5 Quantitative Analysis

With these estimates in hand, we can evaluate the ability of the model to match the cross-industry and firm-level trends outlined in the data section. Figure 1.7 documents the model's qualitative success in reproducing four of the the major cross-industry trends seen in the data. In each pane, the variable of interest is plotted at the industry level against the sequentiality for that industry, which is simply the fraction of patents that are transferred in their lifetime. This measure is used so as to most closely match the facts presented in the data section and varies monotonically with the underlying theoretical sequentiality parameter α. The median profitability, firm growth, and exit rate figures are the result of simulations and so have noise associated with each point, but the magnitude of this noise is not enough to obscure the evident trends.

To get an idea of the quantitative match of the model with regards to these trends, the average level of each variable over all industries in the data and in simulations is given in Table 1.3. The expiry rate and firm growth volatility are both matched well, while the profitability and exit rate figures are both not matched entirely. The profitability figure, though targeted at the economy-wide level, is grouped by industry and can thus be different due to aggregation effects. The exit rate may reflect asymmetries between entry and exit that are not present in the model, which features only entry of and exit by one-product firms. In the data, though entering firms are generally quite small, large firms may exit or cease to exist due to mergers. The framework used here implicitly categorizes innovations in to internal, sequential, and independent types. By targeting patent transfers and internal citation rates, the close alignment of the data here further supports the notion that sequential innovation is reflected in patent transfers, while independent innovation is reflected in patent expiry.

Having analyzed the levels of each of the four variables of interest here, I now study the variation across industries. To assess the amount of variation across industries in the

FIGURE 1.7: CROSS-INDUSTRY PREDICTIONS

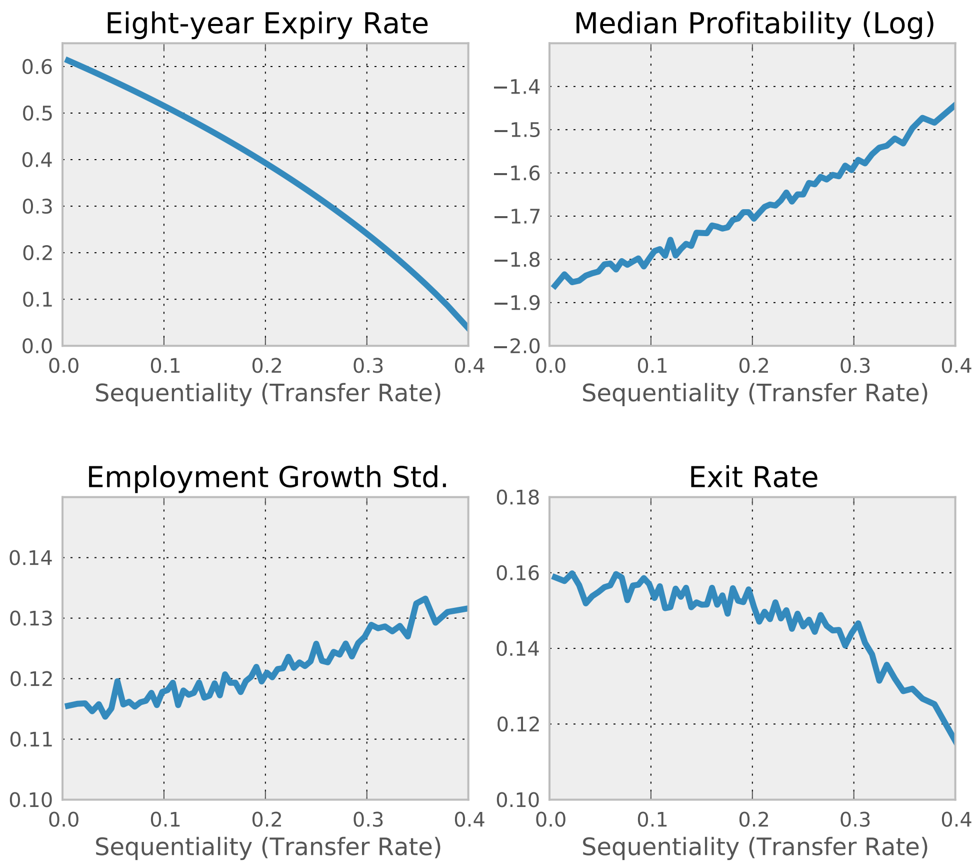

various quantities, I perform a weighted least squares regression on sequentiality (the rate of patent transfer) and look at the predicted effect of moving between the mean plus or minus one standard deviation, as normalized by the mean value for that particular quantity. Comparing the values given by this metric for both the data and the simulated model, I can assess the predictive power of the model. It is important to note that though the variation in sequentiality was targeted, variation in the other quantities has not been, meaning these implications come purely from the model structure. Using this metric, the model can account for approximately 65% of the variation in profitability (return on sales), 26% of the variation in firm growth volatility (employment), and 40% of the variation in exit rates.

TABLE 1.3: AVERAGE INDUSTRY-WIDE VALUES

Variable Name	Data	Model
Eight-year Expiry Rate	34%	39%
Median Profitability	11%	18%
Firm Growth Volatility	13%	12%
Exit Rate	29%	15%

Meanwhile, the level of variation in expiry rates is overpredicted by the model. In Figure 1.8, I plot the proportional variation for each quantity. Both the data and model generated numbers are normalized by their respective mean values, as is the regression line for the data.

FIGURE 1.8: CROSS-INDUSTRY VARIATION

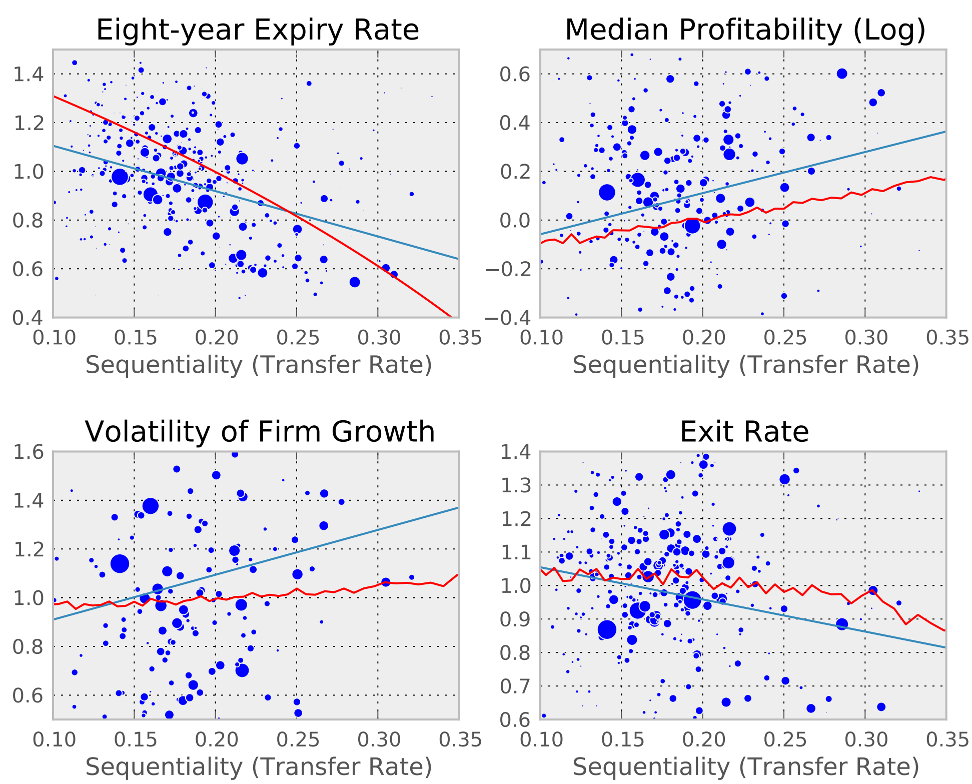

Figure 1.9 depicts a number of equilibrium variables of interest as they vary with in-

dustry sequentiality. Of particular interest in this framework is the effect of industry sequentiality on the incentives to innovate and the resulting innovation rates. The effect on external innovation is *ex ante* ambiguous due to the existence of two opposing forces on the expected product line valuation. The results of the estimate indicate a negative effect of sequentiality on the the rate of external innovation, with the effect being larger for high type firms. To understand this result, consider the time profile of returns delivered by an innovation. In non-sequential industries, the payoff is large early on but is cut off relatively soon due to creative destruction. Meanwhile, in highly sequential industries, the initial payoff is smaller, but there are ongoing payouts from future sequential innovations. Even if the total payout is similar in these two cases, a firm only captures a fraction of this surplus through bargaining, thus back-loaded incentives in highly sequential industries result in lower innovation rates.

As expected theoretically, internal innovation rates rise with sequentiality due to the reduced threat of creative destruction in more sequential industries. The n-shaped dependence of the average firm type on sequentiality is somewhat unexpected. One would anticipate that high type firms would face the largest drops in product loss rates moving to more sequential industries, which is indeed the case. However, this effect is eventually overwhelmed by the disproportionate drop in high type external innovation rates. Finally, the level of monopoly distortion (Δ) on a per-industry basis is shown. As would be expected, this rises with sequentiality, as a result of the agglomeration of larger patent portfolios by firms.

Now consider the aggregate trends in the economy. The following table summarizes the growth contributions of entrants, external innovation, and internal innovation.

FIGURE 1.9: CROSS-INDUSTRY EQUILIBRIUM VARIABLES

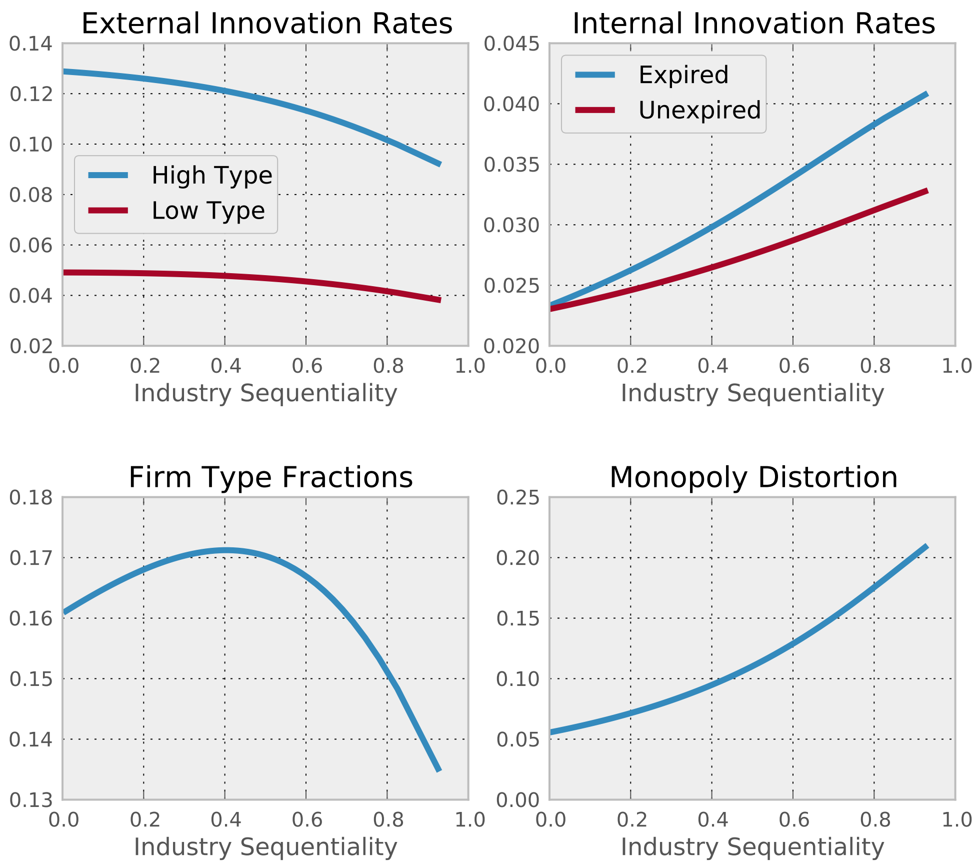

EQUILIBRIUM DECOMPOSITIONS (PERCENTAGES)

	Entrant	**Inc. External**	**Inc. Internal**	**Total**
Growth	0.49 (14.3%)	2.02 (58.9%)	0.92 (26.8%)	3.43 (100.0%)
Labor	1.37 (16.0%)	5.38 (62.8%)	1.81 (21.1%)	8.56 (100.0%)

External innovation by incumbents still plays a large role in innovation, followed by incumbent internal innovation, then entrants. Additionally, internal innovation achieves notable performance in terms of innovations per unit labor. Considering both the cost and capacity are similar to that of external innovation, this is largely a sign that it is simply utilized less.

As noted previously, introducing firm types into such a model naturally allows us to consider the potential effects of sequentiality on reallocation. Entering firms start out with a particular fraction of high-type firms ($\zeta = 58\%$). Over time the surviving firms decay at rate $\nu = 13\%$. High type firms have a higher survival probability, so the distribution over type by firm age exceeds the simple case of exponential decay. The transfer direction parameter dictates that when firms of like type transfer patents, the innovator becomes the eventual owner in 73% of cases. Using the respective shares high type and low type products, this implies that overall a high type innovator assumes final ownership in 96% of cases, while the same number for low types is 61%, and the overall number is 76%.

FIGURE 1.10: EQUILIBRIUM DISTRIBUTIONS

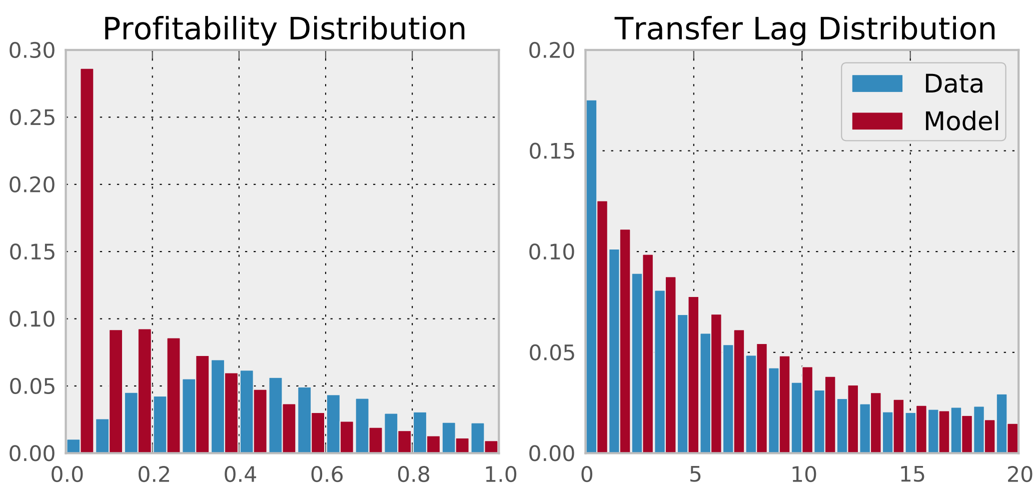

Looking at Figure 1.10, we can see that the match between the predicted and observed lag from patenting filing and time of first transfer is quite good. It should be noted, though, that the model delivers an exponential form for this function by construction, and relative mass at zero is indeed targeted for the purposes of estimating the direction of transfer parameter (q). As for the profitability distribution, the match between data and model is less exact. The median profit is targeted as a moment, however, the model generated clearly shows excess mass near zero. This could arise from compositional issues. Firm level profitability data is available only for Compustat firms, which are larger than the average firm, and hence the observed data could be "overaggregated", leading to greater weight in the middle of the distribution and less at the very bottom.

1.5.1 Mechanism Investigation

To better understand the means through which sequentiality affects the incentives for innovation, in this section I consider various modification to the estimated model. First, I investigate the effects of varying the bargaining power parameter, which was previously set to the symmetric value of $1/2$. Because payoffs are back-loaded in more sequential industries and hence only partially delivered to the incumbent through the bargaining process, this parameter will be an important determinant of the incentives for external innovation.

FIGURE 1.11: BARGAINING PROCESS ALTERNATIVES

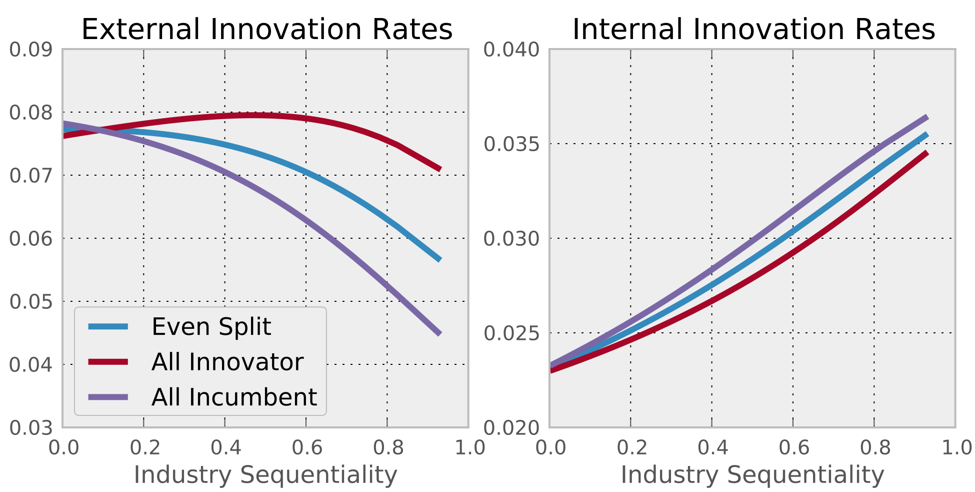

In Figure 1.11, the baseline case, as well as the two extremes of giving all the bargaining power to the innovator ($p = 0$) and to the incumbent ($p = 1$), are plotted. Here we see that though the rate of internal innovation is largely invariant to the bargaining parameter, the decrease in external innovation with sequentiality is larger the more bargaining power is vested with the incumbent. The intuition here is that when none of the incentives are transferred intertemporally through bargaining, as is the case when the innovator has all the bargaining power, the backloading of payoffs that comes with increased sequentiality does not interact with the bargaining distortions, thus the profile is relatively flat.

A second modification I consider is simply eliminating sequential innovation across the economy, that is, setting $\alpha = 0$. This provides insight into the aggregate effects of sequentiality. Moving from the baseline case to the no sequential case, the growth rate falls from 3.42% to 3.27%. Most of this comes from changes in internal innovation rates, which fall from 2.7% to 2.2% annually. Distortion from production labor misallocation (Δ) fall substantially from 9.4% to 5.5%. Recall that because of patent expiry dynamics and internal innovation, there will still be labor utilization heterogeneity, though it is much smaller here.

The final model modification performed is the elimination of firms type. Consider an economy where instead of having multiple firm types, there is a single firm type whose cost is equal to the expected cost of a new entrant, that is $\bar{C} = \zeta c^H + (1 - \zeta)c^L$. By studying the difference between our benchmark and this economy, we can illuminate the effects of firm types. The primary motivation for introducing firm types was to capture the disproportionate flow of patent transfers from older and larger firms to smaller and younger firms. Moving to the homogeneous settings, the fraction of patents transfers that are directed towards younger firms falls from 61% in the baseline to 39%. Thus firm types are critical for producing the disproportionate flows we see towards younger firms, which in either case constitute a small fraction of the overall patent stock (15% in the baseline and 9% in the homogeneous case).

There are also important effects on the firm size distribution. Moving from the homogeneous case to the heterogeneous case, we see mean firm size increase by 40%. However, this effect is concentrated mostly amongst the smaller firms. The skewness of the firm distribution is actually much larger in the heterogeneous case, at 20.5 compared to only 3.2 in the homogeneous case. For comparison, this value is 26.0 when looking at the patent stock data. Thus the baseline model captures nearly all the skewness in the data, and the introduction of firm types is partially responsible for this. In particular, the fact that certain firms (high type firms) undergo sustained periods of abnormally high growth is an important factor in generating realistic levels of variation in the firm size distribution.

1.5.2 Social Optima

As discussed previously, I consider both a constrained social planner, who can make innovation decisions but is still subject to patent policy and the resulting monopoly distortions, and an unconstrained planner who can make both innovation and production decisions at will. The constrained optimum yields innovations rates by type on a per-industry basis, while the unconstrained optimum will feature uniform innovation rates across industry. Note that the constrained planner is still also subject to firm type dynamics as in the decentralized case, while the unconstrained planner can reassign product lines to different firms at will but is still subject to exogenous heterogeneity in entrant types. In the constrained setting, the share of labor allocated to research rises from 8.6% to 15.3%, while the aggregate growth rate rises to 4.6%.

In Figure 1.12, both the internal and external innovation rates are plotted for the equilibrium and the constrained and unconstrained planner. Because monopoly distortions arising from innovation are much more severe in sequential industries, the planner reduces innovation of both types in these industries. As a result the level of monopoly distortion in the economy goes down by 0.50 percentage points. The net welfare gains from moving to the constrained planner's allocation are 2.5%. The aggregate growth and labor decompositions

FIGURE 1.12: SOCIAL OPTIMUM INNOVATION RATES

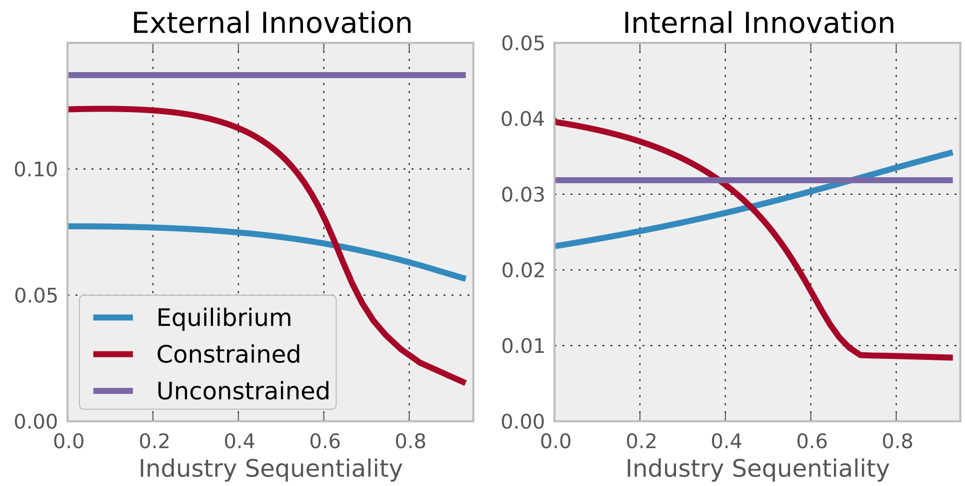

are show in the table below

CONSTRAINED OPTIMUM DECOMPOSITIONS

	Entrant	**Inc. External**	**Inc. Internal**	**Total**
Growth	0.64 (13.9%)	2.96 (64.0%)	1.02 (22.1%)	4.62 (100.0%)
Labor	2.41 (15.8%)	10.54 (68.9%)	2.33 (15.3%)	15.29 (100.0%)

In the aggregate, in addition to the increase in the level of research labor overall, there is a shift from internal to external innovation. This is consistent with the intuition presented in Aghion and Howitt (1992) by which externally oriented innovation can be either under or over-invested in, dependent on the step size.

In the unconstrained optimum, we actually see a partial reversal of the trends that occurred when going from the equilibrium to the constrained optimum. The external innovation rate is uniformly higher than in the equilibrium case. However, due to the shift to all high type research, this is done using proportionately less labor. Meanwhile, the uniform innovation rate of the unconstrained planner is quite close to the equilibrium rate seen in highly sequential industries, where one would expect firms to internalize much of the social gains from innovation. The following table summarizes the overall allocation by research type

UNCONSTRAINED OPTIMUM DECOMPOSITIONS

	Entrant	**Inc. External**	**Inc. Internal**	**Total**
Growth	0.51 (8.9%)	4.14 (72.2%)	1.08 (18.8%)	5.72 (100.0%)
Labor	1.47 (8.9%)	11.92 (72.2%)	3.11 (18.8%)	16.49 (100.0%)

Here we see a further shift towards external innovation, with accompanying gains due to the increase in the quality of incumbent firms. The total amount of research labor rises only slightly above the constrained planners case, while the growth rate rises considerably to 5.72%. It is interesting to note the alignment of the share of growth and labor from each source. This arises from the fact that from the perspective of the unconstrained planner, each type of innovation has the same effect, namely increasing the overall productivity level by a common factor in perpetuity, and there is a common cost elasticity across each type. The consumption equivalent welfare gains associated with moving to the unconstrained planner's allocation are 16%. Much of this, however, comes from the total elimination of monopoly distortions, while the rest is from growth related effects.

1.5.3 Policy Implications

The primary policy lever that will be considered here is the strength of patent protection, as embodied in the rate of patent expiry b. This will not be directly analogous to the real world notion of patent length, as patent expiry is stochastic in the model and fixed-length in reality. However, we will map policies between the two based in the mean patent length.

In highly sequential industries, a greater fraction of the benefits from a particular innovation are internalized by the original inventor, in the form of payments for patent sales to subsequent innovators. Additionally, these industries also feature a greater concentration of patent portfolios, leading to larger monopoly distortions. The fundamental trade-off of patent policy, as articulated by Arrow (1962), is between increasing the incentives to innovate so as to more closely align private and public returns to innovation and the deleterious effects of granting monopolies in production. Looked at through this lens, both of the above mentioned features of sequential industries lead us to expect that an optimal policy will include lower patent protection in more sequential industries. The first implies less of a need for the realignment of incentives, while the second implies that the inherent costs associated with granting patents are more severe in these industries.

Motivated by these considerations, I study both policies that are constant across industries and those that vary linearly with industry sequentiality. In terms of implementation, though sequentiality cannot be directly observed, one can infer its value for each industry using the equilibrium relationship between sequentiality and the probability of patent transfer (or any other observable that varies monotonically with sequentiality).

FIGURE 1.13: OPTIMAL LINEAR PATENT POLICY

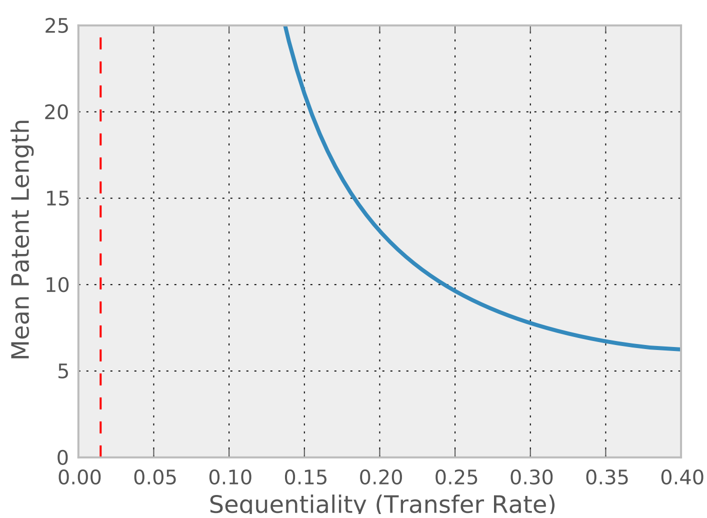

The optimal constant policy calls for a patent expiry rate of 8.6% annually, implying a mean patent length of 11.6 years. This causes the growth rate to fall to 2.94%. However, since the fraction of production labor rises to 93.4% and monopoly distortions fall to 6.0%, there is a net welfare increase of 0.9%. The optimal linear patent length decreases sharply with sequentiality. For industries with extremely low sequentiality, the policy calls for an infinite patent. The patent policy eventually reaches a minimal mean length of 6 years in the most sequential industries. The full path of the patent expiry rate is depicted in Figure 1.13. It is interesting to note that the optimal patent length in the median industry (sequentiality 17%) is about 18 years, almost exactly what current law prescribes. The resulting growth rate is 3.25% and the share of production labor is 91.9%, which are quite similar to the equilibrium values. However, the monopoly distortions fall noticeably, having been curbed in the most offending industries, to 6.7%, resulting in a net welfare increase of 1.7%.

1.6 Conclusion

In this paper, I propose a notion of sequentiality in innovation and argue that it is an important determinant of a firm's incentives to innovate and of firm dynamics. Specifically, though externally oriented innovation has generally been assumed to result in creative destruction (or new products), I emphasize the notion that patent protection applies not just to contemporaneous imitators but to future cumulative innovators. The result is that innovating firms must in some cases come to an agreement with incumbent firms regarding ownership of the underlying portfolio. This ultimately has strong effects upon firms overall incentive to innovate. Not only that, these patent sales allows us to use data on the transfer of patent ownership as a window into the nature of the innovation process.

When looking at the cross-industry trends in patenting, a number of notable trends regarding transfer rates, expiry rates, profitability, and firm dynamics are apparent. Additionally, patent transfers flow disproportionately towards smaller and younger firms. To capture these trends, I take the basic model of sequential innovation described above and introduce heterogeneity both across industries and across firms within industry. The resulting estimated model is able to match these trends qualitatively and account for a large fraction of the cross-industry variation.

The quantitative estimates point to an underallocation in the quantity of labor devoted to research. This result is not surprising given the existing theoretical and empirical literature. However, even fixing the quantity of labor devoted to research, there is a misallocation of research labor towards those industries with the highest sequentiality. I find that monopoly distortions are quantitatively important in this setting. In highly sequential industries, firms accumulate large patent portfolios, allowing them acquire a substantial technological lead over their nearest legal competitor. To remedy these misallocations in both research and production, I introduce both constant and industry dependent patent policies and evaluate their welfare effects.

I find that a patent policy featuring weaker patent protection in more sequential industries can generate large welfare gains over both current policy and the optimal uniform patent policy. The implications of this finding are not out of line with certain existing proposals by policymakers. In particular, there have been numerous calls by interest groups to either eliminate or severely restrict patenting of software, an industry which I find to be highly sequential. Other authors, such as Boldrin and Levine (2008), who advocate the elimination of patents frequently cite the software industry as an example of patenting gone wrong. Similarly, those concerned about "patent trolls" cite software as an industry which has been hit particularly hard by costs associated with intellectual property litigation.[9] Though I do not address these costs directly, reducing patent protection in these industries would limit the potential damage that patent trolls could cause.

Though this paper takes a very detailed approach to modeling patenting dynamics and the incentives to innovate, this setting is undoubtedly extremely complex, featuring a large

[9] http://www.schumer.senate.gov/record.cfm?id=341612&

quantity of heterogeneity both at the firm and industry level. There are many potential avenues of research that remain to explored. For instance, the notion of sequentiality may be related to patent breadth, meaning there are patent policy levers in addition to length that could be explored in this context. Additionally, not all innovation is necessarily protected with patenting. Incorporating an endogenous decision to patent (as opposed to using secrecy, for instance) could have interesting implications, though observability is naturally an issue.

Chapter 2

Back to Basics

Joint work with Ufuk Akcigit (University of Pennsylvania and NBER) and Nicolas Serrano-Velarde (Bocconi University and IGIER).

2.1 Introduction

Fostering economic growth is one of the primary objectives of economists and policy-makers. The amount of resources invested in research is often at the heart of the debate regarding how to best achieve this. The level of research investment plays an important role in the pace of long-term technological progress and economic growth, and countries allocate a significant share of their GDP to researching new products and technologies in this spirit (see Figure 2.1). Less well known, however, is what role the composition of this research plays in determining growth, particularly when considering the breakdown between basic and applied research. In this paper, we aim to fill this gap by studying the differential effects of basic versus applied research on economic growth.

According to the NSF, basic research investment refers to a "systematic study to gain more comprehensive knowledge or understanding of the subject under study without specific applications in mind." Conversely, applied research is defined as a "systematic study to gain knowledge or understanding to meet a specific, recognized need."[1] This distinction is empirically important since almost half of total research investment is allocated to basic research in countries such as France and the US (see Figure 2.2).

FIGURE 2.1: TOTAL RESEARCH TO GDP RATIO IN FRANCE AND THE US

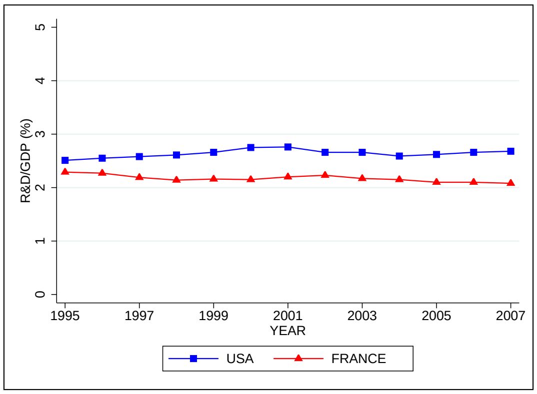

FIGURE 2.2: COMPOSITION OF RESEARCH INVESTMENT IN FRANCE AND THE US

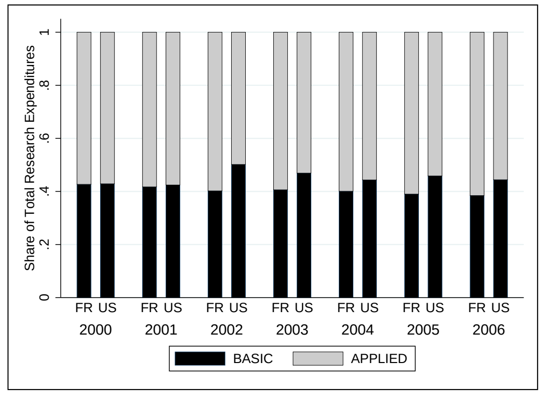

Notes: Figure 2.1 plots the time series of total research spending as a fraction of GDP for the US and France. The data come from the OECD, Main Science and Technology Indicators, volume 2009/1. Figure 2.2 plots the composition of total research spending for the US and France. The data for the US come from the National Science Foundation, Division of Science Resources Statistics (NSF/SRS), while the data for France come from the French Ministry of Research.

The issue of investment in basic research also received fresh policy interest in a recent report by the US Congress's Joint Economic Committee, where it is argued that despite its value to society as a whole, basic research is underfunded by private firms precisely because it is performed with no specific commercial applications in mind. The report states that the level of federal funding for basic research is "worrisome" and should be increased (**(alias?)**).

[1]Although basic research may not have specific applications as its goal, it can be directed to fields of current or potential interest. This focus is often the case when performed by industry or mission-driven federal agencies. In industry, applied research includes investigations to discover new scientific knowledge that has specific commercial objectives with respect to products, processes, or services. http://www.nsf.gov/statistics/seind10/c4/c4s.htm#sb2

Despite clear empirical importance and considerable policy interest, the differential nature of the roles played by basic and applied research in the growth process is still relatively unexplored, and many related questions remain to be answered: What are the key roles of basic and applied research for productivity growth? What are the incentives of private firms to do basic research? How does publicly funded basic research contribute to innovation and productivity growth? How sizable are the spillovers from basic research? What are the potential inefficiencies in a competitive economy, and what are the appropriate government policies to mitigate them and promote economic growth?

In order to understand the potential inefficiencies in research investment and to design appropriate industrial policies to address them, it is necessary to adopt a structural framework that explicitly models the incentives for different types of research investments by private firms. Our goal in this project is to take an important step toward developing this theoretical framework, identify the potential spillovers, and study their macroeconomic implications for innovation policy.

We follow the influential literature on basic science and consider the possibility that basic research not only generates large spillovers within an industry, but it can also be applicable to many different industries.[2] The historical example of Du Pont de Nemours' financing of William Carothers' research serves as a fine showcase of these spillovers. As Nelson (1959) describes it: "Carothers' work in linear super-polymers began as an unrestricted foray into the unknown, with no practical objective in mind. But the research was in a new field in chemistry and Du Pont believed that any new chemical breakthrough would likely be of value to the company. In the course of research Carothers obtained some super-polymers that became viscous solids at high temperatures, and the observation was made that filaments could be made from this material if a rod were dipped in the

[2] Upstream technologies having multiple downstream applications have been mentioned by various important papers, such as Bresnahan and Trajtenberg (1995), Rosenberg and Trajtenberg (2004), Nelson (1959), Rosenberg (1990) and Dasgupta and David (1994).

molten polymer and withdrawn. At this discovery the focus of the project shifted to these filaments and Nylon was the result." Nylon is now used in many industries such as textiles, automobiles, and military hardware, three industries in which Du Pont had operations.

Ideally, in order to capture the full return from this new scientific knowledge in industries where it could have an application but in which the innovating firm is not present, the innovator would first patent and then license or sell the innovation to other firms in those industries. However, the applications of basic scientific advances are often not immediate and firms are often only able to transform them into patentable applications in their own industries. This is the well-known *appropriability problem* of basic research that has been discussed in a vast literature.[3] Hence, firms operating in more industries will be able to utilize more facets of a given basic innovation. As Nelson hypothesized it: "It is for this reason that firms which support research toward the basic-science end of the spectrum are firms that have *fingers in many pies*." Note that the key concept that is being emphasized here is not firm size per se, but the diversity of its operations. This interesting argument (which we will refer to as Nelson's hypothesis) will be the central building block of our analysis in this paper.

Our analysis proceeds in three steps. Using micro-level data on French firms, we first document stylized facts on their investment in basic and applied research. Second, motivated by those empirical facts, we propose a general equilibrium, multi-industry framework with private firms and a public research sector. Firms conduct both basic and applied research, whereas the public sector focuses exclusively on basic research. In our model, basic research generates fundamental technological innovations and generates spillovers, both within and across industries, that affect subsequent applied innovations.[4] In line with

[3]Among many others, see for instance Nelson (1959), Rosenberg (1990) and Dasgupta and David (1994).

[4]By fundamental innovation, we mean the major technological improvements that generate larger than average contributions to the aggregate knowledge stock of society. In addition, these will have long-lasting spillover effects on the size of subsequent innovations within the same field.

the "Ivory Tower" theory of academic research, basic research by private firms in our model will turn into consumer products faster than that undertaken by public research labs. Applied research, on the other hand, will be done only by private firms and will generate follow-on innovations building on the existing basic knowledge stock.

To highlight the key economic forces, we will first consider a benchmark economy with tractable functional forms, characterize the dynamic equilibrium analytically, and discuss the resulting dynamics and inefficiencies. Our ultimate goal in this paper is to undertake a quantitative investigation of the impacts of various innovation policies on the aggregate economy. As such, we then generalize our benchmark framework to allow for greater quantitative flexibility and estimate the structural parameters.[5] Finally, we use the estimated model to assess the extent of inefficiencies in basic and applied research and to study the implications of several important innovation policies.

Our main result is the quantification of the inefficiencies due to dynamic misallocation in research. We find that 89% of spillovers from basic research across industries are not internalized and that basic research makes applied innovation 60% more productive. As a result, there is a dynamic misallocation of research efforts, which reduces welfare by 4.7 percentage points in consumption equivalent terms. One striking feature of the solution to the social planner's problem is that the fraction of resources devoted to research activities is not substantially greater than in the decentralized equilibrium. Indeed, the dominant misallocation here is not that between production and research, as is common in this class of models, but among the various types of research activities, in this case, applied and basic innovation. Another striking feature is that in the case of applied innovation, there is actually an *overinvestment* in the decentralized economy due to the strategic complementarity

[5]For instance, in his famous book *Pasteur's Quadrant* (1997), Stokes (1997) describes how some path-breaking innovations can emerge from applied research (as opposed to emerging only through basic research). In the theoretical analysis, we will allow for the possibility of applied research generating radical innovations as well.

between basic research spillovers and the returns to applied research.

This raises an important question: to what extent can public policies address this inefficiency? The first policy we analyze is a uniform research subsidy to private firms. In this environment, subsidizing overall private research is ineffective since this policy oversubsidizes applied research, which is already overinvested in due to competition. Therefore, the welfare improvement from such a subsidy is limited unless the policymaker is able to discriminate between types of research projects at the firm level. We thus consider a hypothetical type-dependent research subsidy and find that the optimal policy is to subsidize basic research by 50% and applied research by 14%.

However, given that this type-dependent research subsidy is difficult to implement, we analyze a further policy tool: the level of funding for public research labs. We show that due to the Ivory Tower nature of public basic research, allocating more money to the academic sector without giving property rights to the researchers is not necessarily a good idea. To demonstrate this, we mimic a policy exercise similar to the Bayh-Dole Act enacted in the US in 1980. We consider alternative scenarios in which public researchers have no property rights, then 50% and 100% property rights. We find a complementarity between the level of property rights and the optimal allocation of resources to academic research. The optimal combination turns out to grant full property rights to the academic researcher and allocating 3.7% of GDP to public research. This reduces the welfare gap from 4.7 to 1.7% in consumption equivalent terms.

Related Literature

Our paper contributes to a number of different branches of the literature. Our first contribution is to the growing literature on the role of industrial policies in productivity and welfare (see, for instance, Acemoglu, Akcigit, Bloom, and Kerr (2013), Aghion, Blundell, Griffith, Howitt, and Prantl (2004, 2009), Aghion, Bloom, Blundell, Griffith, and Howitt

(2005), Cozzi and Impullitti (2010), Garicano, Lelarge, and Van Reenen (2013), Hsieh and Klenow (2009, 2012), Impullitti (2010), Peters (2013), Restuccia and Rogerson (2008), Song, Storesletten, and Zilibotti (2011), among many others). We first evaluate the impact of innovation policies considered by policymakers in many OECD countries. Various papers have empirically shown that R&D subsidy policies have been ineffective at boosting aggregate productivity.[6] This result is theoretically puzzling, as standard endogenous growth models typically predict that growth rises when private R&D is subsidized at the expense of lower initial consumption.[7] However, these models feature only a single type of research. Once the distinction between basic and applied research is introduced, the results can differ greatly, shedding light on the aformentioned puzzle.

Second, our paper contributes to the extensive literature on technology and R&D spillovers (Griliches (1992), Jaffe, Trajtenberg, and Henderson (1993), Eeckhout and Jovanovic (2002), Jones (2005), Acemoglu, Aghion, Lelarge, Van Reenen, and Zilibotti (2007), König, Lorenz, and Zilibotti (2012), Lucas and Moll (2013), Perla and Tonetti (2013), Benhabib, Perla, and Tonetti (2013), Bloom, Schankerman, and Van Reenen (2013))[8] Moreover, the reduced-form analysis of our paper contributes to the empirical innovation literature by introducing two new ways of identifying spillovers. First, we use the variation in the level of basic research spending between firms operating in different numbers of industries to infer the magnitude of cross-industry spillovers. Second, we use heterogeneous citation patterns across public and private patents in order to identify within-industry spillovers.

Our third contribution is to the macro literature on endogenous technical change. Although the different characteristics of basic and applied research and public and private research have been widely recognized to be of first-order importance by policymakers,

[6] See, for instance, Romer (2001), Goolsbee (1998), and Wilson (2009).

[7] See Aghion and Howitt (1998) p. 486 and Acemoglu (2008) p. 478 for more detailed discussions.

[8] See also the earlier literature Griliches (1986), Audretsch and Feldman (1996), Anselin, Varga, and Acs (1997).

these issues have received insufficient attention from the economic growth literature. In particular, models of endogenous technological change (see **?** for a recent survey) mainly considered a uniform type of (applied) research and ignored basic research investment in the economy. A few exceptions, such as Aghion and Howitt (1996, 2009), Cozzi and Galli (2009, 2014), Morales (2004), and Mansfield (1995), have considered theoretical models with both basic and applied research investment. We contribute to this literature by building a model with rich firm dynamics as in Klette and Kortum (2004) and Lentz and Mortensen (2008) that is estimated with new firm-level micro data on firms' research investment. In addition to including the private investment in basic research, we enrich the analysis of the distinct features of basic research by identifying within- and cross-industry spillovers.

Finally, our analysis contributes to the long-standing and highly influential empirical literature on basic science (some of the earliest papers include Nelson (1959), Rosenberg (1990), Dasgupta and David (1994))[9]. In this paper, we take a macro approach and use the micro evidence on basic science to discipline and guide the general equilibrium model. This approach allows us to assess the very important effects of industrial policies on the reallocation of resources across firms and industries. Considering these policies in a general equilibrium framework is particularly helpful in understanding the potential inefficiencies and their quantitative magnitudes.

The rest of the paper is organized as follows. Section 2.2 introduces some new empirical facts on basic research spillovers to motivate our modeling approach. The discussion of our theoretical framework consists of two parts: In Section 2.3 we provide a benchmark version of the main model, characterize its dynamic equilibrium in an intuitive manner, and discuss the main mechanisms. In Section 2.4 we describe a generalization of the benchmark model that we bring to the data and estimate. Section 2.5 describes our quantitative analysis,

[9] See also Griliches (1986), Henderson, Jaffe, and Trajtenberg (1998) Link (1981), Mansfield (1980, 1981), Murray, Aghion, Dewatripont, Kolev, and Stern (2009), Rosenberg and Nelson (1994), Trajtenberg, Henderson, and Jaffe (1992).

including the welfare properties of the estimated model. Section 2.6 provides a detailed discussion of the welfare effects of various policies on the decentralized economy. Section 2.7 concludes. The Appendix contains omitted proofs and derivations (Appendix B.1), the data description (Appendix B.2), further details on within-industry spillovers (B.3), robustness checks on the stylized facts (Appendix B.4), and further details on the quantitative analysis, in particular, target moments and identification (Appendix B.5).

2.2 Empirical Facts

To understand firms' incentives to invest in basic and applied research we use unique data on the French economy combining information not only on product market and R&D investment characteristics of individual firms but also on plant and ownership information for the period 2000-2006. The R&D information comes from an annual survey conducted by the Ministry of Research that covers a large, representative cross-section of innovating firms. In this survey firms are asked to report their expenditures for basic and applied research.[10] Details regarding data sources and the policy environment are provided in Appendix B.2.

The next section presents the main empirical facts emerging from these data.

2.2.1 Basic Versus Applied Research

First, we document that private firms' investment in basic research forms a non-negligible fraction of both total private research spending and total basic research spending. Table 2.1 reports official statistics from the French Ministry of Research on public investment in basic research and private investment in applied and basic research for the period 2000-2006.

[10]The definition used by the French authorities for basic and applied research is based on the Frascati manual. It is therefore similar to the NSF definition presented in the introduction.

TABLE 2.1: EXPENDITURES ON BASIC RESEARCH

Year	Private Basic	Private Applied	Public Basic	Private $\left(\frac{\text{Basic}}{\text{Applied}}\right)$	Basic $\left(\frac{\text{Private}}{\text{Public}}\right)$
2000	802	7005	6425	.11	.13
2001	795	7748	6786	.1	.12
2002	959	8899	7037	.11	.14
2003	1092	8928	7133	.12	.15
2004	1175	9482	7338	.12	.16
2005	1227	9469	7331	.13	.17
2006	1213	10278	7755	.12	.16

Notes: Expenditures on basic and applied research in millions of euros for the period 2000-2006. Source: French Ministry of Research.

Private spending on basic research amounted to an average of 1 billion euros per year as opposed to 8.3 billion on applied research for the period 2000-2006. During the same period, public expenditures on basic research represented an average of 7 billion euros per year in France.[11] This implies that more than 11% of private research is spent on basic research. More important, almost 15% of total basic research in the economy is undertaken by private entities.[12]

The picture that emerges therefore hints at a significant involvement of the private sector in undertaking basic scientific research. Thus, ignoring the private incentives behind basic research might prevent economists and policymakers from designing more effective policies for productivity growth.

2.2.2 Multi-Industry Distributions

Another stylized fact emerging from the data is the extent of the multi-industry presence of firms. Figure 2.3 uses our micro-level data on companies between 2000 and 2006 in order to plot their empirical distribution into multiple industries.

