1-1-2015

Essays on Multiple Selves and Temporal Framing

Rudolph Henkel
University of Pennsylvania, rudyhenkel@gmail.com

Essays on Multiple Selves and Temporal Framing

Abstract
There exist numerous documented behavioral deviations from standard discounted utility maximizing behavior. These include time inconsistency, violations of the axiom of independence of irrelevant alternatives, preferences for commitment, and preferences over the timing of information. I develop two novel models which provide new insights into, and plausible explanations for, several of these deviations. In the first model, Multiself Bargaining, I propose a dual-self model in which two selves have conflicting preferences over the action to be taken by an agent. The selves have identical payoff utility, and only differ in their time preference factor. The default action of the agent is modeled as the outcome of a Tullock contest among the selves, where the self who wins chooses their preferred action. Viewing the outcome of this contest as the point of disagreement, the selves are allowed to negotiate to a mutually preferred outcome, and this negotiation is modeled as a Nash bargaining problem. I show that many of the deviations of interest are generated by this model, including time inconsistent behavior, such as diminishing impatience, as well as violations of independence of irrelevant alternatives in choice problems. Notably the preference reversals from time inconsistency are "smooth", as opposed to the singular reversal in quasi-hyperbolic discounting, the standard model used in the literature. In the second model, Temporal Reference Points, I develop an alternate form of prospect evaluation in which agents form a set of subjectively important points in time in the life of a prospect, termed "temporal reference points." When determining the present value of a prospect, agents discount based on the time between each of these temporal reference points, as opposed to based on the entire time between the present and the payout of the prospect. Under restrictions on the formation of temporal reference points, diminishing impatience in an agent is shown to be equivalent to a preference for informational updates occurring at the same time. Finally, the model is shown to allow the novel resolution of an apparent conflict in the experimental evidence on diminishing impatience.

Degree Type
Dissertation

Degree Name
Doctor of Philosophy (PhD)

Graduate Group
Economics

First Advisor
David Dillenberger

Second Advisor
Mallesh Pai

Keywords
Behavioral, Decision, dual-self, Economics, irrationality

Subject Categories
Economics | Economic Theory

ESSAYS ON MULTIPLE SELVES AND TEMPORAL FRAMING

Rudy Henkel

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

2015

Supervisor of Dissertation

____

David Dillenberger
Assistant Professor of Economics

Co-Supervisor of Dissertation

____

Mallesh Pai
Assistant Professor of Economics

Graduate Group Chairperson

____

Jesus Fernandez-Villaverde, Professor of Economics

Dissertation Committee

David Dillenberger, Assistant Professor of Economics
Mallesh Pai, Assistant Professor of Economics
Andrew Postlewaite, Professor of Economics; Professor of Finance

ACKNOWLEDGEMENT

I would first like to thank my co-advisors, Drs. David Dillenberger and Mallesh Pai. They have provided invaluable support for me in the development of this work, in my professional development, and in the face of personal developments. Cliché as it may be, I could not have done this without their motivation and input, and it has instilled in me a lifelong gratitude.

I would second like to thank my committee member, Dr. Andrew Postlewaite, for his helpful input and support in the development of the flagship paper of this thesis, and for his availability and advice in general.

I would further like to extend thanks to Drs. George Mailath, Aislinn Bohren, Hanming Fang and Garth Baughman for their critiques and suggestions.

Finally, I would like to extend my thanks and my love to my wife, Stacey, for her patience and forbearance, and for her tireless support and love over the course of my Ph.D. work.

ABSTRACT

ESSAYS ON MULTIPLE SELVES AND TEMPORAL FRAMING

Rudy Henkel

David Dillenberger

Mallesh Pai

There exist numerous documented behavioral deviations from standard discounted utility maximizing behavior. These include time inconsistency, violations of the axiom of independence of irrelevant alternatives, preferences for commitment, and preferences over the timing of information. I develop two novel models which provide new insights into, and plausible explanations for, several of these deviations. In the first model, Multiself Bargaining, I propose a dual-self model in which two selves have conflicting preferences over the action to be taken by an agent. The selves have identical payoff utility, and only differ in their time preference factor. The default action of the agent is modeled as the outcome of a Tullock contest among the selves, where the self who wins chooses their preferred action. Viewing the outcome of this contest as the point of disagreement, the selves are allowed to negotiate to a mutually preferred outcome, and this negotiation is modeled as a Nash bargaining problem. I show that many of the deviations of interest are generated by this model, including time inconsistent behavior, such as diminishing impatience, as well as violations of independence of irrelevant alternatives in choice problems. Notably the preference reversals from time inconsistency are "smooth", as opposed to the singular reversal in quasi-hyperbolic discounting, the standard model used in the literature. In the second model, Temporal Reference Points, I develop an alternate form of prospect evaluation in which agents form a set of subjectively important points in time in the life of a prospect, termed "temporal reference points." When determining the present value of a prospect, agents discount based on the time between each of these temporal reference points, as opposed to based on the entire time between the present and the payout of the prospect. Under restrictions on the formation of temporal reference points, diminishing impatience in an agent is shown to be equivalent to a preference for informational updates occurring at the same time. Finally, the model is shown to allow the novel resolution of an apparent conflict in the experimental evidence on diminishing impatience.

Contents

Part I

Multiself Bargaining

"The idea of self-control is paradoxical unless it is assumed that the psyche contains more than one energy system, and that these energy systems have some degree of independence from each other." (Donald McIntosh, The Foundations of Human Society, 1969)

1 Introduction

The observable behavior of decision making agents includes a number of ubiquitous effects not consistent with the predictions of standard utility maximization with geometric time discounting. Among these, agents exhibit time inconsistency; specifically, they have reversals of preference when outcomes are mutually delayed. Agents exhibit violations of independence of irrelevant alternatives in choice problems; in general, their decision may depend on the entire choice set, and they may be tempted "toward" an option that they yet do not choose. Related to both of these, agents seek out commitment devices, whether to prevent preference reversals, or to remove temptations, even when such devices are costly.

That an individual might contain conflicting internal preferences is certainly one potential cause of these phenomena. One strand of work in neurology views the brain as operating with a "team of rivals" architecture, wherein different sections of the brain compete with each other directly for control over the actions of the individual, e.g. Eagleman (2011). Some MRI evidence is consistent with the notion that decisions made with different time horizons engage very different areas of the brain, e.g. McClure et al. (2004). Threads of research in psychology also address the idea of conflict between multiple selves, e.g. Ainslie (1986).

Without laying claim to being a model of the brain, which it is not, this work has the goal of formalizing this neurological inspiration in order to account for the described empirical regularities in behavior. I do this through the use of a novel "dual-self" model. The model gives rise to a smooth form of time inconsistency, which has the same qualitative implications as hyperbolic time discounting. It further generates a two-sided temptation effect, in which unchosen alternatives alter the decisions of the agent. The model additionally predicts the use of costly commitment devices by the agent, provides strong intuition about the nature of these phenomena, and provides some novel insight into welfare evaluation.

I model an individual decision maker, or agent, as consisting of two selves: one patient and one impatient. The selves are taken to share the same payoff utility, and each discounts time geometrically, but they differ in their time discount factor. Thus, while they agree on the immediate utility granted by a decision, such as which flavor of ice cream is the most delicious, they will in general disagree with regard to decisions which have consequences over time, such as whether it is a good idea to eat the ice cream. If the selves cannot agree on what action to take, it is assumed that they engage in a costly conflict over control of the action, modeled as a variant on a Tullock (1980) rent-seeking game.[1] The self that wins the conflict then chooses their most preferred action.

I allow the selves to negotiate in order to select an action mutually preferred to this costly conflict; this negotiation is modeled as a Nash bargaining problem. The bargaining set is the set of possible utility vectors created by available actions, or lotteries over actions, and the outcome of the costly conflict is treated as the disagreement point for the bargaining. Thus, in equilibrium, the costly conflict will not occur, as the selves will negotiate to a better outcome. This bargaining can result in either a deterministic choice in some applications, such as consumption-savings, or an agreed upon mixing between discrete choices in some discrete menu choice applications; both of these are addressed. Section 2 details the model in full.