[11] Public research has three major components: public research labs, universities, and the French National Science Foundation (CNRS). Their relative shares within public research expenditures are around 20%, 40% and 40%, respectively. To simplify the terminology we will refer to them as public research labs.

[12] Similarly, Howitt (2000), using an NSF survey, finds that around 22% of all basic research in the US during the period 1993-1997 was performed by private enterprises.

To measure multi-industry presence, we count the number of distinct SIC codes in which a firm is present. Our data allow us to identify a firm's links to different industries not only through product lines within the same firm but also through its majority ownership links. To avoid misclassification of related industries, we consider as our benchmark case the number of distinct 1-digit SIC codes (10 industries). The final sample is composed of 13,708 firm-year observations, for which we provide descriptive statistics in Table B.1 of Appendix B.2.

FIGURE 2.3: DISTRIBUTION OF FIRMS BY NUMBER OF INDUSTRIES

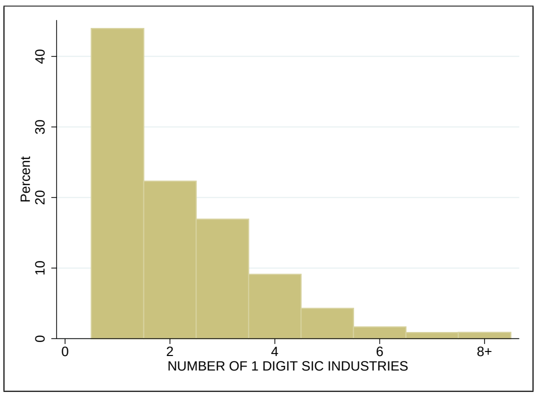

FIGURE 2.4: BASIC RESEARCH INTENSITY AND NUMBER OF INDUSTRIES

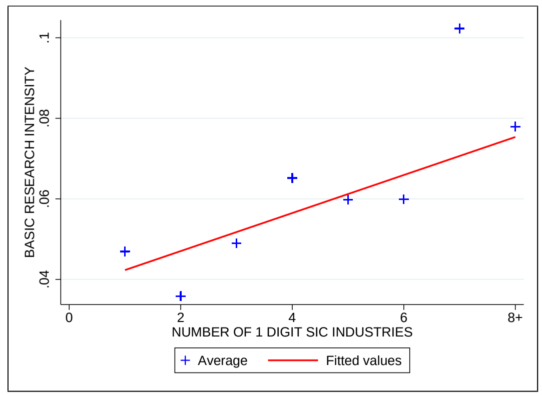

Notes: Both panels use 13,708 firm-year observations from the pooled data for the period 2000-2006. The left panel plots the share of firms as a function of the *number of 1-Digit SIC industries*, which is the number of distinct SIC codes in which a firm is present. The right panel plots *Basic Research Intensity* as a function of *Number of 1-Digit SIC Industries*. For each firm-year observation *Basic Research Intensity* is defined as the ratio of total firm investment in basic research divided by total firm investment in applied research. *Average* is the average basic research intensity for firms conditional on the number of their 1-Digit SIC industries, while the red line represents a linear fit of the firm-year observations.

On average firms are present in 2 distinct industries as defined by 1-digit SIC codes. Although nearly 44% of the firms are operating in only one industry, the remaining firms occupy a large spectrum of industries. The full distribution is plotted in Figure 2.3, and the results are very similar when using more disaggregate SIC classifications (up to the 4-digit

SIC level) or when changing the definition of an industry link.[13]

2.2.3 Basic Research and Cross-Industry Spillovers

Next we discuss the link between multi-industry presence and the private incentives for basic research. More specifically, we test Nelson's hypothesis that the largest investors in basic research should be those firms that have *fingers in many pies*. According to this argument, as the range of a firm's products and industries becomes more diversified, its incentive for investing in basic research relative to applied research should increase due to better appropriability of potential knowledge spillovers.

Figure 2.4 plots average basic research intensity against the total number of distinct 1-digit SIC codes in which the firm is present, together with a simple linear fit of the data. Basic research intensity is defined as the ratio of total firm investment in basic research to total firm investment in applied research. The figure suggests a positive and statistically significant relationship between the two variables.

Table 2.2 provides further evidence about the relationship between multi-industry presence and basic research intensity. To account for zeros in the basic research intensity data, we estimate a Tobit model. In all specifications basic research intensity is increasing in the number of industries. According to the benchmark estimation, presence in an additional industry increases a firm's basic research intensity by 3 percentage points on average. In terms of magnitude, this corresponds to a 50% increase in the average research intensity of a single industry firm.

Table B.2 in Appendix B.4 of the Appendix provides a rich set of robustness checks in terms of control variables, alternative measures of multi-industry presence, and estimation methods. Most important, it exploits historical ownership structures and changes in gov-

[13] Figures available upon request.

ernment policies as instrumental variables. The IV estimates are larger in magnitude and seem to suggest that the positive correlation is not driven by omitted variables.

TABLE 2.2: BASIC RESEARCH INTENSITY AND MULTI-MARKET ACTIVITY

	1-Digit SIC	2-Digit SIC	3-Digit SIC	4-Digit SIC
Log # of Industries	0.032***	0.027***	0.024***	0.021***
	(0.01)	(0.00)	(0.00)	(0.00)
Log Employment	0.003**	0.002	0.001	0.001
	(0.00)	(0.00)	(0.00)	(0.00)
Year & Organization Fixed Effects	YES	YES	YES	YES
N	13708	13708	13708	13708

Notes: Pooled data for the period 2000-2006. Estimates are obtained using Tobit models and relate to the marginal effect of the regressors at the sample mean. *Basic Research Intensity* is defined as the ratio of total firm investment in basic research divided by total firm investment in applied research. *Log # of Industries* is the number of distinct SIC codes in which a firm is present. *Year FE* denotes year fixed effects, and *Organization FE* denotes whether the firm operates its activity as a conglomerate or as a business group. See the Appendix for the definition of variables. Robust standard errors clustered at the firm level are in parentheses. One star denotes significance at the 10% level, two stars denote significance at the 5% level, and three stars denote significance at the 1% level.

2.2.4 Basic Research and Within-Industry Spillovers

Basic research can also contribute to economic growth through its impact on subsequent innovations within the same industry. This is because applied research can potentially build on the latest technological knowledge in the product line. However, the returns from building on the original breakthrough innovation diminish as more and more firms exploit it.

To empirically capture these types of spillovers, we turn to patent data. The idea is to pin down the age at which a patent derived from basic research cannot be distinguished from a patent derived from applied research in terms of its importance for follow-up innovations. Two empirical issues need to be addressed: (i) distinguishing patents derived from basic and applied research, and (ii) capturing the idea of successively less original contributions. We address the first point by distinguishing between patents applied for by corporations from patents applied for by public institutions.[14] We address the second point by computing

[14]While this proxy is simple to measure in the data, it potentially misclassifies the contribution of private basic patents. However, given that our interest lies in the relative difference between those two groups of

a citation-based measure of the marginal contribution of citing patents over time.

To accomplish this, we use the NBER patent data set covering the period 1974-2006.[15] The analysis of our final data set will focus exclusively on French patenters but the construction of the different variables uses information from the entire data set. For each patent we first identify citing patents across time. The age of a patent is given by the difference between the grant year of the patent and the current year. For each of the citing patents we compute the cumulative 10-years-forward citations these citing patents receive. For each originally cited patent we are then able to compute across age the mean of the citing patents' cumulative 10-years-forward citations. Our measure, *Average Citations of Citing Patents*, captures the marginal importance of each successive citing patent. Appendix B.3 provides a more detailed explanation of the construction of this variable.

Figure 2.5 presents graphical evidence for the 15,383 French patents granted between 1975 and 1985, and follows their citation patterns until 2005. More specifically, it plots *Average Citations of Citing Patents* for French public patents (blue line) and French private patents (red line).

Patents citing private patents receive on average 1.6 citations within the first 10 years. The relative importance of patents citing private patents remains stable and slightly increasing through the age of the private patent. Patents citing public patents receive on average two citations within their first 10 years. The importance of citing patents is stable until the original public patent is 8 years old, at which point there is a significant drop in citations of citing patents from 2.4 to 1.7. This is when the difference between private and public, in terms of citings' citations, becomes non-significant. Although public citations of citing patents slightly increase again after this drop, the difference remains smaller and statistically non-significant, as indicated in Table 2.3. The results are similar when using the

patents across time, time invariant errors in the classification should not impact our conclusions.

[15] The use of US patent data was linked to the availability for a long time horizon of publicly available data on patents granted, depositor classification, and the associated citations.

FIGURE 2.5: CITATION PATTERNS FOR FRENCH PUBLIC AND PRIVATE PATENTS

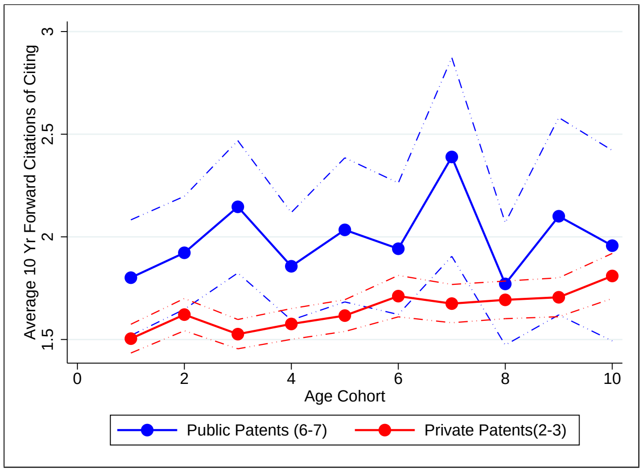

Notes: Panel plots *Average Citations of Citing Patents* for French public patents (blue line) and French private patents (red line) across patent age. *Average Citations of Citing Patents* is computed as the 10-years-forward citations of the citing patents and is measured for 15,383 patents granted in the period 1975-1985.

Wilcoxon-Mann-Whitney test. Appendix Appendix B.3 provides further robustness checks related to computation of the citations variable and public and private patent classification.

TABLE 2.3: CITATION DIFFERENCES FOR FRENCH PUBLIC AND PRIVATE PATENTS

Age	1	2	3	4	5	6	7	8	9	10
Difference	.3**	.3**	.62***	.28**	.41**	.23	.71***	.08	.39	.14
	(0.15)	(0.15)	(0.17)	(0.14)	(0.18)	(0.17)	(0.25)	(0.16)	(0.25)	(0.24)

Notes: Differences in citation patterns of 15383 patents granted by the USPTO to French private (92%) and public (8%) depositors. The difference is computed in terms of *Average Citations of Citing Patents* across patent age. *Average Citations of Citing Patents* is computed as the 10-years-forward citations of the citing patents and is measured for patents granted in the period 1975-1985. Two sample t-test with unequal variances were used. One star denotes significance at the 10% level, two stars denote significance at the 5% level, and three stars denote significance at the 1% level.

Our results are consistent with previous stylized facts related to citations of private and academic patents. Henderson, Jaffe, and Trajtenberg (1998) and Trajtenberg, Henderson, and Jaffe (1997) show that corporate patents tend to be cited less relative to academic patents and are less general in terms of the technological fields that cite them subsequently. Evidence on European patent data is more scarce, but Bacchiocchi and Montobbio (2009) show that patents of universities and public research organizations are more highly cited during the first five years but then become similar in terms of citations.

2.3 Theory: Growth with Basic and Applied Research

We first discuss important choices in modeling before going into specific details. Our theoretical framework will depart from standard endogenous growth models in a number of ways. First, we will allow private firms to invest in both basic and applied research. Second, firms will be able to operate in multiple industries. Third, our analysis relies on the appropriation of spillovers from basic research by multi-industry firms. Hence there will be cross-industry spillovers from basic research. Fourth, there will also be within-industry spillovers from basic research. Finally, we will also introduce a public research sector, which can be thought of as universities or publicly funded research labs.

The key distinction between private basic research and public basic research will be that an outcome of the former will turn immediately into a consumer product of the innovating firm, while the latter will contribute to the general pool of basic knowledge and will not turn into a consumer product until a firm uses that knowledge. This will induce a delay in the effect of public basic research, as argued in the introduction. The social trade-off will be that while private firms are better at turning abstract basic research into consumer products, they do not internalize all the spillovers associated with it. Hence, there will be room for meaningful policy interventions, which we will investigate after our quantitative analysis.

For ease of exposition and intuition, in this section, we will first outline a simplified baseline framework with myopic (one-period-ahead maximizing) firms that highlights the key elements of the main model. After deriving the theoretical results and discussing the main economic forces at play, in Section 2.4 we will describe the generalizations we make to the benchmark model.

2.3.1 Baseline Model

We consider a representative household economy in continuous time. The household consists of a measure M of workers. Each worker has one unit of labor that is supplied inelastically in the labor market. There is a unique final good $Z(t)$. The economy is a closed economy, there is no physical capital investment and all expenses are in terms of the labor units. Therefore, $Z(t)$ will also be equal to household consumption at time t.

Production

Production is divided into three major sectors: *downstream*, *midstream,* and *upstream*. The upstream sector produces intermediate goods (y_{ij}) that are used to produce industry aggregates (Y_i) in the midstream sector. Finally, the downstream sector combines these industry aggregates into the final good (Z). We will now describe them in detail.

Downstream Sector The final good $Z(t)$ is produced in the downstream sector by infinitely many competitive firms that combine inputs from M different industries according to the following constant elasticity of substitution (CES) production function

$$Z(t) = \left[\frac{1}{M}\sum\nolimits_{i=1}^{M} Y_i(t)^{\frac{\gamma-1}{\gamma}}\right]^{\frac{\gamma}{\gamma-1}}. \tag{2.1}$$

In this production function, $Y_i(t)$ is the aggregate output from industry $i \in \{1, ..., M\}$. The economy consists of $M \in \mathbb{Z}_+$ industries. In the context of firm-level data, each industry i can be thought of as a different 1-digit Standard Industrial Classification (SIC) code and $Z(t)$ is simply the aggregate GDP of the economy.[16] We normalize the price of the final good to 1 at every instant t without any loss of generality. For notational simplicity, time

[16]Note that we introduce this multi-industry structure in order to model cross-industry spillovers. To avoid any additional theoretical complications, we will focus on symmetric equilibria in which industry aggregates assume a common value.

subscripts will henceforth be suppressed.

Midstream Sector Each industry aggregate Y_i is produced competitively, combining inputs from a continuum of product lines. Let y_{ij} denote the production of upstream good j in industry i by the firm that has the best technology in that product line. Industry aggregate i is produced according to the following CES production function

$$Y_i = \left[\int_0^1 y_{ij}^{\frac{\varepsilon-1}{\varepsilon}} dj\right]^{\frac{\varepsilon}{\varepsilon-1}}. \tag{2.2}$$

Upstream Sector In product line j, the firm that has the latest (and also the best) technology produces as a monopolist according to the following linear production technology that takes only labor as an input

$$y_{ij} = q_{ij} l_{ij} \tag{2.3}$$

where $q_{ij} > 0$ is the labor productivity associated with product line j and l_{ij} is the number of production workers employed. Let us denote the wage rate in the economy by w in terms of the final good. The specification in Equation 2.3 implies that each product y_{ij} has a constant marginal cost of production $w/q_{ij} > 0$. We denote the productivity index of industry i by

$$\bar{q}_i \equiv \left(\int q_{ij}^{\varepsilon-1} dj\right)^{\frac{1}{\varepsilon-1}}. \tag{2.4}$$

Definition of a Firm In this model, as in Klette and Kortum (2004), a firm is defined as a collection of product lines in which it is the lead producer. These product lines can come from multiple industries. In what follows, $m_f \in \{1, .., M\}$ will denote the number of industries in which the firm actively operates, $n_{if} \in \mathbb{Z}_+$ will denote the number of product lines firm f owns in a given industry i, and finally n_f will stand for the total number of product lines owned by the firm and will satisfy $n_f \equiv \sum_{i \in m} n_{if}$. For notational tractability, henceforth we will drop the firm index f, when it creates no confusion.

A firm's payoff in a given product line j in industry i depends on its productivity level q_{ij}. Therefore, the payoff-relevant state of a firm is denoted by

$$\mathbf{q} = (\mathbf{q}_1, \mathbf{q}_2, \ldots, \mathbf{q}_m)$$

where $\mathbf{q}_i = \{q_{i,1}, q_{i,2}, \ldots, q_{i,n_i}\}$ is a multi-set keeping track of all the productivity levels of the firm in industry i where it has the best technology.[17] Working with such a large and complex state space proves burdensome in practice. Later on, we will impose sufficient assumptions that allow us to use a much simpler equivalent representation for a firm's product portfolio.

Example 1. *An example is helpful to summarize the description so far. Figure 2.6 illustrates an example of an economy that consists of $M = 3$ industries. It also shows an example of a firm (f) that operates in $m = 2$ industries ($i = 1$ and $i = 3$) and has $n_1 = 3$ product lines in industry $i = 1$ and $n_3 = 2$ product lines in $i = 3$. This firm does not currently operate in industry $i = 2$.*

FIGURE 2.6: EXAMPLE OF A FIRM

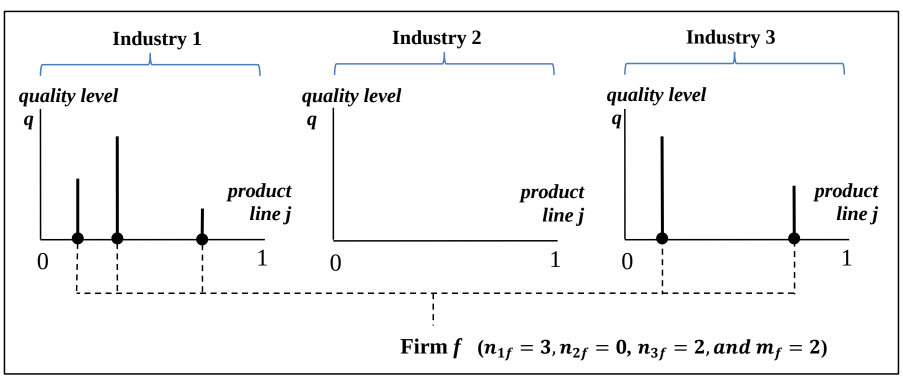

A firm's portfolio of products will expand through successful innovation. Likewise, it

[17]A multi-set is a generalization of a set that can contain more than one instances of the same member. For instance, given $j \neq j'$, a multiset $\mathbf{q}_{if}$ can contain $q_{if}(j)$ and $q_{if}(j')$ regardless of whether $q_{if}(j) = q_{if}(j')$.

will lose product lines when other firms or potential entrants successfully innovate on one of its product lines (thus stealing it). These innovations will be the source of economic growth in this economy. The next subsection will describe the details of the innovation technology.

Innovation and Technological Progress

In this economy, there are two types of innovations (basic and applied) and two different groups of agents (private and public sectors) generating productivity growth.

Firms invest in both *basic* and *applied* research, thus generating innovations that drive productivity growth. As Nelson (1959) and Aghion and Howitt (1996) describe it, fundamental advances in technological knowledge come through basic innovation and open up windows of opportunity for future research. Applied innovation builds on these existing basic innovations, thus realizing these opportunities. That being said, innovations eventually run into diminishing returns. If the latest basic innovation in a product line becomes outdated, applied innovations in that product line become less productive until a new basic innovation introduces additional fundamental knowledge that can make the applied innovation more productive again. Therefore, there will be complementarity between the two types of innovation at the aggregate level.

These innovations come from two sources: First, the private sector invests in both basic and applied innovation with the goal of increasing their market share. Second, the government uses tax revenues to fund public research labs to produce basic innovations. In what follows, we are going to describe firms' research technology and the distinction between basic and applied research. Then we will describe the public research technology.

Research by Private Firms Firms choose their flow rate of innovation and pay a labor cost that is increasing and convex in this rate. Basic and applied research levels are chosen separately, and there is no complementarity between them in terms of research costs. For the innovation production function, we will follow the literature (see Klette and Kortum (2004), Lentz and Mortensen (2008), and Acemoglu, Akcigit, Bloom, and Kerr (2013)). Firms undertake innovation by combining their existing, non-tradable intangible capital with researchers (hired at wage rate w, as with production workers) in a Cobb-Douglas production function. In our model, the intangible capital stock in a particular industry i is proxied by the number of product lines n_i that a firm owns in that industry. The production function for applied and basic research then takes the following form

$$A_i = n_i^{1-\frac{1}{\nu_a}} H_{ai}^{\frac{1}{\nu_a}} \Omega_a \quad \text{and} \quad B_i = n_i^{1-\frac{1}{\nu_b}} H_{bi}^{\frac{1}{\nu_b}} \Omega_b$$

where $\Omega_a, \Omega_b > 0$ are scale parameters, $\nu_a, \nu_b > 1$ are the inverse of the innovation production function elasticities with respect to researchers and H_{ai} and H_{bi} denote the number of researchers that firm f needs to hire in order to generate the Poisson flow rates for applied (A_i) and basic research (B_i) in industry i.

The above specifications, which are standard in this class of models, capture the idea that a firm's knowledge capital facilitates innovation.[18] Let us define $a_i \equiv A_i/n_i$ and $b_i \equiv B_i/n_i$ as the applied and basic *innovation intensities*. Similarly, let $h_a(a_i) \equiv H_{ai}/n_i$ and $h_b(b_i) \equiv H_{bi}/n_i$ be defined as the number of researchers per product line hired for applied and basic research. As a result, we can summarize the cost of doing applied and basic research as

$$C_a(a_i \mid n_i) = wn_i a_i^{\nu_a} \xi_a \quad \text{and} \quad C_b(b_i \mid n_i) = wn_i b_i^{\nu_b} \xi_b \tag{2.5}$$

where w is the wage rate, $\xi_a \equiv \Omega_a^{-\nu_a}$, and $\xi_b \equiv \Omega_b^{-\nu_b}$. Notice that total cost is directly proportional to the number of product lines.

[18]It also simplifies the analysis by making the problem proportional to the number of product lines.

Similar to Klette and Kortum (2004), Lentz and Mortensen (2008), Acemoglu, Akcigit, Hanley, and Kerr (2012), and Acemoglu, Akcigit, Bloom, and Kerr (2013), both applied and basic research are *directed* toward particular industries but *undirected* within those industries. In other words, once a firm chooses A_i and B_i, the realization of innovations will take place on a random product within industry i.

Innovation through *basic research* introduces a new generation of fundamental technical knowledge. The utilization of this fundamental knowledge for production requires what we call industry-specific *working knowledge*. This translates into our model one of the main insights on basic research presented in the introduction. Although the knowledge generated by basic research is often applicable to many industries, the ability to turn it into patents and capture its full economic value critically depends on the spectrum of activities and technologies operated by the firm. In the model, each firm has therefore some working knowledge in the industries where it has undertaken production (m). For now, we take the joint distribution $\Gamma_{m,n}$ over m and n as given but we will endogenize it in the generalized model in Section 2.4.

Let $q_{ij}(t)$ be the highest productivity technology for producing j in industry i. When a firm that has working knowledge in i produces a basic innovation that has a direct application in industry i and product line j, the same firm uses this basic knowledge for production and patents this new high-value technology. As a result, the firm improves $q_{ij}(t)$ by $\eta\bar{q}_i(t)$

$$q_{ij}\left(t+\Delta t\right)=q_{ij}\left(t\right)+\eta\bar{q}_i(t) \tag{2.6}$$

where $\eta > 0$ is the step size, and $\bar{q}_i$ is the productivity index defined in Equation 2.4. When the firm produces this new innovation, it adds this product line with the productivity improvement into its portfolio $\mathbf{q}(t+\Delta t) = \mathbf{q}(t) \cup \{q_{ij}(t+\Delta t)\}$, which generates per-period profit of $\pi\left(q_{ij}(t+\Delta t)\right)$. Going back to Example 1, firm f would increase its total number of product lines from 5 to 6 with this basic innovation.

Moreover, basic research features two potential spillovers:

- *within-industry spillover (Section 2.2.4)*: Each new basic innovation changes the evolution of the product line by introducing a radically new technology. The introduction of this new basic technology causes subsequent applied innovations to be larger until the latest basic technology becomes outdated through some random process. We refer to product lines just hit by basic innovation as *hot* product lines, as opposed to *cold* product lines, whose latest basic innovation has become outdated.

- *cross-industry spillover (Section 2.2.3)*: Each new basic innovation has the potential for spillovers into other industries. With some probability a basic innovation will generate an additional basic innovation in some other industry. If the firm has working knowledge in this other industry, it can use the innovation for production. Otherwise, the new technology contributes to the pool of existing basic knowledge and will eventually contribute to a new consumer product made by some other producer.

These two types of spillovers lie at the heart of our analysis; therefore, we will now discuss each in more detail.

Within-Industry Spillover from Basic Research *Applied research* makes use of the *within-industry* spillover from basic research and builds on the existing latest basic technological knowledge in a product line. The productivity of each applied innovation is a function of how depreciated the latest basic technology is. If the latest basic knowledge in j is undepreciated (i.e., still hot), a successful applied innovation will benefit from it and improve the latest productivity $q_{ij}(t)$ of that product line by $\eta\bar{q}_i(t)$, as in expression (2.6): $q_{ij}(t+\Delta t) = q_{ij}(t) + \eta\bar{q}_i(t)$. If the latest basic technology of the product line is depreciated (i.e., cold), a successful applied innovation will improve the latest productivity only by a magnitude proportional to $\lambda < \eta$ so that

$$q_{ij}(t + \Delta t) = q_{ij}(t) + \lambda \bar{q}_i(t). \tag{2.7}$$

We assume that a new basic technology depreciates (innovations run into diminishing returns) at a Poisson rate $\zeta > 0$. On the other hand, a new basic innovation reactivates the product line until the next time it cools down again. Let us denote the arrival rate of basic innovations to product lines by τ_b. Then during a small time interval Δt, each product line will be subject to the transition rates denoted in Table 2.4:

TABLE 2.4: TRANSITION MATRIX FOR WITHIN-INDUSTRY SPILLOVERS

	hot	**cold**
hot	$1 - \zeta\Delta t$	$\zeta\Delta t$
cold	$\tau_b\Delta t$	$1 - \tau_b\Delta t$

Figure 2.7 illustrates the implications of within-industry spillovers. In every industry, at any point in time, some product lines will be hot (red dotted lines) and some product lines will be cold, in cases where the latest technology is outdated (black solid lines). We will denote the share of hot product lines by $\alpha_i \in [0, 1]$. In a balanced-growth-path equilibrium, the share of the hot product lines will be determined through the transition rates in Table 2.4 and will remain invariant. An applied innovation is more productive if the latest basic knowledge in that product line is still "hot" and improves the productivity by $\eta\bar{q}_i$; otherwise, the contribution is only $\lambda\bar{q}_j$ where $\eta > \lambda$. This highlights the complementarity between basic and applied research.

FIGURE 2.7: WITHIN-INDUSTRY SPILLOVER

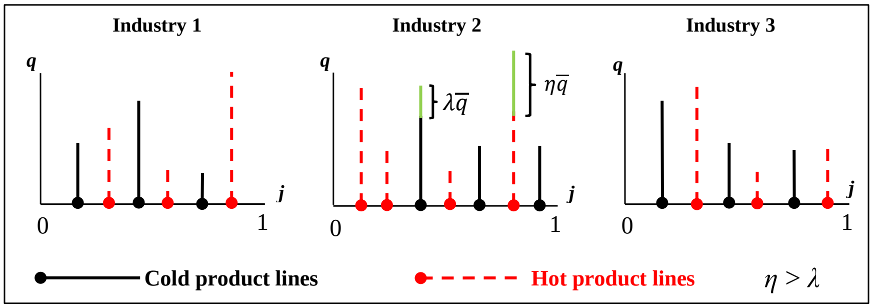

Cross-Industry Spillover from Basic Research Basic research features an additional element of uncertainty arising from random spillovers into other industries. When a firm successfully innovates through basic research, the resulting new fundamental knowledge will be applied first by that firm to increment the productivity of a random product in the target industry.

The characteristic feature of basic research we wish to capture is that it often has applications in many industries other than the one for which it was originally intended. Therefore, we will assume that when a basic innovation occurs, it applies with probability one to the target industry, and with probability $p \in (0, 1)$, it generates an additional basic innovation in another industry determined by nature at random. Thus, p is our measure of the intensity of cross-industry spillovers. Let $\mathbf{1}_{i,i'}$ be an indicator function that takes a value of one if a basic innovation in industry i has an application in industry i' and zero otherwise. Then the unconditional probabilities satisfy

$$\Pr\left[\mathbf{1}_{i,i'} = 1\right] = \begin{cases} \frac{p}{M-1} & \text{if } i' \neq i \\ 1 & \text{if } i' = i \end{cases}. \tag{2.8}$$

The spillover innovation in industry i' will be of step size η as well but will not generate additional cross-industry spillovers. This new innovation will be used by the same firm f if it has working knowledge in i'. Otherwise the production potential of this innovation will be used by the next inventor in that product line.

This structure captures Nelson's hypothesis. When a firm generates some basic knowledge, it can turn this into an immediate application only in the sectors in which it has working knowledge. In order to capture the full return from new basic scientific knowledge in industries where a firm is not present but the knowledge could have an application, the innovating firm must first patent and then license or sell the innovation to other firms in those industries. However, the applications of significant scientific advances are often *not immediate* and firms can turn them into patentable applications mostly in their own industries due to their expertise in the field.

Example 2. *Cross-industry spillovers are depicted in Figure 2.8. Firm f from Example 1 now produces a basic innovation in industry 3. This adds a new product line to the firm's portfolio and hence the number of product lines of the firm goes from 2 to 3 in $i = 3$. In addition, this basic knowledge has a potential application in industries $i = 1$ and $i = 2$ with probability p. The spillover in industry $i = 1$ is used by the firm since it has working knowledge there. However, the application in $i = 2$ is not immediate to the firm due to lack of working knowledge and therefore it is not used by the current firm but contributes to the pool of basic knowledge in $i = 2$, which can be used by another firm in the future.*

FIGURE 2.8: CROSS-INDUSTRY SPILLOVER

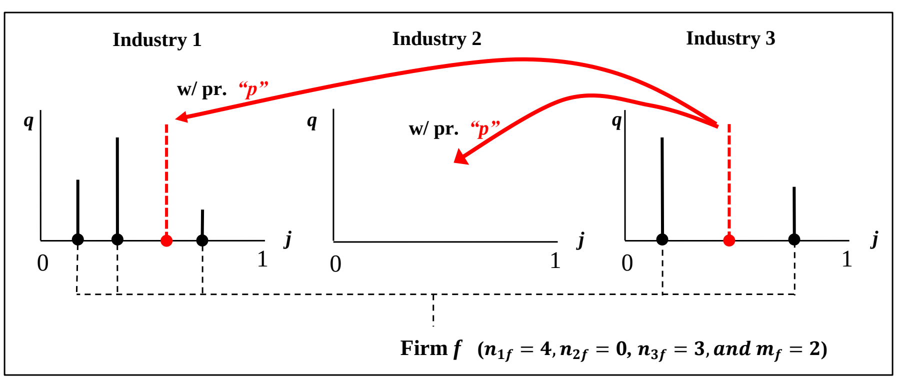

Recall that m denotes the number of industries in which a firm has working knowledge. Then the probability of a used spillover for the firm is

$$\rho_m \equiv \frac{p\,(m-1)}{M-1} \in [0,1).$$

This highlights the well-known *appropriability problem* of basic research. There is a significant chance that the new basic knowledge will be relevant to multiple industries, but it is not always clear that a firm will be in a position to exploit all of these avenues of production and patenting. However, firms operating in more industries will have a greater probability of being able to directly use all facets of a basic innovation. As Nelson puts it, firms that have *fingers in many pies* have a higher probability of using the results of basic research. A broad technological base increases the probability of benefiting from successful basic research.

Public Basic Research In our model, the academic sector will be the other source of basic knowledge creation. One of the main tasks of public research labs in an economy is to produce the necessary basic scientific knowledge that will be part of the engine for subsequent applied innovations and growth. We assume that the public research sector consists of a measure U of research labs *per industry*. Each lab receives the same transfer $\bar{R}$ from the government to finance its research which results in an overall funding level of $R = \bar{R} \times U \times M$.

We assume that each public research lab generates a flow rate of u by hiring h_u researchers with the same basic research technology as a one-product firm in Equation 2.5, so that $u = \Omega_b h_u^{\frac{1}{\nu_b}}$.[19] This specification implies that the government can affect the basic knowledge pool in the economy through the amount of funds R allocated to the academic

[19]In reality, public research labs may have a different research technology than private labs. However, obtaining data on both the inputs and outputs of individual public labs is difficult. The separate estimation of public and private innovation production functions is left for future research.

sector. The flow rate of basic innovation from the academic sector will satisfy

$$u = \left(\bar{R}/w\right)^{\frac{1}{\nu_B}} \Omega_b \tag{2.9}$$

where u is the academic basic innovation flow *per lab*. In this economy, R is a policy lever controlled by the policymaker. As with private firms, each basic innovation generated by the academic sector applies to industry i and a random product line j and makes that product line hot. However, this innovation by public labs will turn into output only upon a subsequent private applied innovation. In addition to i, the same basic knowledge will contribute to the basic knowledge pool in another industry $i' \neq i$ and line j' with probability $p \in (0,1)$. Note that the equilibrium fraction of hot product lines α will be determined by the aggregate rates of public (u) and private (b_m) basic research as well as the cool-down rate (ζ).

Remark It is important to note that we follow the empirical Ivory Tower nature of basic research and assume that innovation done by public labs is turned into consumer products only upon subsequent innovation by private firms. The lag between the creation of publicly funded innovations and actual goods production is empirically shown in a large literature.[20] This important issue is generally overlooked in the theoretical growth literature. Inclusion of this feature generates some new and interesting dynamics, such as the importance of involvement of the private sector in basic research.

Entry and Exit The research technology for a single outside entrant is assumed to be the same as that of applied innovation for a firm with a single product line. Thus if an outside entrant hires h_e researchers, it produces a flow probability of entry of $a_e = h_e^{\frac{1}{\nu_a}} \Omega_a$.

[20] Trajtenberg, Henderson, and Jaffe (1992), Rosenberg and Nelson (1994), Henderson, Jaffe, and Trajtenberg (1998), and Mowery, Nelson, Sampat, and Ziedonis (2004), among several others.

There is a mass E of outside entrants per industry. Varying this parameter will control the relative importance of outside entry in the economy. This will imply that creative destruction arising from new entrants will be equal to $E \times a_e$.

In our model, there will be both endogenous and exogenous channels for firm exit. First, a firm that loses all of its product lines to other competitors will have a value of zero and thus will exit. Second, each firm has an exogenous death rate $\kappa > 0$. When this occurs, the firm sells all of its product lines to random firms at a "fire sale" price $\mathcal{P}$.[21] On the flip side, firms will receive a buyout option with a probability that is proportional to their number of products.

Labor Market Labor is split between production (L_p) and research labor. Research labor can be further subdivided into that devoted to private basic (L_b), public basic (L_u), private applied research (L_a) and firm entry (L_e). Since the total labor supply is M workers, the labor market clearing condition is given by

$$M = L_p + L_b + L_a + L_e + L_u.$$

The labor utilization from each component can be expressed in a more concise form when we investigate the properties of the dynamic equilibrium in the next section.

Household Problem Finally, we close the model by describing the household problem that determines the equilibrium interest rate in this model. The household consumes the final good and maximizes the following lifetime utility

$$W_0 = \int_0^\infty \exp\left(-\delta t\right) \frac{C(t)^{1-\gamma} - 1}{1-\gamma} dt \tag{2.10}$$

[21] The exact value of this price will not play any role for the equilibrium determination.

where $C(t)$ is consumption at time t, γ is the constant relative risk aversion parameter, and δ is the discount rate. The household owns all the firms in the economy, which generates a risk-free flow return of r in aggregate. The household also supplies labor in the economy, through which it earns wage rate $w(t)$. Finally, the household pays a lump-sum tax $T(t) \geqslant 0$ every instant. Thus, the household's intertemporal maximization is simply to maximize Equation 2.10 subject to the following budget constraint

$$C(t) + \dot{A}(t) \leqslant r(t) A(t) + Mw(t) - T(t)$$

where $A(t)$ is the asset holdings of the household.

2.3.2 Equilibrium

In this section, we characterize the dynamic equilibrium of our model. Our focus is on a symmetric balanced-growth-path (SBGP) equilibrium where all industries start with the same initial conditions at time $t = 0$ and all aggregate variables grow at the same endogenous rate g.

In this model, three variables affect the payoff of the firm: the number of product lines n, the number of industries m, and the relative productivity

$$\widehat{q}_{ij} \equiv q_{ij}/\bar{q}_i \tag{2.11}$$

of its product lines, which is the absolute productivity in line j normalized by the productivity index $\bar{q}_i$ in industry i. Thus, each incumbent firm is characterized by its state $k \equiv (\mathbf{q}, n, m)$.

More specifically, given a government policy sequence $[T(t)]_{t=0}^{\infty}$, an SBGP equilibrium is composed of a sequence of intermediate good quantities, prices, the basic and applied innovation rates of private firms and entrants, the wage rate and interest rate, the joint distribution of multi-industry presence and product count, hot and cold product line productivity distributions, the fraction of hot product lines, i.e., $[y_k(t), p_k(t), b_k(t), a_k(t), a_e(t), w(t), r(t), \Gamma_{m,n}(t), \mathcal{F}_H(t), \mathcal{F}_L(t), \alpha(t)]_{t=0}^{\infty}$, such that all firms choose quantity and price to maximize their profits, incumbent and entrant firms invest in research to maximize their firm value, the labor market clears, the household maximizes its discounted sum of future utilities, and the distributions satisfy the relevant flow equations.

Solution of the Model The standard monopoly profit maximization delivers the following familiar equilibrium price and quantities (interested readers are referred to the Appendix B.1 for the detailed derivations)

$$y_j = \widehat{q}_j^{\epsilon} Z \qquad \text{and} \qquad p_j = \frac{1}{M\widehat{q}_j}. \tag{2.12}$$

Clearly, a monopolist's quantity is increasing and price decreasing in the relative productivity $\widehat{q}$ of the product line. Finally, the equilibrium profits of the monopolist are again increasing in its relative productivity $\widehat{q}$ and the average market size Z/M:

$$\pi(\widehat{q}) = \frac{\widehat{q}^{\varepsilon-1}}{\varepsilon} \frac{Z}{M}. \tag{2.13}$$

Next, only in this section, we focus on myopic firms that maximize their one-period-ahead returns (as opposed to forward-looking firms that maximize the discounted sum of future profits). This will allow us to provide some useful analytical results and highlight the key economic forces of our model. In our quantitative analysis (Section 2.5), we will generalize this and focus on forward-looking firms.

Myopic Firms Consider now a firm that has n product lines in m industries. Moreover, in an SBGP, an α fraction of product lines are hot. Then the maximization problem when deciding for the amount of basic research can be written as

$$\max_{b_m} \left\{ nb_m \left(1+\rho_m\right) V^H - \widetilde{w} n b^{\nu_b} \xi_b \right\}$$

where $V^H \equiv \mathbb{E}^H_{\widehat{q}} \pi \left(\widehat{q} + \eta\right)$ is the expected return to a *successful* basic innovation and $\widetilde{w} \equiv \frac{w}{Z/M}$ is the normalized wage rate. Several observations are in order. First, the expected return from basic research investment is increasing as the firm has fingers in more pies as Nelson argued (higher ρ_m). Second, the innovations are undirected within industries; therefore, the firm has to form an expectation for the expected profit $\mathbb{E}^H_{\widehat{q}} \pi \left(\widehat{q} + \eta\right)$, which means that we have to keep track of the invariant relative productivity distribution to compute V^H. Finally, both the returns and the costs are proportional to the number of product lines n, which makes the problem much more tractable and the quantitative solution manageable. Now we can express the first-order condition as

$$b_m = \left[\frac{\left(1+\rho_m\right) V^H}{\nu_b \xi_b \widetilde{w}} \right]^{\frac{1}{\nu_b - 1}}$$

The most important result here is the fact that basic research investment is increasing in the multi-industry presence of the firm. The strength of this positive relationship will be mainly governed by the probability of the cross-industry spillover parameter p, which will help us match Figure 2.4.

Fact 8. *A firm's basic research investment is increasing in its multi-industry presence.*

Both private firms and public research labs are generating basic research in this economy. It is useful to break down total basic research into its embodied and disembodied components. The distinction is based on whether the basic knowledge is immediately turned into a consumer product (embodied) or simply added to the stock of knowledge available for future innovators (disembodied). We obtain the following aggregates

$$\begin{aligned}\text{Embodied:} \quad \tau_b^e &\equiv \sum_{m=1}^{M} \mu_m(1+\rho_m)b_m \\ \text{Disembodied:} \quad \tau_b^d &\equiv \sum_{m=1}^{M} \mu_m(p-\rho_m)b_m + (1+p)u \\ \text{Total:} \quad \tau_b &\equiv \tau_b^e + \tau_b^d\end{aligned} \tag{2.14}$$

where we define the mass of product lines owned by firms in m industries by μ_m, which can be computed from the joint distribution using $\mu_m \equiv \sum_{n=1}^{\infty} n \cdot \Gamma_{m,n}$. Then τ_b^e and τ_b^d correspond respectively to the embodied and disembodied components of basic research. Note that the disembodied component includes both private spillovers that are unused and the results of public basic innovation. Finally, τ_b is simply the overall flow of basic innovation, including all spillovers.

Using this aggregate rate and the cool-down rate ζ, we can express the steady-state flow equation: the number of product lines that become hot must be equal to the number of product lines that cool down. In other words, we must have $\alpha\zeta = (1-\alpha)\tau_b$. As a result, the steady-state fraction of hot product lines is

$$\alpha = \frac{\tau_b}{\zeta + \tau_b}. \tag{2.15}$$

The share of hot product lines, those having basic knowledge that can be turned into better consumer products (α), is increasing in the amount of basic research flow. This expression highlights the role of public policy in affecting the knowledge stock. The more money is allocated to public basic research, the higher will be the basic research flow from public research labs (u), which will then increase the fraction of hot product lines through τ_b, as in Equation 2.14 and Equation 2.15.

However, a bigger α is meaningful only when there is subsequent applied research that turns this existing basic knowledge stock into consumer products. Therefore, we now turn to the applied research decision of the firms. Their maximization problem is simply

$$\max_{a} \left\{ na \left[\alpha V^H + (1-\alpha) V^C \right] - \widetilde{w} n a^{\nu_a} \xi_a \right\}$$

where $V^H \equiv \mathbb{E}^H_{\widehat{q}} \pi \left(\widehat{q} + \eta \right)$ is the expected returns from hot product lines and $V^C \equiv \mathbb{E}^C_{\widehat{q}} \pi \left(\widehat{q} + \lambda \right)$ is that from cold ones and $\widetilde{w} = \frac{w}{Z/M}$ is the normalized wage rate. When investing in applied research, firms form two types of expectations. The first one is due to the undirected nature of research: firms have to form expectations over the relative productivity $\widehat{q}$ that they are going to land on. The second, and more important one, is due to the complementarity between basic and applied research: firms take into account the fraction of hot product lines. Firms invest in applied research according to

$$a = \left[\frac{\alpha V^H + (1-\alpha) V^C}{\nu_a \xi_a \widetilde{w}} \right]^{\frac{1}{\nu_a - 1}}.$$

The crucial observation here is the complementarity between basic and applied research. In equilibrium $V^H > V^C$ since hot product lines are associated with a larger step size η. Hence, if there are more hot product lines (a higher α), each firm increases its investment in applied research.

Fact 9. *Basic and applied research investments are complementary. In particular, higher public basic research investment encourages firms to invest more in applied research.*

However, the fraction of hot product lines α is not sufficient to determine the incentives for applied research alone due to the correlation between this product state and productivity. The incentives will be a function of the fraction of hot and cold product lines and the average qualities within those types. In particular, firms must know the values of $\mathbb{E}_{\widehat{q}}^{H}(\widehat{q}+\eta)^{\varepsilon-1}$ and $\mathbb{E}_{\widehat{q}}^{C}(\widehat{q}+\lambda)^{\varepsilon-1}$ due to the exact form of the profit function in Equation 2.13. Therefore, Theorem 4 describes the laws of motion for the type-specific productivity distributions.

Let us denote the aggregate rate of applied innovation by τ_a such that

$$\tau_a = \sum_{m=1}^{M} \mu_m a_m + E a_e. \tag{2.16}$$

Note that in the baseline model, $a_m = a$ for all m, but this will not necessarily be the case in the general model in Section 2.4. Recall that τ_b^e denotes the arrival rate of embodied basic research, as defined in Equation 2.14. Now we can denote the aggregate rate of *creative destruction* (the rate at which firms lose product lines to other firms) by τ:

$$\tau \equiv \tau_a + \tau_b^e. \tag{2.17}$$

Creative destruction is determined by the rate at which incumbents produce basic innovations which can be embodied into production immediately (τ_b^e), and by the rate at which incumbents and entrants produce applied innovations (τ_a). Now we are ready to state the following lemma.

Lemma 4. *Let $\mathcal{F}_H(\cdot,t)$ and $\mathcal{F}_C(\cdot,t)$ be the aggregate product cumulative measures by type (hot or cold). The flow equations for these objects are, respectively,*

$$\dot{\mathcal{F}}_H(\widehat{q}) = -\tau\left[\mathcal{F}_H(\widehat{q}) - \mathcal{F}_H(\widehat{q}-\eta)\right] + \tau_b^e \mathcal{F}_C(\widehat{q}-\eta) - \zeta \mathcal{F}_H(\widehat{q}) + \tau_b^d \mathcal{F}_C(\widehat{q}) + g\widehat{q}[\partial \mathcal{F}_H(\widehat{q})/\partial \widehat{q}]$$
$$\dot{\mathcal{F}}_C(\widehat{q}) = -\tau_a\left[\mathcal{F}_C(\widehat{q}) - \mathcal{F}_C(\widehat{q}-\lambda)\right] - \tau_b \mathcal{F}_C(\widehat{q}) + \zeta \mathcal{F}_H(\widehat{q}) + g\widehat{q}[\partial \mathcal{F}_C(\widehat{q})/\partial \widehat{q}]$$

Proof. See Appendix B.1. □

The labor market clearing condition can now be expressed in terms of the above endogenous variables. One additional relationship we will exploit is that between the mass of labor devoted to production and the normalized wage rate. This can be derived from the goods production specification (see Appendix B.1 in the Appendix for its detailed derivation)

$$L_p = \frac{Z}{w}\left(\frac{\varepsilon - 1}{\varepsilon}\right)$$

Using this and the symmetric nature of the equilibrium, we express the labor market clearing condition as an average over industries

$$1 = \frac{1}{\widetilde{w}}\left(\frac{\varepsilon - 1}{\varepsilon}\right) + \xi_b \left(\sum_m \mu_m b_m^{\nu_b} + U u^{\nu_b}\right) + \xi_a (a^{\nu_a} + E a_e^{\nu_a}) \tag{2.18}$$

This expression equates the labor supply per industry ($= 1$ since the total labor supply is M) to labor demand for production workers; private basic researchers, which is a function of the multi-industry presence of the firms; public basic researchers, which is determined by public policy; incumbent applied researchers; and entrant basic researchers.

Finally, plugging the equilibrium intermediate good quantity from Equation 2.12 into the aggregate production functions from Equation 2.2 and Equation 2.1, we find that the aggregate output is

$$Z = \bar{q} L_p / M \tag{2.19}$$

This expression simply says that the aggregate output is equal to the product of the number of workers employed for production and the aggregate productivity index of the economy. In an SBGP equilibrium, the labor allocated for production is constant. Therefore the growth rate of aggregate output (and also output per worker) will be equal to the growth rate of the productivity index $\bar{q}$. The following proposition provides the exact growth rate of the productivity index.

Proposition 5. *In an SBGP, the growth rate of the productivity index is*

$$g = \frac{\tau_a\left[\alpha \mathbb{E}^H_{\widehat{q}}\left(\widehat{q}+\eta\right)^{\varepsilon-1} + (1-\alpha)\mathbb{E}^C_{\widehat{q}}\left(\widehat{q}+\lambda\right)^{\varepsilon-1} - 1\right] + \tau^e_b\left[\mathbb{E}_{\widehat{q}}\left(\widehat{q}+\eta\right)^{\varepsilon-1} - 1\right]}{\varepsilon - 1} \quad (2.20)$$

Proof. See Appendix B.1 □

This growth expression shows that the engines of economic progress include both applied and basic innovation. More important, the basic knowledge stock in the economy, represented by α, makes each applied innovation more valuable and contributes more to growth (since $\eta > \lambda$). This expression shows how public funding can contribute to growth through its indirect impact on private research.

2.3.3 Discussion of the Model

In this section, we briefly discuss sources of inefficiency and what policy can achieve in this model. First, as in standard quality ladder models, there are intertemporal spillovers within each product line. Second, firms simply enjoy the expected duration of monopoly power due to the competition channel of creative destruction. As a result, the private value of innovation differs from the social value of innovation. It is also worth highlighting that in this model, there could be either over- or underinvestment in R&D. In addition to the standard channels, our model features additional spillovers due to basic research, both within and across industries. Finally, there are additional static distortions due to monopoly power. However, since we are primarily interested in the dynamic inefficiencies associated with innovation and basic research, we will consider the case of a social planner who is still subject to monopoly distortions on the production side.

All of these inefficiencies will generate room for innovation policy, and our estimated model will govern whether there is over- or underinvestment in the various types of research expenditures in the decentralized equilibrium. It will also provide a framework within which to evaluate the effects of these innovation policies.

2.4 Generalizations of the Model

The previous section introduced a simplified version of the main model to highlight the key economic forces in analytical forms. Our ultimate goal in this paper is to bring this general equilibrium framework to the data. Therefore, this section generalizes the baseline model to provide richer and more realistic dynamics (with forward-looking firms and heterogeneous innovation qualities, for instance) for the economy and its agents and to give the model some more flexibility to match the data (e.g., introducing the fixed cost of doing basic research). Those not interested in the technical details can skip directly to the quantitative Section 2.5.

Stochastic Innovation Step Sizes Stokes (1997) argued that technological breakthroughs do not necessarily derive from basic research. According to the "Pasteur Quadrant" hypothesis, applied research efforts can potentially also lead to important technological changes. Our first generalization takes this possibility into account by introducing stochastic innovation step sizes into the model. We assume, as in Klette and Kortum (2004) and Lentz and Mortensen (2008), that these step sizes are drawn from exponential distributions. For basic research, the mean of the distribution is always η. For applied research, the distribution mean is η if the product line is hot and λ if it is cold. It is important to note that we do not take any stand on the comparison of the average step sizes ($\eta > \lambda$ or vice versa) and let them be determined by the data.

Fixed Cost of Basic Research In our sample, some firms do not invest in basic research. To capture this fact, we generalize the basic research technology by introducing a fixed cost of doing basic research. At each instant, a firm with n product lines draws a fixed labor cost of doing basic research $n\phi \geqslant 0$, where ϕ is distributed according to the distribution $\mathcal{B}(\cdot)$. Then a firm that operates in n product lines and has a fixed cost of basic research ϕ this period has the following cost function $C_b(b_m \mid n, \phi) = nc_b(b_m \mid \phi)$. This implies that firms will follow a cutoff rule as a function of their multi-industry presence ϕ_m^* such that they will not invest in basic research if $\phi > \phi_m^*$. Otherwise, in addition to the variable cost, they will also pay the fixed cost.

Industry Expansion (Start-up Buy-outs) In the baseline model, we took the working knowledge of the firms (m) as exogenously given. We now endogenize m by introducing the possibility of buy-out offers for new entrants. The economy features $E \times a_e$ flow of entry at any instant. We will assume that a ς fraction of new entrants will meet a randomly selected incumbent firm. Thus, an incumbent will have a flow rate of incoming buy-out offers

$$x \equiv \varsigma E a_e / F.$$

where F is the equilibrium measure of firms. If $\bar{n}$ denotes the average number of product lines per firm, then $F = 1/\bar{n}$. Clearly this new company will be from a new industry with probability $(1 - m/M)$ or from an industry that already exists in the incumbent's portfolio with probability m/M. Our goal is to keep the M&A margin as tractable as possible, and we will achieve this by assuming that the M&A price that the incumbent firm has to pay is equal to the full surplus of the new merger. The resulting invariant joint distribution $\Gamma_{m,n}$ over multi-industry presence m and firm product count n is described in Appendix B.1.

Forward-Looking Firms For expositional purposes, in the previous section we described the model with myopic firms that maximize their one-period-ahead returns. For the rest of our analysis, we relax this assumption and consider firms that maximize the discounted sum of future returns. The analysis of this new model is very similar to that of the previous model except that the returns to innovation take the form of a value function that takes into account all future contingencies. The following proposition provides the exact forms of the value of a firm that has a productivity portfolio $\hat{\mathbf{q}}$ and operates in m industries.

Proposition 6. *Let the value of a firm with a productivity portfolio $\hat{\mathbf{q}}$ in m industries be denoted by $\mathcal{V}(\hat{\mathbf{q}}, m)$. This value is equal to*

$$\mathcal{V}(\hat{\mathbf{q}}, m) = \frac{Z}{M}\left[\sum_{\hat{q}\in\hat{\mathbf{q}}} V(\hat{q}) + nV_m\right]$$

where

$$V(\hat{q}) = \frac{\hat{q}^{\epsilon-1}}{\varepsilon\left[r + \tau + \kappa + g\left(\epsilon - 2\right)\right]}$$

and

$$(r-g)\,V_m = \max_{a,b}\left\{\begin{array}{c} -\widetilde{w}\left[h_a(a) + h_b(b) + \mathbf{1}_{(b>0)}\phi\right] \\ +a\left[\alpha V^H + (1-\alpha)\,V^C + V_m\right] + b\,(1+\rho_m)\left[V^H + V_m\right] \\ +x\left(1 - \frac{m}{M}\right)\left[V_{m+1} - V_m\right] - \tau V_m + \kappa\mathbb{E}_{\hat{q}}V(\hat{q}_t) \end{array}\right\}. \tag{2.21}$$

The analogous production values are defined as $V^H \equiv \mathbb{E}^H_{\hat{q},\eta}V(\hat{q}+\eta)$ and $V^C \equiv \mathbb{E}^C_{\hat{q},\lambda}V(\hat{q}+\lambda)$.

Proof. See Appendix B.1 □

This important result has a number of implications. First, the value of a firm has a tractable additive form across product lines. Moreover, the firm value has two major components: the first component is the production value $V(\hat{q})$, which simply computes the sum of the future discounted profits where the effective discount rate takes into account the rate of creative destruction τ, the exogenous destruction rate κ, and the obsolescence of the relative productivity $\hat{q}$ due to the growth of $\bar{q}$. The second component is the R&D option value

V_m, which is a direct function of the multi-industry presence due to the associated internalization of spillovers. Finally, because of the stochastic nature of step sizes, the expectations now integrate over the productivity (which are type specific) and step size.

Welfare Finally, we close this section by describing the SBGP equilibrium welfare. In an SBGP equilibrium that has an initial consumption C_0 and a growth rate of g, welfare is computed as

$$W(C_0, g)^{SBGP} = \int_0^\infty \exp(-\delta t) \frac{(C_0 e^{gt})^{1-\gamma}}{1-\gamma} dt = \frac{1}{1-\gamma} \left(\frac{C_0^{\frac{1-\gamma}{\varepsilon-1}}}{\rho - (1-\gamma) g} - \frac{1}{\rho} \right)$$

We will report our results in consumption-equivalent terms. In particular, when two different public policies T_1 and T_2 generate different SBGP equilibrium welfare values as $W(C_0^{T_1}, g^{T_1})$ and $W(C_0^{T_2}, g^{T_2})$, we will report β such that

$$W(\beta C_0^{T_1}, g^{T_1}) = W(C_0^{T_2}, g^{T_2}).$$

In other words, β constitutes the compensating differential in initial consumption that equalizes the welfare of the two proposed policy environments. It therefore provides an intuitive measure for evaluating policy tools. This completes the description of the theoretical environment. Now we are ready to move on to the quantitative analysis.

2.5 Quantitative Analysis

In this section we describe the estimation strategy used. We will assume that the fixed costs are drawn from a lognormal distribution $\mathcal{B}(\phi)$ with mean $\bar{\phi}$ and variance σ^2. As a result, the set of parameters of the model is

$$\theta = \{\delta, \gamma, \varepsilon, p, \eta, \lambda, E, U, \zeta, \nu_a, \nu_b, \xi_a, \xi_b, \kappa, \bar{\phi}, \sigma\} \in \Theta.$$

During the period we consider, there was existing government support for R&D activities in France. Our data set contains information on the publicly funded portion of private R&D. On average, 10% of private R&D was funded publicly. Therefore in our estimation, we introduce a uniform subsidy to the total R&D spending of the firm $\psi = 0.10$. The government has a balanced budget every period, so that the sum of total subsidies (S) and public research funding (R) must be equal to tax revenues, that is

$$T = S + R = \psi \left[\sum_{m=1}^{M} \mu_m C_B(b_m \mid \phi) + C_A(a) \right] + UC_B(u \mid \bar{\phi})$$

where T is a lump-sum tax on consumers. In France, during 2000-2006, the fraction of GDP devoted to public research labs and academic institutions was approximately 0.5%. Therefore, we pick R/Z, which is the share of GDP devoted to public basic research, to be 0.5%.

2.5.1 Estimation Method

In our data set, for each firm f and each time period t, we have a vector of N observables from the actual data $\mathbf{y}_{ft} \equiv \left[y_{ft}^{1} \ \ . \ . \ . \ \ y_{ft}^{N} \right]'_{N \times 1}$ including the number of industries in which the firm is present, sales, profits, and labor costs. Let the entire data set be denoted by $\mathbf{y}$.

We use the simulated method of moments (SMM) for the estimation.[22] Define $\Lambda(\mathbf{y})$ and $\Lambda(\theta)$ to be, respectively, the vectors of real data moments (generated from $\mathbf{y}$) and equilibrium model moments (generated for some vector of parameters θ). Since certain mo-

[22]See Bloom (2009) and Lentz and Mortensen (2008) for further description and usage information on SMM.

ments require knowledge of the joint distribution of firms over the number of products and industries (m, n) and the portfolio of product qualities $\mathbf{q}$, which has no apparent analytic form, we simulate a large panel of firms to calculate $\Lambda(\theta)$ to a high degree of accuracy.[23]

Our proposed estimator minimizes a quadratic form of the difference between these two vectors

$$\widehat{\theta} = \arg\min_{\theta \in \Theta} [\Lambda(\theta) - \Lambda(\mathbf{y})] \cdot W \cdot [\Lambda(\theta) - \Lambda(\mathbf{y})]$$

where W is the weighting matrix. We use a diagonal weighting matrix with entries equal to the inverse square of the data moment value, or in notational terms $W_{ii} = 1/\Lambda_i(\mathbf{y})^2$ and $W_{ij} = 0$ for $i \neq j$. In our estimation, we use 26 moments. We pick moments that are most informative for the unique features of our model. In particular, we target both the intensive and extensive margins of basic research intensity as it varies with multi-industry presence. Since multi-industry presence is one of the key determinants of innovation, we target both the mean and the variance of that quantity. In addition, we include aggregate and conditional firm-level growth rates. Since our model is macro growth model and household's welfare (and accordingly the policy analysis) depends crucially on the level of aggregate growth, hitting that moment is of particular importance. For that purpose, we boost the weighting on the aggregate growth moment.[24] To capture the within-industry spillover, we target the spillover differentials described in Section 2.2.4. Finally, to further inform the model parameters on firm dynamics, we include the mean return on sales, the R&D/production labor ratio, the exit rate, and mean firm age by size. The details of the moments and identification are described in Appendix Appendix B.5.

[23]For our results, we simulate 32K firms with a burn-in time of 100 years and 100 time steps per year.

[24]Increasing the weighthing factor to 3 was sufficient to align the aggregate growth rate in the data and the model.

2.5.2 Computer Algorithm Outline

An equilibrium of this model is described by a system of five equations in the five variables $(\tau_a, \tau_b^e, \tau_b^d, \widetilde{w}, g)$. This system can be evaluated using the following procedure:

1. Calculate α and the distribution of $\widehat{q}$ using τ_a, τ_b^e, τ_b^d, and g according to Theorem 4 and Equation 2.15.

2. Calculate g using, τ_a, τ_b, and the distribution over $\widehat{q}$ with Equation 2.20.

4. Calculate $V^H = \mathbb{E}^H_{\widehat{q},\eta} V(\widehat{q} + \eta)$ and $V^C = \mathbb{E}^C_{\widehat{q},\lambda} V(\widehat{q} + \lambda)$ using the relevant step size distribution and the type-specific productivity distributions.

5. Find a_m and b_m using first-order conditions with $\widetilde{w}$ from Equation 2.21.

6. Impose an upper bound on n and find the steady state $\Gamma_{m,n}$ using the flow rates in Appendix B.1.

7. Compute the updated values of τ_a, τ_b^e, and τ_b^d using Equation 2.16 and Equation 2.14.

8. The difference between the conjectured and updated values of τ_a, τ_b^e, τ_b^d, and g in conjunction with the labor market clearing differential from Equation 2.18 constitute the five desired equations.

We use Powell's (Powell (1970)) hybrid equation solver to solve this set of equations for a given set of parameters. To minimize the SMM objective function, we perform a search over the parameter space using a combination of a naive simulated annealing algorithm and a Nelder-Mead simplex (Nelder and Mead (1965)) algorithm. See Zangwill and Garcia (1981) for more information on solving systems of nonlinear equations.

TABLE 2.5: PARAMETER ESTIMATES

#	Description	Sym	Value	#	Description	Sym	Value
1.	Discount Rate	δ	0.038	9.	Applied Cost Curvature	ν_a	1.367
2.	CRRA Utility Parameter	γ	2.933	10.	Basic Cost Curvature	ν_b	1.538
3.	Elasticity of Substitution	ε	5.800	11.	Applied Cost Scale	ξ_a	1.228
4.	Cross-industry Spillover	p	0.113	12.	Basic Cost Scale	ξ_b	5.437
5.	Basic Step Size	η	0.079	13.	Exogenous Exit Rate	κ	0.006
6.	Applied Step Size	λ	0.049	14.	Basic Fixed Mean	$\bar{\phi}$	-4.761
7.	Mass of Entrants	E	0.502	15.	Basic Fixed Std. Dev.	σ	0.327
8.	Mass of Academic Labs	U	0.491	16.	Product Cooldown Rate	ζ	0.116

2.5.3 Estimation Results

Table 2.5 reports the values of the estimated structural parameters. The estimated values of the discount rate and CRRA utility parameters are within their standard macro ranges. The elasticity of substitution parameter generates $17\%(= 1/\varepsilon)$ gross profits, resulting in 7.9% net profits after subtracting R&D expenses as a share of sales.

One of the most important parameters of our model is the cross-industry spillover parameter $p = 0.11$, which measures the probability that a basic innovation will have an additional immediate application. This estimate affects the extent to which basic innovations contribute to cross-sectional growth. In equilibrium, firms operate in two industries out of 10 on average. Therefore, any given spillover is not embodied with probability $89\%(= 8/9)$. Given that the probability of having a spillover is 11%, the probability of having a disembodied spillover is $10\%(= 0.11 * 0.89)$.

The estimated innovation size of basic research is $\eta = 7.9\%$ and the innovation size of each new applied innovation is $\lambda = 4.9\%$. This implies that basic research (hot product lines) makes applied innovation $60\%(= 7.9/4.9 - 1)$ more productive.

Additionally, each basic innovation has a within-industry spillover. The cool-down rate of hot product lines is estimated to be $\zeta = 0.12$, which indicates that a basic innovation affects the subsequent innovations in the same product line for almost $8.3(= 1/0.12)$ years on average.

The elasticity of applied innovation counts with respect to the research dollars spent is estimated to be 0.73 ($= 1/\nu_a$) and similarly the elasticity of basic innovation with respect to the basic research investment is 0.65 ($= 1/\nu_b$). These values are close to the elasticity estimates in the literature, which typically finds a value around 0.5 (Blundell, Griffith, and Windmeijer (2002), Griliches (1990), Pakes and Griliches (1984) and Kortum (1992, 1993)).

2.5.4 Goodness of Fit

In this section, we will first focus on the moments that we targeted in our estimation and then turn to the moments that we did not directly target but still find useful in understanding the model's performance.

Targeted Moments Table 2.6 contains the moments from the actual data and our estimated model.

TABLE 2.6: MOMENTS USED IN ESTIMATION

#	Description	Model	Data	#	Description	Model	Data
1-8	Basic Research Extensive	See Figure 2.10		21	R&D/Labor	0.284	0.260
9-16	Basic Research Intensive	See Figure 2.9		22	Employment Growth	0.111	0.103
17	Mean Industries	2.217	2.203	23	Aggregate Growth	0.013	0.015
18	Mean Square Industries	7.213	6.975	24	Spillover Differential	8.378	8.000
19	Return on Sales	0.032	0.032	25	Age, Small Firms	11.53	14.99
20	Exit Rate	0.082	0.091	26	Age, Large Firms	18.69	24.87

The results indicate that the model performs very well in generating firm and industry dynamics similar to those in the data. As documented in Section 2.2.1, a significant fraction of innovating firms invest in basic research. In particular, 29% of firms are investing in basic research, which was 27% in the data. We also capture the positive relationship between the extensive margin of basic research and multi-industry presence, as evidenced in Table 2.6 and Figure 2.9.

The positive correlation between multi-industry presence of a firm and its basic research intensity was one of the primary motivations for introducing multi-industry presence into

FIGURE 2.9: FRACTION POSITIVE BASIC BY # INDUSTRIES

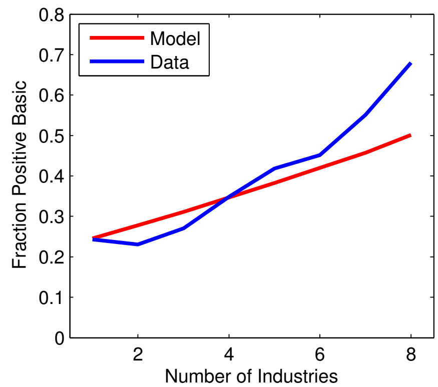

FIGURE 2.10: BASIC RESEARCH INTENSITY BY # INDUSTRIES

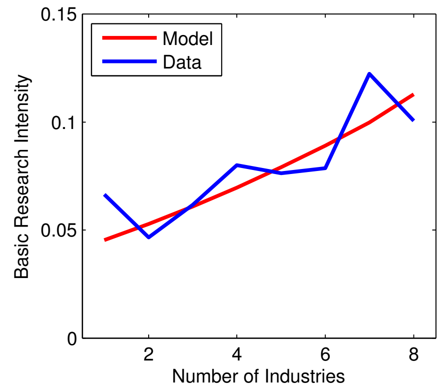

our model. As explained previously in the text, multi-industry presence plays an important role in increasing basic research incentives, by allowing a greater potential to internalize the positive spillovers from basic research. In our reduced-form analysis, we found a significant and positive correlation between multi-industry presence and basic research intensity. This has been the key moment to identify the cross-industry spillover parameter. Our model successfully generates this positive correlation.

In the data, firms operate on average in 2.2 industries, and the same is true in the model. Furthermore, we find the mean squared in the model to be 7.2, compared to 7.0 observed in the data.

The table above reports some additional moments that are not captured by the stylized facts. For instance, the mean profitability is 3.2% in the model and in the data. The prime determinants of profitability are the step sizes for basic and applied innovation. However, these also affect the investment levels for both types of research, since this increases the return to successful innovation. Therefore, the step size parameters are set to be a compromise between hitting the profitability moment and the research investment and growth moments.

We are targeting additional moments regarding research investments. The first is the average ratio of total research labor to production labor by incumbent firms. The model comes very close to hitting this ratio exactly (28.4% vs 26.0%), largely in order to hit the aggregate growth and return on sales.

All of these components of the economy determine the aggregate growth rate. Our model matches the observed growth rate closely. Our model economy grows at a rate of 1.3%, while the French economy grew at an average rate of 1.5% during the period studied (2000-2006).

Untargeted Moments In this part, we discuss our model's prediction about some of the moments that we did not directly target.

Interestingly in the data the correlation between profitability and basic research intensity is not significantly different from zero. The same implication emerges from our model. In the baseline model, the correlation between profitability and basic research intensity is only 0.04. This result emerges because basic research investment is determined through the multi-industry presence of the firms, whereas profitability is determined by the share of hot and cold product lines, type of research investment, and the productivity distribution $\mathcal{F}(\hat{q})$ in the economy.

Our model naturally generates a positive correlation between multi-industry presence and firm size, which is also empirically true in the data. This arises since both of these moments are strongly correlated with firm survival. In the model, we find a correlation of 0.52 between the log employment and multi-industry presence. In the data, this value is 0.76.

Another stylized fact in our data is that the firm size distribution is highly skewed. This is a well-known feature that is documented extensively in the literature. For detailed references, see **?**. In our model, we capture this fact with a skewness of the firm size distribution of 4.12. This value is 3.07 in the data.

Our estimates indicate that entrants play an important direct role in overall growth.

The innovation rate from entrants is 0.43%, whereas that number is 0.92% for incumbents. That implies that entrants account for 32% of growth. Though our number is for the French economy, our number is in line with Foster, Haltiwanger, and Krizan (2001) who find that 25% of productivity growth in the US comes from new entry.

We will now focus on the details of the equilibrium and the social planner's problem to study the efficiency properties of this economy. Then we will turn to various policies that could address this inefficiency.

2.5.5 Endogenous Variables of the Baseline Economy

The following table provides equilibrium values for some of the important variables in the model:

TABLE 2.7: DECENTRALIZED ECONOMY: ENDOGENOUS VARIABLES (IN PERCENTAGES)

ψ	τ_a	τ_b^e	τ_b^d	L_p	L_b	L_u	L_e	L_a	α	g	β
10	22.0	0.58	0.28	85.6	0.53	0.52	4.5	8.9	6.9	1.34	95.3

In this table, τ_a denotes the aggregate rate of applied innovation by incumbents and entrants, whereas τ_b^e and τ_b^d denote the aggregate rates of embodied and disembodied basic innovation, respectively. The next five columns report the labor allocations into production, private basic, public basic, entry, and applied research. The remaining columns report the fraction of hot product lines α, the ratio of consumption to that for the social planner's economy, the growth rate g, and the welfare in consumption equivalent terms β.

The model highlights the dynamic misallocation of research effort and its welfare consequences. In our benchmark economy, 85.6% of labor is used for production, and 14.4% is employed for innovation activities. Among researchers, roughly 7% are engaged in basic research activities. Note that this composition within innovation activities will be the central focus of the policy analysis, since uninternalized spillovers are one of the main sources of inefficiency. As a consequence, the arrival rate of basic innovation in our baseline economy is significantly smaller (25 times) than that of applied innovation. This translates into significant welfare losses, with the economy achieving only 95.3% of welfare with respect to the social planner's optimum, which we will analyze next.

2.5.6 Quantifying the Social Planner's Optimum

In this section, we are going to provide the solution to the social planner's problem under two scenarios. In the first, the planner, as in the Ivory Tower approach of the baseline case, cannot appropriate public basic research returns. This is illustrated in Panel A. In the second, the planner can appropriate and use new basic inventions for production immediately. This case is reported in Panel B. We will set the welfare to 100% in this case and report the remaining welfare numbers relative to this baseline. Finally, recall that we are considering a planner who controls firm's research labs but not their production decisions. The following table summarizes these results.

TABLE 2.8: SOCIAL PLANNER'S OPTIMUM (IN PERCENTAGES)

τ_a^{SP}	$\tau_b^{e,SP}$	$\tau_b^{d,SP}$	L_p^{SP}	L_b^{SP}	L_u^{SP}	L_e^{SP}	L_a^{SP}	α	g	β
A. NON-APPROPRIATED PUBLIC RESEARCH										
19.1	5.1	0.2	82.9	5.6	0.5	3.7	7.3	31.1	1.80	98.7
B. APPROPRIATED PUBLIC RESEARCH										
18.5	6.5	0.0	82.6	4.6	2.3	3.5	7.0	35.8	1.93	100

One striking feature of the solution to the social planner's problem under both scenarios is that the fraction of labor devoted to research activities is not substantially greater than in the decentralized equilibrium. In particular, in Panel A the total labor allocated to research activities was 14% in the decentralized economy, while it is only 17% when set by the social planner.

Indeed, the dominant misallocation here is not that between production and research, as is common in this class of models, but among the various types of research activities, in this case, applied and basic innovation. In the decentralized economy, only 1.05% of the total labor force is devoted to basic research, whereas in the social planner's economy, this number rises to 6.1%. In other words, the social planner devotes 36% of research labor to basic research, whereas this fraction was only 7% in the decentralized economy. This happens on both the intensive and the extensive margins of basic research. In fact, the planner finds it optimal to employ nearly all private research labs, regardless of their fixed cost draw.

Another interesting and important finding is that in the case of applied innovation, there is actually an *overinvestment* in the baseline economy. The applied research labor utilization (including entrants) is 13.6% in the decentralized case. This figure drops to 11% in the social planner's solution. This is in spite of the fact that the fraction of hot product lines rises from 7% to 31%, meaning the average step size of an applied innovation rises by almost a third.

The net result of the above changes is that growth rises to 1.8% from 1.34%. Overall, the decentralized economy's welfare corresponds to a decrease of $3.4\%(= 1 - 95.3/98.7)$ in consumption-equivalent terms from the social planner's optimum. The following policy experiments will try to bridge this gap.

Panel B reports these numbers for the case of appropriated public research. The main difference is the sizable increase in the labor devoted to public basic research, which rises to 2.3% relative to 0.5% in both the decentralized economy and Panel A. When public basic research turns into production immediately, this contributes to aggregate growth by an additional 0.13 percentage point and increases welfare by an additional 1.3 percentage points in consumption equivalent terms. Policies such as the Bayh-Dole Act allow academic researchers to appropriate their innovations through patenting. In our setting, this would correspond to an increase in the rate of appropriation of innovation by public researchers. We will consider this as a policy tool in Section 2.6.3.

2.6 Policy Analysis

In this section, we analyze the impact of different types of research subsidies. Given our distinction between basic and applied research, it seems natural to propose different subsidy policies for different types of research spending. However, this could potentially generate a moral hazard problem, since firms would have an incentive to misreport the type of research they undertake, which is very difficult for a policymaker to verify. However, it is still useful to consider this hypothetical case to form a benchmark.

This section is organized as follows: Section 2.6.1 starts with this hypothetical case, Section 2.6.2 considers a uniform research subsidy as in the real world, Section 2.6.3 considers only optimal funding of public research labs, and finally Section 2.6.4 combines both uniform subsidy and public research funding using feasible policy tools.

2.6.1 Type-Dependent Research Subsidy

Assume first that the policymaker can distinguish between different types of research efforts and accordingly provide differentiated subsidy rates. Let ψ_a and ψ_b denote the applied research and basic research subsidy rates, respectively. The share of GDP allocated to public research (R/Z) is kept constant by the policymaker. Note that an increase in the subsidy rate (ψ_a or ψ_b) reduces research costs for the firm and leads to an increase in research effort as a result. The following table reports the optimal subsidy rates and resulting equilibrium variables.

Since the underinvestment is mainly in basic research, the optimal type-dependent sub-

TABLE 2.9: TYPE-DEPENDENT RESEARCH SUBSIDY (IN PERCENTAGES)

ψ_a^{TD}	ψ_b^{TD}	τ_a^{TD}	$\tau_b^{e,TD}$	$\tau_b^{d,TD}$	L_p^{TD}	L_b^{TD}	L_u^{TD}	L_e^{TD}	L_a^{TD}	α^{TD}	g^{TD}	β^{TD}
14	50	19.3	4.50	0.38	83.1	5.3	0.50	3.7	7.5	29.6	1.75	98.2

sidy dictates a much larger subsidy rate for it, namely, $\psi_b = 50\%$ and $\psi_a = 14\%$. Here, the major component of policy is a fivefold increase in the subsidy rate for basic research, whereas the subsidy rate on applied innovation remains roughly the same.

The large value for the basic research subsidy is straightforward to understand. There are spillovers associated with basic innovation that are not internalized by firms. By subsidizing this type of innovation, we can mitigate this effect. This policy can almost achieve the level of welfare seen in the relevant social planner's case in Panel A of Table 2.8 (98.2% vs 98.7%).

As discussed above, this policy is hard to implement in the real world due to the moral hazard problem. Therefore, we focus on a policy providing a uniform subsidy across the economy.

2.6.2 Uniform Private Research Subsidy

With this policy, the government subsidizes a fraction ψ of each firm's total research investment, keeping the share of funds allocated to academic research constant. Note that such a policy subsidizes not only basic research but also applied research. The following table summarizes the results of the optimal uniform subsidy rate.

TABLE 2.10: UNIFORM RESEARCH SUBSIDY (IN PERCENTAGES)

ψ^{UP}	τ_a^{UP}	$\tau_b^{e,UP}$	$\tau_b^{d,UP}$	L_p^{UP}	L_b^{UP}	L_u^{UP}	L_e^{UP}	L_a^{UP}	α^{UP}	g^{UP}	β^{UP}
31	25.4	1.52	0.26	81.8	1.54	0.49	5.41	10.8	13.2	1.70	96.1

Our analysis of the baseline economy and the planner's economy documented a slight underinvestment in research overall and a large misallocation between the different types of research. A uniform subsidy is therefore ill suited to address these issues as it cannot directly affect the allocation between research types, and any attempt to subsidize basic research will only worsen the overinvestment in applied research. Although the optimal type-dependent basic subsidy is 50%, the optimal uniform subsidy is only 31%, due to cross-subsidization of applied research whose optimal level was 14%.

Under this policy, we are allocating a larger fraction of the labor force to research relative to the social planner's economy. Overall, the researcher's share goes up to 18% from 14%. As a result, we have too few workers devoted to production of the consumption good (81.8%) relative to the social planner's allocation (82.9%), which reduces the initial consumption of the baseline economy. Even though we have more labor working for research, the economy grows at a lower rate (1.7%) than the social planner's (1.8%). This interesting result emerges due to the misallocation of researchers between basic and applied innovation. The welfare gain from this policy is 0.8 percentage points, which is only $28\%(= 0.8/2.9)$ as large as that for the type-dependent policy.

Although the underinvestment in basic research is sizable, the uniform policy partially makes up for this at the cost of worsening the overinvestment in applied research. The main lesson to be drawn from this is that a uniform research subsidy should take into account its negative welfare consequences through its oversubsidization of applied research. Finding a feasible method to differentiate basic and applied research is essential for better innovation policies.

2.6.3 Optimal Academic Fraction of GDP

In this section we will look for the optimal public funding level for academic research as a fraction of GDP (R/Z) keeping the baseline subsidies fixed. This is particularly important because the rate of academic innovation is a major factor in determining the share of hot product lines, which determines the effectiveness of applied innovation.

TABLE 2.11: OPTIMAL ACADEMIC FUNDING (IN PERCENTAGES)

ψ	R/Z	τ_a	τ_b^e	τ_b^d	L_p	L_b	L_u	L_e	L_a	α	g	β
10	0.8	22.0	0.54	0.76	85.4	0.50	0.80	4.5	8.9	10.0	1.37	95.4

The results indicate that welfare can be improved by allocating a larger fraction of GDP to academic research. In particular, when we consider only this as the policy tool, the optimal funding rate is 0.8% of GDP. This figure is only 0.5% in France (and in the benchmark case). Such a policy increases the fraction of hot product lines from 6.9% to 10%. However, this policy makes a limited contribution to growth and welfare due to the Ivory Tower nature of academic research.

So far we have assumed that public innovations have no immediate effect on productivity in a particular product line. However, one can argue that policies such as the Bayh-Dole Act, which was adopted in the US in 1980, enhance the applicability of academic innovations by allowing universities to retain ownership of inventions made using federal funds. This is an interesting policy question, which we can analyze in our setting. We will study this appropriability problem by considering a scenario where academic research is focused on immediately applicable innovations half of the time (Panel A), and one where all innovations are immediately applicable (Panel B). In the latter extreme, academic research functions much like private corporate research. It should be noted that academic research, as we have all experienced, has a much wider set of objectives than purely generating consumer products (such as education, to say the least). Our analysis will abstract from those considerations.

Table 12 summarizes the results of the optimal academic policy.

Under these alternative cases, the optimal level of academic funding is increasing in the applicability of academic research. While the optimal fraction is 1.9% when half of the

TABLE 2.12: ACADEMIC FUNDING WITH ALTERNATIVE BAYH-DOLE SCENARIOS (IN PERCENTAGES)

PANEL A: BAYH-DOLE=50%												
ψ	R/Z	τ_a	τ_b^e	τ_b^d	L_p	L_b	L_u	L_e	L_a	α	g	β
10	1.9	21.7	1.35	0.94	84.7	0.38	1.9	4.4	8.7	16.4	1.49	96.3
PANEL B: BAYH-DOLE=100%												
ψ	R/Z	τ_a	τ_b^e	τ_b^d	L_p	L_b	L_u	L_e	L_a	α	g	β
10	3.7	20.8	3.4	0.01	83.7	0.24	3.7	4.1	8.2	22.7	1.69	98.1

innovations are immediately applicable, this number rises to 3.7% when all innovations are immediately applicable.

These optimal allocations bring with them large welfare gains, between 2 and 3 percentage points. Some of this is simply due to the increase in the Bayh-Dole factor, while the rest can be attributed to the optimal allocation of academic funding. The growth rate rises as well, attaining levels seen in the social planner's optimum in the last case.

Our results highlight the special role of academic research in overall growth and show the complementarities present between public and private research. Allocating resources to academic research has not only a direct effect on growth but also an indirect effect by making private research more productive. However, one should also note that this particular policy alone cannot make up for the underinvestment in research on the part of the private sector. Therefore, the next policy experiment is of particular importance.

2.6.4 Optimal Feasible Policy: Uniform Subsidy and Academic Budget

Our final policy experiment combines both of the feasible policies that have been considered thus far individually. We will allow both the uniform subsidy rate and the academic funding rate to be chosen by the policymaker. The advantage of considering both types of policies is to introduce more freedom to control the incentives for both types of research in a largely separate way. In particular, ψ and R/Z are going to be the choice variables in this exercise. The following table contains the results of this experiment

TABLE 2.13: OPTIMAL ACADEMIC AND UNIFORM POLICY (IN PERCENTAGES)

ψ	R/Z	τ_a	τ_b^e	τ_b^d	L_p	L_b	L_u	L_e	L_a	α	g	β
31	0.7	25.4	1.5	0.6	81.6	1.5	0.7	5.4	10.8	15.5	1.72	96.1

When considered jointly, the optimal uniform R&D subsidy is 31% and the optimal fraction of GDP allocated for public research is 0.7%. This combination generates a limited improvement, however. The growth rate increases 0.02 percentage points relative to the optimal uniform policy of Table 2.10 and 0.35 percentage point relative to the optimal public funding of Table 2.11. These improvements are mitigated by the limited applicability of the academic research. Next, we consider both uniform policy and academic funding jointly under the scenario where academic innovations have immediate applications for production.

TABLE 2.14: ACADEMIC AND UNIFORM POLICY WITH 100% APPLICABILITY (IN PERCENTAGES)

ψ	R/Z	τ_a	τ_b^e	τ_b^d	L_p	L_b	L_u	L_e	L_a	α	g	β
26	3.3	23.5	3.8	0.0	81.2	0.9	3.3	4.9	9.7	24.7	1.92	98.3

When academic innovations are geared toward consumer needs (i.e., have immediate application for production), the share of academic funding becomes much more effective. In this case, the optimal fraction of GDP allocated for academic research is 3.3%. Under this policy, the growth rate increases to 1.92% and achieves the highest welfare result among all policies considered. By using the level of academic funding to reach the proper share of researchers, the policymaker is able to lower the uniform subsidy, thus reducing needless cross-subsidization of applied research. Under the current policy 19% of the labor force is allocated to research, an increase over that of the baseline case. This time around, the composition of workers between applied and basic research is closer to the social optimum.

To summarize our findings, we first considered the most widely discussed policy, which is a uniform subsidy. Using this tool optimally yielded limited improvement in welfare due to oversubsidization of applied research since the policy could not distinguish between the research types with different spillover and productivity implications. Considering a policy combination that governs both private and academic research in which the researchers can appropriate the returns to their innovations could generate a significant improvement. The first main conclusion to be drawn for innovation policy is the importance of recognizing different types of innovations and the impact of policies on these types of research. The second is that it is important to take into account both the direct and indirect effects of academic research on productivity growth and the role of researchers' appropriability of their outcomes when considering growth and innovation policies.

2.7 Conclusion

In this paper, we distinguished between basic and applied research and identified spillovers associated with each. Our quantitative analysis highlighted the importance of this distinction. Indeed, in the competitive equilibrium, applied research is overinvested and basic research is underinvested. As a result, imposing a uniform research subsidy does not generate the expected welfare improvement due to inefficient cross-subsidization of applied research. An increase in the uniform subsidy improves the underinvestment in basic research by worsening the overinvestment in applied research.

The key message of our paper is that standard R&D policies can accentuate the dynamic misallocation in the economy. Our findings point to the need for policies that target basic research more directly. One method of achieving this is by increasing the intellectual property rights granted to academic researchers. Alternatively, one can reward collaboration between universities and the private sector, which would encourage focusing on research that can more directly lead to tangible gains in production technologies.

Our paper took a first step in trying to quantify the inefficiencies regarding different types of research and innovation efforts. There are still important open questions awaiting further study. In particular, the effect of university licensing and collaboration opportunities between universities and the private sector are two examples. We hope further structural work will be undertaken to enhance our understanding of the aforementioned issues, which can then guide the relevant innovation policies.

Chapter 3

Transition to Clean Technology

Joint work with Daron Acemoglu (MIT), Ufuk Akcigit (University of Pennsylvania and NBER), and William Kerr (Harvard Business School)

3.1 Introduction

Recent economic research has recognized the importance of transition to clean technology in controlling and reducing fossil fuel emissions and potentially limiting climate change.[1] Recent empirical work has also shown that innovation may switch away from dirty to clean

[1]On climate change, see, e.g., Stott, Peter, D.A. Stone, and M.R. Allen (2004) on the contribution of human activity to the European heatwave of 2003, Emanuel, Kerry (2005) and Landsea, Christopher (2005) on the increased impact and destructiveness of tropical cyclones and Atlantic hurricanes over the last decades; and Nicholls, Robert, and Jason Lowe (2006) on sea-level rise. On economic costs of climate change, see Mendelsohn, Robert, William Nordhaus, and Daigee Shaw (1994), Pizer, William (1999), and Weitzman, Martin (2009). On economic analyses of climate change, see, e.g., Golosov, Mikhail, John Hassler, Per Krusell, and Aleh Tsyvinski (2011), Hassler, John, and Per Krusell (2012), Krusell, Per, and Anthony Smith (2009), MacCracken, Christopher, James Edmonds, Son Kim, and Ronald Sands (1999), Nordhaus, William (1994), Nordhaus, William, and Joseph Boyer (2000), Nordhaus, William (2008), and Stern, Nicholas (2007). On endogenous technology and climate change, see, Acemoglu, Daron, Philippe Aghion, Leonardo Bursztyn, and David Hemous (2012), Bovenberg, Lans, and Sjak Smulders (1995), Bovenberg, Lans, and Sjak Smulders (1996), Goulder, Lawrence, and Koshy Mathai (2000), Goulder, Lawrence, and Stephen Schneider (1999), Grimaud, Andre, Gilles Lafforgue, and Bertrand Magné (2011), Hartley, Peter, Kenneth Medlock, Ted Temzelides, and Xinya Zhang (2011), Hassler, John, Per Krusell, and Conny Olovsson (2011), Popp, David (2002), Popp, David (2004), and Van der Zwaan, Robert, Reyer Gerlagh, G. Klaassen, and L. Schrattenholzer (2002).

technologies in response to changes in prices and policies. For example, Newell, Richard, Adam Jaffe, and Robert Stavins (1999) show that following the oil price hikes, innovation in air-conditioners turned towards producing more energy-efficient units compared to the previous focus on price reduction; Popp, David (2002) finds that higher energy prices are associated with a significant increase in energy-saving innovations; Hassler, John, Per Krusell, and Conny Olovsson (2011) estimate a trend break in factor productivities in the energy-saving direction following the era of higher oil prices; and Aghion, Philippe, Antoin Dechezlepretre, David Hemous, Ralf Martin, and John Van Reenen (2012) find a significant impact of carbon taxes on the direction of innovation in the automobile industry and further provide evidence that clean innovation has a self-perpetuating nature feeding on its own past success. Based on this type of evidence, Acemoglu, Daron, Philippe Aghion, Leonardo Bursztyn, and David Hemous (2012) suggest that a combination of (temporary) research subsidies and carbon taxes can successfully redirect technological change towards cleaner technologies. Several conceptual and quantitative questions remain, however. The first is whether, in the context of a micro-founded quantitative model, reasonable policies can secure a transition to clean technology. The second is whether, in the presence of carbon taxes, there is still any role for significant research subsidies. The third concerns how rapidly the transition to clean technology should take place under optimal policy.

A systematic investigation of these questions necessitates a micro model of innovation and production where clean and dirty technologies can compete given the prevailing policies and research incentives (and the direction of technological change) are also endogenously determined as a function of these policies.[2] It also necessitates a combination of micro data for the modeling of competition in production and innovation, and a quantitative model flexible enough to represent realistic dynamics of carbon emissions and potential

[2] Acemoglu, Daron, Philippe Aghion, Leonardo Bursztyn, and David Hemous (2012) assume that clean and dirty inputs are combined with a constant elasticity of substitution, which allows for limited form of competition between clean and dirty technologies.

climate change. This paper is an attempt in this direction.

Our first contribution is to develop a tractable and parsimonious microeconomic model for this purpose. In our model, which we view as an abstract representation of the energy production and delivery sectors, each one of a continuum of intermediate goods can be produced either using a dirty or clean technology, each of which has a knowledge stock represented by a (separate) quality ladder. Given production taxes (which are differential by type of technology), profit-maximizing final good producers choose which technology to utilize. Profit-maximizing firms also decide whether to conduct research to improve clean or dirty technologies. Clean research, for example, leads to an improvement over an existing clean technology, though there is also a small probability of a breakthrough which will build on and surpass the dirty technology when the dirty technology is the frontier in the relevant product line. Research and innovation decisions are impacted both by policies and the current state of technology in the two sectors. For example, when clean technology is far behind, most research directed to that sector will generate incremental innovations that cannot be profitably produced (unless there are very high levels of carbon taxes). However, if clean research can be successfully maintained for a while, it slowly becomes self-sustaining as the range of clean technologies that can compete with dirty ones expands as a result of a series of incremental innovations.

Our second contribution is to estimate parameters of this model using microdata on R&D expenditures, patents, sales, employment and firm entry and exit from a sample of US firms in the energy sector. The data we use for this exercise are from the Census Bureau's Longitudinal Business Database and Economic Censuses, the National Science Foundation's Survey of Industrial Research and Development, and the NBER Patent Database. We design our sample around innovative firms in the energy sector that are in operation during the 1975-2004 period. We use our sample to directly estimate some key parameters of the model and the initial distributions of dirty- and clean-energy product lines.[3] In particular, we estimate two of the key parameters of the model with regression analysis using R&D and patents. We also estimate the initial distribution of productivity gaps between clean and dirty technologies in the economy by allocating the patent stocks of firms innovating in these technology areas across the three-digit industries in which the firms are operating. The remaining four crucial parameters are estimated using simulated method of moments (we impose the discount rate and the fraction of scientists in the labor force from the data rather than estimating these from the model). We show that, despite its parsimony, the fit of the model to a rich and diverse set of moments not targeted in the estimation is fairly good.

We then combine this structure with a parsimonious model of the carbon cycle. Our modeling of the carbon cycle follows Golosov, Mikhail, John Hassler, Per Krusell, and Aleh Tsyvinski (2011) and is fairly flexible despite its simplicity. Our final contribution is to use this estimated quantitative model for the analysis of optimal policy, in particular optimal carbon taxes and research subsidies,[4] and a range of counterfactual policy experiments.

Our main results are as follows. Though it is intuitive to expect that carbon taxes should do most of the work in the optimal allocation—because they both reduce current emissions and encourage R&D directed to clean technologies—quantitatively we find a major role for research subsidies. For example, with an annual discount rate of 1% (similar to the number favored by Nordhaus, William (2007)) and focusing on constant policies, the optimal research subsidy is 61% (meaning that the government pays for 61 cents out of every dollar of R&D expenditure for clean technology) while the carbon tax is 16%. The numbers are more extreme with a discount rate of 0.1% for the social planner (similar to the number favored by Stern, Nicholas (2007)) but with a similarly major role for research subsidies: a

[3]See Popp, David (2006) and Jaffe, Adam, David Popp, and Richard Newell (2010) for background on technology, R&D and innovation in the energy sector.

[4]We do not allow additional tax instruments to remove the monopoly distortions in the economy.

research subsidy of 95% and a carbon tax of 44%. When we allow time-varying policies, the overall pattern is broadly similar and still heavily relies on research subsidies, but with some notable differences: first, the research subsidy is initially slightly more aggressive and then declines somewhat over time; second, with a discount rate of 1%, carbon taxes are backloaded (low, in fact zero, for an extended period of time and then high); and third, with a discount rate of 0.1%, carbon taxes are frontloaded (starting out higher and declining over time).[5] Despite the differences between the models, the reason for the major role for research subsidies is related to the one emphasized in Acemoglu, Daron, Philippe Aghion, Leonardo Bursztyn, and David Hemous (2012). [6] Research subsidies are powerful in redirecting technological change, and given this, it is not worth distorting the initial production too much by introducing heavy carbon taxes. It is important to emphasize that research subsidies are not being used just because there is a market failure (and an uninternalized externality) in research. In fact, in our model, in the absence of externalities from carbon, or in the special case in which there is only a dirty or a clean sector, the social planner would have no reason to use research subsidies—because a scarce factor, skilled labor, is being used for research and no other purpose, and thus the social planner cannot increase the growth rate by subsidizing research. The reason why the social planner heavily uses research subsidies is because when carbon creates negative externalities, inducing a transition to clean technology is an effective way of reducing future carbon emissions by changing

[5] Our time-varying optimal policy results need to be interpreted with caution, since the resulting optimal policy sequence is not time consistent.