The first empirical regularity, addressed in Section 3, is diminishing impatience. This is a particular form of time inconsistent preference reversal, where as the consequences of a decision are pushed into the future, an individual's choices exhibit less impatience regarding the outcome. As a simple example, an individual given the choice between \$100 today and \$120 tomorrow may choose the \$100 today, but when presented with the choice between \$100 in seven days and \$120 in eight days may "switch", and now desire to wait the extra day for the greater reward. As another example, an individual may desire to save a large portion of his next paycheck for retirement; however, when payday arrives, he may change his mind when faced with the immediate reality of reducing his consumption. Proposition 1 establishes that the multiself bargaining model reflects a smooth form of diminishing impatience. In discrete decisions, as payoffs at different times are both pushed into the future, the observed patience of the agent increases continuously, in that they choose latter rewards with continuously increasing probability. For continuous decisions, such as consumption-savings, the model predicts that their decision reflects greater patience as the decision is made farther in advance, such as through a higher savings rate.

The empirical evidence for diminishing impatience is broad, and spans disciplines, seeing particular interest in psychology and economics; e.g. Thaler (1981), Loewenstein and Prelec (1992), Kirby and Herrenstein (1995), Frederick et al. (2002), Fang and Silverman (2009). Many models attempt to incorporate the phenomenon by utilizing forms of time discounting other than geometric. Evidence, e.g. Ainslie (1992) and Myerson and Green (1995), suggests that the form of time discounting reflected in the decisions of agents is well fit by hyperbolic discounting, in which the

[1] Also used by Benabou and Pycia (2002).

discount factor falls steeply at first, but then less rapidly.[2] Consumption data is shown by some work, e.g. Angeletos et al. (2001), to be much better fit by hyperbolic models than exponential. The most commonly used model of diminishing impatience, quasi-hyperbolic discounting, also referred to as "beta-delta", was introduced by Laibson (1997). Quasi-hyperbolic discounting amounts to having one time discount factor between the current period and the next period, and a second, higher, time discount factor for evaluating between all periods thereafter. The difficulty with such an approach is that it allows only a discontinuous form of diminishing impatience: a stark division between "now" and "later" extremely sensitive to the precise definition of the length of a period. One advantage of the multiself bargaining model is that the form of diminishing impatience it creates more closely reflects hyperbolic than quasi-hyperbolic discounting, avoiding this discontinuity and period sensitivity.

The tendency of agents to utilize costly commitment actions is addressed in Section 4. Evidence for the use of commitment options includes voluntary intermediate deadlines, e.g. Ariely and Wertenbroch (2002), Kaur et al. (2009), and savings pre-commitment, e.g. Thaler and Benartzi (2004). In many cases these options are *a priori* inefficient, as the same choices can be made without the costly commitment device. Individuals now make use of smart phone apps which impose financial penalties for not going to the gym, for example. As we are dealing with preference aggregation, and not a single set of preferences, the model cannot easily address a true "preference" for commitment, and commitment options cannot benefit both selves in this model. However, we can speak of what the model implies about the behavior of the agent. Proposition 2 establishes that for sufficiently low, but positive, commitment cost, agents will make use of commitment devices to commit to specific actions in advance with strictly positive probability. Further, under a broad set of conditions, this probability will be 1. Intuitively, this arises from the fact that the patient self has greater foresight when making decisions further in advance, and this foresight places that self in a better position for bargaining; the commitment device then allows the patient self the ability to "lock in" that superior position. Proposition 3 establishes conditions under which the same result will hold when commitment devices for minimal (or maximal) actions are available, such as a minimum amount of savings.