[6] Major differences between the models include: (1) here the damage from atmospheric carbon is modeled as impacting production along the lines of previous literature rather than directly utility; (2) here there is no "environmental disaster" threshold, making it possible for us to calibrate the parameters more closely to data and without taking a position on carbon emissions in the rest of the world; (3) in contrast to the constant elasticity of substitution formulation, dirty and clean sectors are not complements in our model, but explicitly compete in each product line. This last one is the most important distinction, enabling us to use microdata on innovation and production. It also implies a different pattern of production distortions from carbon taxes. In Acemoglu, Daron, Philippe Aghion, Leonardo Bursztyn, and David Hemous (2012), carbon taxes are particularly distortionary when the dirty sector is behind (and thus its relative prices high because of the imperfect substitutability). In contrast, in our model the carbon tax is least distortionary when the clean technology has already taken over or is about to take over almost all product lines.

the path of technological progress.

Another useful comparison is to current US policies. We estimate the effective research subsidy from the differential between clean and dirty firms in our sample in the use of federally funded R&D expenditure. Utilizing this estimate and different values of effective carbon tax at the moment and its likely values in the future, our estimated optimal policies are quite different from their US counterparts, and we show that under US policies, climate change dynamics will be significantly different (and worse).

In terms of counterfactual policies, we investigate the welfare costs of just relying on carbon taxes and delaying intervention. The most notable result here is that the welfare costs of delaying the optimal policy by 50 years (laissez faire) is very significant. With a discount rate of 1%, delaying optimal policy by 50 years has a welfare cost equivalent to a permanent 8% drop in consumption. With a discount rate of 0.1%, the consumption-equivalent welfare cost is 16.6%. The costs of relying just on carbon tax (without any research subsidy) are more modest but still significant, 4.2% and 3.4%, with the same two discount rates, respectively.

We also consider several variations and robustness checks to show which aspects of the model are important for our main theoretical and quantitative results. In particular, we investigate the implications of using different discount rates and estimates of the damage of carbon concentration on economic activity, allowing different degrees of distortions from research subsidies, different estimates of the microeconomic elasticities in the R&D technology, and different distributions of productivity gaps between clean and dirty technologies. Overall, most of the main qualitative and quantitative features of optimal policy appear to be fairly robust to a range of plausible variations.

Our model combines elements from four different lines of research (and is thus related to each of these four lines). First, we build on the growing literature on quantitative general equilibrium models of climate change, such as Golosov, Mikhail, John Hassler, Per

Krusell, and Aleh Tsyvinski (2011), Hassler, John, and Per Krusell (2012), Krusell, Per, and Anthony Smith (2009), Nordhaus, William (1994), Nordhaus, William, and Joseph Boyer (2000), Nordhaus, William (2008), and Stern, Nicholas (2007). We follow these papers in introducing a simple model of the carbon cycle and the economic costs of carbon emissions in a general equilibrium model, and then characterizing optimal policy. Second, we introduce endogenous and directed technological change along the lines of Acemoglu, Daron (1998) and Acemoglu, Daron (2002) in a model where producers have a choice between clean and dirty production methods. In combining these two first lines of research, we are following Acemoglu, Daron, Philippe Aghion, Leonardo Bursztyn, and David Hemous (2012) as well as several other papers listed in footnote 1 above. Third, we develop a tractable but rich model of competition between dirty and clean technologies building on the literature on step-by-step competition as in Harris, Christopher, and John Vickers (1995), Aghion, Philippe, Christopher Harris, Peter Howitt, and John Vickers (2001), and Acemoglu, Daron and Ufuk Akcigit (2012). Fourth, we model the microeconomics of innovation, employment and output dynamics building on Klette and Kortum (2004), where each firm consists of a number of products and technologies (different from other applications, technologies here are different from products because of the competition between clean and dirty sectors).

In estimating a general equilibrium model of firm-level innovation and employment dynamics, we follow Lentz and Mortensen (2008) and Acemoglu, Daron, Ufuk Akcigit, Nick Bloom, and William Kerr (2012). We differ from existing work in this area in three important respects, however. First, we combine this type of estimation strategy with a model of clean and dirty technologies and estimate some of the parameters of the R&D technology directly from microdata. Second, rather than focusing on steady-state comparisons, we study non-steady-state dynamics, which is crucial for the question of transitioning to clean technology. Third, we characterize optimal policies in such a framework.

The remainder of the paper is organized as follows. Section 3.2 introduces our model and characterizes the equilibrium. Section 3.3 describes the dataset we will use for estimation and quantitative evaluation, outlines the different components of our estimation strategy, and presents the estimates of some of the parameters we obtain from our micro data. Section 3.4 presents the simulated method of moments estimates of our parameters and discusses the fit of the model. Section 3.5 quantitatively characterizes the structure of optimal environmental policy. In this section, we also conduct a range of counterfactual exercises. Section 3.6 discusses a range of robustness exercises intended to convey which sorts of assumptions and parameters are important for the qualitative and quantitative results of the paper. Section 3.7 concludes.

3.2 Model

In this section, we present our baseline model. This is a simple dynamic general equilibrium model, where final output combines intermediates produced either using a clean or dirty technology. The productivity of the dirty and clean technology for each intermediate is represented by a quality ladder. Production is also subject to taxes, so profit-maximizing final good producers choose whether to use clean or dirty intermediates as a function of the productivity gap between the two and taxes. Research is directed towards clean or dirty technology, and progresses both with incremental research increasing productivity by one rung on the quality ladder and with occasional breakthrough research which enables the firm to surpass the current frontier technology. Research is conducted both by entrants and incumbent firms which already hold a portfolio of products and technologies. Finally, dirty technology contributes to carbon emissions, which create potential economic damage. We next describe each module of the model in turn.

3.2.1 Preferences and Endowments

We model an infinite-horizon closed economy in continuous time. Since the consumer side is not our focus, we simplify the discussion by modeling it with a representative household with a logarithmic instantaneous utility function. The lifetime utility is then

$$U_0 = \int_0^\infty e^{-\rho t} \ln C_t dt, \tag{3.1}$$

where C_t is the household's consumption at time t and $\rho > 0$ is the discount rate. We assume that the representative household consists of mass one of production workers and mass L^s of "scientists" who will be employed in R&D activities. All workers supply one unit of labor inelastically. The representative household owns all the firms in the economy, so its problem will be to maximize Equation 3.1 subject to the following budget constraint

$$w_t^u + w_t^s L^s + \Pi_t \geqslant C_t,$$

and the usual no Ponzi-game condition. Here Π_t is the total sum of corporate profits net of R&D expenses, w_t^u and w_t^s are the wage rates (and thus wage incomes) of the production and R&D workers.

Since the economy is closed, there is no physical capital, and intermediates and the R&D sector use labor, aggregate consumption is equal to the production of the final good:

$$C_t = Y_t.$$

3.2.2 Final Good Technology, Intermediate Production and Pricing

The final good is produced by combining a measure one of intermediates with an elasticity of substitution equal to one. In addition, its production is negatively affected by the amount of atmospheric carbon concentration, which we denote by S_t. We follow the formulation suggested by Golosov, Mikhail, John Hassler, Per Krusell, and Aleh Tsyvinski (2011), which builds on earlier work by Mendelsohn, Robert, William Nordhaus, and Daigee Shaw (1994), Nordhaus, William (1994), and Nordhaus, William (2008), and assume

$$\ln Y_t = -\gamma \left(S_t - \bar{S}\right) + \int_0^1 \ln y_{i,t} di, \tag{3.2}$$

where $\bar{S} > 0$ is the pre-industrial level of the atmospheric carbon concentration, $\gamma \geqslant 0$ is a scale parameter, and $y_{i,t}$ is the quantity of intermediate good i. When $\gamma = 0$, Equation 3.2 gives the standard (unitary elasticity of substitution) production function for combining intermediates to produce a final good. When $\gamma > 0$, levels of atmospheric carbon concentration above the pre-industrial level reduce productivity with elasticity γ, for reasons discussed in Mendelsohn, Robert, William Nordhaus, and Daigee Shaw (1994), Nordhaus, William (1994), Nordhaus, William (2008), and Stern, Nicholas (2007).

A feature of Equation 3.2, which will play a central role in our quantitative exercise, is worth noting: the proportional cost of a unit increase in atmospheric carbon concentration is independent of its current level. Though nonlinearities, or even major threshold effects, are likely to be present in the impact of atmospheric concentration on economic activity, this functional form is not only in line with assumptions made by other economic approaches to climate change (e.g., Nordhaus, William (1994), Nordhaus, William (2007), Nordhaus, William, and Joseph Boyer (2000), Stern, Nicholas (2007), Golosov, Mikhail, John Hassler, Per Krusell, and Aleh Tsyvinski (2011), but also enables us to study the implications of carbon emissions from our economy, with parameters estimated from the US and calibrated to US aggregates, without taking a position on the path of carbon emissions on the rest of the world. Without this assumption, the marginal cost of carbon emissions, and thus optimal policy, would strongly depend on assumptions on the evolution of emissions from other countries.

Each intermediate $i \in [0,1]$ can be produced with either a dirty or a clean technology, and when it is produced with the clean (dirty) technology we denote it by $y_{i,t}^{c}$ ($y_{i,t}^{d}$). We will sometimes refer to clean and dirty technologies as clean and dirty "sectors," and we also use the terms "intermediaries" and "product lines" interchangeably.

Firm f can produce intermediate i with either a clean or dirty technology ($j \in \{c,d\}$) with the following production function $y_{i,t}^{j}(f) = q_{i,t}^{j}(f)\, l_{i,t}^{j}(f)$, where $l_{i,t}^{j}(f)$ is the employment of (production workers) by this firm and $q_{i,t}^{j}(f)$ is the labor productivity of the technology that this firm has access to for producing with clean or dirty technology j in product line i. In equilibrium, only firms with the highest technology either in the clean or dirty sector will produce, so we simplify this equation by suppressing firm indices and with the implicit convention that the labor productivity q always refers to the most advanced clean or dirty technology, thus writing:

$$y_{i,t}^{j} = q_{i,t}^{j} l_{i,t}^{j}.$$

Though only firms with the most advanced technology for intermediate i within the clean or dirty sector can ever produce it, because of taxes it is not necessarily the most advanced technology between these two sectors that will always be active. In particular, there is a tax at the rate τ_t^{j} on sector (technology) j at time t, which implies that the marginal cost of production is

$$MC_{i,t}^{j} = \frac{\left(1+\tau_t^{j}\right) w_t^{u}}{q_{i,t}^{j}},\ j \in \{c,d\} \text{ and } i \in [0,1],$$

where w_t^{u} is the wage rate of production workers. We define tax-adjusted labor productivity as

$$\widetilde{q}_{i,t}^{j} \equiv \frac{q_{i,t}^{j}}{1+\tau_{t}^{j}}.$$

In equilibrium, only the technology with the lower marginal cost (inclusive of taxes)—or equivalently the one with the higher tax-adjusted labor productivity—will produce. Summarizing this, we have

$$\text{produce intermediate } i \text{ with technology } j \text{ if } \widetilde{q}_{i,t}^{j} > \widetilde{q}_{i,t}^{-j} \text{ where } j \neq -j \in \{c, d\}.$$

We assume that if clean and dirty technologies have equal tax-adjusted labor productivities, each produces with probability 50% at any point in time.[7] Thus, the tax-adjusted technology level use in the production of intermediate i at time t can be written as

$$\bar{q}_{i,t} = \begin{cases} \widetilde{q}_{i,t}^{d} & \text{if } \widetilde{q}_{i,t}^{d} \geqslant \widetilde{q}_{i,t}^{c} \\ \widetilde{q}_{i,t}^{c} & \text{otherwise} \end{cases}.$$

Finally, we also assume that at the initial date $t = 0$, for each leading technology of quality $q_{i,0}^{j}$, there also exists an intermediate good of quality $q_{i,0}^{j}/\lambda$, which ensures that markups in the initial date will not exceed λ (this will be guaranteed endogenously in subsequent dates).

3.2.3 Innovation, the Quality Ladder and Dynamics

Labor productivity for each intermediate (for each technology) evolves as a result of innovation. Research is directed towards clean or dirty technologies. A successful innovation leads to one of two types of innovation. The first is an *incremental* innovation, which takes

[7]In other models of this type, e.g., Acemoglu, Daron and Ufuk Akcigit (2012), which of two firms produces is immaterial. But here, since one of them uses the dirty technology and thus will contribute to carbon emissions, we need to specify exactly who produces in this case.

place with probability $1 - \alpha$; and the second is a *breakthrough* innovation, which takes place with probability α (independently of all other events).

If research directed to sector $j \in \{c, d\}$ leads to an incremental innovation, then the innovator improves over the sector j technology of a randomly chosen intermediate. This is incremental innovation in the sense that it enables the innovator to go up by one rung in the quality ladder over producing technology, and we assume that each rung corresponds to an improvement of $\lambda > 1$. Consequently, labor productivity of technology j in intermediate i at time t can be written as

$$q_{i,t}^{j} = \lambda^{n_{i,t}^{j}},$$

where $n_{i,t}^{j} \in \mathbb{Z}_{+}$ is the effective number of steps that this technology has taken since time $t = 0$ (when all technologies are, by assumption, normalized to $q_{i,0}^{j} = 1$).

Relative productivity of dirty to clean technology in intermediate i at time t can be written as

$$\frac{q_{i,t}^{d}}{q_{i,t}^{c}} = \lambda^{n_{i,t}}$$

where

$$n_{i,t} \equiv n_{i,t}^{d} - n_{i,t}^{c} \in \mathbb{Z}$$

is defined as the technology gap between dirty and clean sectors in product line i at time t. In what follows we will need to keep track of the share of intermediates with technology gap $n \in \mathbb{Z}$, and we denote this by $\mu_{n,t} \in [0, 1]$ at time t.

Breakthrough innovations, on the other hand, enable the successful innovator to improve by one rung over the frontier technology, even if this frontier is set by the alternative technology—i.e., a breakthrough clean innovation will improve over the dirty technology even if the latter is far ahead of the clean sector, thus allowing the clean sector to leapfrog the dirty one.

Therefore, conditional on an innovation in technology j for intermediate i between times t and $t + \Delta t$, the evolution of $q_{i,t}^{j}$ can be written as

$$q_{i,t+\Delta t}^{j} = \begin{cases} \lambda q_{i,t}^{j} & \text{with probability} & 1 & \text{if } q_{i,t}^{j} \geqslant q_{i,t}^{-j} & \text{(incremental)} \\ \lambda q_{i,t}^{j} & \text{with probability} & 1-\alpha & \text{if } q_{i,t}^{j} < q_{i,t}^{-j} & \text{(incremental)} \\ \lambda q_{i,t}^{-j} & \text{with probability} & \alpha & \text{if } q_{i,t}^{j} < q_{i,t}^{-j} & \text{(breakthrough)} \end{cases} .$$

Let z_t^j denote the aggregate innovation rate which is the sum of incumbents' and entrants' innovation rates in technology j. The law of motion for the technology gap $n_{i,t}$ can then be expressed as follows:

$$n_{i,t+\Delta t} = \begin{cases} n_{i,t} - 1 & \text{with probability} & (1-\alpha)\, z_t^c \Delta t & \forall n_{i,t} \\ n_{i,t} + 1 & \text{with probability} & (1-\alpha)\, z_t^d \Delta t & \forall n_{i,t} \\ -1 & \text{with probability} & \alpha z_t^c \Delta t & \text{if } n_{i,t} > 0 \\ n_{i,t} - 1 & \text{with probability} & \alpha z_t^c \Delta t & \text{if } n_{i,t} \leqslant 0 \\ 1 & \text{with probability} & \alpha z_t^d \Delta t & \text{if } n_{i,t} \leqslant 0 \\ n_{i,t} + 1 & \text{with probability} & \alpha z_t^d \Delta t & \text{if } n_{i,t} > 0 \\ n_{i,t} & \text{otherwise} & & \end{cases}$$

Note that innovations here have a creative destruction element (e.g., Aghion, Philippe, and Peter Howitt (1992), Grossman, Gene, and Elhanan Helpman (1991)) because, by improving over an existing product typically operated by another firm, they transfer the leading-edge technology to the current innovator.

In what follows, for notational and computational tractability, we assume that the gross tax rates are multiples of λ such that $1 + \tau_t^j = \lambda^{m_t^j}$. Since taxes are chosen by the social planner, especially when λ is not too large, this is without much loss of generality. Given this assumption, we can write

$$\frac{1 + \tau_t^d}{1 + \tau_t^c} = \lambda^{m_t},$$

where

$$m_t \equiv m_t^d - m_t^c,$$

and thus tax-adjusted technologies can be written as

$$\frac{\widetilde{q}_{i,t}^d}{\widetilde{q}_{i,t}^c} = \frac{q_{i,t}^d}{1+\tau_t^d}\frac{1+\tau_t^c}{q_{it}^c} = \lambda^{n_{i,t}-m_t}.$$

We will say that dirty is the leading (tax-adjusted) technology if $n_{i,t} > m_t$; the two technologies are neck and neck if $n_{i,t} = m_t$; and clean is the leading technology otherwise.

3.2.4 Firms, R&D and Free Entry

Following **?**, we define a firm as a collection of leading-edge technologies. Let u_f^j denote the number of intermediates where firm f has the leading-edge technologies in sector $j \in \{c, d\}$ (but these are not necessarily more advanced than the technologies available in the other sector $-j$). Again following **?**, we assume that u_f^j captures the stock of knowledge of the firm for further innovations with technology $j \in \{c, d\}$. In particular, we assume that firms combine their knowledge stock u_f^j with scientists (R&D workers) H^j in order to generate a Poisson flow rate of X^j new innovations (in continuous time) according to the following production function

$$X^j = \theta\left(H^j\right)^\eta\left(u^j\right)^{1-\eta}, \tag{3.3}$$

where $\eta \in (0,1)$ is the R&D elasticity with respect to scientists and $\theta > 0$ is a scale parameter. Thus the *variable* cost of generating a flow rate of X^j is simply $w_t^s u\left(x^j\right)^{\frac{1}{\eta}}\theta^{-\frac{1}{\eta}}$ where $x^j \equiv X^j/u^j$ is the innovation intensity per product line and w_t^s is the wage rate of scientists. In addition, R&D activities also require each firm to hire a number of scientists per product line (as fixed cost). We assume that, per product line, firm f will need to hire

$F_{I,i}u$ scientists where $F_{I,i,t} \in [(1-\xi) F_I, (1+\xi) F_I]$ is an iid (across firms and over time) draw with mean F_I and $\xi \in (0,1)$.[8] Hence, the total cost of R&D for firm i performing R&D directed at technology $j \in \{c,d\}$ at time t is

$$\begin{aligned} C_t\left(u, x^j\right) &= w_t^s u \left(h^j + F_{I,i,t}\right) \\ &= w_t^s u \left(\left(x^j\right)^{\frac{1}{\eta}} \theta^{-\frac{1}{\eta}} + F_{I,i,t}\right), \end{aligned}$$

where $h^j \equiv H^j/u$ is the average scientists hired per product line and the cost function is indexed by time because of the wage rate of scientists.

Entrants can also undertake R&D directed to either sector. We assume that to do this they need to hire $F_E \geqslant F_I$ scientists, and this will lead to a flow rate of innovation equal to one. We denote the endogenously determined mass of entrants performing R&D directed to technology j at time t by E_t^j.

On the policy side, incumbents performing R&D for sector j receive a proportional government subsidy at the rate $s_{I,t}^j \in [0,1]$, and entrants performing R&D for sector j receive a subsidy at the rate $s_{E,t}^j \in [0,1]$.

3.2.5 The Carbon Cycle

While clean intermediate production $y_{i,t}^c$ creates no carbon emission, dirty production $y_{i,t}^d$ emits κ units of carbon per intermediate output. This implies that total amount of carbon emission at time t is

$$K_t = \kappa \int_0^1 y_{i,t}^d di. \tag{3.4}$$

[8] This heterogeneity in fixed costs is necessary to make the dynamics (computationally) well behaved. Because of "creative destruction" in these types of models, equilibrium path in which some types of firms stop doing R&D (as clean firms will do without policy and dirty firms under our optimal policy), there will be a discontinuous behavior shortly before this point because creative destruction is expected to cease. Heterogeneity in fixed costs smooths this transition.

We follow Golosov, Mikhail, John Hassler, Per Krusell, and Aleh Tsyvinski (2011) in assuming that the atmospheric carbon concentration S_t is determined as follows

$$S_t = \int_0^{t-T} (1 - d_l) K_{t-l} dl, \tag{3.5}$$

where $t = T$ is the first date when emission started and

$$d_l = (1 - \varphi_P) \left[1 - \varphi_0 e^{-\varphi l}\right]$$

is the amount of carbon emitted l years ago still left in the atmosphere. In addition, $\varphi_P \in (0, 1)$ is the share of emission that remains permanently in the atmosphere, $(1 - \varphi_P) \varphi_0 \in (0, 1)$ is the fraction of the transitory component that remains in the first period, and $\varphi \in (0, 1)$ is the rate of decay of carbon concentration over time. As explained in Golosov, Mikhail, John Hassler, Per Krusell, and Aleh Tsyvinski (2011), this is a flexible specification that approximates the more complex dynamics of carbon concentration in the atmosphere used by Nordhaus, William (2008). Though considerably simpler, this specification fairly closely approximates the observed dynamics of atmospheric carbon concentration as we show below.

3.2.6 Equilibrium

In this section, we characterize certain properties of the equilibrium path of this economy.

The economy at time $t = 0$ is characterized by a distribution of technology gaps between clean and dirty sectors $\mu_{n,0}$ for $n \in \mathbb{Z}$, and the equilibrium path will be defined for a given sequence of taxes and subsidies. Then a dynamic equilibrium path is a sequence of intermediate outputs, prices, innovation rates by incumbents and entrants, skilled and unskilled wages, measures of entrants, growth rate of aggregate output, interest rate, and atmospheric concentration, i.e., $\left[y^j_{i,t}, p^j_{i,t}, x^j_{I,t}, x^j_{E,t}, w^s_t, w^u_t, E^j_t, r_t, S_t\right]^{\infty}_{t=0}$, such that, given sequences of policies, all firms maximize profits, skilled and unskilled labor markets clear, free entry conditions hold (with complementary slackness), consumers optimize dynamically, and atmospheric carbon evolves according to the carbon cycle model presented above (i.e., Equation 3.5). To determine this dynamic time path we also have to keep track of the distribution of sectors by technology gaps, $\{\mu_{n,t}\}^{\infty}_{n=-\infty}$.

3.2.7 Prices and Profits

Given the aggregate production function in Equation 3.2, which implies unit elastic demand for intermediates, the demand for intermediates at time t is

$$y_{i,t} = \frac{\widetilde{Y}_t}{p_{i,t}}, \quad \forall i \in [0,1], \tag{3.6}$$

where $\widetilde{Y}_t \equiv Y_t \exp\left(\gamma\left(S_t - \bar{S}\right)\right)$ is net aggregate output (net of environmental damage).

We now characterize equilibrium prices. As explained in the previous section, if the leading technology for intermediate i at time t is $q^j_{i,t}$, another firm will have access to technology $q^j_{i,t}/\lambda$ for free. This clearly also applies to tax adjusted labor productivity, which is what is relevant for production decisions: when the leading technology is $\widetilde{q}^j_{i,t}$, there is always a follower with technology with $\widetilde{q}^j_{i,t}/\lambda$. Thus, equilibrium markups can never exceed λ. However, in equilibrium, there may not be any markup in some of the intermediates because of the competition between clean and dirty technologies. In particular, intermediate i will be produced using technology $j \in \{c, d\}$ only if there $\widetilde{q}^{-j}_{i,t} \leqslant \widetilde{q}^j_{i,t}$. If $\widetilde{q}^{-j}_{i,t} < \widetilde{q}^j_{i,t}$, the equilibrium markup will be λ, and if $\widetilde{q}^{-j}_{i,t} = \widetilde{q}^j_{i,t}$, there will be zero markup. Therefore:

$$p^j_{i,t} = \begin{cases} \frac{w^u_t}{\widetilde{q}^j_{i,t}} & \text{if} \quad \widetilde{q}^j_{i,t} = \widetilde{q}^{-j}_{i,t} \\ \frac{\lambda w^u_t}{\widetilde{q}^j_{i,t}} & \text{if} \quad \widetilde{q}^j_{i,t} > \widetilde{q}^{-j}_{i,t} \end{cases}. \tag{3.7}$$

Now, combining Equation 3.6 and Equation 3.7, the output of intermediate i as a function of tax-adjusted labor productivities output can be written as

$$y_{i,t}^{j} = \begin{cases} \frac{\widetilde{Y}\widetilde{q}_{i,t}^{j}}{w_t^u} & \text{if} \quad \widetilde{q}_{i,t}^{c} = \widetilde{q}_{i,t}^{d} \\ \frac{\widetilde{Y}\widetilde{q}_{i,t}^{j}}{\lambda w_t^u} & \text{if} \quad \widetilde{q}_{i,t}^{j} > \widetilde{q}_{i,t}^{-j} \end{cases} . \tag{3.8}$$

Then, the equilibrium profits (gross of R&D expenditures), as a function of m and n, can be expressed follows

$$\begin{array}{llll} \pi_{n,t}^{c} = \widetilde{Y}_t \frac{\lambda-1}{\lambda} & \pi_{n,t}^{d} = 0 & \text{if} & n < m \\ \pi_{n,t}^{c} = 0 & \pi_{n,t}^{d} = \widetilde{Y}_t \frac{\lambda-1}{\lambda} & \text{if} & n > m \\ \pi_{n,t}^{c} = 0 & \pi_{n,t}^{d} = 0 & \text{if} & n = m \end{array} . \tag{3.9}$$

3.2.8 Innovation Incentives

We now characterize innovation incentives, which are the only forward-looking part of firm behavior in our model. To simplify the exposition, we first assume that firms are myopic and maximize instantaneous (one-step ahead profits) rather than discounted sum of profits. This enables us to provide analytical expressions for R&D decisions, clarifying the basic economic forces. We will then turn to forward-looking maximization by firms and show that exactly the same expressions and intuitions apply, with the only exception that one term will then be replaced by a solution to a Hamilton-Jacobi-Bellman (HJB) equation rather than being explicitly given as in this subsection.

Not every successful innovation leads to profitable production for two reasons. First, the innovation might be in technology j which is behind technology $-j$, and thus may still not be active even after the improvement in labor productivity. Second, even if it leads to production, this might happen at zero markup if the tax-adjusted labor productivities are the same with the two technologies. Clearly, innovation incentives will be determined by the probability of generating positive profits following innovation. We denote this probability for innovation directed at sector j by $\Gamma_t^j \in [0,1]$. For dirty sector, this is

$$\begin{aligned} \Gamma_t^d &\equiv (1-\alpha)\sum\nolimits_{n\geqslant m}\mu_{n,t} + \alpha\left(\mathbb{I}_{(m\leqslant 0)} + \mathbb{I}_{(m>0)}\sum\nolimits_{n\geqslant m}\mu_{n,t}\right) \\ &= \sum\nolimits_{n\geqslant m}\mu_{n,t} + \alpha\left(1-\sum\nolimits_{n\geqslant m}\mu_{nt}\right)\mathbb{I}_{(m\leqslant 0)}, \end{aligned}$$

where $\mathbb{I}_{(m\leqslant 0)}$ is the indicator function for the event $m \leqslant 0$.

The interpretation of this expression is as follows: If the innovation is incremental (which, conditional on successful innovation, has probability $1-\alpha$), then it will only be profitable if it builds on an intermediate technology where the dirty sector is ahead or neck and neck with the clean sector which, given uniform random draws from the set of all intermediates, has probability $\sum_{n\geqslant m}\mu_{n,t}$. Alternatively, with probability α, the innovating firm will necessarily be at least one step ahead of the competing technology (either the dirty sector is ahead or with the breakthrough technology, it leapfrogs the clean sector). However, in this case, it may leapfrog the clean technology but still not compete with it on the basis of tax-adjusted productivity because of the higher tax on dirty production (i.e., because of a "carbon tax"). In particular, if $m \leqslant 0$ (so that $\mathbb{I}_{(m\leqslant 0)} = 1$), then there is no carbon tax (if anything there might be a carbon subsidy), then it will certainly be at least one step ahead of the clean technology and will be able to charge a markup. If, on the other hand, $m > 0$, then the innovation will be profitable only for intermediates where the technology gap is already sufficiently large for the dirty sector to have higher tax-adjusted technology, which is in the sectors with $n \geqslant m$.

A similar reasoning leads to a probability of positive profit following clean innovation of

$$\begin{aligned}\Gamma_t^c &\equiv (1-\alpha)\sum\nolimits_{n\leqslant m}\mu_{n,t}+\alpha\left(\mathbb{I}_{(m\geqslant 0)}+\mathbb{I}_{(m<0)}\sum\nolimits_{n\leqslant m}\mu_{n,t}\right)\\ &= \sum\nolimits_{n\leqslant m}\mu_{n,t}+\alpha\left(1-\sum\nolimits_{n\leqslant m}\mu_{n,t}\right)\mathbb{I}_{(m\geqslant 0)}. \end{aligned} \tag{3.10}$$

Let us denote the expected value of a successful innovation in technology j by $\bar{v}_t^j$. Since in this subsection we are assuming myopic behavior on the sides of firms, this is equal to the expected immediate (rather than discounted) profits from a successful innovation given by

$$\bar{v}_t^j = \frac{\Gamma_t^j (\lambda-1)\widetilde{Y}_t}{\lambda}. \tag{3.11}$$

Then, dropping the firm subscript i (in $F_{I,i,t}$), the maximization problem of a firm with the leading-edge technology in u^j intermediates in sector $j \in \{c, d\}$ can be written as:

$$\max_{X_{I,t}^j\geqslant 0}\left\{X_{I,t}^j\bar{v}_t^j-\left(1-s_{I,t}^j\right)w_t^s\left[H\left(X_{I,t}^j,u^j\right)+\mathbb{I}_{(X_t^j>0)}u^jF_{I,t}\right]\right\}, \tag{3.12}$$

where $H\left(X_{I,t}^j, u^j\right)$ denotes the number of scientist hired by a firm that has u^j product lines and innovates at the rate $X_{I,t}^j$. In this expression, the indicator function allows us to turn off the fixed costs of R&D when the firm chooses not to perform any R&D activities. Dividing this objective function by u^j, the maximization problem of a firm "per leading-edge technology" (i.e., Equation 3.12 divided by the number of products in which the firm has the leading-edge technology in sector j) is

$$\max_{x_{I,t}^j\geqslant 0}\left\{x_{I,t}^j\bar{v}_t^j-\left(1-s_{I,t}^j\right)w_t^s\left[h\left(x_{I,t}^j\right)+\mathbb{I}_{(x_{I,t}^j>0)}F_{I,t}\right]\right\}. \tag{3.13}$$

where $h\left(x_{I,t}^j\right) \equiv H\left(X_{I,t}^j, u^j\right)/u^j$ is defined as the average number of scientists hired and $x_{I,t}^j \equiv X_{I,t}^j/u^j$ is the average innovation intensity. Using the R&D production function defined in Equation 3.3, equilibrium innovation rate for $j \in \{c, d\}$ can be expressed as

$$x_{I,t}^{j} = \mathbb{I}_{(x_{I,t}^{j} > 0)} \left(\frac{\bar{v}_{t}^{j} \eta \theta^{\frac{1}{\eta}}}{\left(1 - s_{I,t}^{j}\right) w_{t}^{s}} \right)^{\frac{\eta}{1-\eta}}. \tag{3.14}$$

A number of important conclusions follow from Equation 3.14:

1. Higher net output, higher markups and lower scientist wages increase research effort as should be expected.

2. Subsidies to research increase research effort. This will be important in encouraging clean innovation by means of research subsidies.

3. Through the Γ_t^j's, carbon taxes increase clean research effort (and reduce dirty research effort). This can be seen by considering higher values of m in Equation 3.10, which given the distribution of technology gaps, increases Γ_t^c, because production shifts from dirty to clean technologies (and neck-and-neck sectors shift to positive markups for clean technologies). This shows that just carbon taxes may be sufficient to encourage clean innovation and thus a transition to clean technology. Whether they will in fact be sufficient is an empirical and quantitative question we will try to address below.

4. Again through the Γ_t^j's, we can also see the path-dependent nature of innovation. When there are large technology gaps between dirty and clean, $\sum_{n \leqslant m} \mu_{n,t}$ will be very small, and thus Γ_t^c will be small (and Γ_t^d will be high), discouraging clean innovation and encouraging dirty innovation. But if clean innovation can be maintained for a while, then $\sum_{n \leqslant m} \mu_{nt}$ will increase, and so will Γ_t^c (while Γ_t^d declines). Thus clean innovation will naturally self-reinforce over time. To the extent that $\sum_{n \leqslant m} \mu_{n,t}$ changes only slowly, this will also be a slow process.

3.2.9 Free Entry and Labor Market Clearing

The previous subsection characterized the R&D decisions of the incumbents (as a function of the state of the economy and policies). The other component of R&D, creating demand for scientists, is from entrants. With a similar reasoning to the profitability of the R&D of incumbents, the free entry condition for entrants for technology $j \in \{c, d\}$ can be written as

$$\max_{x^j_{E,t} \geqslant 0} \left\{ x^j_{E,t} \bar{v}^j_t - \left(1 - s^j_{E,t}\right) w^s_t \left[h\left(x^j_{E,t}\right) + F_E \right] \right\} \leqslant 0, \tag{3.15}$$

with this condition holding as equality if $E^j_t > 0$. Hence, the innovation rate by entrants is

$$x^j_{E,t} = \mathbb{I}_{\left(x^j_{E,t}>0\right)} \left(\frac{\bar{v}^j_t \eta \theta^{\frac{1}{\eta}}}{\left(1 - s^j_{E,t}\right) w^s_t} \right)^{\frac{\eta}{1-\eta}} \text{ for } j \in \{c, d\}. \tag{3.16}$$

Inspection of Equation 3.15 establishes that at time t, there can be positive entry into technology j only if the "policy-adjusted" value of innovation is higher in sector j than in sector $-j$. In other words, entrants will direct their R&D to the clean technology if $\bar{v}^c_t / \left(1 - s^c_{E,t}\right) > \bar{v}^d_t / \left(1 - s^d_{E,t}\right)$ and to the dirty technology if the reverse inequality holds. We also adopt the tiebreaking rule that if $\bar{v}^c_t / \left(1 - s^c_{E,t}\right) = \bar{v}^d_t / \left(1 - s^d_{E,t}\right)$, then half of the entrants will direct R&D to each sector. Therefore, denoting the total number (measure) of entrants at time t by E_t, we have that the number of entrants with technology directed to sector j is given by

$$E^j_t = \begin{cases} E_t & \text{if } \bar{v}^j_t / \left(1 - s^j_{E,t}\right) > \bar{v}^{-j}_t / \left(1 - s^{-j}_{E,t}\right) \\ 0 & \text{if } \bar{v}^j_t / \left(1 - s^j_{E,t}\right) < \bar{v}^{-j}_t / \left(1 - s^{-j}_{E,t}\right) \\ E_t/2 & \text{if } \bar{v}^j_t / \left(1 - s^j_{E,t}\right) = \bar{v}^{-j}_t / \left(1 - s^{-j}_{E,t}\right) \end{cases}.$$

A comparison of Equation 3.14 and Equation 3.16 shows, conditional on entry an entrant's innovation rate (directed to sector $j \in \{c, d\}$) will only be different from an incumbent's in

Equation 3.14 because of differential subsidies.

It is also useful to inspect the R&D to sales relationship implied by our model. Suppose that free entry condition holds for entry directed at technology $k \in \{c, d\}$. Conditional on investing in R&D, $x_{I,t}^{j} > 0$, the equilibrium R&D to sales ratio (for $j \in \{c, d\}$) would be:

$$\frac{R\&D_{i,t}^{j}}{Sales_{i,t}^{j}} = \eta^{\eta}(1-\eta)^{1-\eta}\frac{\lambda-1}{\lambda}\frac{\Gamma_t^k}{1-s_{E,t}^k}\theta F_E^{-(1-\eta)}\left[\left(\frac{\Gamma_t^j}{\Gamma_t^k}\frac{\left(1-s_{E,t}^k\right)}{\left(1-s_{I,t}^j\right)}\right)^{\frac{1}{1-\eta}}\frac{\eta}{1-\eta}F_E + F_{I,i,t}\right].$$

Note that higher profitability of R&D in the sector for which the free entry condition holds increases the R&D to sales ratio of that sector, but may reduce it in the other sector. The impact of fixed cost requirements of incumbents on R&D to sales ratio result is positive. However, the impact of the fixed cost of entry is ambiguous. On the one hand, it reduces the equilibrium wage, and thus R&D expenditure. On the other, it increases labor requirements, increasing R&D expenditures. The interplay of these two forces makes R&D to sales ratio non-monotonic in the fixed cost for entrance. These different impacts of the fixed cost for incumbents and entrants will enable us to identify both parameters in the estimation.

The labor market clearing condition for scientists can be written as

$$L^s = \sum_{j\in\{c,d\}}\left[\begin{array}{c}\left(\left(\frac{\bar{v}_t^j\theta\eta}{\left(1-s_{E,t}^j\right)w_t^s}\right)^{\frac{1}{1-\eta}} + F_E\right)E_t^j \\ +\int_0^1 \mathbb{I}_{\left(x_{it}^j>0\right)}\left(\left(\frac{\bar{v}_t^j\theta\eta}{\left(1-s_{E,t}^j\right)w_t^s}\right)^{\frac{1}{1-\eta}} + F_{I,i,t}\right)di\end{array}\right]. \tag{3.17}$$

This equation shows that the demand for scientists is decreasing in the skilled wage w_t^s and will be higher when R&D is more profitable and is subsidized more heavily.

We next characterize labor market clearing for production workers. From the equilibrium production decision in Equation 3.8 the unskilled labor demand is

$$l_{i,t} = \begin{cases} \frac{\widetilde{Y}_t}{\left(1+\tau_i^j\right) w_t^u} & \text{if } \widetilde{q}_{i,t}^j = \widetilde{q}_{i,t}^{-j} \\ \frac{\widetilde{Y}_t}{\left(1+\tau_i^j\right) \lambda w_t^u} & \text{if } \widetilde{q}_{i,t}^j \neq \widetilde{q}_{i,t}^{-j} \end{cases}$$

Substituting the optimal quantities from Equation 3.8 into the final good production function in Equation 3.2,

$$w_t^u = \bar{Q}_t \Lambda_t^{\mu}, \tag{3.18}$$

where

$$\bar{Q}_t \equiv \exp\left(\int \ln \widetilde{q}_{it} di\right)$$

is the quality index of active tax-adjusted labor productivities, and

$$\Lambda_t^{\mu} = \lambda^{-(1-\mu_{m,t})}$$

is an inverse function of equilibrium markups (where $\mu_{m,t}$ refers to the fraction of product lines where the lead of dirty is exactly equal to m steps, so that clean and dirty are neck and neck in tax-adjusted productivity). In particular, Λ_t^{μ} takes the value λ^{-1} when all intermediates charge a markup (which is the case when $\mu_{m,t} = 0$) and the value 1 when no intermediates charger markup (which is the case when $\mu_{m,t} = 1$). The labor market clearing for production workers can then be expressed as

$$1 = \frac{\widetilde{Y}_t}{w_t^u} \left\{ \frac{\mu_{m,t}}{2} \left(\frac{1}{1+\tau_t^d} + \frac{1}{1+\tau_t^c} \right) + \frac{1}{\lambda} \left(\frac{\sum_{n<m} \mu_{n,t}}{1+\tau_t^c} + \frac{\sum_{n>m} \mu_{n,t}}{1+\tau_t^d} \right) \right\}.$$

This equation shows both the impact of taxes on labor demand (both types of taxes reduce labor demand and thus wages) and the distribution of technology gaps (because these affect markups). It also shows that if there were only one type of technology, an increase in the tax rate would have no impact on production, just reducing the unskilled wage rate. This is no longer true, however, with two types of technologies, because a tax on dirty technology, for example, would also change the prices of intermediates produced by dirty technology relative to those produced by clean technology, thus impacting production.

This equation also enables us to express aggregate output as a function of the quality index of active tax-adjusted labor productivities as follows

$$Y_t = \exp\left(-\gamma\left(S_t - \bar{S}\right)\right)\tilde{Y}_t = \frac{\bar{Q}_t \Lambda_t^{\mu}}{\Omega_t^{\mu}\exp\left(\gamma\left(S_t - \bar{S}\right)\right)}, \tag{3.19}$$

where

$$\Omega_t^{\mu} \equiv \frac{\mu_{m,t}}{2}\left(\frac{1}{1+\tau_t^d} + \frac{1}{1+\tau_t^c}\right) + \frac{1}{\lambda}\left(\frac{\sum_{n<m}\mu_{n,t}}{1+\tau_t^c} + \frac{\sum_{n>m}\mu_{n,t}}{1+\tau_t^d}\right)$$

is an adjustment for labor demand coming both from taxes and markups.

3.2.10 Dynamics and Equilibrium Redux

Equilibrium dynamics are determined by changes in the interest rate and the evolution of technologies and technology gaps. Household maximization leads to the usual Euler equation

$$g_t = r_t - \rho, \tag{3.20}$$

where g_t is the growth rate of consumption and r_t is the interest rate at time t (and in addition we impose the usual transversality condition).

The evolution of technology gaps $\mu_{n,t}$ can be derived as follows. Let us denote the aggregate innovation rate in technology j as $z_t^j \equiv \left(1 + E_t^j\right) x_t^j$ and the total innovation rate as $z_t \equiv z_t^d + z_t^c$. Then, the flow equations for the distribution of technology gap $n > 1$ can be expressed as

$$\dot{\mu}_{n>1,t} = z_t^d \mu_{n-1,t} + (1-\alpha)\, z_t^c \mu_{n+1,t} - z_t \mu_{n,t}.$$

The change in the share depends on the difference between inflows and outflows. There will be inflows into state n from $n-1$ when a dirty innovation occurs and from $n+1$ when a clean innovation occurs without leapfrogging. On the other hand, an outflow will happen with both clean or dirty innovation as it will bring the state into $n+1$, $n-1$ or -1 depending on the innovation type. We repeat the same reasoning for $n \leqslant 1$ below:

$$\begin{aligned}
\dot{\mu}_{1,t} &= z_t^d \mu_{0,t} + (1-\alpha) z_t^c \mu_{2,t} + \alpha z^d \mu_{-t}^c - z\mu_{1,t} \\
\dot{\mu}_{0,t} &= (1-\alpha) z^d \mu_{-1,t} + (1-\alpha) z_t^c \mu_{1,t} - z\mu_{0,t} \\
\dot{\mu}_{-1,t} &= z^c \mu_{0,t} + (1-\alpha) z_t^d \mu_{-2,t} + \alpha z_t^c \sum_{n>0} \mu_{n,t} - z\mu_{-1,t} \\
\dot{\mu}_{n<-1t} &= z^c \mu_{n+1,t} + (1-\alpha) z_t^d \mu_{n-1,t} - z\mu_{n,t}.
\end{aligned} \tag{3.21}$$

Total dirty intermediate production at time t, Y_t^d, which creates pollution is given as

$$\begin{aligned}
Y_t^d &= \int y_{i,t}^d di = \int_{i\in\mu_m} \frac{y_{i,t}^d}{2} di + \sum_{n>m} \int_{i\in\mu_n} y_{i,t}^d di \\
&= \frac{\widetilde{Y}_t}{\left(1+\tau_t^d\right) w_t^u} \left[\frac{1}{2} Q_{m,t}^d + \frac{1}{\lambda} \sum_{n>m} Q_{n,t}^d \right],
\end{aligned} \tag{3.22}$$

where we break up the productivity aggregates by step size differential n defining (with a slight abuse of notation where $i \in \mu_n$ denotes intermediates where the technology gap is n steps):

$$Q_{n,t}^d \equiv \int_{i\in\mu_n} q_{i,t}^d di.$$

We now summarize the dynamic equilibrium path using the equations we have derived in this section. For any given time path of policies $\left[\tau_t^j, s_{I,t}^j, s_{E,t}^j\right]_{t=0}^{\infty}$, a dynamic equilibrium path is characterized by time path of

$$\left[y_{i,t}^j, p_{i,t}^j, x_{I,t}^j, x_{E,t}^j, w_t^s, w_t^u, E_t^j, \{\mu_{n,t}\}_{n=-\infty}^{\infty}, \{Q_{n,t}^d\}_{n=-\infty}^{\infty}, g_t, r_t, S_t\right]_{t=0}^{\infty}$$

such that $[i]$ $y^j_{i,t}$ and $p^j_{i,t}$ maximize profits as in Equation 3.7 and Equation 3.8; $[ii]$ $x^j_{I,t}$ and $x^j_{E,t}$ solve incumbent's and entrant's R&D decision as in Equation 3.14 and Equation 3.16; $[iii]$ w^u_t clears unskilled labor market as in Equation 3.18; $[iv]$ w^s is determined from the free entry condition Equation 3.15 when there is positive entry and from skilled labor market clearing Equation 3.17 when there is no positive entry; $[v]$ E^j_t is determined from the skilled labor market clearing Equation 3.17 when there is positive entry; $[vi]$ technology gap shares $\{\mu_{n,t}\}^{\infty}_{n=-\infty}$ satisfy the set of flow equations 3.21; $[vii]$ total productivity of the sectors with n-step gap $Q^d_{n,t}$ evolves according to the innovation rates in Equation 3.14 and Equation 3.16, $[viii]$ the growth rate is consistent with the innovation rates $x^j_{I,t}$ and $x^j_{E,t}$; and $[ix]$ the interest rate satisfies the Euler equation Equation 3.20, and $[x]$ S_t is given by Equation 3.5.

3.2.11 Full Model

We now relax the assumption of myopic firms and assume that firms maximize their discounted profits (and this full model will be used in our quantitative analysis also).

Let $\vec{n}^j \equiv \left[n^j_1, ..., n^j_u\right]$ denote the vector of product lines where the firm in question holds the leading-edge technology (a total of $u = u^j_t$ of them for this firm) and n^j_i the technology gap compared to technology $-j$ within the same product line. Let $\vec{n}^j_{-i}$ denote the same vector $\vec{n}^j$ without its ith element n^j_i. Then the value of a firm with a portfolio of products given by $\vec{n}^j$ then satisfies the HJB equation:

$$
\begin{aligned}
& rV^j_{\vec{n}^j,t} - \dot{V}^j_{\vec{n}^j,t} \qquad (3.23) \\
= \;& \sum_{i=1}^{u} \left[\begin{array}{l} \pi^j_{n_i,t} + z^j_t \left(V^j_{\vec{n}^j_{-i}} - V^j_{\vec{n}^j,t} \right) \\ + z^{-j}_t (1-\alpha) \left(V^j_{\vec{n}^j_{-i} \cup \{n^j_i - 1\},t} - V^j_{\vec{n}^j_i,t} \right) + z^{-j}_t \alpha \left(V^j_{\vec{n}^j_{-i},t} - V^j_{\vec{n}^j_i,t} \right) \end{array} \right] \\
& + \int \max_{x^j_t \geqslant 0} \left[u^j_t x^j_t \left(V^j_{\vec{n}^j \cup \{n^j_{u+1}\},t} - V^j_{\vec{n}^j,t} \right) - \left(1 - s^j_{I,t}\right) u^j_t w^s_t \left(\left(x^j_t\right)^{\frac{1}{\eta}} \theta^{-\frac{1}{\eta}} + \mathbb{I}_{(x^d_{n,t}>0)} F_{I,t} \right) \right] dF_{I,t}.
\end{aligned}
$$

The interpretation is straightforward. The right-hand side includes the profits generated from u product lines, which is given by the first term. In addition, at the flow rate rate z_t^j, each product line i will experience an innovation by another firm from the same technology j in which case i is taken out of firm's portfolio (so that the firm's portfolio becomes $\vec{n}_{-i}^j$). If instead production line i experiences an innovation from the alternative technology $-j$, which happens at the rate z_t^{-j}, then there are two possibilities: either the innovation is incremental (probability $(1-\alpha)$) and the current incumbent will still continue with its production in which case the technology gap will be narrowed by one step (so that $n_i^j = n_i^j - 1$) or the innovation might be drastic (probability α), in which case the firm will lose this product line (again reducing its portfolio to $\vec{n}_{-i}^j$). Finally, the firm invests in R&D itself and innovates at the flow rate $X_t^j = u_t^j x_t^j$, and the option value of this R&D (inclusive of costs) is added as the second line of the right-hand side, with the integral taking care of the fact that fixed costs are stochastic. (Note in particular that when it is successful, the firm adds a new product line so that its portfolio becomes $\vec{n}^j \cup \{n_{u+1}^j\}$).