The effects of tempting options on choices is addressed in Section 5. Temptation here is used to refer to cases where the addition of an option to a choice set alters the decision made by the individual, even when the new option is not chosen; in other words, a violation of independence of irrelevant alternatives. This is documented in choices from discrete sets, as well as in consumption-savings decisions, e.g. Hanks et al. (2012), Ashraf et al. (2006), Huang et al. (2013). Existing literature on temptation primarily relies on choices being assigned some explicit temptation, or self-

[2] Specifically, under generalized hyperbolic discounting, a reward in t time is discounted by $(1+\alpha t)^{-(\gamma/\alpha)}$.

control value, e.g. Gul and Pesendorfer (2001). The agent is taken to have a limited, or costly, capacity to resist such temptations. In contrast with this literature, in this model temptation effects will arise endogenously from the differing utility evaluations of the selves, which in turn result from their differing time preference. Proposition 4 establishes that when the utility granted to a self by their most preferred action increases, the agent will choose an action, or lottery, that grants that self a higher expected utility, even if this most preferred action is an "irrelevant" alternative.

However, the overall temptation effect of an action will be shown to depend on how much the option is valued by *both* selves. If the preferred action of the short-term self is regarded as particularly awful by the long-term self, such as a very unhealthy dessert choice, then the anticipated effort the long-term self exerts in conflict increases. Proposition 5 examines how the utility granted to the long-term self by the short-term self's preferred action influences the outcome, establishing a cutoff which determines the way in which changes to this utility move the action chosen by the agent. The implications of Propositions 4 and 5 are then combined to characterize the net effect of the addition of a new option to an existing choice set.

Section 6 investigates the implications of the model for welfare evaluation, with the primary insight being that while the model is more limiting than the classic model for welfare evaluation, we can recapture much that is lost by models that use non-geometric time discounting. This is accomplished by viewing the selves as individuals for the purposes of examining welfare, so that we can evaluate policies from the standpoint of how they affect the welfare of both selves. Utility aggregation methods are considered, and several results are established for a weighted utility welfare function. Notably, commitment devices are argued to be welfare improving for sufficiently long time horizons.

Section 7 discusses extensions and alterations to the model, and concludes. Extensions discussed include variations on the conflict game, the robustness of the model to other methods of bargaining between selves, and varying the preferences of the selves along dimensions other than time preference.

2 The Model

2.1 Decision Problem and Notation

Time is continuous, but there are a finite number of discrete times $t_1 = 0,\ t_2,\ t_3...\ t_N$, where decisions are made and payoffs received, with $\Delta_n = t_{n+1} - t_n \geq 0$. The decision made at time t_n is referred to as decision n. An agent is assumed to consist of two selves. Both selves share an identical utility function over payoffs, $u(\cdot)$; they differ only in their time discount rate, ρ. One self is referred to as "long-term" with ρ_l, and one is referred to as "short-term" with ρ_s. It is assumed that these time discount rates satisfy $\rho_s > \rho_l > 0$; note that this means the discount factor of the long-term self is larger.[3]

The choice of continuous time here allows for ease in comparative statics exercises to follow, but as utility is not evaluated continuously, the model should conceptually be thought of as consisting of discrete periods with varying period lengths. In particular, for fixed intervals, $\Delta_1 = \Delta_2 = ... = \Delta_{N-1} = \Delta$, the model becomes analogous to a standard discrete model, with time discount factors given by

$$\beta_l = e^{-\Delta\rho_l} > \beta_s = e^{-\Delta\rho_s}.$$

The actions available to the agent at decision n are given by action set A_n; A_n is assumed to either be a discrete choice set, or a compact Euclidean space. The set of lotteries over actions is given by $\mathcal{A}_n$. At each decision point the agent must select one lottery over actions $\alpha_n \in \mathcal{A}_n$. Let the history of realized actions up to decision n be given by $h_n = \{a_1,\ a_2,\ ...\ a_{n-1}\}$. In general, A_n, and therefore $\mathcal{A}_n$, may depend on this history. In addition to its effect on future action sets, each realized action grants some payoff vector, and the corresponding payoff utility is denoted $u(a_t)$.[4] Each self evaluates their own welfare using standard expected utility with geometric time discounting, so that the time 0 discounted utility of a realized action stream $a_1,\ a_2, ...\ a_n$ to self i is given by

$$\sum_{j=1}^{n} e^{-\rho_i t_j} u(a_j).$$

Every lottery α_n at decision n creates an expectation over realized actions, and thus an expected payoff utility, given by

$$E_{\alpha_n}[u(a_n)].$$

[3]For $\rho_s = \rho_l$ the model coincides with standard expected utility with geometric discounting.