The next lemma provides a convenient re-expression of this Bellman equation in per product terms:

Lemma 7. *Equation 3.23 can be re-expressed as* $V_{\vec{n}^j,t}^j = \widetilde{Y}_t \sum_{i=1}^u v_{n_i,t}^j$ *where*

$$\begin{aligned}
&\rho v_{n_i,t}^j - \dot{v}_{n_i,t}^j = \pi_{n_i}^j - z_t^j v_{n_i,t}^j + z_t^{-j}(1-\alpha)\left(v_{n_i-1,t}^j - v_{n_i,t}^j\right) + z_t^{-j}\alpha\left(v_{-1,t}^j - v_{n_i,t}^j\right) \\
&+ \int \max_{x_t^j \geqslant 0}\left[x_t^j \bar{v}_t^j - \left(1 - s_{I,t}^j\right)\widetilde{w}_t^s\left(\left(x_t^j\right)^{\frac{1}{\eta}}\theta^{-\frac{1}{\eta}} + \mathbb{I}_{(x_{n,t}^d > 0)}F_{I,t}\right)\right]dF_{I,t}.
\end{aligned} \tag{3.24}$$

and $\bar{v}_t^j \equiv \mathbb{E}_i v_{n_i,t}^j$.

Proof of Theorem 7. Substituting Equation 3.24 into Equation 3.23, we obtain

$$\begin{aligned}
& r\widetilde{Y}_t \sum_{i=1}^u v_{n_i,t}^j - \frac{d}{dt}\widetilde{Y}_t \sum_{i=1}^u v_{n_i,t}^j - \widetilde{Y}_t \sum_{i=1}^u \dot{v}_{n_i,t}^j \\
= \;& \sum_{i=1}^u \left[\widetilde{Y}_t \pi_{n_i}^j - z_t^j \widetilde{Y}_t v_{n_i,t}^j + z_t^{-j}(1-\alpha)\left(\widetilde{Y}_t v_{n_i-1,t}^j - \widetilde{Y}_t v_{n_i,t}^j\right) + z_t^{-j}\alpha\left(\widetilde{Y}_t v_{-1,t}^j - \widetilde{Y}_t v_{n_i,t}^j\right)\right] \\
& + \int \max_{x_t^j \geqslant 0}\left[u x_t^j \widetilde{Y}_t \bar{v}_t^j - \left(1 - s_{I,t}^j\right) u w_t^s\left(\left(x_t^j\right)^{\frac{1}{\eta}}\theta^{-\frac{1}{\eta}} + \mathbb{I}_{(x_{n_i,t}^d > 0)}F_{I,t}\right)\right]dF_{I,t}.
\end{aligned}$$

where

$$\widetilde{w}_t^s \equiv \frac{w_t^s}{\widetilde{Y}_t}$$

and

$$\begin{array}{llll} \pi_n^c = \frac{\lambda-1}{\lambda} & \pi_n^d = 0 & \text{if} & n < m \\ \pi_n^c = 0 & \pi_n^d = \frac{\lambda-1}{\lambda} & \text{if} & n > m \\ \pi_n^c = 0 & \pi_n^d = 0 & \text{if} & n = m \end{array} .$$

Then this can further be simplified to

$$\begin{aligned} & r \sum_{i=1}^{u} v_{n,t}^j - g_t \sum_{i=1}^{u} v_{n,t}^j - \sum_{i=1}^{u} \dot{v}_{n,t}^j \\ = \; & \sum_{i=1}^{u} \left[\pi_{n_i}^j - z_t^j v_{n,t}^j + z_t^{-j} (1-\alpha) \left(v_{n-1,t}^j - v_{n,t}^j \right) + z_t^{-j} \alpha \left(v_{-1,t}^j - v_{n,t}^j \right) \right] \\ & + u \int \max_{x_t^j \geqslant 0} \left[x_t^j \bar{v}_t^j - \left(1 - s_{I,t}^j\right) \widetilde{w}_t^s \left(\left(x_t^j\right)^{\frac{1}{\eta}} \theta^{-\frac{1}{\eta}} + \mathbb{I}_{(x_{n,t}^d > 0)} F_{I,t} \right) \right] dF_{I,t}. \end{aligned}$$

where we used the fact that $\frac{d}{dt}\widetilde{Y}_t = g_t \widetilde{Y}_t$ in the first line. Next, we eliminate $\widetilde{Y}_t$ throughout and use Equation 3.20, $r_t - g_t = \rho$, to establish the desired result. □

An important implication is that incumbent innovation rates per product line will be independent of the portfolio of the incumbent and given by

$$x_t^j = \mathbb{I}_{(x_{I,t}^j > 0)} \left(\frac{\bar{v}_t^j \eta \theta^{\frac{1}{\eta}}}{\left(1 - s_{I,t}^j\right) \widetilde{w}_t^s} \right)^{\frac{\eta}{1-\eta}} \text{ for } j \in \{c, d\},$$

which can be easily verified to be identical to Equation 3.14 except that now $\bar{v}_t^j$ is given as a solution to Equation 3.24.

We can now describe the full dynamic equilibrium path of this economy, which will be essentially identical to the equilibrium path with myopic firms, with $\bar{v}_t^j$ given as the solution to the HJB equation 3.24.

For any given time path of policies $\left[\tau_t^j, s_{I,t}^j, s_{E,t}^j\right]_{t=0}^{\infty}$, a dynamic equilibrium path is characterized by time path of

$$\left[\bar{v}_t^j, y_{i,t}^j, p_{i,t}^j, x_{I,t}^j, x_{E,t}^j, w_t^s, w_t^u, E_t^j, \{\mu_{n,t}\}_{n=-\infty}^{\infty}, \{Q_{n,t}^d\}_{n=-\infty}^{\infty}, g_t, r_t, S_t\right]_{t=0}^{\infty}$$

such that $\bar{v}_t^j \equiv \mathbb{E}_i v_{n_i,t}^j$, and each $v_{n,t}^j$ satisfies Equation 3.24. In addition: $[i]$ $y_{i,t}^j$ and $p_{i,t}^j$ maximize profits as in Equation 3.7 and Equation 3.8; $[ii]$ $x_{I,t}^j$ and $x_{E,t}^j$ solve incumbent's and entrant's R&D decision as in Equation 3.14 and Equation 3.16; $[iii]$ w_t^u clears unskilled labor market as in Equation 3.18; $[iv]$ w^s is determined from the free entry condition in Equation 3.15 when there is positive entry and from skilled labor market clearing Equation 3.17 when there is no positive entry; $[v]$ E_t^j is determined from the skilled labor market clearing Equation 3.17 when there is positive entry; $[vi]$ technology gap shares $\{\mu_{n,t}\}_{n=-\infty}^{\infty}$ satisfy the set of flow equations in 3.21; $[vii]$ total productivity of the sectors with n-step gap $Q_{n,t}^d$ evolves according to the innovation rates in Equation 3.14 and Equation 3.16, $[viii]$ the growth rate is consistent with the innovation rates $x_{I,t}^j$ and $x_{E,t}^j$; and $[ix]$ the interest rate satisfies Equation 3.20, and $[x]$ S_t is given by Equation 3.5.

3.3 Empirical Strategy and Data

Our model has 14 parameters/variables to be determined:

$$\left\{\rho, \bar{S}, \gamma, \varphi, \varphi_0, \varphi_P, \kappa, L^s, \alpha, \eta, \theta, \lambda, F_I, F_E\right\}.$$

In addition, the initial distribution of technology gaps between clean and dirty technologies, $\{\mu_{0t}\}_{n=-\infty}^{\infty}$, needs to be determined. It is useful to note that, as will become clearer below, given $\{\mu_{nt}\}_{n=-\infty}^{\infty}$, estimation of the remaining parameters can be done without knowledge of taxes and subsidies, and also without any information on $\bar{S}, \gamma, \varphi, \varphi_0$, and φ_P. These become relevant only for our policy analysis. Nevertheless, here we specify our choices for all these parameters.

We proceed in four steps. First, we externally calibrate some of the parameters, in particular the parameters of the carbon cycle and the discount rate. In all, the parameters $\rho, \bar{S}, \gamma, \varphi, \varphi_0, \varphi_P$, and κ are taken from external sources. Second, we directly estimate $L^s, \alpha,$ and η from microdata. Third, we choose the initial distribution of technology gaps to match the distribution of patents between firms innovating mostly with clean and mostly with dirty technologies as we explain below. Finally, we estimate the remaining parameters θ, λ, F_I and F_E using simulated method of moments, with moments being selected to model the firm-level R&D behavior, growth rates, and entry/exit rates for the energy sector as we describe below. The model performs well and is able to replicate these moments reasonably closely.

Throughout our focus is on the energy sector, the behavior of which has motivated our theoretical model. The energy sector is defined as firms involved in the sourcing, refinement and delivery of energy inputs for residential and industrial applications (e.g., oil and gas, electricity), firms that provide complementary inputs and equipment into this energy-production process (e.g., drilling equipment, power plant technologies), and firms that interface with the energy inputs for residential and industrial use (e.g., motor manufacturers). As such, our group of 1576 firms that make up our sample includes oil and gas producers, mining and exploration firms, engine manufacturers, power companies building upon multiple techniques, energy equipment manufacturers, and similar.[9]

The data we use for estimation comes from the Census Bureau's Longitudinal Business Database and Economic Censuses, the National Science Foundation's Survey of Industrial Research and Development, and the NBER Patent Database. We design our sample around innovative firms in the energy sector that are in operation during the 1975-2004 period.

[9]We exclude approximately 50 non-profit research centers and similar entities to match our model's focus on profit-seeking firms. Our estimations below are robust to retaining these entities.

3.3.1 External Calibration

We choose $\bar{S}, \gamma, \varphi, \varphi_0, \varphi_P$, and κ to link our model to the carbon cycle and its impact on aggregate output following Golosov, Mikhail, John Hassler, Per Krusell, and Aleh Tsyvinski (2011) approach. This approach takes into account the current level of carbon stock and its increase since pre-industrial times; the rate at which new emissions enter the atmosphere, the terrestrial biosphere or shallow oceans, and the deep oceans; how that movement and the various reservoirs of carbon influence the earth's temperature; and how higher temperatures and environmental damage hurt the economy. This work builds upon prior work in environmental economics (e.g., Nordhaus, William (2008), Nordhaus, William, and Joseph Boyer (2000), but is more flexible in allowing non-linear absorption of atmospheric carbon, but does not allow any delay on the impact of this carbon content on economic outcomes and temperature changes (which result from different rates at which oceans change temperature, for example) and does not separately keep track of the dynamics of the atmospheric concentration of carbon dioxide (CO_2), methane (CH_4) and nitrous oxide (N_2O).

Our value of the pre-industrial stock of carbon dioxide in the atmosphere $\bar{S}$ is 581 GtC (gigatons of carbon). To model how emission increases the atmospheric stock of CO2, we define the three parameters φ, φ_0, and φ_P as follows. First, φ_P is the portion of new emissions that will remain in the atmosphere for a very long time, likely for thousands of years, and estimates of this parameter from Solomon, S. D. Qin, M. Manning, Z. Chen, M. Marquis, K.B. Averyt, M.Tignor and H.L. Miller (2007) and Archer, David (2005) are about 20%. The other two parameters, φ and φ_0, govern the short- and medium-term movement of the emitted carbon that will not become part of this very long duration stock in the atmosphere. These emissions influence the earth's temperature over short horizons, but they are ultimately absorbed into the deep oceans. To identify these parameters, we utilize the 30 year half-life of carbon and match the carbon stock evolution under emissions during the 1900-2008 period. We use the following formula to determine the atmospheric carbon concentration S_t in every year during 1900-2008 period

$$S_t = \int_0^{t-1900} (1 - d_l)\, K_{t-l} dl + S_{1900}, \tag{3.25}$$

where

$$d_l = (1 - \varphi_P) \left[1 - \varphi_0 e^{-\varphi l}\right]$$

The emission data for $\{K_t\}_{t=1900}^{2008}$ is shown in Figure 1 below.

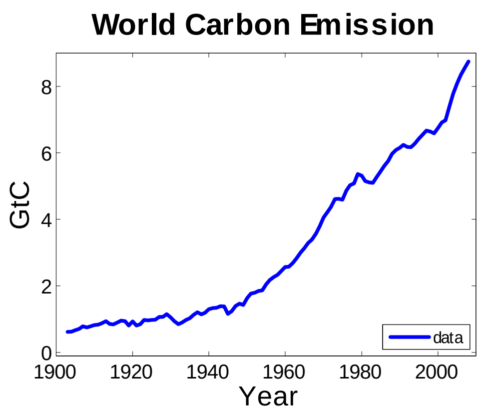

FIGURE 1

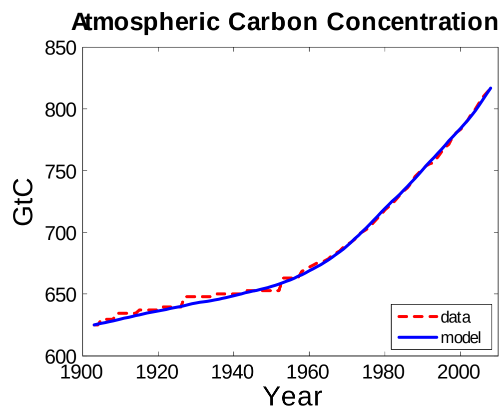

FIGURE 2

Figure 2 shows the fit of Equation 3.25 which yields $\varphi = 0.0313$ and $\varphi_0 = 0.7661$. The dynamics implied by Equation 3.5 at these parameter values match the actual evolution of atmospheric carbon over the past century very well as shown by the close correspondence between the solid blue line representing the data and the dashed red line corresponding to the model-implied atmospheric concentration in Figure 2.

Our damage function also follows Golosov, Mikhail, John Hassler, Per Krusell, and Aleh Tsyvinski (2011) and we choose the γ parameter in our baseline policy analysis to be the same as theirs, 5.3 x 10^{-5} GtC^{-1},[10] though this number may be too low, particularly

[10]This approach provides a fairly good approximation of the damage function developed in Nordhaus, William (2007), who incorporates a typical estimate that a doubling of the stock of atmospheric carbon leads

because, in contrast to their paper, we are not allowing policy to adjust to new information about damages as this becomes available, so the certainty equivalent average of estimates rather than arithmetic average might be more appropriate. Section Section 3.6.1 provides robustness checks with higher values of γ.

The κ parameter is chosen to link current emissions levels to the baseline output level of the model. In doing so, we are making a simplifying assumption that the emission of our economy can proxy for the emission of the entire US economy. As suggested in Golosov, Mikhail, John Hassler, Per Krusell, and Aleh Tsyvinski (2011), the modeling of the carbon cycle and its impact on production has the attractive feature that the social value of marginal emissions is the same (relative to output). This implies that our results would be essentially unchanged if we take a future path of emissions from the rest of the world, with the only difference being that the implied temperature changes we depict below would no longer apply (and we would need to talk about incremental temperature changes due to the US energy sector). It is also worth noting that our model and this modeling strategy certainly abstract from several important aspects of international cooperation or competition that impact climate change outcomes (e.g., Hassler, John, and Per Krusell (2012).

Finally, we report all of our results for a single private discount rate $\rho = 1\%$ and two values of the social discount rate 1% and 0.1%. The first is $\rho = 1\%$, which is close to the 1.5% chosen by Nordhaus in his models, and the second is $\rho = 0.1\%$ used by Stern, Nicholas (2007), on the basis that with a higher discount rate there is almost no weight put on the welfare of future generations.

---

to a 3°C increase and then a proportional damage function of how global temperature increases affect the economy. Golosov, Mikhail, John Hassler, Per Krusell, and Aleh Tsyvinski (2011) show a close correspondance of these functions over relevant temperature ranges.

3.3.2 Sample Construction and Data Sources

We combine several datasets for this study. The NBER Patent Database and the NSF Survey of Industrial Research and Development are the backbones for our study, with additional information and details being collected from the Longitudinal Business Database and the various Economic Censuses conducted by the Census Bureau. We introduce each dataset as we describe the steps in our sample construction.

Patent Data for Energy Sector

Our first data source is the individual records of all patents granted by the United States Patent and Trademark Office (USPTO) from January 1975 to May 2009. This database was first developed by the NBER and was subsequently extended by HBS Research to include patenting in recent years. Each patent record provides information about the invention and the inventors submitting the application. Hall, Bronwyn, Adam Jaffe, and Manuel Trajtenberg (2001) provide extensive details about these data, and Griliches (1990) surveys the use of patents as economic indicators of technology advancement. We collect from this database the patents that are 1) filed by inventors living in the United States at the time of the patent application, and 2) assigned to industrial firms. In a representative year, 1997, this group comprised about 77 thousand patents (about 40% of the total USPTO patent count in 1997).

We then identify patents that are related to the energy sector. This is a key step for our study, and we outline our approach in detail. We use patent technology codes to identify patents related to the energy sector. The technology codes are the most disaggregated level of the USPTO's classification scheme and number over 150,000. This is important as energy-sector patents are spread out over multiple patent classes (the next higher level of the classification system with about 450 groups). Two examples related to solar energy are

“Power Plants/utilizing natural heat/Solar” and “Stoves and Furnaces/Solar heat collector”. Moreover, we describe later how these patenting technologies are also used to classify firms as being primarily clean- or dirty-energy firms. This separation can only be done at the technology level as the patent class level includes both types (e.g., “Power Plants” includes technologies for coal-powered plants too).

We identify relevant technology codes through four steps. First, we adopt the prior classifications developed in a study of alternative energy by Popp, David (2002) and Popp, David, and Richard Newell (2012). The work of Popp, David, and Richard Newell (2012) is particularly helpful in that they provide classifications into various types of energy technologies that we discuss in greater detail below. Given this authoritative prior work, we also report results below for our key parameters that just use their classification system.

We are interested, however, in several technologies (e.g., nuclear power) not considered by Popp, David, and Richard Newell (2012). We thus extend their list through three additional steps. Our second step utilizes resources from the OECD’s work to identify environment-related technologies. OECD Environment Directorate (2011) lists such technologies using the International Patent Classification (IPC) scheme, which some observers believe is better designed to identify and group environment-related technologies than the USPTO classification framework. We use concordances between the IPC and USPTO framework to identify additional technologies to be included.

The third step utilizes information on energy-related R&D data from the NSF Survey of Industrial Research and Development that we describe in greater detail next. For the firms identified in this survey to be conducting energy-related R&D, we list their patent technology codes and frequency. We then manually search the 600 most-frequent codes to identify if they are energy related. In a related fourth step, we also manually search the USPTO database using key words like “Coal” and “Solar” to determine relevant technologies. This identification process constructs a pool of patents related to the energy sector.

As a representative year, 1997, our energy-related patents comprised 7.6% of the total US patent count.[11]

Operating Data for Energy Sector

Our next step links our energy patent data to firm-level operating data collected by the US Census Bureau. The Longitudinal Business Database is a business registry that contains annual employment levels for every private-sector establishment in the United States with payroll from 1976 onward. We also employ Economic Censuses that are conducted every five years by sector of the economy; these censuses provide additional plant and firm operation data (e.g., sales). Sourced from US tax records and Census Bureau surveys, these micro-records document the universe of establishments and firms, making them an unparalleled laboratory for studying our model of firm dynamics in the US energy sector.

We match the patent data to these operating data using firm-name and geographic-location matching algorithms. The basic concept in these algorithms is to identify Census Bureau firms that have similar names to USPTO patent assignees and that have establishments in the same geographic area as where inventors of the patents are located.[12] This linkage also accomplishes a related step of aggregating patent assignees to firms, as some firms file patents through multiple patent assignee codes. This aggregation is due to the Census Bureau's establishment-firm hierarchy, as we observe establishment-level names within multi-unit firms that help identify subsidiaries and major corporate restructurings like mergers and acquisitions, and through the name-matching process that consolidates

[11] Energy-related patents account for 5%-15% of US patents over our sample period, with a declining share in recent years; in absolute terms, patent counts for the energy sector are stable or growing throughout the period. The declining share is partly due to the sector not growing as fast as others, and partly due to external changes over time to allow for patents to be made in sectors that traditionally did not patent, especially software patents.

[12] The algorithms are described in detail in an internal Census Bureau report by Ghosh, Kaushik, and William Kerr (2010). This patent matching builds upon the prior work of Balasubramanian, Natarajan, and Jagadeesh Sivadasan (2011) and Kerr, William, and Shihe Fu (2008).

slight name variants over patent assignees.

We focus our sample on the years in which Economic Censuses are conducted, specifically every five years starting in 1977 and ending in 2002. We adopt this focus for several reasons: 1) the operating data are often best measured around these years due to heightened Census Bureau resources, 2) some specific variables from the Economic Censuses are only available at those five-year marks, and 3) our innovation data are most appropriately considered over short time periods. The third rationale is due, for example, to lumpiness in firm applications for patents; as we describe next, our R&D expenditures data are also often collected biannually. We thus measure variables using the average of observed values for firms in five-year windows surrounding these Economic Census years (e.g., 1985-1989 for the 1987-centered period). We have six time periods covering the 30 year interval of 1975-2004.

R&D Data for Energy Sector

We next utilize the Census Bureau's internal linkages to collect information on R&D expenditures from the NSF's Survey of Industrial Research and Development (R&D Survey). The R&D Survey is the US government's primary instrument for surveying the R&D expenditures and innovative efforts of US firms. This is an annual or biannual survey conducted jointly by the Census Bureau and NSF. The survey includes with certainty all public and private firms, as well as foreign-owned firms, undertaking over a minimum threshold of R&D expenditure in the United States. For most of our sample period, this expenditure threshold is one million dollars of R&D within the US. The survey frame also sub-samples firms conducting less than the certain expenditure threshold. These micro-records begin in 1972 and provide the most detailed statistics available on firm-level R&D efforts. In 1997, 3,741 firms reported positive R&D expenditures that sum to $158 billion. Foster, Lucia, and Cheryl Grim (2010) and Akcigit and Kerr (2010) discuss these data in greater detail.

The R&D Survey provides us with information on many firms' R&D expenditures and employments of science and engineering workers. We use the data, along with the patenting of the firm, to calculate the innovation production function for the sector (e.g., the η and α parameters). We describe these calculations below, and for these calculations we only utilize firm observations for which we always observe reported data on R&D expenditures or scientist counts—meaning that these calculations use only firms that conduct more than the minimum threshold of one million dollars in R&D or are sub-sampled completely. While this might present an issue for sectors like consumer internet start-ups, this is not very restrictive for the supply side of the energy sector given the large amounts of R&D expenditures required by many start-ups.

For our broader moments on firm dynamics, this minimum threshold creates a challenge, however, for the consistent calculation of the entry margin and growth rates. Our model requires that firms be innovative from the start of their lives, and thus an innovative firm that falls below threshold value in its first period would be inappropriately dropped if we restricted the sample only to firms for which we always observe R&D expenditure. By contrast, the patent data are universally observed. To ensure a complete distribution, we thus use patents to impute R&D values for firms that are less than the certainty threshold and not sub-sampled. Overall, our moments combine the R&D and patent data into a single measure of innovation (in R&D terms) that accords with the model for the characterization of firm dynamics in the energy sector. As the R&D expenditures in these sub-sampled cases are low (by definition), this imputation choice versus treating unsurveyed R&D expenditures as zero expenditures conditional on patenting is not very important. The firm does not need to conduct R&D or patent in every year of the initial five-year window, but the firm must do one of the two activities at least once.[13]

[13] In a small number of cases where we have scientists counts from the R&D Survey but lack R&D expenditures, we similarly use the scientist counts to impute R&D values for firms below the certainty threshold.

To summarize, the key idea is that our sample requires that a firm either patent or have measured R&D in the first period of its life. If the firm is an incumbent in the initial 1975-1979 period, it must have either a patent or measured R&D. Our sample does not condition on innovative activity before 1975-1979. Thus, these incumbents may have had some point in their past when they did not conduct R&D or patent. Our model allows for firms to transition out of R&D, and thus we include firms that quit being innovative. On the other hand, we do not consider non-innovative incumbent firms starting to do innovation. As the probability that an existing, non-innovative firm commences R&D or patenting over the ensuing five years (conditional on survival) is only about 1%, this exclusion is reasonable.[14] As one exception to this sample construction, we only estimate the key innovation production function over firms that have continually observed R&D expenditures (so that imputation procedures are not required).

Sample Inclusion Rules and Sample Size

These procedures define the base pool of innovative firms in the energy sector. To be retained in our final sample, the firm must meet two additional requirements. The first is that the firm has positive employment and obtains three or more patents in the energy sector during the 1975-2004 period. These are not very high hurdles given our purpose, and we thus exclude entities that only obtain one or two energy-related patents over their lifetime. Second, and more important, we also require that 10% of the firm's total patenting be in energy-related fields. This is an important hurdle as it excludes innovative firms that are not very active in the sector. The 10% bar is more substantial than it may initially appear as we have been fairly conservative in terms of defining energy-sector patents.

[14]Note that it would have been impossible to build a consistent sample that would also include incumbents switching into innovation. To see why, consider keeping all of the past records for incumbent firms that start conducting R&D in 1987. In the prior periods, this approach would induce a mismeasurement of exit propensities and growth dynamics because a portion of the sample will include firms conditioned on survival until 1987.

Thus, our compiled dataset includes innovative firms in the energy sector from 1975-2004. A record in our dataset is a firm-period observation that aggregates over the firm's different establishments. We have 6228 observations from 1576 firms. While focused on a single sector, our firm panel contains 19% of all US R&D industrial expenditures during the 1975-2004 period. The panel accounts for about 70% of industrial patents for the energy sector in the United States. Across all activity in the economy, our sample typically account for 1% of establishments, 5% of employment, and 10% of sales. In the 1997 period, our sample accounts for over a trillion dollars in sales, 3.9 million employees, and 25 billion dollars in R&D expenditures, and the firms obtain 56 thousand patents during 1995-1999.

Our sample is very important for studying emissions in two ways. First, we account for a substantial amount of activity in several of the main sectors responsible for emissions (e.g., Mueller, Nicholas, Robert Mendelsohn, and William Nordhaus (2011)). In the 1992-1997 period, for example, we account for 59% of sales in industries related to coal and oil extraction, refinement, and shipment; 33% of sales in industries related to electricity production; and 21% of manufacturing sales. Among manufacturing industries, our sample contains higher shares in industries more closely linked with emissions (e.g., 64% in petroleum refinement, 31% in primary metals). Second, while our sample does not include many firms directly from two high-emission sectors, agriculture and transportation, our sample does include many of the manufacturers of products that are key inputs to these sectors.

Designation of Firms as Clean or Dirty Energy

Beyond the development of the firm panel, our leap-frogging calculations below require us to identify whether patents are related to "dirty" or "clean" energy. We do so through the field of patent technologies. We identify patents as related to dirty technologies if they are connected to the extraction, refinement or use of fossil-fuel based energy, including oil, coal, natural gas, and shale technologies. We group into clean-energy patents fields that are related to geothermal, hydroelectric, nuclear, solar, and wind energy. We also include in the clean-energy group identified patents for conservation and utilization of energy. The results below are robust to reclassifications of border group types.

For our model's initial conditions, we also need to identify whether firms are primarily operating in dirty- or clean-energy applications. We do so through a simple rule that has two steps. We first classify a firm-period observation as focused on clean energy if 25% or more of its energy-sector patents are devoted to clean-energy fields; otherwise the firm is classified as a dirty-energy firm in the period. We use the 25% threshold as our assignments of clean-energy fields are conservative compared to dirty-energy fields. We then describe the firm overall as a clean-energy firm if half or more of its time periods achieve this clean-energy focus. The distribution between clean and dirty uses at the firm level is fairly bimodal—96% of observations have 75% or more of their patents in one technology—making the exact details of these procedures less important. In total, 11% of our firms are classified in the clean-energy sub-sector; 14% of energy-sector patents are classified as clean energy.

Several points are worth noting at this stage. First, we generally include technologies that are designed to make fossil fuels cleaner in the dirty-energy group. In this one regard, we deviate from the classifications developed by Popp, David, and Richard Newell (2012) where coal liquefaction and gasification are included in alternative energy, for example. When we directly use the technology scheme of Popp, David, and Richard Newell (2012) as a robustness check, we classify the technologies as in their original work. Second, we have not built our sample selection or grouping procedures around technologies related to pollution abatement. We retain all patents for included firms, and thus they are part of our overall technology description, but we only use energy-directed patents to classify patents and firms into dirty- or clean-energy groups. Finally, we also use the more detailed information the R&D Survey collects from selected major R&D producers. We specifically utilize information collected from about 100 firms on R&D expenditures related to specific energy applications like coal or solar energy. We earlier identified one application of this extra information in that we manually searched the major patenting technology codes from these R&D entities to identify energy-sector patenting groups that we should be including. A second application is to verify our data development procedures, for example by assigning firms based upon the types of R&D they conduct rather than observed patents. This group from the R&D Survey also confirms the bimodal nature of our firm groupings. While the group of firms asked these questions is too small and selected to serve as the backbone of our sample, these checks are comforting. While Census Bureau disclosure prevents us from listing firms, we overlap well with producers listed in Popp, David, and Richard Newell (2012) as one example.

3.3.3 Estimation and Choice of Parameters from Microdata

We first estimate the η parameter from our innovation production function, $X_f = \theta(H_f)^\eta (u_f)^{1-\eta}$, which can be rewritten as $\ln(X/u_f) = \ln(\theta) + \eta \cdot \ln(H/u_f)$. We measure X by the firm's count of patents, H by the firm's R&D expenditures or scientist counts, and u by the number of four-digit SIC industries in which a firm is operating. We also check the robustness of our results to using three-digit SIC industry counts, sales and establishment counts as our proxies for firm size u. Our patent count measure is weighted by citations, with citation counts normalized by the average citations achieved by other patents in the same patent class and application year.

To estimate the η elasticity as accurately as possible, we use the panel nature of our data and later return to estimating the θ parameter. As noted earlier, we only use for this exercise firms that have a full panel of reported R&D data. To focus on higher quality data for our differenced estimations, we also require that the firm be present in at least three periods. We first estimate a linear regression with year fixed effects δ_t, yielding

$$\ln(Patents/product_{f,t}) = \delta_t + 0.625\ (0.043) \cdot \ln(R\&D/product_{f,t}) + \epsilon_{f,t}, \quad (3.26)$$

with standard errors clustered by firm. We then extend the estimation to allow for firm fixed effects, and we estimate the panel elasticity in a first-differenced format, yielding

$$\Delta \ln(Patents/product_{f,t}) = \delta_t + 0.353\ (0.057) \cdot \Delta \ln(R\&D/product_{f,t}) + \epsilon_{f,t}, \quad (3.27)$$

The range of these point estimates is representative of a broader set of estimates for the η parameter. Table 1a summarizes eight variants of the OLS levels regressions. The rows indicate four measures of firm size u_f: SIC3 industry counts, SIC4 industry counts, sales, and establishment counts. Column headers indicate whether R&D inputs are being measured through expenditures or counts of scientists. The eight coefficients are from eight separate estimations of the regression in Equation 3.26. The average of these eight estimations is 0.69, and the estimates are consistently within the range of 0.63-0.76.

TABLE 1A. OLS LEVELS ESTIMATES FOR η PARAMETER

	R&D Input Measure H_f	
Firm Size Measure u_f:	**R&D Expenditure**	**Scientist Counts**
SIC3 Counts	0.632 (0.042)	0.653 (0.048)
SIC4 Counts	0.625 (0.043)	0.644 (0.048)
Sales	0.761 (0.053)	0.751 (0.048)
Establishments	0.714 (0.039)	0.732 (0.041)

Notes: Table presents eight variants of the regression in Equation 3.26.

Table 1b similarly summarizes eight estimation variants of the first-differenced regression from Equation 3.27. The average across these variants is lower at 0.37, with a wider range from 0.29 to 0.51.

TABLE 1B. OLS FIRST-DIFFERENCED ESTIMATES FOR η PARAMETER

	R&D Input Measure H	
Firm Size Measure u_f:	**R&D Expenditure**	**Scientist Counts**
SIC3 Counts	0.342 (0.056)	0.286 (0.052)
SIC4 Counts	0.353 (0.057)	0.296 (0.053)
Sales	0.405 (0.075)	0.348 (0.065)
Establishments	0.505 (0.058)	0.455 (0.054)

Notes: Table presents variants of the regression in Equation 3.27.

Our baseline value for η is 0.5, taking a mid point within the range of estimates in Tables 1A and 1B.

We also find comparable η parameters in robustness checks off of this sample platform. For example, restricting the sample to firms with energy patents as more than 30% of their innovations yields levels and first-differences estimates of 0.744 (0.065) and 0.384 (0.100), respectively. Restricting our sample to firms that would have been defined for the sector using codes from Popp, David, and Richard Newell (2012) yields levels and first-differences estimates of 0.704 (0.049) and 0.301 (0.071), respectively. Relaxing our requirement that the firm be present in three periods yields levels and first-differences estimates of 0.614 (0.043) and 0.338 (0.056), respectively. We likewise find similar outcomes when incorporating industry-year fixed effects, using unweighted patent counts, or similar adjustments. In addition, Acemoglu, Daron, Ufuk Akcigit, Nick Bloom, and William Kerr (2012) describe a related instrumental variable elasticity of patenting to science and engineering workers of 0.694 (0.295) across the whole economy developed through H-1B visa reforms estimated by Kerr, William, and William Lincoln (2010).

Finally, Table 1C shows estimates from Poisson models that allow for zero patenting outcomes. We report both random effects and fixed effects formats; to conserve space, we only provide two choices of firm size that mostly bound the other variants. Standard errors are bootstrapped. Using four-digit industry counts to measure size consistently delivers elasticities around 0.33, while using establishment counts delivers elasticities around 0.57. Our baseline estimate of $\eta = 0.5$ falls again within these ranges.

TABLE 1C. POISSON ESTIMATES FOR η PARAMETER

	R&D Input Measure H_f	
Technique, Firm Size Measure u_f:	**R&D Expenditure**	**Scientist Counts**
Random Effects, SIC4 Counts	0.326 (0.122)	0.361 (0.079)
Fixed Effects, SIC4 Counts	0.321 (0.106)	0.357 (0.089)
Random Effects, Establishments	0.567 (0.108)	0.584 (0.064)
Fixed Effects, Establishments	0.565 (0.103)	0.583 (0.076)

Notes: Table presents fixed and random effects Poisson estimates similar to Tables 1A and 1B.

We next turn to the α parameter, which in our model describes technology leap-frogging. This process is challenging to model empirically, and we are unfortunately unable to identify exact races between clean and dirty technologies directly within the patent data. This limitation is due to the extreme narrowness of the technology codes that are entirely clean or dirty in application, while patent class divisions are too broad and few in number. We thus instead identify the rate at which patents with exceptional quality emerge using patent citations. We specifically quantify the rate at which patents enter and establish quickly high levels of citations compared to their incumbent peers.

We start with our dataset of all energy-sector patents granted to US inventors during the post-1975 period. We calculate among these energy-sector patents the citation count distribution among incumbent patents by year, excluding within-firm citations. Incumbent patents are defined to be those that were applied for 5-10 years before the focal year; we cap at 10 years prior so that we can have a stable window across a time period from 1985 onwards for analysis. Citations are coming from patents being applied for in the focal year. By conditioning the citation distribution upon a patent receiving a citation in a given year, we are effectively looking at technologies that are being actively used, with many incumbent patents dropping out as no one is building on them.

We then calculate for new patents the citations they receive by year. We designate a major entrant as any patent that has a citation count that exceeds the 90th percentile of the incumbent distribution in any of its first three years. This evaluation approach is designed to keep the incumbent groups (5-10 years earlier) separate from the entrant groups (max of three years earlier). 4.2% of entrants achieve this level of major entrant. We find a slightly lower estimate at 4.0% using Popp, David, and Richard Newell (2012) definitions, and a rate of 4.1% when making the citation distributions specific to each patent class. Based upon these findings, we set $\alpha = 4\%$.

Finally, for L^s, which is the supply of scientists and engineers involved in R&D-type activities in the model (relative to production workers), we use 5.5%. We calculate this share from Census IPUMS using the 2000 5% sample. We keep non-group quartered workers who are aged 20-65 years old and working in industries closely related to the energy sector. We also require 20 weeks worked within the year and a usual hours worked of 20 or more during each week. 5.5% is the share of these workers with bachelor's educations and higher employed in occupations related to science and engineering.

3.3.4 Initial Technology Gaps

To provide the initial distributions of the model, we develop estimates of the cumulative stock of technologies invented by clean- and dirty-energy firms using three-digit SIC industries as approximations of product lines. We develop these distributions in three steps. The first step is to calculate the sum of patents made by each firm during the 1975-2004 period and the firm's distribution of employment across SIC3 industries in these sectors over the same period. We then apportion the firm's cumulative patent stock across SIC3 industries using the firm's employment distributions. For each SIC3 industry, we finally sum the apportioned patents made by clean- and dirty-energy firms. This sum of patents across all firms, active or inactive, reflects the quality ladders structure of our model.

These calculations provide us over 400 estimates of comparative clean- and dirty-energy stocks. Across these SIC3 industries, clean-energy firms are estimated to have a higher cumulative patent stock in 13.1% of industries. For data quality and Census Bureau disclosure restrictions, we focus on the upper half of the industry distribution in terms of cumulative clean and dirty patent counts, which has 13.0% of industries being led by the clean-energy stock; within manufacturing and energy production specifically, this share is

12.5%. The following table summarizes some details of these lines:

TABLE 2. INITIAL CONDITION DISTRIBUTIONS SIC3

Metric:	**Clean Energy**	**Dirty Energy**
Mean Patent Total	260	1029
Standard Deviation	515	1500
Share: [0,20]	37%	0%
Share: [21,100]	25%	6%
Share: [101,500]	22%	50%
Share: [500+]	16%	44%

The average gap to the frontier for dirty-patents stocks in the 13% of cases where clean patents have the lead is 424 patents, or in relative terms, 39% of the total patenting in that line to date. The average gap to the frontier for clean-patent stocks in the 87% of cases where dirty patents have the lead is 947 patents and 76% in relative terms. To convert the empirical gap into the units of the model, we use the following reasoning. In our model, the annual patent flows of incumbents is 16.1% per product line (the sum of $x^c = 3.9\%$ and $x^d = 12.2\%$). In the data, the median annual flow of patents is approximately 17.1 per line. Hence we divide the empirical patent distribution of clean and dirty (which consists of patents registered between 1975-2004) by a conversion factor $17.1/0.161$ and round it to the closest integer. This gives us the initial number of improvements $n^d_{j,0}$ and $n^c_{j,0}$. Then we compute the initial productivities as $q^d_{j,0} = \lambda^{n^d_{j,0}}$ and $q^c_{j,0} = \lambda^{n^c_{j,0}}$ to provide the initialization values. The following graph plots the density of this distribution with gaps between dirty and clean technologies on the horizontal axis:

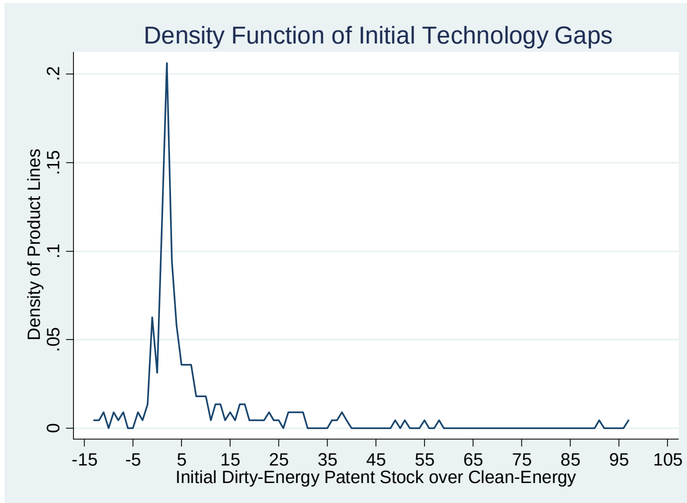

FIGURE 3

This graph shows that in most product lines the dirty technology is only a few steps ahead of clean technology, but there is a long tail of product lines with a large gap between dirty and clean, and a small set in which clean is ahead of dirty. The fraction of product lines with a non-zero gap in terms of step sizes is 90%. Clean energy leads by one or more step sizes in 9% of cases. Dirty energy has a lead of 20 and 50 steps sizes or more in 11% and 2% of technologies, respectively. We later consider an alternative initial distributions that modifies several of the modelling choices made here.

3.3.5 Simulated Method of Moments

The remaining parameters θ, λ, F_I and F_E are estimated using simulated method of moments (SMM). McFadden, Daniel (1989), Pakes, Ariel, and David Pollard (1989), and Gouriéroux, Christian, and Alain Monfort (1996) characterize the statistical properties of the SMM estimator. This quantitative approach takes a set of key moments from our model, and then chooses the parameter vector so as to minimize the distance between these moments as implied by our model and as computed from the Census Bureau data,

$$\min \sum_{i=1}^{4} |\text{model}\,(i) - \text{data}\,(i)| \, ,$$

where we index each moment by i. SMM iteratively searches repeatedly across sets of parameter values for θ, λ, F_I and F_E in the model until the model's moments are as close as possible to the empirical moments (see Adda, Jerome, and Russell Cooper (2003) for further details). We also choose the heterogeneity parameter, ξ, as 10% and verify that our results are not sensitive to this parameter.

We use three moments from the microdata—firm entry rates, firm exit rates, and the average R&D/sales ratio of firms—together with the growth rate of the sector to identify these parameters. The entrant's labor share and exit rates are calculated across the five-year intervals of our Census Bureau data and then transformed into annualized net rates of 1.3% and 1.75%, respectively. We match the construction of these entry and exit rate moments in the model. The weighted average R&D-to-sales ratio is 6.57%, using log sales as weights and capping the R&D/sales ratio at 10x to reduce outliers (approximately the 99.8th percentile of the distribution). The aggregate annual sales growth per worker is 1.23% for the sector across the 1975-2004 period. After identifying these parameters in the estimation section, we investigate the fit of the model by comparing the implications of the model with a battery of other non-targeted moments from the microdata.

3.3.6 Computational Algorithm

Our theoretical analysis shows that microeconomic behavior is independent of climate dynamics. We therefore solve for value functions, innovation rates, and distributions first, then use those to find the time path for the carbon stock, temperature and other variables of interest.

The solution algorithm for this model involves finding the transition dynamics as the fixed point of a forward-backward iteration process, as in Conesa, Juan Carlos and Dirk Krueger (1998). See Zangwill and Garcia (1981) for further references. If the state space were of a more manageable size, one could simply solve the value function over this space and characterize the dynamics given arbitrary initial conditions. However, in this case the state space is the distribution of product lines over the technology gaps between the clean and dirty technologies. For any reasonable approximation, this results in a very high dimensional state space. Therefore, we solve each element of the model as a function of time given the initial conditions from the patenting data. The value function in early periods will thus depend on value functions in later periods. These later period value functions will in turn depend on the later period product distribution, which depends on early period value functions and innovation rates.

To solve for the fixed point of the sequence of value functions, we first discretize time into $N = 2048$ steps and set a terminal period $T = 2000$. Due to the symmetry between technology types inherent in this model, when a single type of technology is dominant—in the sense that the technology gap distribution is heavily skewed to either clean or dirty technology—one can analytically characterize value functions $v^j_{n,\infty}$ and innovation rates x^j_∞ and z^j_∞. We use these values as terminal conditions, though we do not know in advance which technology (clean or dirty) will be dominant. In addition, we set large upper (100) and lower (-100) bounds on the step gap distribution space. The algorithm proceeds as follows:

1. Posit an initial guess for the value function at time zero of the form

$$v^j_{n,t}(0) = \frac{\pi^j_n}{\rho + \bar{z}} \quad \forall t$$

   where $\bar{z}$ represents an estimated rate of creative destruction (we use $\bar{z} = 0.15$). Instantiate the technology gap distribution using the patent data with

$$\mu_{n,t}(0) = \mu_{n,0} \quad \forall t.$$

2. Iterate forward in time from $t = 0$ to $t = T$ by finding innovation rates x_t^j and z_t^j given value function and product distributions guesses at time $t+1$, $v_{n,t+1}^j(k)$ and $\mu_{n,t+1}(k)$. Using these innovation rates, update the time $t+1$ product distribution $\mu_{n,t+1}(k+1)$ using discrete time versions of the flow equations in 3.21.

3. Find the implied dominant technology at the terminal period by determining which technology type has a higher aggregate innovation rate as some late stage period $T - T_P$ (we use $T_P = 200$). Use the corresponding terminal value functions to update $v_{n,T}^j(k+1) = v_{n,\infty}^j$.

4. Iterate backward in time from $t = T$ to $t = 0$ by updating value function $v_{n,t-1}^j(k+1)$ using $v_{n,t}^j(k)$ and $\mu_{n,t}(k)$ according to a discretized version of the HJB equation in Theorem 7, re-solving for innovation rates x_t^j and z_t^j in the process.

5. Repeat steps 2-4 until the convergence criterion

$$\max_{n,t} \left| v_{n,t}^j(k+1) - v_{n,t}^j(k) \right| < \varepsilon$$

is met. We use $\varepsilon = 10^{-6}$.

In order to avoid any instability, particularly when one is close a threshold where the asymptotically dominant technology switches over, we also introduce heterogeneity in incumbent fixed costs as explained above.[15]

[15]Our algorithm introduced a similar heterogeneity in entrant fixed costs, but at the end, entrants are never in the region where this heterogeneity matters.

Using up-to-date computer hardware, the equilibrium solver takes anywhere from five seconds to two minutes, depending on the speed of convergence. The code is written mostly in Python, with core routines written in C/C++.

Estimation To find the moments used in the SMM estimation, we simulate a panel of 16384 firms using equilibrium variables from the above model solving routine. Each firm has a portfolio of product lines with various technology gap values. We cap the maximum number of product lines a firm can have at 200. In order to determine the sales and R&D activity of the firm, we must keep track of both the number of product lines it is currently operating in, as well as the knowledge stock of the firm, which can in general differ. We simulate this panel of firms for 5 years to replicate the data generating process, and discretize time to have 100 subperiods per year, so that the simulations have 500 periods.

Optimal Policy We compute optimal policies for both the constant and time-varying cases. In the constant case, we use a straightforward grid search to find the optimum. In the time-varying case, we parameterize policies using a three stage carbon tax and a three stage research subsidy. Within each stage, whose boundaries are also parameterized and optimized over, we have constant policy levels. We then search over this space of functions using a combination of a simple simulated annealing algorithm (Kirkpatrick, S. C.D. Gelatt, and M.P. Vecchi (1983)) and a Nelder-Mead (simplex) algorithm (Nelder and Mead (1965)).

3.4 Estimation Results

In this section, we provide the results of the simulated method of moments estimation and discuss the fit of our model to non-targeted moments. Finally, we show how atmospheric carbon concentration, temperatures and aggregate output evolve given these parameters in a laissez-faire equilibrium (with no policy intervention) starting from the observed distribution of technology gaps.

3.4.1 Parameter Estimates

Table 3 summarizes our parameter estimates.

TABLE 3. PARAMETER ESTIMATES

Parameter	Description	Value
θ	Innovation productivity	0.500
λ	Innovation step size	1.075
F_I	Fixed cost of incumbent R&D	0.002
F_E	Fixed cost of entry	0.035

Our innovation productivity estimate implies that one unit of labor with a single product line generates an innovation with probability of about 8% a year. We estimate the innovation step size as 1.075 which implies a gross profit margin 7%. Finally our model predicts a sizable fixed cost advantage for incumbent firms. Their fixed cost of operation is equivalent to 6% of the entrants' fixed cost.

3.4.2 Goodness of Fit

Here we describe the goodness of fit of our model, first focusing on the four targeted moments, and then a range of diverse non-targeted moments.

Targeted Moments

Table 4 shows the values of the four moments used for estimation in the data and the model implications.

TABLE 4. MOMENT MATCHING

Moments	Model	Data
Entry Share	0.013	0.013
Exit Rate	0.018	0.018
Average R&D/Sales	0.066	0.066
Aggregate Sales per Employee Growth	0.007	0.012

On the whole, there is a very good match between the model and the data. The entry share, exit rate and R&D intensity are identical between the data and the model. Moreover the aggregate sales per worker growth is also very close.

Non-targeted Moments

Our main method of evaluating the quantitative fit of our model is to look at a range of non-targeted moments, which are presented (together with the model implications) in Tables 5A-5C.

We choose the non-targeted moments to represent aspects of the firm size distribution and its growth properties, which are quite different from the moments we targeted in our estimation. Our first non-targeted moment compares the size ratio of the median entrant to the median incumbent firm. Our targeted moments on entry/exit rates, overall sector growth, and R&D intensity do not directly impose strong constraints on this size distribution. Table 5A contrasts the size ratios in the model and data with respect to employment, sales, and sales per employee, and shows that our model implications match the data very closely with respect to the latter two metrics, though not as well for employment.[16]

TABLE 5A. ENTRANT SIZE RATIO TO INCUMBENTS

	Ratio of Median Sizes	
Size Measure:	**Model**	**Data**
Employment	0.17	0.03
Sales	0.18	0.20
Sales per Employee	1.12	1.05

Notes: Table compares non-targeted moments in model and data.

Our next point of comparison is for the structure of the growth distribution. We first calculate the unconditional growth rate of employment for each firm in the model and data defined as $(Emp_t - Emp_{t-1})/((Emp_t + Emp_{t-1})/2)$. This formula divides the employment change across the period by the average of the start and end values. As argued by Davis, Steven, John Haltiwanger, and Scott Schuh (1996), this approach has attractive properties like a symmetric treatment of positive and negative growth and bounded values that minimize outliers. We calculate growth over five-year intervals. We then calculate the probability that firms experience substantial movements in either positive or negative directions. Comparing these movements in the model to the data is meaningful as it provides insights into how well the innovation step sizes and associated firm dynamics mirror the

[16]To pass Census Bureau disclosure restrictions, the empirical medians are "fuzzy" median estimates that use the average values over the 45th to 55th percentiles.

sector's true performance. Table 5B summarizes this comparison:

TABLE 5B. COMPARISON OF GROWTH DISTRIBUTION

	Employment Growth Probability	
Change over 5-Years:	**Model**	**Data**
Decrease 75% or more	0.17	0.11
Decrease 50% or more	0.20	0.15
Decrease 25% or more	0.27	0.26
Increase 25% or more	0.24	0.31
Increase 50% or more	0.17	0.20
Increase 75% or more	0.15	0.14
Increase 100% or more	0.08	0.11

Notes: Table compares non-targeted moments in model and data.

On this dimension, the model matches the data quite well. Compared to the data, the model somewhat over-predicts major downward employment declines of 50% or more. Recall that we matched the exit rate itself almost exactly, so this indicates an over-prediction of large employment declines conditional on survival. On the other hand, the model matches the data well for predicting employment decreases or increases of 25% or more. Likewise, the model and data are in very close agreement on the relative probabilities and sizes of large employment increases for firms.

Our final comparison is the variation in conditional growth rates for employment across the firm size distribution, again over a five-year period. We divide firms into quantiles based upon their initial size in each period. We then compute the growth rates using the above formula, and the following table provides the comparison:

TABLE 5C. COMPARISON OF GROWTH OVER SIZE DISTRIBUTION

	5-Year Conditional Growth Rates	
Quantile of Sizes:	**Model**	**Data**
Smallest	18%	31%
2nd	25%	14%
3rd	18%	11%
4th	-5%	-1%
Largest	-0%	-10%

Notes: Table compares non-targeted moments in model and data.

The comparison is again reasonable. We match the general feature in the data that conditional growth rates are highest for small firms. The model's employment distribution is a little less fine-grained than the data, as about 50% of our firms have one product and employment is partially proportional to product counts. In consequence, there is limited variation across the smaller quantiles in the model compared to the more regular distribution in the data. The model and data then match quite well in identifying lack of growth for the top two quantiles compared to the bottom three (though the model does not predict employment declines in the largest firms that are present in the data).

Overall, our reading of the evidence is that for this range of diverse moments, which we did not target in our estimation, the model performs reasonably well, and this bolsters our confidence that our quantitative model is able to capture production and innovation dynamics in the energy sector.

3.4.3 Climate Dynamics in the Laissez-faire Economy

We next describe the implied future equilibrium and atmospheric carbon paths of the model given our estimates with the case of no carbon taxes and research subsidies. Given the initial distribution of technology gaps, dirty innovation is more productive and with no policy intervention, most R&D is initially targeted to the dirty technology as shown in Figure 4. Moreover, at these innovation rates, technology gaps and the profitability of dirty technologies increases relative to those of clean technologies, and clean R&D converges to zero. Consequently, in the long-run clean technologies disappear completely and dirty technologies take over the whole economy.

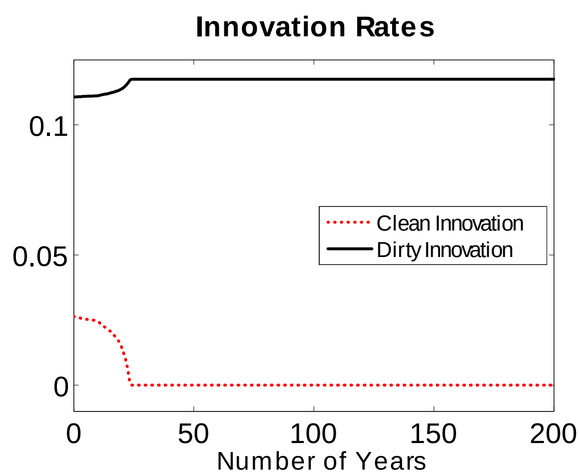

FIGURE 4: INNOVATION RATES, x_t^c AND x_t^d

The obvious implication of this time path of innovations is a steady increase in dirty energy production and carbon emissions. There are two ways of ascertaining the implications of these growing carbon emissions from our economy estimated and calibrated to US data. The first is to ignore emission growth from the rest of the world (i.e., keeping their emissions at a constant level). This is done in Figure 5A, which shows an increase in temperature of an additional 2.5°C in the next 200 years.[17] The alternative is shown in Figure 5B and assumes that emissions from the rest of the world will grow at the same rate as the US.

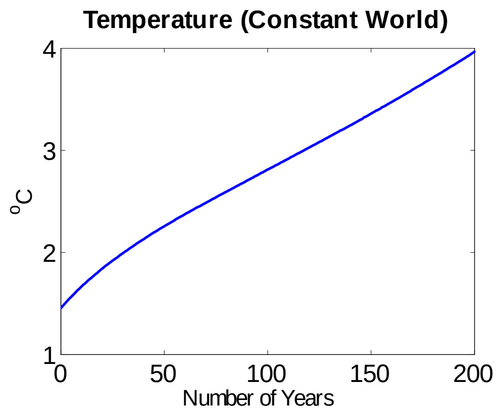

FIGURE 5A

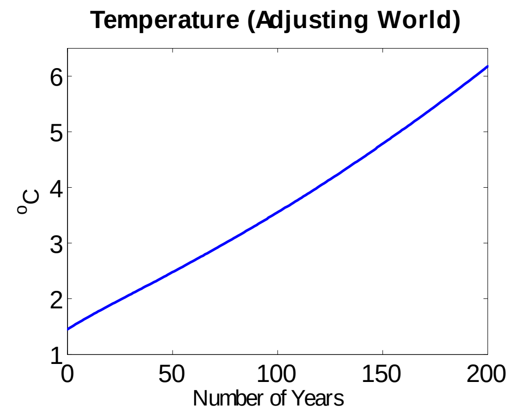

FIGURE 5B

The impact on global temperature is considerably exacerbated in view of the fact that we are now showing the increase in global temperature resulting from growth in global emissions, not just US emissions. It is important to recall that, as noted above, given the functional form in Equation 3.2, the exact path of emissions from the rest of the world has no impact on optimal policy (since the marginal damage created by US emissions are independent of the level of global emissions). This is the reason why we did not have to take a position on the time path of emissions in the rest of the world until this point (and now we are only taking a position in order to translate its implied path of emissions into changes in global temperatures).

[17]We use the following formula to compute the temperature changes:

$$\Delta temperature = \frac{\lambda \left(\ln S_t - \ln \bar{S}\right)}{\ln 2}.$$

3.5 Policy Analysis

In this section, we derive the policy sequence that maximizes discounted welfare. Throughout, we do not allow the social planner to correct for monopoly distortions, thus limiting ourselves to the policy instruments discussed above—a carbon tax and subsidy to clean research.[18] In addition, our theoretical analysis makes it clear that what is relevant is the differential tax and subsidy rates for clean vs. dirty energy, thus we just look at taxes on dirty production, which we refer to as "carbon taxes," and subsidies to clean innovation. Finally, we restrict the subsidies to entrants and incumbents to be the same, i.e., $s_{Et} = s_{It}$ for all t (based on early results which suggest that when both instruments are allowed to vary they are often equal or very close to each other). Throughout, we consider a private discount rate of $\rho = 1\%$ and present results for two different social discount rates: $\rho_{sp} = 1\%$ and $\rho_{sp} = 0.1\%$.[19] In both cases, paths that involve no switch to clean technology will lead to unbounded atmospheric carbon and temperature increases and consumption limiting to zero because of the economic distortions created by the unbounded increase in atmospheric carbon (see Equation 3.2), and we assign minus infinite social welfare to such paths, so that when feasible, a switch to clean technology is preferred.[20] Finally, we first focus on optimal constant policies (where carbon taxes and research subsidies are constant over time), which have several advantages: they are simpler and more transparent and the optimal time-varying policies, which we also characterize below, are time-inconsistent, raising some caveats about interpretation.

[18] As mentioned above, in the one-sector version of our model (either with only dirty or only clean technology), taxes or subsidies to research would only affect relative wages of skilled workers (employed in the research sector), and crucially not the aggregate rate of innovation. For this reason, subsidies to clean research or taxes on dirty research are identical in our model.

[19] The reasoning here is that, following Stern, Nicholas (2007), the social planner—society—may have a lower discount rate than that implied by market interest rates. Thus 0.1% is what should be applied to the welfare analysis of the social planner, while still keeping the discount rate that firms use in their decisions at 1%.

[20] This is implied whenever the growth rate g is greater than ρ_{sp}, which is always satisfied when $\rho_{sp} = 0.1\%$, but may or may not be satisfied when $\rho_{sp} = 1\%$ depending on the growth rate.

3.5.1 Optimal Constant Policies

Table 6 shows optimal constant policies for the two values of social discounting, $\rho_{sp} = 1\%$ (high), $\rho_{sp} = 0.1\%$ (low).

TABLE 6. OPTIMAL CONSTANT POLICY

	$\rho_{sp} = 1\%$	$\rho_{sp} = 0.1\%$
τ	16%	44%
s	61%	95%

In both cases, there is a very aggressive research subsidy for clean technology. With $\rho_{sp} = 1\%$, the carbon tax is fairly low, at 16%, while research directed at clean technologies receives a 61% government subsidy (meaning that for every dollar of R&D spending, there is a 61 cents subsidy). With the social discount rate of $\rho_{sp} = 0.1\%$, carbon taxes are raised to 44%, but now clean research subsidies are even more aggressive, at 95%.

The intuition for why optimal policy relies so much on subsidies to clean research is instructive. The social planner would like to induce a switch from R&D directed at carbon intensive dirty technologies towards clean technologies. She can do so by choosing a sufficiently high carbon tax rate today and in the future, because this would reduce the profitability of production using dirty technologies and secure both a switch to clean production and, on the basis of this, to research directed at clean technologies. However, this is socially costly because given the current state of technology, switching most production to clean technology has a high consumption cost (because the marginal costs of production of clean technologies are initially significantly higher than those of dirty technologies). Hence it is a better strategy for the social planner to choose the carbon tax to only deal with the carbon emission externality and rely on the research subsidy to induce the switch to clean technologies in the long run. Figure 6 in fact shows that the social planner is able to do this, particularly with $\rho_{sp} = 0.1\%$, where the optimal policy involves a very rapid ramp up of clean innovation rates and the disappearance of all research directed to dirty technologies in about 130 years. Interestingly, however, with $\rho_{sp} = 1\%$, the social planner chooses not to completely replace dirty research with clean research in the first several hundred years. Instead, she subsidizes clean research just enough to make sure that clean research and thus clean technologies also survive for a long time, but not so much that they overtake the dirty sector completely. As a result, dirty innovation survives for several hundreds years (clean innovation exceeds dirty innovation only around year 500), but throughout, innovation rates in the clean technology are significantly higher than in the laissez-faire equilibrium shown in Figure 4 where they converged to zero in about 25 years.

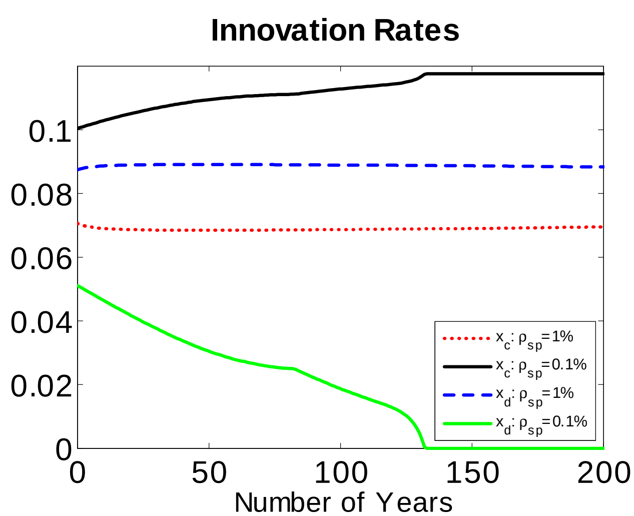

FIGURE 6: INNOVATION RATES UNDER OPTIMAL POLICIES

Figure 7 depicts the implied path of temperature increases under the optimal policies, in the same two ways as we have done in Figure 5—assuming either that emissions from the rest of the world are constant or that they grow at the same rate as US emissions. In both cases, global temperature increases less, and in fact significantly less with $\rho_{sp} = 0.1\%$, than in

Figure 5.

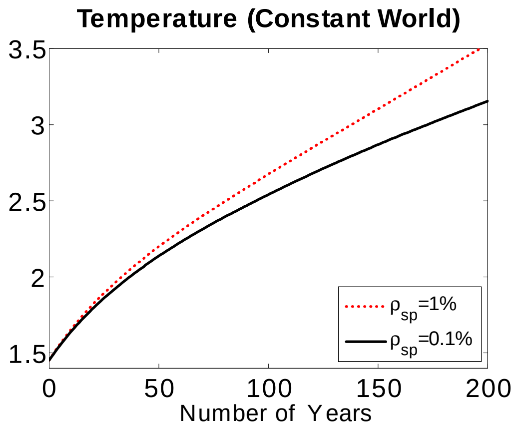

FIGURE 7A

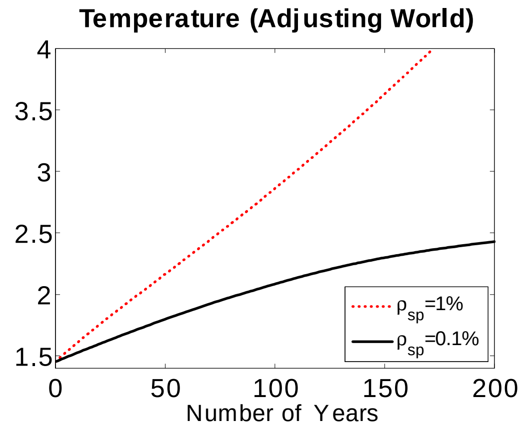

FIGURE 7B

3.5.2 Optimal Time-Varying Policies

We now return to optimal time-varying policies and characterize the welfare gains from using time-varying rather than constant policies. For computational reasons, we look for policies that take a simple "step function form" with three endogenously determined switch points.[21]

The resulting optimal policies are shown in Figure 8. A couple of features are worth noting. First, the subsidy rate for clean research is very similar to the constant policies in both cases. With a social discount rate of $\rho_{sp} = 0.1\%$, this subsidy rate is roughly constant (starting at 95% and declining to 93% at year 54).[22] Moreover, in this case, the carbon tax starts somewhat higher than with the optimal constant policy but then declines from 54% to 34% at year 92). With $\rho_{sp} = 1\%$, the pattern is different: the subsidy rate starts and

[21]For our baseline results, we experimented with increasing the number of switch points and with alternative formulations of time variation, with broadly similar results.

[22]Note that research subsidies in the far future may not matter very much because early research subsidies may have already induced a large reallocation of research from dirty to clean. Nevertheless, research subsidies in the future are not undetermined because there is always some positive fraction of firms with a dirty portfolio of product lines which will then have incentives to undertake research in the dirty technology (this fraction declines to zero asymptotically in the optimal allocation with $\rho_{sp} = 0.1\%$).

remains at 25%, again very close to its constant optimal policy value, but the carbon tax starts at zero and stay there for quite a while, and then increases dramatically to 650% in year 328.

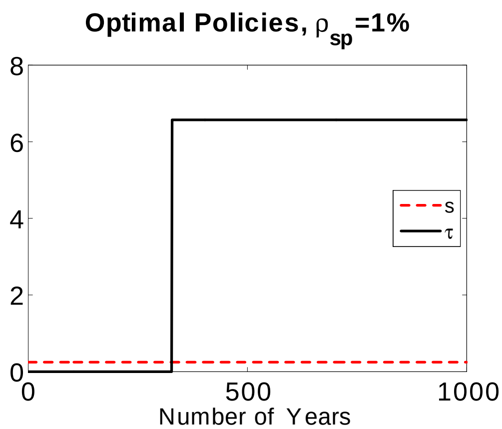

FIGURE 8A

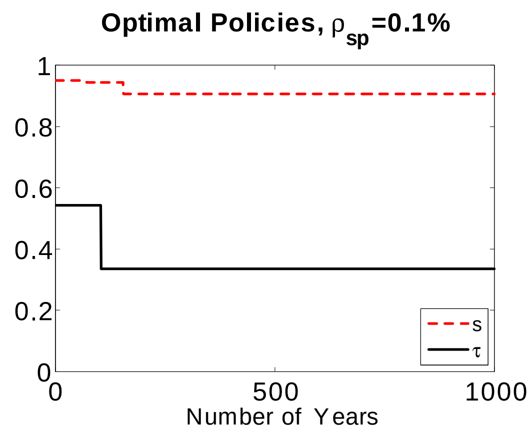

FIGURE 8B

The patterns shown in Figure 8 result from the interplay of two counteracting forces. First, all else equal, the social planner would like to delay as much as possible the consumption loss from switching to clean technologies.[23] Second, carbon taxes early on are more effective in both switching production and reducing emissions (given the long half life of carbon in the atmosphere imposed in our model of the carbon cycle). With $\rho_{sp} = 1\%$, the first effect is dominant because with this relatively high social discount rate, high consumption during the early years is highly valued, encouraging the planner to delay the start of high carbon taxes for quite a while. In consequence, in this case carbon taxes are sharply backloaded, and in fact, as in the constant policy case, dirty innovation disappears only very slowly—over several hundreds of years. With $\rho_{sp} = 0.1\%$, the future is less heavily discounted, strengthening the second effect and making carbon taxes frontloaded and the complete switch to clean innovation much more rapid. In both cases, however, the average

[23]This effect in part reflects the fact that the social planner is committing to a policy path and is not "time consistent".

values of the carbon tax in the first 200 years is in the ballpark of the constant optimal policies (16% with $\rho_{sp} = 1\%$ and 44% with $\rho_{sp} = 0.1\%$).

Table 7 shows that the welfare loss from using constant policies is quite small, 0.3% with $\rho_{sp} = 0.1\%$, but sizable, about 16%, with $\rho_{sp} = 1\%$, which reflects the benefits for social planner's utility resulting from high consumption growth at the expense of high emissions in the first 300 years. This pattern suggests that the results with the low social discount rate, $\rho_{sp} = 0.1\%$, are more plausible in a range of dimensions.[24]

TABLE 7. WELFARE COSTS

$\rho_{sp} = 1\%$	$\rho_{sp} = 0.1\%$
16%	0.3%

3.5.3 Counterfactual Policy Analysis

Our model enables us to investigate the welfare and climatic implications of a range of counterfactual policies. Here we focus on two counterfactuals. The first is relying just on a carbon tax (i.e., no research subsidy) as the policy tool, and the second is delaying intervention for 50 years and then choosing the optimal policy from that point onwards. We focus on time-varying optimal policies, which are shown for these two counterfactual

[24]We have also verified that our main results with an intermediate social discount rate of $\rho_{sp} = 0.5\%$ are very similar to those with $\rho_{sp} = 0.1\%$, again making us trust these results more than the ones based on $\rho_{sp} = 1\%$.

policies in Figure 9.

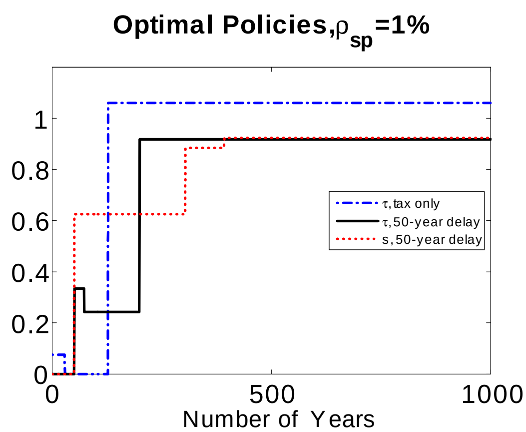

FIGURE 9A

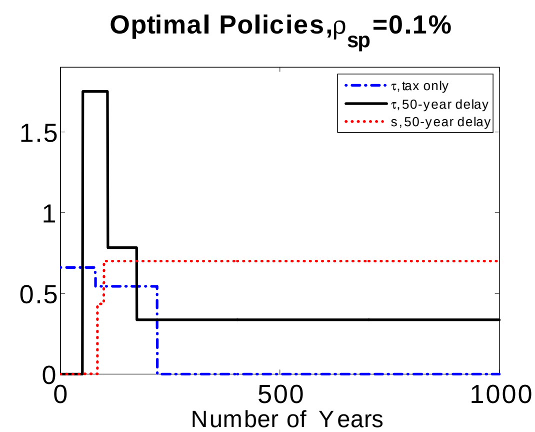

FIGURE 9B

Optimal policies following the 50-year delay are more aggressive than the baseline optimal policies, and when the only policy tool is the carbon tax, this tax also is typically higher.[25]

For example, in the carbon tax only counterfactual, this tax is higher because it is also being used to redirect innovation towards clean technologies. As a result, with $\rho_{sp} = 1\%$, aggregate temperature increases less at long horizons under this constrained optimal policy than our actual optimal policy, but this is at the expense of slower output growth, especially early on. As a result, the cost of just relying on carbon tax for optimal policy, shown in Table 8, is 4.2% with $\rho_{sp} = 1\%$ and 3.4% with $\rho_{sp} = 0.1\%$. Delaying the start of optimal policies by 50 years leads to greater losses—a consumption equivalent welfare cost of 8% with $\rho_{sp} = 1\%$, and 16.6% with $\rho_{sp} = 0.1\%$. These numbers indicate that delaying policy interventions to combat carbon emissions could have very significant welfare costs, especially when the social discount rate is low. Moreover, just relying on carbon taxes—

[25]When just relying on the carbon tax and with $\rho_{sp} = 0.1\%$, the carbon tax reaches zero earlier than in the baseline shown in Figure 8. Nevertheless, it is effectively more aggressive than the baseline since it starts at a higher level (66% instead of 54%) and remains at a higher level (54% instead of 34%) for the first 220 years, and this induces both a much more rapid switch to clean production and also encourages a switch to clean innovation despite the absence of research subsidies in this case.

eschewing research subsidies—could also have sizable welfare costs.

TABLE 8. WELFARE COSTS

Carbon Tax Only		50-year Delay	
$\rho_{sp} = 1\%$	$\rho_{sp} = 0.1\%$	$\rho_{sp} = 1\%$	$\rho_{sp} = 0.1\%$
4.2%	3.4%	8.0%	16.6 %

Finally, we also evaluate what the climatic and welfare implications of maintaining current US policies (here interpreted for the whole world) would be relative to adopting an optimal policy moving forward. For this purpose, we have tried to estimate the carbon taxes implied by US policies and the current subsidies to clean innovation (relative to dirty R&D) in our sample of firms. There is much uncertainty about what the carbon tax in the United States will be moving forward. A cap-and-trade program is likely to be implemented, but it is unclear what the implied carbon tax rate will be. On the other hand, Greenstone, Michael, Elizabeth Kopits, and Anne Wolverton (2011) estimate a social cost of carbon equal to about \$21 in 2010, expressed in 2007 dollars, and this number is currently being used for cost-benefit analysis by US agencies. This social cost estimate is the central tendency across a number of models and scenarios considered. The social cost increases in real 2007 terms to \$45 in 2050 as a consequence of future marginal emissions becoming ever more harmful. We therefore use two values for the "business-as-usual" carbon tax, 0% consistent with the current situation, and 24% (approximately implied by \$45 social cost of carbon in 2050, a relatively early point in the transition path).[26] We estimate the current clean

[26]In particular, US carbon emissions are 1.58 billion tons in 2002. One metric ton of carbon is equivalent to 3.667 units of carbon dioxide. Our dirty firms have sales of approximately one trillion dollars in this year. The \$45 social cost is \$39 in 2002 terms. These numbers imply a real tax rate in 2050 of about 23% $((39 \times 3.667 \times 1.58 \times 10^9)/10^{12} \simeq 0.23)$. We approximate this with an 24% tax rate (since our taxes have to be multiples of λ). This carbon tax rate is much less than currently used in Sweden (see Golosov, Mikhail, John Hassler, Per Krusell, and Aleh Tsyvinski (2011) and also less than the numbers suggest that by Nordhaus, William (2008).

research subsidy from our sample as follows: over our full 30 year period, 49% of all R&D expenditures by our clean firms are federally funded, while the same number is 11% for our dirty firms. This implies a 43% $((1-0.49)\,/\,(1-0.11) \simeq 1-0.43)$ subsidy for clean R&D relative to dirty R&D.

The scenario with a zero carbon tax, regardless of the discount rate, involves 100% welfare costs because, in this case, temperature increases rapidly and continues to grow unboundedly. Essentially, 43% R&D subsidy for clean is insufficient to redirect technological change towards clean with no carbon tax. The resulting significant damage to the environment leads to a disastrous welfare result. Interestingly, however, even with this less than optimal subsidy to clean research, it turns out that the temperature increase can be contained if there is a moderate carbon tax at 24%. As a result, with this moderate carbon tax, the welfare costs are still sizable but limited as shown in Table 9.

TABLE 9. WELFARE COSTS

$\tau = 24\%$, $s = 43\%$		$\tau = 0$, $s = 43\%$	
$\rho_{sp} = 1\%$	$\rho_{sp} = 0.1\%$	$\rho_{sp} = 1\%$	$\rho_{sp} = 0.1\%$
18%	8%	100%	100%

3.6 Robustness and Extensions

In this section, we investigate how our estimation, optimal policy and counterfactual results are affected by a range of different modeling assumptions or variations on parameter estimates. Throughout, to economize on space we only report the implied optimal policies (even when the variation in question involves reestimating the parameters of the underlying model).

3.6.1 Alternative Damage Elasticity γ

As noted above, actual damages from atmospheric carbon may be greater than the estimates commonly used in the economics literature. We now show the sensitivity of our results to higher values of these damages, captured by the parameter γ, in our model. Table 10 depicts constant optimal policies for two cases, when γ is twice and 10 times as large as our baseline value, $\gamma = 5.3 \times 10^{-5}$, and Figure 10 shows optimal time-varying policies for the same two cases.

TABLE 10. OPTIMAL CONSTANT POLICIES

	$\gamma = 2\times$			$\gamma = 10\times$	
	$\rho_{sp} = 1\%$	$\rho_{sp} = 0.1\%$		$\rho_{sp} = 1\%$	$\rho_{sp} = 0.1\%$
τ	16%	44%	τ	24%	54%
s	61%	95%	s	95%	95%

Overall, the results are remarkably similar to those in our baseline. Interestingly, with $\rho_{sp} = 1\%$, optimal constant policies are identical when γ is twice as large as the baseline. This result, which at first appears counterintuitive, is because the optimal policy in this case does not eliminate but chooses to contain carbon emissions (and does not even eliminate the dirty sector). When γ is doubled, the social planner still prefers to maintain this containment strategy, making optimal policies very similar to the baseline. When γ is taken to be much larger (10 times as large as the baseline), this is no longer optimal, and we now see a more aggressive carbon tax and a much more aggressive research subsidy, utilized for eliminating the dirty sector.

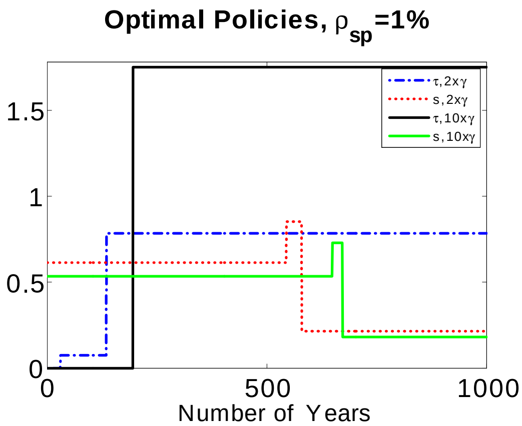

FIGURE 10A

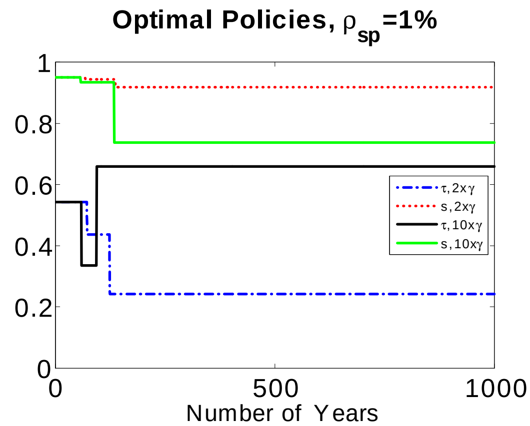

FIGURE 10B

Optimal time-varying policies, which are shown in Figure 10, are also quite similar—but not identical—to the baseline.

Overall, with the exception of the last case mentioned, the results suggest that the qualitative, and even quantitative, messages from our analysis are fairly robust to different economic damages from atmospheric concentration.

3.6.2 Costly Research Subsidy

We next investigate the robustness of our results to assuming that R&D subsidies create direct distortions. In particular, we assume that for every dollar of subsidy, $1 + \chi$ dollars need to be spent, so that χ is a waste, which we subtract from consumption. We consider two values of χ, 50% and 100%, both of which are very aggressive choices on the distortion or implications of research subsidies. The results for constant policies are shown in Table

11.

TABLE 11. OPTIMAL CONSTANT POLICIES

50% Consumption Cost			100% Consumption Cost		
	$\rho_{sp} = 1\%$	$\rho_{sp} = 0.1\%$		$\rho_{sp} = 1\%$	$\rho_{sp} = 0.1\%$
τ	16%	54%	τ	16%	66%
s	61%	53%	s	61%	0%

We find that except in one case, optimal policy still makes use of aggressive research subsidies despite the significant waste that these create. The reason for this is that, as implied by our discussion above, the carbon tax is a poor substitute for research subsidies; it also encourages clean research but does so at the cost of creating more intra-temporal distortions. In consequence, there is ample room for research subsidies even when they are distortionary. In fact, Table 11 shows that with a discount rate of $\rho_{sp} = 1\%$, optimal constant policies are identical to the case without any distortions. Intuitively, the social planner finds it optimal to leave the carbon tax unchanged (recall that the carbon tax can only change in steps), and with unchanged carbon tax, the research subsidy also remains constant. With the lower social discount rate, $\rho_{sp} = 0.1\%$, the carbon tax becomes more aggressive; in fact, with 100% distortions from research subsidies and this lower social discount rate, the optimal constant policy increases the carbon tax significantly and ceases to use research subsidies. However, Figure 11 shows that optimal time-varying policies in this case still involve heavy use of positive research subsidies. Moreover, the qualitative pattern of optimal time-varying policies is quite similar to the baseline, shown in Figure 8, and again involve backloading of carbon taxes for $\rho_{sp} = 1\%$, frontloading of carbon taxes for $\rho_{sp} = 0.1\%$, and fairly aggressive use of research subsidies, especially in the first few hundred years. The fact that research subsidies are now phased out entirely with $\rho_{sp} = 0.1\%$ is also very intuitive: research subsidies early on are sufficient to switch most innovation to clean, and start influencing only a few firms after a while (as most leading-edge technologies are now clean); because they are costly and become largely unnecessary, it is natural for the social planner to rely less on them. With $\rho_{sp} = 1\%$, this does not happen because the transition to clean technology is slower and research subsidies are still useful for several hundreds of years.

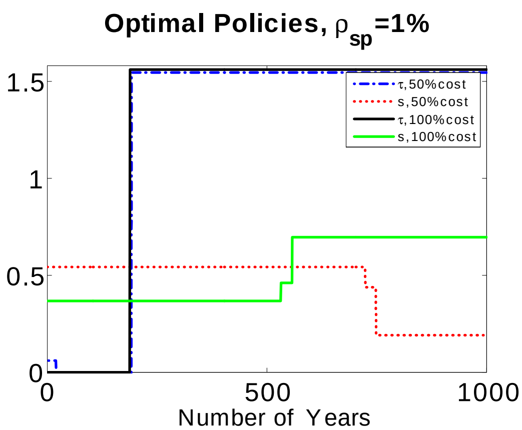

FIGURE 11A

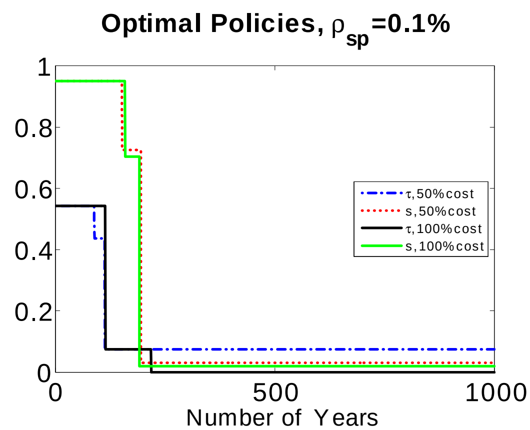

FIGURE 11B

Overall, we conclude that with reasonable values of distortions, and even with certain extreme values of distortions, the broad pattern of optimal policies is quite similar to the baseline case, and research subsidies are still an essential part of the portfolio of optimal policies, even if they may be significantly distortionary.

3.6.3 Alternative R&D Elasticities η

Our baseline results are for $\eta = 0.45$ which averaged across cross-sectional and first-difference estimates. We now reestimate the model using first a value of η in the ballpark of the cross-sectional estimates ($\eta = 0.65$) and then for a value corresponding to the first-difference estimates ($\eta = 0.35$), and investigate the implications of this for the fit of the model and for optimal policy. Overall, the fit of the model is not affected much by the change in η, and the implications for optimal constant policies are shown in Table 12 and optimal time-varying policies are shown in Figure 12.

TABLE 12. OPTIMAL CONSTANT POLICY RATES

$\eta = 0.35$			$\eta = 0.65$		
	$\rho_{sp} = 1\%$	$\rho_{sp} = 0.1\%$		$\rho_{sp} = 1\%$	$\rho_{sp} = 0.1\%$
τ	16%	34%	τ	24%	66%
s	0%	95%	s	84%	84%

When the elasticity of innovation to R&D effort is higher than in our baseline, at $\eta = 0.65$, the results are also remarkably similar to the baseline both with constant and time-varying policies. With the lower elasticity, $\eta = 0.35$ and $\rho_{sp} = 0.1\%$, they are also fairly similar. However, with $\eta = 0.35$ and $\rho_{sp} = 1\%$, the optimal constant policy is quite different. To understand the reason why, recall that in our baseline with $\rho_{sp} = 1\%$, the optimal policy involves positive research effort directed both towards dirty and clean technologies. When the elasticity of innovation to R&D declines, the social planner, restricted to a constant policy and with a reasonably high discount rate, finds it optimal to induce a slow switch to clean technology, which can be achieved with just a carbon tax. This is partly an artifact of constant policies; Figure 12 shows that optimal time-varying policies still heavily rely on research subsidies in this case. We thus conclude that the main message from our baseline results continue to apply with a fairly wideband of elasticities.

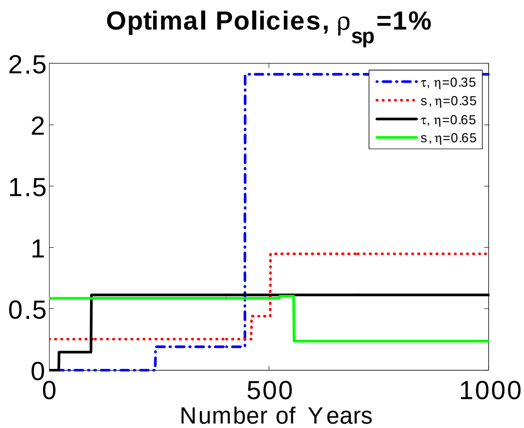

FIGURE 12A

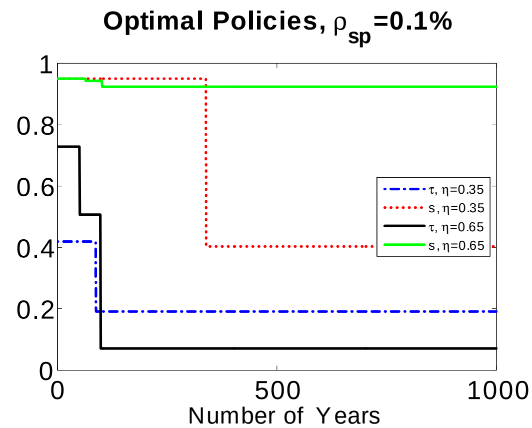

FIGURE 12B

3.6.4 Alternative Leapfrogging Probabilities α

We also investigate the implications of different values of α, in particular, focusing on a lower and higher estimate of α ($\alpha = 0.03$ and $\alpha = 0.05$). The results reported in Table 13 and Figure 13 are quite similar to the baseline results both quantitatively and qualitatively. In particular, optimal constant policies are in the ballpark of our baseline with $\alpha = 0.04$, and optimal time-varying policies have the same backloading and frontloading properties and similar values, though the exact switch points do differ.

TABLE 13. OPTIMAL CONSTANT POLICY RATES

	$\alpha = 0.03$			$\alpha = 0.05$	
	$\rho_{sp} = 1\%$	$\rho_{sp} = 0.1\%$		$\rho_{sp} = 1\%$	$\rho_{sp} = 0.1\%$
τ	24%	54%	τ	16%	44%
s	95%	95%	s	32%	95%

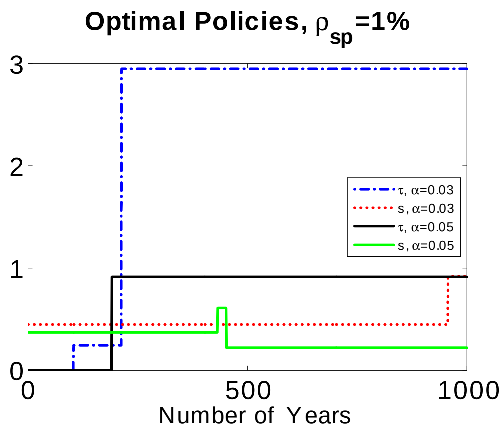

FIGURE 13A

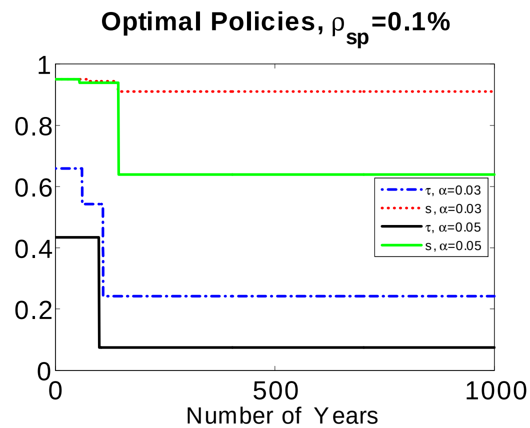

FIGURE 13B

3.6.5 Alternative Initial Technology Distribution

Finally, we also considered an initial technology gap distribution defined with several modifications from our baseline. First, rather than just sum patent counts, we weight patents by the normalized citation counts the patent receives. Second, we consider four-digit industries rather than three-digit industries. And third, we only consider industries within the manufacturing and energy sectors. Using these criteria, there are 332 SIC4 industries that are of sufficient size in terms of innovative firm counts to pass Census Bureau disclosure restrictions. Among these industries, 9.4% are led by the clean-energy stock. Table 14 summarizes some moments of this distribution:

TABLE 14. INITIAL CONDITION DISTRIBUTIONS SIC4

Metric:	**Clean Energy**	**Dirty Energy**
Mean Patent Total	140	663
Standard Deviation	401	1242
Share: [0,20]	53%	2%
Share: [21,100]	23%	18%
Share: [101,500]	17%	48%
Share: [500+]	6%	33%

The average gap to the frontier for dirty-patents stocks in the 9% of cases where clean patents have the lead is 463 patents, or in relative terms, 33% of the total patenting in that line to date. The average gap to the frontier for clean-patent stocks in the 91% of cases where dirty patents have the lead is 624 patents and 82% in relative terms. The conversion factor in this case is $12.6/0.161$. The distribution graph has a broadly similar shape as Figure 3 and we omit it to save space. The fraction of product lines with a non-zero gap in terms of step sizes is 82%. Clean energy leads by one or more step sizes in 7% of cases. Dirty energy has a lead of 20 and 50 steps sizes or more in 8% and 2% of technologies, respectively.

TABLE 15. OPTIMAL CONSTANT POLICY RATES

	$\rho_{sp} = 1\%$	$\rho_{sp} = 0.1\%$
τ	16%	54%
s	74%	95%

Using this alternative distribution of initial technology gaps has fairly limited impact on optimal constant and time-varying policies, which are shown in Table 15 and Figure

14. Both optimal constant and time-varying policies are remarkably similar to the baseline, making us conclude that our qualitative and even quantitative results are fairly robust to reasonable variations in the initial distribution of technology gaps.

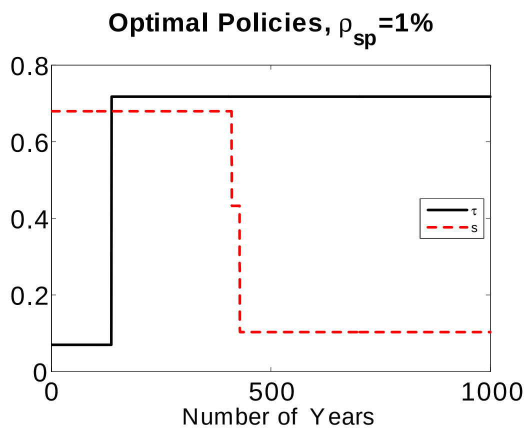

FIGURE 14A

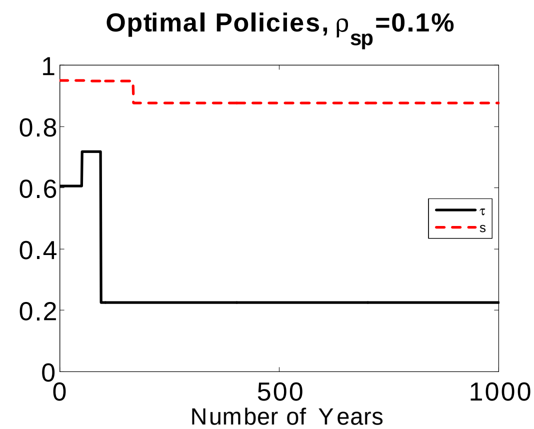

FIGURE 14B

3.7 Conclusion

One of the central challenges facing the world economy is reducing carbon emissions, which appears to be feasible only if a successful transition to clean technology can be induced. This paper has investigated the nature of a transition to clean technology theoretically and empirically. We developed a microeconomic model where clean and dirty technologies compete in production and innovation. If dirty technologies are more advanced to start with, the potential transition to clean technology can be difficult both because clean research must climb several steps to catch up with dirty technology and because this gap discourages research effort directed towards clean technologies. We characterized several properties of the equilibrium in this model and then estimated its key parameters from microdata on production, employment, R&D, patents and entry and exit of firms in the US energy sector, using regression analysis and simulated method of moments. Our model performs fairly well in matching a range of patterns in the data that were not directly targeted in the estimation, giving us confidence that it is potentially useful for the analysis of the transition to clean technology in the US energy sector.

Theoretically, carbon taxes and research subsidies encourage production and innovation in clean technologies. The key question we investigate using our estimated quantitative model is whether optimal policy will indeed secure a transition to clean technology, and if so how rapidly, and whether it will do so using carbon taxes or a combination of carbon taxes and research subsidies. A naïve intuition would be that only carbon taxes should be used because externalities are created by carbon (in the absence of these carbon externalities, the social planner would have no reason to interfere with or subsidize research).

In contrast to this intuition, we find that optimal policy heavily relies on research subsidies, and this result is fairly robust across a range of variations and for different damages and social discount rates. We also use the model to evaluate the welfare consequences of a range of alternative policy structures. For example, just relying on carbon taxes or delaying intervention both have significant welfare costs.