[4]A payoff occurring at a time where there is no decision to be made can be modeled as resulting from a trivial decision from a singleton action set. In this way, the model can accommodate streams of payoffs resulting from single actions, either by modeling the future payoffs as resulting from degenerate decisions, or by altering future action sets to accomodate the changed payoffs. Implicitly, the history of actions acts as a state variable.

The decision procedure employed by the agent will be described recursively. Consider the decision made at time t_N. As there are no future decisions, the concern of the selves is focused entirely on the payoff utility derived from the possible actions. As the selves agree on payoff utility, they agree on rankings of lotteries for this final decision. Thus, the lottery chosen is simply

$$\mathcal{D}(\mathcal{A}_N(h_N)) = arg \max_{\alpha_N \in \mathcal{A}_N(h_N)} E_{\alpha_N}[u(a_N)],$$

where $\mathcal{D}(\mathcal{A}_N(h_N))$ indicates the decision made from lottery set $\mathcal{A}_N(h_N)$.[5] The expected utility self i derives from this decision at time t_N will be denoted $U_{i,N}(\mathcal{D}(\mathcal{A}_N(h_N)), h_N)$.

Now consider decision N-1, at time t_{N-1}, where the lottery is being selected from $\mathcal{A}_{N-1}(h_{N-1})$. Each possible lottery α_{N-1} creates not only an expected payoff utility but also an expectation over the future history, h_N. This future history in turn influences the future action set, which influences the future decision made, and therefore the future utility. Specifically, the discounted utility to self i at time N-1 from lottery α_{N-1} is given by

$$U_{i,N-1}(\alpha_{N-1}, h_{N-1}) =$$
$$E_{\alpha_{N-1}}\left[u(a_{N-1}) + e^{-\rho_i \Delta_{N-1}}\left(U_{i,N}(\mathcal{D}(\mathcal{A}_N(\{h_{N-1}, a_{N-1}\})), \{h_{N-1}, a_{N-1}\})\right)\right].$$

This utility will also be denoted as $U_i(\alpha_{N-1})$ to conserve notation; it consists of the expectation of payoff utility, $u(a_{N-1})$, as well as the expected discounted future utility. Note that $\{h_{N-1}, a_{N-1}\}$ is the future history, and the expectation is over what value a_{N-1} will take. Essentially, the selves are (correctly) projecting the action that will be taken at the final decision point based on the action taken today, and discounting the utility they will receive from that action based on the time difference between the current decision and the latter one, Δ_{N-1}. Note that this implies the assumption of sophistication of the selves. Now, in contrast to the final decision, due to the differing discount rates, the selves do not agree on the ranking for this decision, and generally may prefer different lotteries.

Each lottery α_{N-1} creates a utility vector $(U_s(\alpha_{N-1}),\ U_l(\alpha_{N-1}))$, as determined by the above valuation. Denote by $\mathcal{U}_{N-1}(\mathcal{A}_{N-1},\ h_{N-1})$ the set of all such utility vectors for decision N-1, which will be abbreviated $\mathcal{U}_{N-1}$. Since the choices are lotteries over actions, $\mathcal{U}_{N-1}$ will be a convex set. A conflict/bargaining procedure, described in the next two subsections, is used to select a single utility vector from $\mathcal{U}_{N-1}$, and the lottery chosen by the agent is the one corresponding to that utility vector. These steps can then be applied recursively backward, as the selves can now project the lottery chosen at decision N-1, and so on. The utility of lottery α_n to self i at time t_n is

$$U_{i,n}(\alpha_n, h_n) = E_{\alpha_n}\left(u(a_n) + e^{-\rho_i \Delta_n}\left[U_{i,n+1}(\mathcal{D}(\mathcal{A}_{n+1}(h_{n+1})), h_{n+1})\right]\right).$$

[5] In the non-generic case where arg is multi-valued here, the agent is indifferent between its elements, and the model is agnostic about the lottery chosen from among them. However, as will be seen this indifference does not impact the recursive decision procedure for earlier decisions, which is only dependent on the utility to each self.