Though, to the best of our knowledge, it is the first attempt to develop a microeconomic model of the transition to clean technology and to quantitatively characterize optimal policy in such a setup, our paper has inevitably left several questions unanswered and taken a number of shortcuts, all of which constitute interesting areas for future research and investigation. We list some of these we view as particularly important here:

1. Our damage function enabled us to abstract from emissions in the rest of the world. Though very convenient, this approach left out several interesting considerations. The first is the interaction between US and global emissions.

2. The second is the potential impact of US transition to clean technology on technology choices in the rest of the world. In particular, to the extent that there is such an impact,

optimal policy may be more aggressive (for reasons discussed in Acemoglu, Daron, Philippe Aghion, Leonardo Bursztyn, and David Hemous (2012)) because it might have the power to also induce a switch to clean technology in the rest of the world also.

3. These concerns naturally fit into another important topic: game-theoretic interactions in emissions and technology choice across several countries in the global economy (Harstad, Bard (2012), Dutta, Prajit K and Radner, Roy (2006)).

4. For reasons we have explained, we did not allow for nonlinear threshold effects in the impact of atmospheric carbon on economic efficiency. Such nonlinearities are likely to be important and their exact position might be uncertain. Incorporating such nonlinearities, together with an explicit approach to uncertainty along the lines of Weitzman, Martin (2009), would be an important area for future research. This would also necessitate a much more detailed investigation of the interactions between US emissions and the rest of the world.

5. Our optimal policies are characterized under the assumptions of commitment to the policy sequence by the social planner. In the absence of such commitment, there will be a time inconsistency problem. An obvious important next step is to characterize time-consistent optimal policy.

6. Another interesting area is to investigate the interactions between international trade, technology and emissions (see Hemous, David (2012)).

Appendix A

Technological Interdependence

A.1 Data Construction

Name Matching The following procedure is used to match firm entities by name from both the US patent data and Compustat balance sheet information

1. Remove non-corporate entities

2. Drop corporate name identifiers and common English words

3. Group and standardizing suspected acronyms

4. Construct a similarity score basic on token and positional information for each pair of names

5. Group names by a given cutoff similarity score.

As noted by Hall, Jaffe, and Trajtenberg (2001), weighting tokens by their frequency of appearance would enhance match. The fact that certain uncommon words (such as "Samsung") appear in so many patents may skew this process, so an iterative is needed.

A.2 Two Type Model

I now introduce heterogeneity amongst firms in the form of persistent types in order to explain trends regarding innovation, firm growth, and patent transfers between and amongst large/small and young/old firms. Much of the previous derivations carry through here. The production environment, in particular, is identical.

One major difference that arises when introducing firm-level heterogeneity is that the indeterminacy in the direction of transfer is partially broken. When a high-type and low-type firm have a sequential interaction (in either direction), the patents are ultimately operated entirely by the high-type firm. However, it is still the case that when two firms of common type interact, the direction must be chosen arbitrarily. Moving to a model with a continuum of firm types would entirely eliminate this indeterminacy, but at the cost of tractability.

Equilibrium There are two types of firms: high-intensity and low-intensity innovators, herein referred to as high and low type, that are differentiated by their R&D cost functions. High type firms transition to non-adopting firms at flow rate ν. Being a low type firm is an absorbing state. Denote a generic firm type with $i \in \{H, L\}$. There is a type-specific innovation cost function $c^i(\cdot)$ for external innovation. The internal innovation technology $d(\cdot)$ is the same across types. An entrant firm starts as type H with probability ζ^H and type L with probability $\zeta^L = 1 - \zeta^H$.

Successful innovation yields a present value $\bar{V}^i$. The per-product value of a firm is then

$$\begin{aligned}
\delta_F V^H(\lambda) - \dot{V}^H(\lambda) &= \widetilde{\pi}(\lambda) + \Omega_x^H + \Omega_z^H \\
&+ \alpha\tau p(\mathbb{E}V^H(\beta\lambda) - V^H(\lambda)) + (1-\alpha)\tau(0 - V^H(\lambda)) \\
&+ b(V_0^H - V^H(\lambda)) + \nu(V^L(\lambda) - V^H(\lambda)) \\
\delta_F V^L(\lambda) - \dot{V}^L(\lambda) &= \widetilde{\pi}(\lambda) + \Omega_x^L + \Omega_z^L \\
&+ \alpha\tau^L p(\mathbb{E}V^L(\beta\lambda) - V^L(\lambda)) + \alpha\tau^H p(\mathbb{E}V^H(\beta\lambda) - V^L(\lambda)) \\
&+ (1-\alpha)\tau(0 - V^L(\lambda)) + b(V_0^L - V^L(\lambda))
\end{aligned}$$

where

$$\begin{aligned}
(\delta_F + \tau)V_0^H - \dot{V}_0^H &= \Omega_x^H + \Omega_0^H + \nu(V_0^L - V_0^H) \\
(\delta_F + \tau)V_0^L - \dot{V}_0^L &= \Omega_x^L + \Omega_0^L
\end{aligned}$$

and

$$\begin{aligned}
\Omega_x^i &= \max_{x^i}\left\{-\widetilde{w}c^i(x^i) + x^i\bar{V}^i\right\} \\
\Omega_z^i(\lambda) &= \max_{z^i}\left\{-\widetilde{w}\lambda^{-1}d(z^i) + z^i(\mathbb{E}[V^i(\beta\lambda) - V^i(\lambda))\right\} \\
\Omega_0^i &= \max_{z_0^i}\left\{-\widetilde{w}d(z_0^i) + z_0^i(\mathbb{E}[V^i(\beta\lambda) - V_0^i)\right\}
\end{aligned}$$

As before, posit a linearly separable form

$$\begin{aligned}
V^H(\lambda) &= A^H - B^H\lambda^{-1} \\
V^L(\lambda) &= A^L - B^L\lambda^{-1}
\end{aligned}$$

We then find for the high type

$$
\begin{aligned}
(\delta_F + b + (1-\alpha)\tau)A^H - \dot{A}^H &= 1 + \Omega_x^H + bV_0^H + \nu\left(A^L - A^H\right) \\
(\delta_F + b + (1-\alpha)\tau)B^H - \dot{B}^H &= 1 - \Omega_z^H - \alpha\tau p/(1+\kappa^{-1})B^H + \nu(B^L - B^H)
\end{aligned}
$$

and for the low type

$$
\begin{aligned}
(\delta_F + b + (1-\alpha)\tau)A^L - \dot{A}^L &= 1 + \Omega_x^L + bV_0^L + \alpha\tau^H p(A^H - A^L) \\
(\delta_F + b + (1-\alpha)\tau)B^L - \dot{B}^L &= 1 - \Omega_z^L - \alpha\tau^L p/(1+\kappa^{-1})B^L - \alpha\tau^H p(B^L - B^H/\lambda)
\end{aligned}
$$

The option values of innovation be simplified to

$$
\begin{aligned}
\Omega_x^i &= \max_{x^i}\left\{-\widetilde{w}c^i(x^i) + x^i\bar{V}^i\right\} \\
\Omega_z^i &= \max_{z^i}\left\{-\widetilde{w}d(z^i) + z^i B^i/(1+\kappa^{-1})\right\} \\
\Omega_0^i &= \max_{z^i}\left\{-\widetilde{w}d(z^i) + z^i(A^i - B^i/(1+\kappa) - V_0^i)\right\}
\end{aligned}
$$

Now it can be verified that $B^H = B^L \equiv B$. As such we will also have $\Omega_z^H = \Omega_z^L \equiv \Omega_z$ and $z^H = z^H \equiv z$. As high type firms have a superior innovation technology, they will assume production in the case of sequential innovation. Between firms of a common type, it is ambiguous. The expected gain from innovation is given by

$$
\begin{aligned}
\bar{V}^H &= [(1-\alpha) + \alpha\mu_0]\,\mathbb{E}V^H(\beta) \\
&\quad + \alpha(1-p)\left[\int_1^\infty (\mathbb{E}V^H(\beta\lambda) - V^H(\lambda))d\mu_+^H(\lambda) + \int_1^\infty (\mathbb{E}V^H(\beta\lambda) - V^L(\lambda))d\mu_+^L(\lambda)\right] \\
\bar{V}^L &= [(1-\alpha) + \alpha\mu_0]\,\mathbb{E}V^L(\beta) \\
&\quad + \alpha(1-p)\left[\int_1^\infty (\mathbb{E}V^H(\beta\lambda) - V^L(\lambda))d\mu_+^H(\lambda) + \int_1^\infty (\mathbb{E}V^L(\beta\lambda) - V^L(\lambda))d\mu_+^L(\lambda)\right]
\end{aligned}
$$

These can then be simplified to

$$\bar{V}^H = [(1-\alpha) + \alpha\mu_0]\left(A^H - B^H/(1+\kappa)\right) + \alpha(1-p)\left[(\mu_+\Gamma_+ B^H/(1+\kappa^{-1}) + \mu_+^L(A^H - A^L)\right]$$
$$\bar{V}^L = [(1-\alpha) + \alpha\mu_0]\left(A^L - B^H/(1+\kappa)\right) + \alpha(1-p)\left[(\mu_+\Gamma_+ B^H/(1+\kappa^{-1}) + \mu_+^H(A^H - A^L)\right]$$

Let there be a mass L of researchers. The labor market clearing condition is

$$1 = \frac{\Gamma}{\tilde{w}} + (\mu^H + e\zeta^H)c(x^H) + (\mu^L + e\zeta^L)c(x^L) + \mu_0^H d(z_0^H) + \mu_0^L d(z_0^L) + \mu_+\Gamma_+ d(z)$$

Steady State Imposing stationarity of values and state space elements, the firm value functions simplify to

$$B = \frac{1}{\delta_F + b + (1-\alpha)\tau + \left(\alpha\tau p + z\left(\frac{\delta+\tau+b}{\delta+\tau}\right)(1-1/\eta)\right)/(1+\kappa^{-1})}$$

and

$$A^H = \frac{1 + \Omega_x^H + bV_0^H + \nu A^L}{\delta_F + b + (1-\alpha)\tau + \nu} \qquad A^L = \frac{1 + \Omega_x^L + bV_0^L + \alpha\tau^H p A^H}{\delta_F + b + (1-\alpha)\tau + \alpha\tau^H p}$$

The first part of Theorem 1 regarding the value of Γ_+ still holds in the environment, however the technological lead distribution is no longer tractable. As before the overall inverse technological lead is given by $\Gamma = \mu_0 + \mu_+\Gamma_+$.

The relevant product mass distributions have flow equations

$$\begin{aligned}
\dot{\mu}_0^H &= b\mu_+^H - (\tau + z_0^H + \nu)\mu_0^H \\
\dot{\mu}_0^L &= b\mu_+^L + \nu\mu_0^H - (\tau + z_0^L)\mu_0^L \\
\dot{\mu}_+^H &= \tau^H\mu_0 + z_0^H\mu_0^H + \tau^H\mu_+^L - (1-\alpha)\tau^L\mu_+^H - \nu\mu_+^H - b\mu_+^H \\
\dot{\mu}_+^L &= \tau^L\mu_0 + z_0^L\mu_0^L + (1-\alpha)\tau^L\mu_+^H - \tau^H\mu_+^L + \nu\mu_+^H - b\mu_+^L
\end{aligned}$$

The overall mass distributions by type are

$$\dot{\mu}^H = \tau^H(1-\mu^H) - \tau^L\mu_0^H - (1-\alpha)\tau^L\mu_+^H - \nu\mu^H$$
$$\dot{\mu}^L = \tau^L\mu_0^H - \tau^H\mu^L + (1-\alpha)\tau^L\mu_+^H + \nu\mu^H$$

A bit of algebra reveals that the mass of high type firms can be expressed as

$$\mu^H = \frac{\tau^H}{\tau^H + \tau^L\left[\frac{b+(1-\alpha)(\tau+z_0^H+\nu)}{b+\tau+z_0^H+\nu}\right] + \nu}$$

with $\mu^L = 1 - \mu^H$. The fractions of products that are expired conditional on type are then

$$\mu_0^H = \left[\frac{b}{b+\tau+z_0^H+\nu}\right]\mu^H$$
$$\mu_0^L = \left[\frac{b+\nu\mu_0^H}{b+\tau+z_0^L}\right]\mu^L$$

The average inverse technological lead for patent protected product lines resolves to

$$\Gamma_+ = \frac{(1-\alpha)\tau + (\tau+z)\mu_0}{(1-\alpha)\tau + b + \kappa(\tau+z+b)}$$

These can be used to determine the labor market clearing condition.

For simulations, we also need the conditional step distributions. For high type

$$\dot{\mu}_0^H = b\mu^H - (b+\tau+z_0^H+\nu)\mu_0^H$$
$$\dot{\mu}_1^H = (1-\alpha)\tau^H + \alpha\tau^H\mu_0 + z_0^H\mu_0^H - (b+\tau+z+\nu)\mu_1^H$$
$$\dot{\mu}_n^H = \alpha\tau^H\mu_{n-1} + (\alpha\tau^L+z)\mu_{n-1}^H - (b+\tau+z+\nu)\mu_n^H$$

and for low type

$$\dot{\mu}_0^L = \nu\mu_0^H + b\mu^L - (b + \tau + z_0^L)\mu_0^L$$
$$\dot{\mu}_1^L = \nu\mu_1^H + (1-\alpha)\tau^L + \tau^L\mu_0 + z_0^L\mu_0^L - (b + \tau + z)\mu_1^L$$
$$\dot{\mu}_n^L = \nu\mu_n^H + (\alpha\tau^L + z)\mu_{n-1}^L - (b + \tau + z)\mu_n^L$$

These can be solved for iteratively with foreknowledge of μ^H and μ^L.

A.3 Algorithm

Equilibrium The basic models without industry heterogeneity can be solved in a straightforward fashion by setting up systems of equations consisting of first order conditions and the labor market clearing condition. These will depend on the aggregate innovation rates, so as to allow for the direct computation of product distributions, and the wage rate.

Moving to a setting with industry heterogeneity, the algorithm can be split into two levels. First, a wage rate is proposed, then each industry equilibrium is solved individually as in the basic model. Finally, these solution vectors are aggregated to evaluate the overall labor market clearing condition and the aggregate growth rate. Solving this system for the wage rate and growth rate constitutes solving the equilibrium in its entirety. I use the Powell's hybrid method[1] described in Powell (1970) to solve equations at both the industry level and overall equilibrium level. For more information on solving nonlinear systems, see Zangwill and Garcia (1981).

Simulations Simulating firms in an efficient manner involves a small amount of further derivation. In particular, when a firm undertakes a successful innovation, with probability α it must purchase rights to existing technology from the incumbent, assuming said incum-

[1]The exact code used is from the MINPACK library through the Python wrapper in SciPy.

bent's product has patent protection. In this case, if the innovating firm assumes production (i.e., with probability q is the untyped case), the step size of the product they receive can be drawn from the steady state distribution. In addition there will be one further innovation on top of that.

Interestingly, since the realization of the step size value for a given patent does not affect future patenting dynamics, these two factors will be independent in steady state as well. Assuming step increments are Pareto distributed, the distribution of the technological lead for a product line with n patents will be

$$\log(\lambda) \sim \text{Erlang}(n, 1/\kappa)$$

as the log of a Pareto random variable is exponentially distributed and the sum of i.i.d. exponentials random variables is Erlang distributed. Note that the Erlang distribution is simply the Gamma distribution with an integer curvature parameter.

Social Planner's Problem The social planner's problem in the model with industry heterogeneity can be simplified in a manner similar to that of cost minimization techniques when dealing with multi-product firm problems. In the general setting, a social planner must choose vector $\vec{x}$ of consisting of aggregate innovation rates for each industry, i.e., $\vec{x} = (\tau_1^H, \tau_1^L, z_1, \ldots, \tau_M^H, \tau_M^L, z_M)$. Use the notation $x_{i1} = \tau_i^H$, $x_{i2} = \tau_i^L$, and $x_{i3} = z_i$. Define the following maximization problem for each industry i

$$\Delta_i(\lambda_L, \lambda_g) = \max_{\vec{x}_i} \{-\Delta_i(\vec{x}_i) - \lambda_L L_R(\vec{x}_i) + \lambda_g g(\vec{x}_i)\}$$

where λ_L and λ_g represent the aggregate shadow values of labor and growth. Let the maximands of the above be denoted $\vec{x}_i(\lambda_L, \lambda_g)$. Finally, let $\vec{x}(\lambda_L, \lambda_g) = [\vec{x}_i(\lambda_L, \lambda_g)]_{i=1}^M$

and $\Delta(\lambda_L, \lambda_g) = \prod_{i=1}^{M} \Delta_i(\lambda_L, \lambda_g)$. Define the aggregate maximization

$$S = \max_{\lambda_L, \lambda_g} \left\{ [\delta + (\sigma - 1) g(\vec{x}(\lambda_L, \lambda_g))] \left(\frac{1 - L_R(\vec{x}(\lambda_L, \lambda_g))}{\Delta(\lambda_L, \lambda_g)} \right)^{\sigma - 1} \right\}$$

It can be shown that any maximizer of the above problem is socially optimal. The above formulation has the advantage of having computational complexity that scales linearly with the number of industries, rather than quadratically.

A.4 Proofs

Proof of Theorem 2. It is now necessary to specify a form for the R&D cost function. As is common in the existing literature, I use a constant elasticity function with exponent η

$$c(x) = cx^{\eta}$$

Using the first order condition for innovation intensity, we can find an expression relating the option value of innovation and the expected return from successful innovation

$$\Omega = (1 - 1/\eta) x \bar{V} = \left(\frac{1 - 1/\eta}{1 + e} \right) \tau \bar{V}$$

Then we can construct an equation characterizing the relationship between $\bar{V}$ and τ

$$\bar{V} = \frac{q_0 \left[\left(\frac{1}{\delta + b + (1-\alpha)\tau} \right) - \left(\frac{1}{1+\kappa} \right) B \right] + \alpha(1-p) \left(\Gamma - \mu_0 \right) \left(\frac{\kappa}{1+\kappa} \right) B}{1 - q_0 \left(\frac{1 - 1/\eta}{1+e} \right) \left(\frac{\tau}{\delta + \tau} \right) \left(\frac{\delta + b + \tau}{\delta + b + (1-\alpha)\tau} \right)}$$

where $q_0 = (1 - \alpha) + \alpha \mu_0$ is the probability of not having to purchase patent rights from the existing incumbent upon a successful innovation. This expression can be shown to be well-defined and positive for any $\tau \geqslant 0$. Call this function $\bar{V}(\tau)$. The first order condition for innovation be rearranged to

$$L_R = c(1+e)^{1-\eta}\tau^\eta = \frac{\tau\bar{V}}{\eta\Gamma + \tau\bar{V}}$$

Notice the value on the left is the share of resources devoted to research and the value on the right is strictly less than one, so we are guaranteed to find such a τ satisfying the above equation that results in a mixture of production and research labor. Showing uniqueness, which can be assured with the concavity of $\bar{V}(\cdot)$, is a more difficult matter. However, this can be easily verified for particular sets of parameters.

Proof of Theorem 3. Using the expression for $\log(\Delta)$ in Equation 1.8, we can derive

$$\begin{aligned}\frac{\partial \log(\Delta)}{\partial \tau} = &\left(\frac{1}{b(1+m)+\tau+z_0}\right) - \left(\frac{\tau+z_0}{b(1+m)+\tau+z_0}\right)\left(\frac{1}{1+m}\right)\frac{\partial m}{\partial \tau} \\ &- \frac{1}{b+\tau+z_0} + \frac{mb}{(b+\tau+z_0)^2} + \left(\frac{\tau+z_0}{b+\tau+z_0}\right)\frac{\partial m}{\partial \tau}\end{aligned}$$

Equation 1.7 implies that m increases with both τ and α. Further, we can derive

$$\frac{\partial m}{\partial \tau} = \kappa\left[\frac{\alpha b - (1-\alpha)z}{(b+(1-\alpha)\tau)^2}\right]$$

which implies

$$\frac{\partial}{\partial \alpha}\left[\frac{\partial m}{\partial \tau}\right] > 0$$

The above expression simplifies to

$$\begin{aligned}\frac{\partial \log(\Delta)}{\partial \tau} =& \frac{m}{(b+\tau+z_0)^2}\left(\frac{bm}{bm+b+\tau+z_0}\right) \\ &+ \frac{\partial m}{\partial \tau}\left(\frac{\tau+z_0}{b+\tau+z_0} - \left(\frac{1}{1+m}\right)\left(\frac{\tau+z_0}{b(1+m)+\tau+z_0}\right)\right)\end{aligned}$$

which can be seen to be increasing in α. As τ and z_0 enter into the above expression in the exact same manner, the same logic applies for z_0 as well. However, z_0 does not affect m, the resulting expression is simply

$$\frac{\partial \log(\Delta)}{\partial z_0} = \frac{m}{(b+\tau+z_0)^2}\left(\frac{bm}{bm+b+\tau+z_0}\right)$$

Conversely, z affects m positively but does not change the composition between expired and unexpired product lines, meaning we find

$$\frac{\partial \log(\Delta)}{\partial z} = \frac{\partial m}{\partial z}\left(\frac{\tau+z_0}{b+\tau+z_0} - \left(\frac{1}{1+m}\right)\left(\frac{\tau+z_0}{b(1+m)+\tau+z_0}\right)\right)$$

Varying z and z_0 simultaneously would yield an expression equivalent to that for τ, where both terms are present.

Appendix B

Back to Basics

B.1 Theoretical Proofs

As the downstream production technology is unchanged in the generalized model and we continue to impose symmetry across the industries. This implies that

$$P_i = P = \frac{1}{M} \qquad \text{and} \qquad Y_i = Y = Z. \tag{B.1}$$

Henceforth, we can drop the industry index i. The perfectly competitive firm that produces midstream good Y_i takes equilibrium prices P and p_j as given while maximizing its profit

$$\max_{y_j} \left\{ P \left[\int_0^1 y_j^{\frac{\varepsilon-1}{\varepsilon}} dj \right]^{\frac{\varepsilon}{\varepsilon-1}} - \int_0^1 p_j y_j dj \right\}.$$

This maximization leads to the following inverse demand for upstream good j

$$p_j = P \left(\frac{Y}{y_j} \right)^{\frac{1}{\varepsilon}}.$$

Monopolist in product line j, j has productivity q_j. The firm takes the demand function for its product as given and solves the following maximization problem

$$\pi_j = \max_{y_j} \left\{ PY^{\frac{1}{\varepsilon}} y_j^{\frac{\varepsilon-1}{\varepsilon}} - \frac{w}{q_j} y_j \right\}$$

This delivers the following optimal quantity

$$y_j = \left[\frac{1}{M} \left(\frac{\varepsilon - 1}{\varepsilon} \right) \left(\frac{q_j}{w} \right) \right]^{\varepsilon} Z$$

Plugging this into the production function for midstream goods, we find a relationship between wage w and aggregated productivity $\bar{q} \equiv \left(\int q_j^{\varepsilon-1} dj \right)^{\frac{1}{\varepsilon-1}}$

$$w = \frac{1}{M} \left(\frac{\varepsilon - 1}{\varepsilon} \right) \bar{q} \tag{B.2}$$

With this, we can greatly simplify the expression of the firm's quantity and price choices as a function of its normalized productivity $\widehat{q}_j = q_j / \bar{q}$

$$y_j = \widehat{q}_j^{\varepsilon} Z \qquad \text{and} \qquad p_j = \frac{1}{M \widehat{q}_j}$$

Denote variables normalized by Z/M with a "$\sim$". Then the normalized profit and labor are given by

$$\widetilde{\pi}_j = \frac{\widehat{q}_j^{\varepsilon-1}}{\varepsilon} \qquad \text{and} \qquad l_j = \frac{\widehat{q}_j^{\varepsilon-1}}{\widetilde{w}} \left(\frac{\varepsilon - 1}{\varepsilon} \right). \tag{B.3}$$

where $\widetilde{w}$ is the normalized wage. Note that by construction $\int \widehat{q}_j^{\varepsilon-1} dj = 1$. As a result, we integrate B.3 over j to find profit share and production labor share as

$$\frac{M \int_0^1 \pi_j dj}{Z} = \frac{1}{\varepsilon} \qquad \text{and} \qquad \frac{w L_P}{Z} = \frac{\varepsilon - 1}{\varepsilon}. \tag{B.4}$$

Finally, we combine B.2 and B.4 to find the final output as a function of aggregate productivity $\bar{q}$ and total production labor L_P :

$$Z = \bar{q}L_P/M.$$

Proof of Lemma 4 Let $\mathcal{F}_H(\cdot,t)$ and $\mathcal{F}_C(\cdot,t)$ be the aggregate product cumulative measures by type (hot or hold) at time t. For a small time step Δ, hot distribution $\mathcal{F}_H(\cdot,t)$ will satisfy

$$\begin{aligned}\mathcal{F}_H(\hat{q},t+\Delta) =& \mathcal{F}_H(\hat{q}/(1+\Delta g),t) - \Delta\tau\left[\mathcal{F}_H(\hat{q}/(1+\Delta g),t) - \mathcal{F}_H(\hat{q}/(1+\Delta g)-\eta,t)\right] \\ &+ \Delta\tau_b^e\mathcal{F}_C(\hat{q}/(1+\Delta g)-\eta,t) - \Delta\zeta\mathcal{F}_H(\hat{q}/(1+\Delta g),t) + \Delta\tau_b^d\mathcal{F}_C(\hat{q}/(1+\Delta g),t)\end{aligned}$$

Similarly, the cold distribution $\mathcal{F}_C(\cdot,t)$ will satisfy

$$\begin{aligned}\mathcal{F}_C(\hat{q},t+\Delta) =& \mathcal{F}_C(\hat{q}/(1+\Delta g),t) - \Delta\tau_a\left[\mathcal{F}_C(\hat{q}/(1+\Delta g),t) - \mathcal{F}_C(\hat{q}/(1+\Delta g)-\lambda,t)\right] \\ &- \Delta\tau_b\mathcal{F}_C(\hat{q}/(1+\Delta g),t) + \Delta\zeta\mathcal{F}_H(\hat{q}/(1+\Delta g),t)\end{aligned}$$

Finally, for $i \in \{H, C\}$, calculating

$$\dot{\mathcal{F}}_i(\hat{q}) = \frac{\mathcal{F}_i(\hat{q},t+\Delta) - \mathcal{F}_i(\hat{q},t)}{\Delta}$$

and taking the limit as $\Delta \to 0$ yields the desired flow equations. Note that for this we use

$$\frac{\mathcal{F}_i(\hat{q}/(1+\Delta g),t) - \mathcal{F}_i(\hat{q},t)}{\Delta} = -g\hat{q}[\partial\mathcal{F}_i(\hat{q})/\partial\hat{q}]$$

Proof of Proposition 5. Let $\mathcal{F}(\cdot,t)$ be the distribution over q at time t. Similarly, let $\mathcal{F}_H(\cdot,t)$ and $\mathcal{F}_C(\cdot,t)$ be the product type (hot/cold) conditional distributions. Thus, we will have $\mathcal{F}(q,t) = \alpha\mathcal{F}_H(q,t) + (1-\alpha)\mathcal{F}_C(q,t)$. The evolution of the aggregated productivity index

$\bar{q}$ is then given by

$$
\begin{aligned}
\bar{q}^{\varepsilon-1}(t+\Delta t) =& \int_0^\infty q^{\varepsilon-1} d\mathcal{F}(q, t+\Delta t) \\
=&\alpha \int_0^\infty q^{\varepsilon-1} d\mathcal{F}_H(q, t+\Delta t) + (1-\alpha)\int_0^\infty q^{\varepsilon-1} d\mathcal{F}_C(q, t+\Delta t) \\
=&\alpha \int_0^\infty \left[\Delta\tau\left(q+\eta\bar{q}\right)^{\varepsilon-1} + (1-\Delta\tau)q^{\varepsilon-1}\right] d\mathcal{F}_H(q,t) \\
&+ (1-\alpha)\int_0^\infty \left[\Delta\tau_a\left(q+\lambda\bar{q}\right)^{\varepsilon-1} + \Delta\tau_b^e\left(q+\eta\bar{q}\right)^{\varepsilon-1} + (1-\Delta\tau)q^{\varepsilon-1}\right] d\mathcal{F}_C(q,t)
\end{aligned}
$$

Thus the differential is

$$
\begin{aligned}
\frac{\bar{q}^{\varepsilon-1}(t+\Delta t) - \bar{q}^{\varepsilon-1}(t)}{\Delta} =&\alpha \int_0^\infty \tau\left[\left(q+\eta\bar{q}\right)^{\varepsilon-1} - q^{\varepsilon-1}\right] d\mathcal{F}_H(q,t) \\
&+ (1-\alpha)\int_0^\infty \left(\tau_a\left[\left(q+\lambda\right)^{\varepsilon-1} - q^{\varepsilon-1}\right] + \tau_b^e\left[\left(q+\eta\right)^{\varepsilon-1} - q^{\varepsilon-1}\right]\right) d\mathcal{F}_C(q,t)
\end{aligned}
$$

and the normalized differential is

$$
\begin{aligned}
\frac{\bar{q}^{\varepsilon-1}(t+\Delta t) - \bar{q}^{\varepsilon-1}(t)}{\Delta\bar{q}^{\varepsilon-1}(t)} =&\alpha \int_0^\infty \tau\left[\left(\hat{q}+\eta\right)^{\varepsilon-1} - \hat{q}^{\varepsilon-1}\right] d\mathcal{F}_H(\hat{q},t) \\
&+ (1-\alpha)\int_0^\infty \left(\tau_a\left[\left(\hat{q}+\lambda\right)^{\varepsilon-1} - \hat{q}^{\varepsilon-1}\right] + \tau_b^e\left[\left(\hat{q}+\eta\right)^{\varepsilon-1} - \hat{q}^{\varepsilon-1}\right]\right) d\mathcal{F}_C(\hat{q},t)
\end{aligned}
$$

Finally, the growth can be expressed compactly as

$$
g = \frac{\alpha\tau\mathbb{E}_{\hat{q}}^H\left[\left(\hat{q}+\eta\right)^{\varepsilon-1} - \hat{q}^{\varepsilon-1}\right] + (1-\alpha)\left(\tau_a\mathbb{E}_{\hat{q}}^C\left[\left(\hat{q}+\lambda\right)^{\varepsilon-1} - \hat{q}^{\varepsilon-1}\right] + \tau_b^e\mathbb{E}_{\hat{q}}^C\left[\left(\hat{q}+\eta\right)^{\varepsilon-1} - \hat{q}^{\varepsilon-1}\right]\right)}{\varepsilon-1}
$$

This can also be rearranged into

$$
g = \frac{\tau_a\left(\alpha\mathbb{E}_{\hat{q}}^H\left[\left(\hat{q}+\eta\right)^{\varepsilon-1} - \hat{q}^{\varepsilon-1}\right] + (1-\alpha)\mathbb{E}_{\hat{q}}^C\left[\left(\hat{q}+\lambda\right)^{\varepsilon-1} - \hat{q}^{\varepsilon-1}\right]\right) + \tau_b^e\mathbb{E}_{\hat{q}}\left[\left(\hat{q}+\eta\right)^{\varepsilon-1} - \hat{q}^{\varepsilon-1}\right]}{\varepsilon-1}
$$

$\square$

Proof of Proposition 6. The firm value, in general form, can be expressed as

$$
\begin{aligned}
& r\mathcal{V}_t(\mathbf{q}, m) - \dot{\mathcal{V}}_t(\mathbf{q}, m) \\
= \quad & \max_{a,b} \left\{ \begin{array}{c}
\sum_{\widehat{q}\in\mathbf{q}} \frac{1}{\varepsilon}\widehat{q}^{\epsilon-1}\frac{Z_t}{M} - nw_t\left[h_a(a) + h_b(b) + \mathbf{1}_{(b>0)}\phi\right] \\
+na\left[\alpha\mathbb{E}^H_{\widehat{q}}\mathcal{V}_t\left(\mathbf{q}\cup\{\widehat{q}+\eta\}, m\right) + (1-\alpha)\,\mathbb{E}^C_{\widehat{q}}\mathcal{V}_t\left(\mathbf{q}\cup\{\widehat{q}+\lambda\}, m\right) - \mathcal{V}_t\left(\mathbf{q}, m\right)\right] \\
+nb\left(1+\rho_m\right)\left[\mathcal{V}_t\left(\mathbf{q}\cup\{\widehat{q}+\eta\}, m\right) - \mathcal{V}_t\left(\mathbf{q}, m\right)\right] \\
+\sum_{\widehat{q}\in\mathbf{q}}\tau\left[\sum_{\widehat{q}\in\mathbf{q}}\left[\mathcal{V}_t\left(\mathbf{q}\setminus\{\widehat{q}\}, m\right) - \mathcal{V}_t\left(\mathbf{q}, m\right)\right]\right] \\
+x\frac{m}{M}\left[\alpha\mathbb{E}^H_{\widehat{q}}\mathcal{V}_t\left(\mathbf{q}\cup\{\widehat{q}+\eta\}, m\right) + (1-\alpha)\,\mathbb{E}^C_{\widehat{q}}\mathcal{V}_t\left(\mathbf{q}\cup\{\widehat{q}+\lambda\}, m\right) - p'_m - \mathcal{V}_t\left(\mathbf{q}, m\right)\right] \\
+x\left(1-\frac{m}{M}\right)\left[\alpha\mathbb{E}^H_{\widehat{q}}\mathcal{V}_t\left(\mathbf{q}\cup\{\widehat{q}+\eta\}, m+1\right) + (1-\alpha)\,\mathbb{E}^C_{\widehat{q}}\mathcal{V}_t\left(\mathbf{q}\cup\{\widehat{q}+\lambda\}, m+1\right) - p_m - \mathcal{V}_t\left(\mathbf{q}, m\right)\right] \\
+n\kappa\left[\mathbb{E}_{\widehat{q}}\mathcal{V}_t\left(\mathbf{q}\cup\{\widehat{q}\}, m\right) - \mathcal{P} - \mathcal{V}_t\left(\mathbf{q}, m\right)\right] \\
+\kappa\left[n\mathcal{P} - \mathcal{V}_t\left(\mathbf{q}, m\right)\right]
\end{array} \right\}.
\end{aligned}
$$

Now, conjecture $\mathcal{V}_t(\mathbf{q}) = \frac{Z_t}{M}\left[\sum_{\widehat{q}\in\mathbf{q}} V(\widehat{q}_t) + nV_m\right]$. When we substitute the conjecture into the the above expression and using the prices

$$
\begin{aligned}
p_m &= V_{m+1} + \mathbb{E}_{\widehat{q},s}V\left(\widehat{q}_{t+\Delta t} + \widehat{s}\right) \\
p'_m &= V_m + \mathbb{E}_{\widehat{q},s}V\left(\widehat{q}_{t+\Delta t} + \widehat{s}\right)
\end{aligned}
$$

we find

$$
(r-g)\,V_m = \max_{a,b}\left\{ \begin{array}{c}
-\widetilde{w}\left[h_a(a) + h_b(b) + \mathbf{1}_{(b>0)}\phi\right] \\
+a\left[\alpha V^H + (1-\alpha)\,V^C + V_m\right] \\
+b\left(1+\rho_m\right)\left[V^H + V_m\right] \\
+x\left(1-\frac{m}{M}\right)\left[V_{m+1} - V_m\right] \\
-\tau V_m + \kappa\mathbb{E}_{\widehat{q}}V(\widehat{q}_t)
\end{array}\right\}.
$$

and

$$
V'(\widehat{q}_t)g\widehat{q} + \left[\tau+\kappa+r-g\right]V(\widehat{q}_t) = \frac{1}{\varepsilon}\widehat{q}^{\epsilon-1}.
$$

Note that the last expression is a differential equation as a function of $\widehat{q}$. Then

$$
V(\widehat{q}_t) = \frac{\widehat{q}_t^{\epsilon-1}}{\varepsilon\left[r+\tau+\kappa+g\left(\epsilon-2\right)\right]}.
$$

This completes the proof. $\square$

Derivation of Multi-industry Distribution $\Gamma_{m,n}$. We assume that when a firm loses its last product in a particular industry, it maintains a foothold there, in the sense that it still receives buy-out offers and can still directly use basic research relevant to that industry. When a firm loses all of its products or receives a destructive shock, it ceases to exist. We wish to find the joint distribution over the number of industries a firm is in and how many product lines it owns. For notational convenience, let us denote the basic research flow from m-industry firms by $\widehat{b}_m = \mathcal{B}(\phi_m)b_m$. Let us also denote the expansion rate of a firm into a new industry by e_m. Here the expansion rate comes purely from buy-out offers by entrants. So given a per firm buy-out offer rate of x, a firm in m industries will expand at rate

$$e_m = x\left(\frac{M-m}{M}\right) = \left(\frac{\varsigma E a_e}{F}\right)\left(\frac{M-m}{M}\right)$$

Then the flow equation for firms in m industries with n products is

$$\begin{array}{rcl}
\text{OUTFLOW} & & \text{INFLOW} \\
\left[\begin{array}{c} a_1+\widehat{b}_1+\tau+\kappa \\ +e_1+\kappa \end{array}\right]\Gamma_{1,1} & = & a_e + 2\tau\Gamma_{1,2} \\
\left[\begin{array}{c} a_m+\widehat{b}_m+\tau+\kappa \\ +e_m+\kappa \end{array}\right]\Gamma_{m,1} & = & 2\tau\Gamma_{m,2} + e_{m-1}\Gamma_{m-1,1} \text{ for } m \geqslant 2 \\
\left[\begin{array}{c} n\left(a_m+\widehat{b}_m+\tau+\kappa\right) \\ +e_m+\kappa \end{array}\right]\Gamma_{m,2} & = & \left\{\begin{array}{c} \left(a_m+\widehat{b}_m\left(1-\rho_m\right)+\kappa\right)\Gamma_{m,n-1} \\ +3\tau\Gamma_{m,n+1}+e_{m-1}\Gamma_{m-1,n} \end{array}\right\} \text{ for } m \geqslant 1 \\
\left[\begin{array}{c} n\left(a_m+\widehat{b}_m+\tau+\kappa\right) \\ +e_m+\kappa \end{array}\right]\Gamma_{m,n} & = & \left\{\begin{array}{c} (n-1)\left(a_m+\widehat{b}_m\left(1-\rho_m\right)+\kappa\right)\Gamma_{m,n-1} \\ +(n-2)\rho_m\widehat{b}_m\Gamma_{m,n-2} \\ +(n+1)\tau\Gamma_{m,n+1}+e_{m-1}\Gamma_{m-1,n} \end{array}\right\} \text{ for } n \geqslant 3,\ m \geqslant 1
\end{array}$$

where we use the convention $\Gamma_{m,-1} = \Gamma_{m,0} = 0$ and $e_0 = 0$. The first line equates the outflows from $(m = 1, n = 1)$ that happen once the firm loses its product at the rate $\tau + \kappa$, acquires a new product line at the rate κ, innovates a new good at the rate $a_1 + \widehat{b}_1$ on average or expands into a new industry at the rate e_1. On the other hand, inflow happens from outsiders at the rate a_e and from the firms with 2 products that lose one of their products at the rate 2τ. Similar reasoning applies to the subsequent lines.

Using values for the $\Gamma_{m,n}$ distribution gives us the mass of firms in a given (m, n) state. The total mass of firms is then $F = \sum_{m=1}^{M} \sum_{n=1}^{\infty} \Gamma_{m,n}$. We ultimately want the mass of products in given industry state m. To get this we simply evaluate

$$\mu_m = \sum_{n=1}^{\infty} n \cdot \Gamma_{m,n}$$

$\square$

B.2 Data & Data Organization

Empirical investigation of the relationship between R&D investment and multi-market activity of a firm requires reliable and extensive information not only on product markets and on R&D characteristics of individual firms, but also on firm ownership status. The latter allows us to identify the product markets to which the firm is linked via its business group. We obtain this information from three different data sets.

R&D Information Information about R&D investment comes from the annual R&D Survey conducted by the French Ministry of Research. The R&D survey is available in annual waves of cross-sectional data, where the same firms are not necessarily sampled year after year (Mairesse and Mohnen (2010)). The survey covers a representative sample of French firms of more than 20 employees investing in R&D. However firms with less than .8 million euros of R&D investment fill out a shorter and simplified survey. The survey includes extensive information about the financing of R&D. It not only breaks down R&D investment according to the source of the funds but also provides its allocation to different types of

R&D. More specifically, all firms are asked to report their R&D investment as either basic or applied research.

Multi-Market Activity The identification of business group structures is based on a yearly survey by INSEE called "Enquete Liaisons Financieres" (LIFI). It covers all economic activities but restricts its attention to firms that either employ more than 500 employees or generate more than 60 million euros in revenue, or hold more than 1.2 million euros of traded shares. However, since 1998 the survey is cross-referenced with information from Bureau Van Dijk and thus covers almost the whole economy. The LIFI survey contains information that makes it a unique data set for studying the relationship between multi-market activity and investment in basic research. Besides providing information on direct financial links between firms, it also accounts for indirect stakes and cross-ownership when identifying the head of the group. This is important as it allows us to precisely reconstruct the group structure even in the presence of pyramids. This feature allows us to obtain a reliable account of the structure of business groups in the French economy and, as a consequence, reliable measures of our key variable, the multi-market presence of business groups.

Since each firm can be active in several markets, we cross-reference the data set with an extensive yearly survey by the Ministry of Industry ("Enquete Annuelle des Entreprises"). The survey is filled out by French firms with more than 20 workers and contains information not only on the different markets in which a firm operates but also information on market dedicated sales for each segment. The data cover the vast majority of French firms and span the period 2000-2006.

Balance-Sheet Information We use the firm- and industry-level data sets based on accounting data extracted again from the EAE files. The data also include unique firm identifiers allowing us to match them to the R&D and LIFI data.

Data Organization

We first identify the ownership status of each firm in the economy and the head of the group with which the firm is affiliated. Indeed, our data source (LIFI) defines a group as a set of firms controlled, directly or indirectly, by the same entity (the head of the group). We rely on a formal definition of control, requiring that a firm holds directly or through cross-ownership at least 50% of the voting rights in another firm's general assembly. We do not expect this to be a major source of bias in our sample as most French firms are private and ownership concentration is strong even among listed firms.[1] Firms that do not conform to this definition are classified as stand-alone firms.

We then match the ownership information to our balance-sheet data and to our survey on lines of business within firms. We drop firms that appear in the ownership data but for which we cannot find balance-sheet information. We also delete as outliers firm-year observations whose ROA falls outside a multiple of five of the interquartile range and firms that report 0 employment or which have negative sales. Based on our two sources of information we identify the main line of business from the balance sheets and the different segments of the firm from the survey on lines of business. For computational convenience we create a new firm-group identifier that allows us to aggregate at the same time business groups, business groups with multi-divisional firms, exclusively multi-divisional firms and true stand-alone firms. We then define four measures of multi-market activity. The first measure counts each market in which the firm-group is present either via its ownership links or its multi-

[1] In their overview of ownership structures and voting power in France, Bloch and Kremp (1999) show that ownership concentration is pervasive: for non-listed companies with more than 500 employees the main shareholder's ownership stake is 88%. The degree of ownership concentration is slightly lower for listed companies but still above 50% in most cases.

divisional structure. The second measure counts each market in which the firm-group is present with at least 9 employees via its ownership links or its multi-divisional structure. The third measure counts each market in which the firm-group is present exclusively via its ownership links. The final measure counts each market in which the firm-group is present exclusively via its ownership links and excluding financial activities.

We then define firm characteristics from balance-sheet data. There are three possible organizational types and comparison issues might arise. Taking the firm as the economic unit of interest has the advantage of simplicity since information is directly available in the balance sheets. However, this method has the disadvantage of not being comparable across organizational types. Indeed, most information for multi-divisional firms is aggregated across lines of segment, whereas firms belonging to business groups have market-specific information. Similar to existing studies by the Ministry of Research (Dhont-Peltrault and Pfister (2011)), we decided to aggregate the information to the economic unit at the highest level of control: the firm level for stand-alone and multi-divisional firms, and the business group level for firms affiliated through majority ownership.[2]

In a final step we match the firms' balance-sheet and patent information to information contained in the R&D Survey. We focus on firms for which we have R&D information. Again we aggregate at the highest level of control. As before, one has to be cautious in aggregating on the basis of variables that might be prone to double-counting. When constructing information on the basic R&D intensity of a firm this is not the case as we are focusing exclusively on "internal" research expenditures. Therefore, if a member of the group contracts out research with another member of the group, then one will be counted as "external" research expenditures and the other one as "internal" expenditures. To correct for outliers in the dependent variable, we drop firm-year observations whose basic research

[2]In addition to the economic rationale for constructing the data at the highest level of control there is also a legal argument. Indeed most public administrations and tribunals define the eligibility of firms for subsidy programs with respect to the business groups to which they belong.

intensity, conditional on positive basic research, falls outside a multiple of five of the interquartile range. In addition we exclude firm-year observations whose total R&D to sales ratio falls outside a multiple of five of the interquartile range.[3]

Policy Environment

It is useful to describe the policy environment in France during the period of our data. As shown in Figure 2.1, the share of GDP devoted R&D expenditures in France was on average 2.2% between 2000 and 2006. The public sector, including public universities, accounted for 0.8% of total research and development expenditures to GDP, while firms accounted for the remaining 1.4% of research and development expenditures. Innovation policy during the sample period featured a mix of measures to support R&D investment of firms through public financing. The main instrument to stimulate private innovation activity during that period consisted of approximately 2.5 billion euros of yearly subsidies allocated to firms either through ministries or government agencies such as OSEO-ANVAR. Note that our R&D survey allows us to directly measure this form of public financing in our sample. Finally, the R&D tax credit system was seen by the government as a secondary policy measure until a major reform in 2008 that increased the base and the rate of the the tax credit.

Variable List

All variables are organized and computed according to the method set out in the previous section. To summarize, we decided to aggregate the information to the economic unit at the highest level of control: the firm level for stand-alone and multi-divisional firms, and

[3] Alternatively, we exclude firm-year observations whose basic to applied R&D ratio falls above the 99^{th} percentile of the distribution. The results are qualitatively similar.

the business group level for firms affiliated through majority ownership. In the remainder of the document we will, for the sake of notational convenience, refer generically to firms.

- *Basic Research Intensity:* total basic research by firm i in year t divided by total applied research of firm i in year t. The formulation of the survey questions related to the type of research undertaken is directly derived from the definitions provided by the Frascati Manual;

- *# of Industries:* sum of all distinct SIC codes within firm i in year t irrespective of organizational form (business group or multi-divisional structure). Industries are successively defined at the 4-,3-,2- and 1-digit SIC levels;

- *# of Industries - Weighted Sum:* weighted sum of all distinct bilateral 1-digit SIC links within firm i in year t. Weights are computed on the basis of the empirical frequency of each bilateral SIC link in each year t;

- *# of Patent Classes Applied:* sum of cumulated distinct patent-class applications within firm i in year t. Cumulated patent-class applications are computed for the period leading from 1993 to year t. Patent classes are successively defined at the 5,4,3,2 and 1-digit levels (EPO Classification);

- *# of Patent Classes Granted:* sum of cumulated distinct patent-class grants within firm i in year t. Cumulated patent-class grants are computed for the period leading from 1993 to year t. Patent classes are successively defined at the 5-,4-,3-,2- and 1-digit levels (EPO Classification);

- *Financial Int.:* binary indicator equal to 1 if firm i in year t is present in a financial industry, 0 otherwise;

- *Foreign HQ:* binary indicator equal to 1 if the headquarters of firm i in year t are located outside France, 0 otherwise;

- *Market Share:* weighted average of total sales of firm i, year t in industry k divided by total industry sales in year t. Weights are computed on the basis of the industry share of employment within firm i in year t;

- *Outsourcing to Univ.:* binary indicator equal to 1 if firm i in year t has outsourced R&D to French universities, 0 otherwise;

- *Profitability - ROA:* weighted average of EBIDTA divided by total fixed assets of all subsidiaries within firm i in year t. Weights are computed on the basis of the subsidiaries' share of employment within firm i in year t;

- *Profitability - ROS:* weighted average of EBIDTA divided by total sales of all subsidiaries within firm i in year t. Weights are computed on the basis of the subsidiaries' share of employment within firm i in year t;

- *Public R&D Funds:* binary indicator equal to 1 if firm i in year t has received French public funds, 0 otherwise;

- *Research Area:* weighted average of the share of respectively biotech / software / environmental research in research expenditures in firm i in year t. Weights are computed on the basis of the subsidiaries' share of total R&D within firm i in year t;

- *Total Employment:* total employment of firm i in year t;

- *IV - State Present in 1986:* binary indicator equal to 1 if the French state had a non-zero equity stake in firm i in 1986;

- *IV - SOE in 1986:* binary indicator equal to 1 if the French state had a controlling equity stake in firm i in 1986.

Table B.1 provides the descriptive statistics of the key variables.

TABLE B.1: DESCRIPTIVE STATISTICS

Variable	Mean	25th Percentile	Median	75th Percentile	Standard Deviation	Min	Max	N
R&D Investment								
R&D To Sales	0.11	0.01	0.04	0.14	0.17	0.00	0.86	13708
Basic Research Intensity	0.06	0.00	0.00	0.02	0.19	0.00	1.57	13708
Number of Industries								
1-Digit SIC	2.21	1	2	3	1.48	1	10	13708
4-Digit SIC	4.97	1	2	5	8.87	1	130	13708
Balance Sheet								
Total Employment	1497.88	24	93	506	8445.93	1	195746	13708
Return on Sales	0.032	0.02	0.07	0.13	0.63	-39.39	7.36	13708
Age	21.17	8.79	18.92	30.55	14.97	0	86	13708
Ownership Structure								
Financial Intermediary	0.05	0	0	0	0.22	0	1	13708
Foreign HQ	0.23	0	0	0	0.41	0	1	13708
Public and Private R&D								
Public Subsidy to Private Investment	0.09	0	0	0.04	0.4	0	30.9	13708
Collaboration with Universities	.15	0	0	0	0.36	0	1	13708

Note: Pooled data for the period 2000-2006. *R&D To Sales* is defined as the ratio of total firm research and development expernditures to total firm sales. *Basic Research Intensity* is defined as the ratio of total firm investment in basic research to total firm investment in applied research. *Number of Industries* is the sum of all distinct SIC codes within the firm. *Return on Sales* is the ratio of earnings before interest, taxes, depreciation and amortization to total firm sales. *Total Employment* total employment of the firm. *Age* is the difference between the current year and the year of the firm's incorporation. *Financial Intermediary* binary indicator equal to 1 if the firm is present in a financial industry, 0 otherwise. *Foreign HQ:* binary indicator equal to 1 if the headquarters of the firm are located outside France, 0 otherwise. *Public Subsidy to Private Investment* binary indicator equal to 1 if the firm has received French public funds for innovation expenditures, 0 otherwise. *Collaboration with Universities* binary indicator equal to 1 if the firm has received French public funds for innovation expenditures, 0 otherwise.

B.3 Construction of Within-Industry Spillovers

Figure B.1 provides a graphical intuition for the computation of the citation information.

FIGURE B.1: COMPUTING THE COOL-DOWN RATE

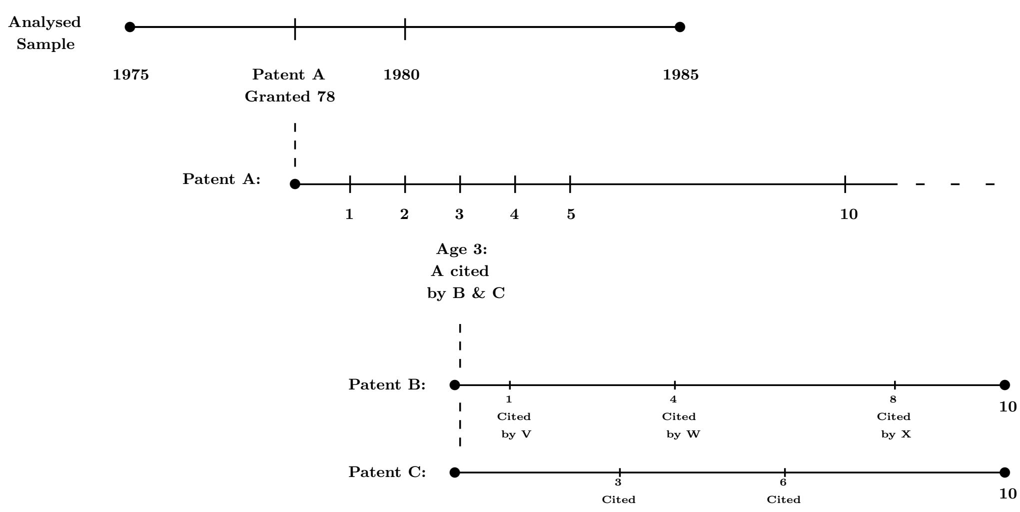

Patent A is granted in 1978. In 1981, when patent A is 3 years old, it receives citations from both patent B and patent C, which was applied for in 1981. Patent B in the following 10 years was cited by patents V, W and X, whereas patent C was only cited by patents Y and Z. The average citation of citing patents for patent A at age 3 is therefore 2.5. The timing of the computation implies that we need to be cautious with respect to possible truncation. We therefore compute our measure for patents between 1975 and 1985. This implies that, inclusive of the 10-years-forward lag, we can observe without truncation all our patents until the age of 10.

Robustness Checks Figure B.2 provides robustness checks for the estimates on the cool-down rate of patents originating from basic and applied research. The left panel of the figure measures *Average Citations of Citing Patents* computing the 5-years-forward citations of the citing patents and is measured for patents granted in the period 1975-1985. The right panel re-classifies university patents that were defined as private depositors. In both cases results are unchanged, with a citation difference between public and private patents that becomes statistically non-significant at year 8. Indeed, in France, most of the academic patents are accounted for in the "public" category. French universities generally manage their patents through public research institutions with which academics are typically affiliated, one example being the CNRS.

B.4 Robustness Checks on Reduced-Form Results

In this section we provide further robustness checks on the correlation between a firm's basic research incentives and its multi-industry presence. Our baseline specification is related to the number of distinct 1-digit SIC activities in which a firm operates but extends to finer SIC classifications. All results are presented in Table B.2.

FIGURE B.2: CITATION PATTERNS FOR FRENCH PUBLIC AND PRIVATE PATENTS

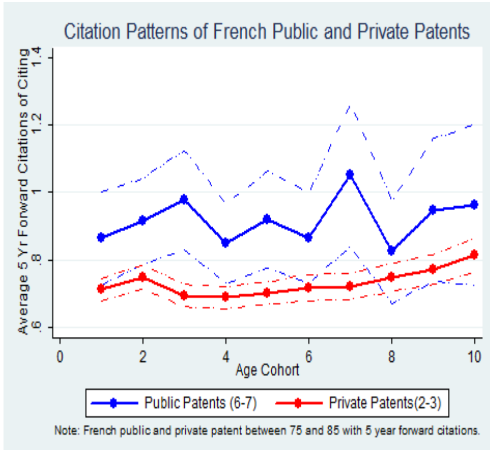

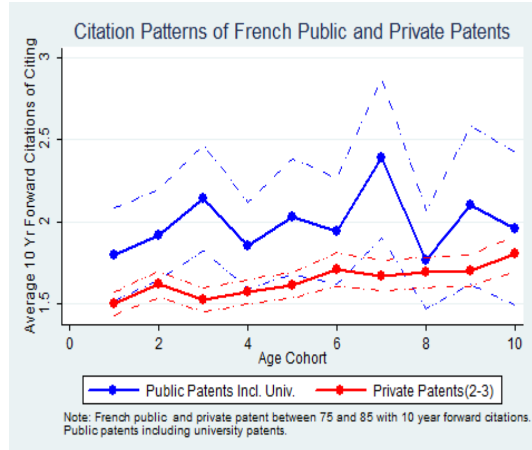

CITATION DIFFERENCES FOR FRENCH PUBLIC AND PRIVATE PATENTS

Age	1	2	3	4	5	6	7	8	9	10
5-Yr-Forward Citations	.15**	.16**	.28***	.16**	.22**	.15**	.33***	.08	.18	.15
	(0.07)	(0.07)	(0.08)	(0.06)	(0.07)	(0.07)	(0.11)	(0.08)	(0.11)	(0.12)
10-Yr-Forward Citations	.3**	.3**	.62***	.28**	.42**	.23	.71***	.08	.39	.15
Including Univ.	(0.15)	(0.15)	(0.17)	(0.14)	(0.18)	(0.17)	(0.25)	(0.16)	(0.25)	(0.24)

*Note:*The figures separately plot *Average Citations of Citing Patents* for French public patents (blue line) and French private patents (red line) across patent age. The left panel computes *Average Citations of Citing Patents* computing the 5-years-forward citations of the citing patents and is measured for patents granted in the period 1975-1985. The bottom panel computes *Average Citations of Citing Patents* computing the 5-years-forward citations of the citing patents and re-classifying university patents as public patents. The table reports differences in citation patterns using two sample t-tests with unequal variances. One star denotes significance at the 10% level, two stars denote significance at the 5% level, and three stars denote significance at the 1% level

Confounding Factors Columns (1) and (2) check robustness of the results with respect to confounding factors. Column (1) estimates the model only allowing for year and organization fixed effects, whereas column (2) includes a set of potential confounding factors. Results in column (1) suggest that presence in an additional industry, not accounting for other variables such as size, is associated on average with a 1.4 percentage-point higher basic research intensity of firms. In column (2) the set of regressors includes controls for size, profitability and headquarter localization. The impact of multi-industry presence is slightly lower but remains statistically significant.[4] Estimates on the localization of headquarters are also statistically significant at the 5% level. Total employment and profitability on the other hand are not.

Measures of Multi-Industry Presence Columns (3) and (4) provide alternative measures for the multi-industry presence of firms. Column (3) defines multi-industry presence on the basis of a firm's technological spectrum. To do so we use EPO patent data for French applicants. We define as the number of technology classes in which a firm is present as the cumulative distinct patent classes granted to the firm between 1993 and t. The coefficient is very similar in magnitude and precision to the one obtained using distinct 1-digit SIC industries. Column (4) measures multi-industry presence as a weighted sum of all distinct bilateral 1-digit SIC links within firm i in year t considering only distinct legal entities linked by majority ownership. Weights are computed on the basis of the empirical frequency of each bilateral SIC link in each year t. Intuitively, if a given bilateral industry link is rare, then industries are more likely to be very different. Multi-industry presence is still positively related to basic research intensity, the different point estimate being linked to the different support of the weighted industry variable.

Causality and Instrumental Variables Columns (5) and (6) address the potential concern of reverse causality, i.e. basic research leading to a larger economic scope of firms. We exploit historical ownership structures that affected a firm's multi-industry presence as instrumental variables. The two instruments are defined as *State Ownership 1985-1987* and *State Owned between 1985-1987*.

The rationale behind our identification strategy is as follows. In 1981 Francois Mitterrand was elected president of the Republic and implemented a vast nationalization pro-

[4] Further checks on control variables included market shares, R&D subsidies, collaborations with universities, the presence of financial intermediaries, state in the capital of the firm, industry fixed effects and the use of a mean patent scaling method.

gram across industries. Even before that period the tradition of French state intervention resulted in a significant fraction of the economy being under state control. Consistent with Colbertist policies, the state also modified the economic scope of its firms by merging unrelated firms into large conglomerates of national champions. In 1987, however, Jacques Chirac was elected prime minister on a liberal platform and this marked the beginning of privatizations, which continued into the 90s. The embedded exclusion restriction therefore requires that state control in the 80 be associated today with a greater basic research intensity of firms only because of politically motivated mergers. The implicit assumption is that when these firms became private they adjusted their research spending from the social to the private optimum but did not adjust their multi-industry presence. First-stage estimates show that state ownership in the 80s is associated on average with 1.2 more industry links for firms between 2000 and 2006. The associated F-test are well above the critical levels related to weak instruments tables.[5] Columns 5 and 6 present the instrumented LATE coefficients related to multi-industry presence of the second stage. The coefficients are nearly twice as large in magnitude with respect to the non-instrumented coefficients of columns 1 and 2.

Estimation Columns (7) and (8) use alternative estimation methods for the baseline model with covariates. Column (7) presents estimates of the Heckman selection model, whereas column (8) presents estimates from a negative binomial model. In both cases estimates suggest a positive and statistically significant relation between basic research intensity and multi-industry presence.

[5]The tables are available upon request.

B.5 Target Moments and Identification

In this section we explain the moments that are used to identify our parameters. For convenience, define expressions for the per product line R&D employment

$$h_a^m = \xi_b a_m^{\nu_b} \qquad \text{and} \qquad \bar{h}_b^m = \mathbb{E}_\phi \left[(\xi_b b_m^{\nu_b} + \phi) \cdot \mathbf{1}_{(\phi < \phi_m^*)} \right]$$

for applied and basic research, respectively. Note that these are functions of m, the number of industries a firm has working knowledge in.

Below, expectations are assumed to be over the distribution of firm characteristics $(m, n, \hat{\mathbf{q}})$. Note that here $\hat{m}$ denotes the number of industries in which a firm has one or more products, rather than the number of industries in which the firm has working knowledge. Since the latter is unobservable, we must compute the former to in order to match the data.

Basic Research Intensity by Number of Industries We define basic research intensity as the ratio of spending on basic research to spending on applied research. Since the effect of multi-industry presence on this quantity is of critical importance to our model, we have one moment for each $\hat{m} \in \{1, \ldots, M\}$. Given a set of parameters and an equilibrium of the model, this moment's value for a given $\hat{m}$ is

$$\Lambda(1-8) = \mathbb{E}_m \left[\frac{\bar{h}_b^m}{h_a^m} \middle| \hat{m} \right]$$

In our estimation, we use $M = 10$. However, in the data there are only a handful of firms with $\hat{m} > 8$, so we have one moment for each $\hat{m} \in \{1, \ldots, 7\}$ and a final moment which is averaged over $\hat{m} \in \{8, 9, 10\}$. The way in which this moment increases with $\hat{m}$ identifies the cross-industry spillover parameter p in our model. Additionally, it provides us with some identification power for the basic research cost parameters (ξ_b, ν_b) .

Extensive Margin of Basic Research Investment by Number of Industries We use the share of positive basic research spending by each $\widehat{m}$ to identify the mean μ and variance σ^2 of the fixed cost distribution basic research. This is simply the probability that the idiosyncratic fixed cost draw is less than the cutoff for a certain $\widehat{m}$

$$\Lambda(9-16) = \mathbb{E}_{m,\phi}\left[\mathbf{1}_{(\phi<\phi^*_m)}\middle|\,\widehat{m}\right].$$

Distribution of m We track two moments relating to the distribution of $\widehat{m}$, the mean and mean squared. They are given by

$$\Lambda(17) = \mathbf{E}_{\widehat{m}}\left[\widehat{m}\right] \qquad \text{and} \qquad \Lambda(18) = \mathbf{E}_{\widehat{m}}\left[\widehat{m}^2\right]$$

These moments identify the merger probability parameter governing the rate of expansion z as well as the mass of potential outside entrants (E). Together, these factors determine the equilibrium distribution of multi-industry presence.

Profitability Firm profitability is defined as the ratio of profits to sales. For a given panel of firms, this moment is given by

$$\Lambda(19) = \frac{1}{\varepsilon} - \mathbb{E}_{m,n,\hat{\mathbf{q}}}\left[\frac{\widetilde{w}\left[h_a^m + \bar{h}_b^m\right]}{\frac{1}{n}\sum_i \widehat{q}_i^{\varepsilon-1}}\right]$$

Notice that there is one fixed component from static production side that yields information on the value of ε and another from dynamic R&D expenditures that yields information on R&D cost and step size parameters.

Exit Rate As exit occurs when firms either receive the exogenous destruction shock or lose their last product, the predicted exit rate will be

$$\Lambda(20) = \kappa + \tau \cdot \sum_m \Gamma_{m,1}$$

However, for consistency, we simply use the value from the simulated firm sample. This moment serves primarily to determine the value of the rate of exogenous destruction κ, as well as the mass of outside entrants E, since the size of the pool of entrants affects the rate creative destruction and hence the exit rate of single-product firms.

Total Research Intensity We have two moments to track levels of R&D: the ratio of total research labor expenditures spending to total production labor expenditures. Since research spending is proportional to n, R&D expenditures per product will be the same across firms with the same m, while employment will be a function of the portfolio of product qualities. Because the wage is common to both types of labor, this will simply be the ratio of R&D employment to production employment given by

$$\Lambda(21) = \mathbb{E}_{m,n,\hat{\mathbf{q}}} \left[\frac{\widetilde{w} \left[h_a^m + \bar{h}_b^m \right]}{\left(\frac{\varepsilon-1}{\varepsilon} \right) \frac{1}{n} \sum_i \widehat{q}_i^{\varepsilon-1}} \right]$$

Conditional on innovation rates, this moment give us information on the research production function parameters.

Firm Growth We have a moment for employment growth amongst firms. This is calculated conditional on the firm not exiting, since we do not observe the last period's growth rate for exiting firms. The moment is calculated by looking at the one-year growth rate of total employment by a firm. It is labeled $\Lambda(22)$. The employment growth primarily informs on the rate of exogenous destruction κ and the R&D cost function parameters.

Aggregate Growth The growth rate gives information on the effectiveness of research spending absent effects coming from the distribution of firm size and its relation to firm growth, particularly on innovation step sizes. This is moment $\Lambda(23)$.

Spillover Differential In order to quantify the spillovers associated with basic research, we turn to patent citation data. The model predicts that innovations that build off of previous basic research should have a larger step size on average. If we take citations as a proxy for step size, then patents that cite basic research should themselves have more citations.

This effect will diminish with the age of the patent due to product line cooldown. Thus the average time after which a public innovation is indistinguishable from a private innovation should be

$$\Lambda(24) = \frac{1}{\zeta}\left(\frac{\tau_a}{\tau}\right)$$

This yields direct information on the value of the cooldown rate ζ.

Firm Age Firm age is highly correlated with firm size. We track the average age of firms for those above and below the median firm size. This yields information entry and exit patterns, as well as on the rate of creative destruction. Moment $\Lambda(25)$ is the average age of firms below the median firm size, while moment $\Lambda(26)$ is the average age for those firms above it.

TABLE B.2: BASIC RESEARCH INTENSITY AND MULTI-MARKET ACTIVITY, ROBUSTNESS CHECKS

	Covariates		**Alternative Measures**		**Instrumental Variables**		**Estimation**	
	(1) No	(2) Yes	(3) Patent Based	(4) Weighted Links	(5) State Present in 1986	(6) SOE in 1986	(7) Heckman	(8) Negative Binomial
# of Industries	0.014***	0.011***	0.012***	0.028***	0.023***	0.020**	0.045***	0.007***
	(0.00)	(0.00)	(0.00)	(0.01)	(0.01)	(0.01)	(0.01)	(0.00)
Log Employment		0.002	0.003	0.018***	-0.003	-0.002	0.012***	-0.001
		(0.00)	(0.00)	(0.01)	(0.00)	(0.00)	(0.00)	(0.00)
Foreign HQ		-0.010**	-0.010	-0.048**	-0.004	-0.005	-0.041***	-0.006
		(0.01)	(0.01)	(0.02)	(0.01)	(0.01)	(0.01)	(0.00)
Profitability		0.006	0.016	0.005	0.005	0.005	0.025**	0.005
		(0.00)	(0.01)	(0.01)	(0.00)	(0.00)	(0.01)	(0.00)
Year & Organization FE	YES	YES	YES	YES	YES	YES	YES	YES
N	13708	13706	3709	14823	13707	13707	13707	13707

Notes: Pooled data for the period 2000-2006. *Basic Research Intensity* is defined as the ratio of total firm investment in basic research divided by total firm investment in applied research. Columns 1 and 2 re-estimate the Tobit model with different sets of regressors. Columns 3 and 4 modify the measure of a firm's multi-industry presence. Column 3 uses patent applications of French firms to the European patent office (1993-2003) to count the number of distinct technological fields in which they are present (1-digit IPC classification). Column 4 weights each bilateral industry link of a firm by the empirical frequency of this link in the French economy, thus giving more weight to less related industries. Columns 5 and 6 re-estimate the model by instrumenting contemporary multi-industry presence by historical ownership structures. More specifically, we exploit the nationalization wave of the Mitterrand era that preceded the privatization of the 90s. The idea is that state ownership effectively increased the scope of a firm's economic activities. Column 5 uses state participation in the capital of a firm in 1986 as an instrument. Column 6 uses state ownership of a company in 1986 as an instrument. Both instruments accurately predict an increased multi-industry presence nowadays. Columns 7 and 8 estimate the relationship between multi-industry presence and basic research intensity by using a Heckman model and a negative binomial model. Tobit estimates relate to the marginal effect of the regressors with respect to the uncensored variable mean and are evaluated at the sample mean of covariates (except for categorical variables evaluated for firms that are present in 1 industry, non-foreign owned, in 2002). Robust standard errors clustered at the firm level in parentheses. See Appendix B.2 for the definition of variables.

Bibliography

ACEMOGLU, D. (2008): *Introduction to Modern Economic Growth*. Princeton University Press.

ACEMOGLU, D., P. AGHION, C. LELARGE, J. VAN REENEN, AND F. ZILIBOTTI (2007): "Technology, Information, and the Decentralization of the Firm," *Quarterly Journal of Economics*, 122(4), 1759–1799.

ACEMOGLU, D., U. AKCIGIT, N. BLOOM, AND W. R. KERR (2013): "Innovation, Reallocation and Growth," Working Paper 18993, NBER.

ACEMOGLU, D., U. AKCIGIT, D. HANLEY, AND W. R. KERR (2012): "Transition to Clean Technology," University of Pennsylvania mimeo.

ACEMOGLU, DARON (1998): "Why do New Technologies Complement Skills? Directed Technical Change And Wage Inequality," *Quarterly Journal of Economics*.

——— (2002): "Directed Technical Change," *Review of Economic Studies*.

ACEMOGLU, DARON AND UFUK AKCIGIT (2012): "Intellectual Property Rights Policy, Competition and Innovation," *Journal of the European Economic Association*.

ACEMOGLU, DARON, PHILIPPE AGHION, LEONARDO BURSZTYN, AND DAVID HEMOUS (2012): "The Environment and Directed Technical Change," *American Economic Review*.

ACEMOGLU, DARON, UFUK AKCIGIT, NICK BLOOM, AND WILLIAM KERR (2012): "Innovation, Reallocation, and Growth," *Working Paper*.

ADDA, JEROME, AND RUSSELL COOPER (2003): "Dynamic Economics: Quantitative Methods and Applications," *MIT Press*.

AGHION, P., U. AKCIGIT, AND P. HOWITT (2014): "What Do We Learn From Schumpeterian Growth Theory?," in *Handbook of Economic Growth*, ed. by P. Aghion, and S. N. Durlauf, vol. 2 of *Handbook of Economic Growth*, pp. 515–563. Elsevier.

AGHION, P., N. BLOOM, R. BLUNDELL, R. GRIFFITH, AND P. HOWITT (2005): "Competition and Innovation: An Inverted-U Relationship," *Quarterly Journal of Economics*, 120(2), 701–728.

AGHION, P., R. BLUNDELL, R. GRIFFITH, P. HOWITT, AND S. PRANTL (2004): "Entry and Productivity Growth: Evidence From Microlevel Panel Data," *Journal of the European Economic Association*, 2(2-3), 265–276.

——— (2009): "The Effects of Entry on Incumbent Innovation and Productivity," *Review of Economics and Statistics*, 91(1), 20–32.

AGHION, P., AND P. HOWITT (1992): "A Model of Growth Through Creative Destruction," *Econometrica*, 60(2), 323–351.

——— (1996): "Research and Development in the Growth Process," *Journal of Economic Growth*, 1(1), pp. 49–73.

——— (1998): *Endogenous Growth Theory*. MIT Press.

——— (2009): *The Economics of Growth*. MIT Press.

AGHION, PHILIPPE, AND PETER HOWITT (1992): "A Model of Growth Through Creative Destruction," *Econometrica*.

AGHION, PHILIPPE, ANTOIN DECHEZLEPRETRE, DAVID HEMOUS, RALF MARTIN, AND JOHN VAN REENEN (2012): "Carbon Taxes, Path Dependency and Directed Technical Change: Evidence from the Auto Industry," *Harvard University and LSE mimeo*.

AGHION, PHILIPPE, CHRISTOPHER HARRIS, PETER HOWITT, AND JOHN VICKERS (2001): "Competition, Imitation and Growth with Step-by-Step Innovation," *Review of Economic Studies*.

AKCIGIT, U., M. A. CELIK, AND J. GREENWOOD (2013): "Buy, Keep or Sell: Economic Growth and the Market for Ideas," .

AKCIGIT, U., AND W. R. KERR (2010): "Growth Through Heterogeneous Innovations," Working Paper 16443, NBER.

ANSELIN, L., A. VARGA, AND Z. ACS (1997): "Local Geographic Spillovers between University Research and High Technology Innovations," *Journal of Urban Economics*, 42(3), 422–448.

ARCHER, DAVID (2005): "Fate of Fossil Fuel CO2 in Geologic Time," *Journal of Geophysical Research*.

ARROW, K. (1962): "Economic Welfare and the Allocation of Resources for Invention," in *The Rate and Direction of Inventive Activity: Economic and Social Factors*, pp. 609–626. NBER.

ATKESON, A., AND A. T. BURSTEIN (2010): "Innovation, Firm Dynamics, and International Trade," *Journal of Political Economy*, 118(3), 433–484.

AUDRETSCH, D. B., AND M. P. FELDMAN (1996): “R&D Spillovers and the Geography of Innovation and Production,” *American Economic Review*, 86(3), 630–640.

BACCHIOCCHI, E., AND F. MONTOBBIO (2009): “Knowledge Diffusion from University and Public Research. A Comparison between US, Japan and Europe Using Patent Citations,” *Journal of Technology Transfer*, 34(2), 169–181.

BALASUBRAMANIAN, NATARAJAN, AND JAGADEESH SIVADASAN (2011): “What Happens When Firms Patent? New Evidence from US Economic Census Data,” *Review of Economics and Statistics*.

BENHABIB, J., J. PERLA, AND C. TONETTI (2013): “Catch-up and Fall-back through Innovation and Imitation,” *Journal of Economic Growth*, pp. 1–35.

BESSEN, J., AND E. MASKIN (2009): “Sequential Innovation, Patents, and Imitation,” *RAND Journal of Economics*, 40(4), 611–635.

BESSEN, J., AND M. J. MEURER (2008): *Patent Failure: How Judges, Bureaucrats, and Lawyers Put Innovators at Risk*. Princeton University Press.

BLOCH, L., AND E. KREMP (1999): “Ownership and Voting Power in France,” FEEM WP 62.

BLOOM, N. (2009): “The Impact of Uncertainty Shocks,” *Econometrica*, 77(3), 623–685.

BLOOM, N., M. SCHANKERMAN, AND J. VAN REENEN (2013): “Identifying Technology Spillovers and Product Market Rivalry,” *Econometrica*, 81(4), 1347–1393.

BLUNDELL, R., R. GRIFFITH, AND F. WINDMEIJER (2002): “Individual Effects and Dynamics in Count Data Models,” *Journal of Econometrics*, 108(1), 113–131.

BOLDRIN, M., AND D. K. LEVINE (2008): *Against Intellectual Monopoly*. Cambridge University Press.

BOVENBERG, LANS, AND SJAK SMULDERS (1995): "Environmental Quality and Pollution-Augmenting Technological Change in a Two-Sector Endogenous Growth Model," *Journal of Public Economics*.

——— (1996): "Transitional Impacts of Environmental Policy in an Endogenous Growth Model," *International Economic Review*.

BRESNAHAN, T. F., AND M. TRAJTENBERG (1995): "General Purpose Technologies Engines of Growth?," *Journal of econometrics*, 65(1), 83–108.

CONESA, JUAN CARLOS AND DIRK KRUEGER (1998): "Social Security with Heterogeneous Agents," *Review of Economic Dynamics*.

COZZI, G., AND S. GALLI (2009): "Science-based R&D in Schumpeterian Growth," *Scottish Journal of Political Economy*, 56(4), 474–491.

——— (2014): "Sequential R&D and Blocking Patents in the Dynamics of Growth," *Journal of Economic Growth*, forthcoming.

COZZI, G., AND G. IMPULLITTI (2010): "Government Spending Composition, Technical Change, and Wage Inequality," *Journal of the European Economic Association*, 8(6), 1325–1358.

DASGUPTA, P., AND P. A. DAVID (1994): "Toward a New Economics of Science," *Research Policy*, 23(5), 487–521.

DAVIS, S. J., J. HALTIWANGER, AND S. SCHUH (1998): *Job Creation and Destruction*. The MIT Press.

DAVIS, STEVEN, JOHN HALTIWANGER, AND SCOTT SCHUH (1996): "Job Creation and Destruction," *MIT Press*.

DHONT-PELTRAULT, E., AND E. PFISTER (2011): "R&D Cooperation versus R&D Subcontracting: Empirical Evidence from French Survey Data," *Economics of Innovation and New Technology*, 20(4), 309–341.

DUTTA, PRAJIT K AND RADNER, ROY (2006): "A Game-Theoretic Approach to Global Warming," *Advances in Mathematical Economics*.

DYKEMAN, D., AND D. KOPKO (2004): "Recording Patent License Agreements in the USPTO," *Intellectual Property Today, August*, 18, 19.

EECKHOUT, J., AND B. JOVANOVIC (2002): "Knowledge Spillovers and Inequality," *American Economic Review*, 92(5), 1290–1307.

EMANUEL, KERRY (2005): "Increased Destructiveness of Tropical Cyclones Over the Past 30 Years," *Nature*.

FIGUEROA, N., AND C. J. SERRANO (2013): "Patent Trading Flows of Small and Large Firms," Working Paper 18982, NBER.

FOSTER, L., J. C. HALTIWANGER, AND C. J. KRIZAN (2001): "Aggregate Productivity Growth: Lessons From Microeconomic Evidence," in *New Developments in Productivity Analysis*, pp. 303–372. University of Chicago Press.

FOSTER, LUCIA, AND CHERYL GRIM (2010): "Characteristics of the Top R&D Performing Firms in the US: Evidence from the Survey of Industrial R&D," *CES Working Paper 10-33*.

GARICANO, L., C. LELARGE, AND J. VAN REENEN (2013): "Firm Size Distortions and the Productivity Distribution: Evidence from France," NBER Working Paper 18841.

GHOSH, KAUSHIK, AND WILLIAM KERR (2010): "Patent-SSEL Matching," *Census Bureau Technical Paper*.

GOLOSOV, MIKHAIL, JOHN HASSLER, PER KRUSELL, AND ALEH TSYVINSKI (2011): "Optimal Taxes on Fossil Fuel in General Equilibrium," *NBER Working Paper 17348*.

GOOLSBEE, A. (1998): "Does Government R&D Policy Mainly Benefit Scientists and Engineers?," *American Economic Review*, 88(2), 298–302.

GOULDER, LAWRENCE, AND KOSHY MATHAI (2000): "Optimal CO2 Abatement in the Presence of Induced Technological Change," *Journal of Environmental Economics and Management*.

GOULDER, LAWRENCE, AND STEPHEN SCHNEIDER (1999): "Induced Technical Change and the Attractiveness of CO2 Abatement Policies," *Resource and Energy Economics*.

GOURIÉROUX, CHRISTIAN, AND ALAIN MONFORT (1996): "Simulation-Based Econometric Methods," *Oxford University Press*.

GRAHAM, S., AND S. VISHNUBHAKAT (2013): "Of Smart Phone Wars and Software Patents," *Journal of Economic Perspectives*, 27(1), 67–85.

GREEN, J. R., AND S. SCOTCHMER (1995): "On the Division of Profit in Sequential Innovation," *RAND Journal of Economics*, 26(1), 20–33.

GREENSTONE, MICHAEL, ELIZABETH KOPITS, AND ANNE WOLVERTON (2011): "Estimating the Social Cost of Carbon for Use in US Federal Rulemakings: A Summary and Interpretation," *MIT Working Paper 11-04*.

GRILICHES, Z. (1986): "Productivity, R&D and Basic Research at the Firm Level in the 1970s," *American Economic Review*, 76(1), 141–154.

——— (1990): "Patent Statistics as Economic Indicators: A Survey," *Journal of Economic Literature*, 28(4), 1661–1707.

——— (1992): "The Search for R&D Spillovers," *Scandinavian Journal of Economics*, pp. S29–S47.

GRIMAUD, ANDRE, GILLES LAFFORGUE, AND BERTRAND MAGNÉ (2011): "Climate Change, Mitigation Options and Directed Technical Change: A Decentralized Equilibrium Analysis," *Resource and Energy Economics*.

GROSSMAN, G. M., AND E. HELPMAN (1991): "Quality Ladders in the Theory of Growth," *Review of Economic Studies*, 58(1), 43–61.

GROSSMAN, GENE, AND ELHANAN HELPMAN (1991): "Innovation and Growth in the Global Economy," *MIT Press*.

HALL, B. H., A. B. JAFFE, AND M. TRAJTENBERG (2001): "The NBER Patent Citation Data File: Lessons, Insights and Methodological Tools," Discussion paper, NBER.

HALL, BRONWYN, ADAM JAFFE, AND MANUEL TRAJTENBERG (2001): "The NBER Patent Citation Data File: Lessons, Insights and Methodological Tools," *NBER Working Paper 8498*.

HARRIS, CHRISTOPHER, AND JOHN VICKERS (1995): "Innovation and Natural Resources: A Dynamic Game with Uncertainty," *RAND Journal of Economics*.

HARSTAD, BARD (2012): "Climate Contracts: A Game of Emissions, Investments, Negotiations, and Renegotiations," *Review of Economic Studies*.

HARTLEY, PETER, KENNETH MEDLOCK, TED TEMZELIDES, AND XINYA ZHANG (2011): "Energy Sector Innovation and Growth: An Optimal Energy Crisis," *Rice University Working Paper*.

HASSLER, JOHN, AND PER KRUSELL (2012): "Economics and Climate Change: Integrated Assessment in a Multi-Region World," *Journal of the European Economic Association*.

HASSLER, JOHN, PER KRUSELL, AND CONNY OLOVSSON (2011): "Energy-Saving Technical Change," *NBER Working Paper 18456*.

HEMOUS, DAVID (2012): "Environmental Policy and Directed Technical Change in a Global Economy: The Dynamic Impact of Unilateral Environmental Policies?," *INSEAD Mimeo*.

HENDERSON, R., A. B. JAFFE, AND M. TRAJTENBERG (1998): "Universities as a Source of Commercial Technology: A Detailed Analysis of University Patenting, 1965–1988," *Review of Economics and Statistics*, 80(1), 119–127.

HOPENHAYN, H., G. LLOBET, AND M. MITCHELL (2006): "Rewarding Sequential Innovators: Prizes, Patents, and Buyouts," *Journal of Political Economy*, 114(6), 1041–1068.

HOWITT, P. (2000): *The Economics of Science and the Future of Universities*. Citeseer.

HSIEH, C.-T., AND P. J. KLENOW (2009): "Misallocation and Manufacturing TFP in China and India," *Quarterly Journal of Economics*, 124(4), 1403–1448.

——— (2012): "The Life Cycle of Plants in India and Mexico," NBER Working Paper 18133.

IMPULLITTI, G. (2010): "International Competition and US R&D Subsidies: A Quantitative Welfare Analysis," *International Economic Review*, 51(4), 1127–1158.

JAFFE, A. B., M. TRAJTENBERG, AND R. HENDERSON (1993): "Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations," *Quarterly Journal of Economics*, 108(3), 577–598.

JAFFE, ADAM, DAVID POPP, AND RICHARD NEWELL (2010): "Energy, the Environment, and Technological Change," Academic Press/Elsevier.

JOINT ECONOMIC COMMITTEE (JEC) (2010): "The 2010 Joint Economic Report," *Washington, DC: Government Printing Office*.

JONES, C. I. (2005): "Growth and Ideas," in *Handbook of Economic Growth*, ed. by P. Aghion, and S. N. Durlauf, vol. 1 of *Handbook of Economic Growth*, pp. 1063–1111. Elsevier.

KERR, WILLIAM, AND SHIHE FU (2008): "The Survey of Industrial R&D—Patent Database Link Project," *Journal of Technology Transfer*.

KERR, WILLIAM, AND WILLIAM LINCOLN (2010): "The Supply Side of Innovation: H-1B Visa Reforms and US Ethnic Invention," *Journal of Labor Economics*.

KIRKPATRICK, S. C.D. GELATT, AND M.P. VECCHI (1983): "Optimization by Simulated Annealing," *Science*.

KLETTE, T. J., AND S. S. KORTUM (2004): "Innovating Firms and Aggregate Innovation," *Journal of Political Economy*, 112(5), 986–1018.

KÖNIG, M., J. LORENZ, AND F. ZILIBOTTI (2012): "Innovation vs Imitation and the Evolution of Productivity Distributions," .

KORTUM, S. (1993): "Equilibrium R&D and the Patent–R&D Ratio: US Evidence," *American Economic Review*, 83(2), 450–457.

——— (1997): "Research, Patenting, and Technological Change," *Econometrica*, 65(6), 1389–1420.

KORTUM, S. S. (1992): "Inventions, R&D and Industry Growth," Ph.D. Dissertation, Yale Univ., New Haven, CT.

KRUSELL, PER, AND ANTHONY SMITH (2009): "Macroeconomics and Global Climate Change: Transition for a Many-Region Economy," *Working Paper*.

LAITNER, J., AND D. STOLYAROV (2013): "Derivative Ideas and the Value of Intangible Assets," *International Economic Review*, 54(1), 59–95.

LANDSEA, CHRISTOPHER (2005): "Atlantic Hurricanes and Global Warming," *Nature*.

LENTZ, R., AND D. T. MORTENSEN (2008): "An Empirical Model of Growth Through Product Innovation," *Econometrica*, 76(6), 1317–1373.

LINK, A. N. (1981): "Basic Research and Productivity Increase in Manufacturing: Additional Evidence," *American Economic Review*, 71(5), 1111–12.

LUCAS, R. E., AND B. MOLL (2013): "Knowledge Growth and the Allocation of Time," *Journal of Political Economy*, forthcoming.

MACCRACKEN, CHRISTOPHER, JAMES EDMONDS, SON KIM, AND RONALD SANDS (1999): "The Economics of the Kyoto Protocol," *The Energy Journal*.

MAIRESSE, J., AND P. MOHNEN (2010): "Using Innovations Surveys for Econometric Analysis," National Bureau of Economic Research WP 15857.

MANSFIELD, E. (1980): "Basic Research and Productivity Increase in Manufacturing," *American Economic Review*, 70(5), 863–873.

——— (1981): "Composition of R and D Expenditures: Relationship to Size of Firm, Concentration, and Innovative Output," *Review of Economics and Statistics*, 63(4), 610–615.

——— (1995): "Academic Research Underlying Industrial Innovations: Sources, Characteristics, and Financing," *Review of Economics and Statistics*, 77(1), 55–65.

MCFADDEN, DANIEL (1989): "A Method of Simulated Moments for Estimation of Discrete Response Models Without Numerical Integration," *Econometrica*.

MENDELSOHN, ROBERT, WILLIAM NORDHAUS, AND DAIGEE SHAW (1994): "The Impact of Global Warming on Agriculture: A Ricardian Analysis," *American Economic Review*.

MORALES, M. F. (2004): "Research Policy and Endogenous Growth," *Spanish Economic Review*, 6(3), 179–209.

MOWERY, D. C., R. R. NELSON, B. SAMPAT, AND A. ZIEDONIS (2004): *Ivory Tower and Industrial Innovation: University-Industry Technology Transfer before and after the Bayh-Dole Act in the United States*. Stanford University Press.

MUELLER, NICHOLAS, ROBERT MENDELSOHN, AND WILLIAM NORDHAUS (2011): "Environmental Accounting for Pollution in the United States Economy," *American Economic Review*.

MURRAY, F., P. AGHION, M. DEWATRIPONT, J. KOLEV, AND S. STERN (2009): "Of Mice and Academics: Examining the Effect of Openness on Innovation," Discussion paper, National Bureau of Economic Research Working Paper 14819.

NELDER, J. A., AND R. MEAD (1965): "A Simplex Method For Function Minimization," *The Computer Journal*, 7(4), 308–313.

NELSON, R. R. (1959): "The Simple Economics of Basic Research," *Journal of Political Economy*, 67(3), 297–306.

NEWELL, RICHARD, ADAM JAFFE, AND ROBERT STAVINS (1999): "The Induced Innovation Hypothesis and Energy-Saving Technological Change," *Quarterly Journal of Economics*.

NICHOLLS, ROBERT, AND JASON LOWE (2006): "Climate Stabilisation and Impacts of Sea-level Rise," .

NORDHAUS, WILLIAM (1994): "Managing the Global Commons: The Economics of Climate Change," *MIT Press*.

——— (2007): "To Tax or Not to Tax: The Case for a Carbon Tax," *Review of Environmental Economics and Policy*.

——— (2008): "A Question of Balance: Weighing the Options on Global Warming Policies," *Yale University Press*.

NORDHAUS, WILLIAM, AND JOSEPH BOYER (2000): "Warming the World: Economic Modeling of Global Warming," *MIT Press*.

OECD ENVIRONMENT DIRECTORATE (2011): "Patent Search Strategies for the Identification of Selected Environment-Related Technologies," .

PAKES, A., AND Z. GRILICHES (1980): "Patents and R&D at the Firm Level: A First Report," *Economics Letters*, 5(4), 377–381.

——— (1984): "Estimating Distributed Lags in Short Panels with an Application to the Specification of Depreciation Patterns and Capital Stock Constructs," *Review of Economic Studies*, 51(2), 243–262.

PAKES, A., AND M. SCHANKERMAN (1984): "The Rate of Obsolescence of Patents, Research Gestation Lags, and the Private Rate of Return to Research Resources," in *R & D, Patents, and Productivity*, pp. 73–88. University of Chicago Press.

PAKES, A. S. (1986): "Patents as Options: Some Estimates of the Value of Holding European Patent Stocks," *Econometrica*, 54(4), 755–84.

PAKES, ARIEL, AND DAVID POLLARD (1989): "Simulation and the Asymptotics of Optimization Estimators," *Econometrica*.

PERLA, J., AND C. TONETTI (2013): "Equilibrium Imitation and Growth," *Journal of Political Economy*, forthcoming.

PETERS, M. (2013): "Heterogeneous Mark-ups and Endogenous Misallocation," LSE Working Paper.

PHILLIPS, G. M., AND A. ZHDANOV (2013): "R&D and the Incentives from Merger and Acquisition Activity," *Review of Financial Studies*, 26(1), 34–78.

PIZER, WILLIAM (1999): "The Optimal Choice of Climate Change Policy in the Presence of Uncertainty," *Resource and Energy Economics*.

POPP, DAVID (2002): "Induced Innovation and Energy Prices," *American Economic Review*.

——— (2004): "ENTICE: Endogenous Technological Change in the DICE Model of Global Warming," *Journal of Environmental Economics and Management*.

——— (2006): "Innovation in Climate Policy Models: Implementing Lessons from the Economics of R&D," *Energy Economics*.

POPP, DAVID, AND RICHARD NEWELL (2012): "Where Does Energy R&D Come From? Examining Crowding Out From Energy R&D," *Energy Economics*.

POWELL, M. J. (1970): "A Hybrid Method For Nonlinear Equations," *Numerical Methods for Nonlinear Algebraic Equations*, 7, 87–114.

RESTUCCIA, D., AND R. ROGERSON (2008): "Policy Distortions and Aggregate Productivity with Heterogeneous Establishments," *Review of Economic Dynamics*, 11(4), 707–720.

ROMER, P. M. (1990): "Endogenous Technological Change," *Journal of Political Economy*, 98(5), S71–102.

——— (2001): "Should the Government Subsidize Supply or Demand in the Market for Scientists and Engineers?," in *Innovation Policy and the Economy, Volume 1*, pp. 221–252. MIT Press.

ROSENBERG, N. (1990): "Why Do Firms Do Basic Research (with Their Own Money)?," *Research Policy*, 19(2), 165–174.

ROSENBERG, N., AND R. R. NELSON (1994): "American Universities and Technical Advance in Industry," *Research Policy*, 23(3), 323–348.

ROSENBERG, N., AND M. TRAJTENBERG (2004): "A General-Purpose Technology at Work: The Corliss Steam Engine in the Late-Nineteenth-Century United States," *Journal of Economic History*, 64(1), 1–39.

SCHERER, F. M. (1965): "Firm Size, Market Structure, Opportunity, and the Output of Patented Inventions," *American Economic Review*, 55(5), 1097–1125.

SOLOMON, S. D. QIN, M. MANNING, Z. CHEN, M. MARQUIS, K.B. AVERYT, M.TIGNOR AND H.L. MILLER (2007): "IPCC (Intergovernmental Panel on Climate Change). Climate Change 2007: The Physical Science Basis," *Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*.

SONG, Z., K. STORESLETTEN, AND F. ZILIBOTTI (2011): “Growing Like China,” *American Economic Review*, 101(1), 196–233.

STERN, NICHOLAS (2007): “The Economics of Climate Change: The Stern Review,” *Cambridge University Press*.

STOKES, D. E. (1997): *Pasteur’s Quadrant: Basic Science and Technological Innovation*. Brookings Institution Press.

STOTT, PETER, D.A. STONE, AND M.R. ALLEN (2004): “Human Contribution to the European Heatwave of 2003,” *Nature*.

TRAJTENBERG, M., R. HENDERSON, AND A. JAFFE (1992): “Ivory Tower versus Corporate Lab: An Empirical Study of Basic Research and Appropriability,” NBER Working Paper 4146.

——— (1997): “University versus Corporate Patents: A Window on the Basicness of Invention,” *Economics of Innovation and New Technology*, 5(1), 19–50.

VAN DER ZWAAN, ROBERT, REYER GERLAGH, G. KLAASSEN, AND L. SCHRATTENHOLZER (2002): “Endogenous Technical Change in Climate Change Modelling,” *Energy Economics*.

WEITZMAN, MARTIN (2009): “On Modeling and Interpreting the Economics of Catastrophic Climate Change,” *Review of Economics and Statistics*.

WILSON, D. J. (2009): “Beggar Thy Neighbor? The In-state, Out-of-state, and Aggregate Effects of R&D Tax Credits,” *Review of Economics and Statistics*, 91(2), 431–436.

ZANGWILL, W. I., AND C. GARCIA (1981): *Pathways to Solutions, Fixed Points, and Equilibria*. Prentice-Hall, Englewood Cliffs, NJ.