This is not a Bellman equation. $\mathcal{D}$ is a decision process depending on the utility values of both selves, not just self i, thus it cannot be expressed as an optimization decision made by self i.

2.2 Conflict

Now we turn to describing the procedure by which the agent selects an action when the selves do not agree.[6] Consider decision n made at time t_n, and corresponding set of utility vectors, $\mathcal{U}_n$, created by the lotteries in $\mathcal{A}_n$. First, note that each self will have a bliss action in A_n: the action which grants them the highest discounted utility.[7] This bliss action is given by

$$b_n^i = arg\, \max_{a_n \in A_n} U_i(a_n).$$

These bliss actions induce a pair of bliss points in $\mathcal{U}_n$, given by:

$$(U_s(b_n^s),\ U_l(b_n^s)) \equiv (X_n^s,\ Y_n^s),\ (U_s(b_n^l),\ U_l(b_n^l)) \equiv (X_n^l,\ Y_n^l).$$

Thus, X_n^l represents the discounted expected utility granted to the short-term self by the bliss action of the long-term self. Finally, we define

$$S_n \equiv X_n^s - X_n^l;\ L_n \equiv Y_n^l - Y_n^s.$$

S_n, for example, is the non-negative difference in utilities that the short-term self will receive from the two bliss points. Figure 1 illustrates these terms.

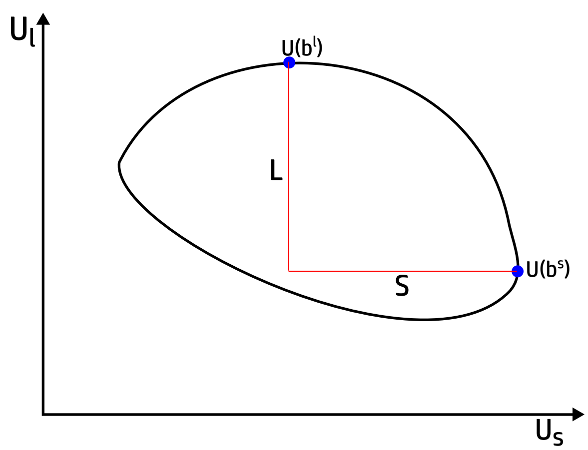

Figure 1: Utility vector set $\mathcal{U}$ with bliss points, b^l and b^s.

[6] If they do agree, the result of the procedure coincides with the mutually preferred lottery, and so need not be viewed as a special case.

[7] Usually, there will only be one such action. It is possible, however, that a self may have multiple "best" actions, granting the same maximum discounted expected utility; in most applications this will be a non-generic occurrence. In this case, the model requires a single bliss action to be selected from among these actions, but is agnostic about which; while the specific outcome that will result in this special case will depend on which bliss point is selected, the general results of the model do not depend on the selection method.

If the selves cannot agree on what action or lottery to take, it is assumed that the selves will engage in a conflict modeled as a slight variant of the Tullock (1980) rent-seeking game.[8] First, the selves simultaneously commit to an effort choice e_i. Second, the winner is determined based on the effort choices. The probability that the short-term self is the winner is given by

$$p \equiv \begin{cases} \frac{1}{2}, & if\ e_s = e_l = 0 \\ \frac{e_s^\gamma}{e_s^\gamma + e_l^\gamma}, & otherwise, \end{cases}$$

with $0 \leq \gamma$. Third, the winner selects the action taken by the agent. Naturally, the self that wins will select their own bliss action. Considering potential equilibria of this conflict game, note that $e_s = e_l = 0$ is not one, as both selves would have incentive to exert marginal effort. Thus, the short-term self selects e_s to maximize:

$$\frac{e_s^\gamma}{e_s^\gamma + e_l^\gamma} X_n^s + \frac{e_l^\gamma}{e_s^\gamma + e_l^\gamma} X_n^l - e_s.$$

Lemma 1. *Given $S_n > 0$, $L_n > 0$, $0 \leq \gamma \leq 1$, the conflict game at decision n has a unique Nash equilibrium given by:*

$$e^*_{s,n} = \frac{\gamma L_n^\gamma S_n^{\gamma+1}}{(S_n^\gamma + L_n^\gamma)^2},\ e^*_{l,n} = \frac{\gamma L_n^{\gamma+1} S_n^\gamma}{(S_n^\gamma + L_n^\gamma)^2},\ p^*_n = \frac{S_n^\gamma}{S_n^\gamma + L_n^\gamma}.$$

This creates an expected utility vector for the selves given by

$$\left(X_n^l + \frac{S_n^{\gamma+1}\left(S_n^\gamma + (1-\gamma)L_n^\gamma\right)}{(S_n^\gamma + L_n^\gamma)^2},\ Y_n^s + \frac{L_n^{\gamma+1}\left(L_n^\gamma + (1-\gamma)S_n^\gamma\right)}{(S_n^\gamma + L_n^\gamma)^2} \right).$$

Proof: Proofs of all results not given in text are in Appendix A.

To ensure uniqueness, we will restrict attention to $0 \leq \gamma \leq 1$ in the conflict game.[9]

2.3 Bargaining

The model allows the selves to negotiate to an option mutually preferred by the selves to the outcome of the conflict. This is modeled as a Nash bargaining problem, using the set of utility vectors, $\mathcal{U}_n$, as the bargaining set, and the equilibrium of the conflict game as the disagreement point.

[8]The variance is that in the classic Tullock game the players are competing for an equal prize; in this application the prize - control over the action of the agent - has different values for the two players.

[9]This is different from the standard Tullock requirement of $0 \leq \gamma \leq 2$ precisely because of the homogeneous prize value

For the remainder of the paper, I will add the cost of effort from the conflict game back into the disagreement point, leaving the disagreement point as a simple mixing between the utility vectors induced by the two bliss actions. This step will allow for cleaner intuition and understanding: the actual payment of the cost of effort in the conflict game is not the source of any interesting behavior in the model. The qualitative nature of the results are not changed by this step, and where relevant both disagreement points are shown to yield the same result in the proofs.[10] This gives a simpler form to the disagreement point derived from the conflict game, now given by

$$d_n = (d_{s,n},\ d_{l,n}) = \left(p_n^* X_n^s + (1-p_n^*)X_n^l,\ p_n^* Y_n^s + (1-p_n^*)Y_n^l\right) = \left(X_n^l + \frac{S_n^{\gamma+1}}{S_n^\gamma + L_n^\gamma},\ Y_n^s + \frac{L_n^{\gamma+1}}{S_n^\gamma + L_n^\gamma}\right).$$

The standard Nash bargaining solution maximizes the product of the gains of the two selves relative to a disagreement point; this product is denoted the Nash product. Thus, the action taken by the agent is given by

$$\mathcal{D}(\mathcal{A}_n(h_n)) = arg\ \underset{\alpha_n \in \mathcal{A}_n}{max} \left(U_s(\alpha_n) - X_n^l - \frac{S_n^{\gamma+1}}{S_n^\gamma + L_n^\gamma}\right)\left(U_l(\alpha_n) - Y_n^s - \frac{L_n^{\gamma+1}}{S_n^\gamma + L_n^\gamma}\right).$$

The chosen utility vector is illustrated as α in figure 2.[11] Note that this procedure implies that conflict does not occur in equilibrium. Rather, anticipation of conflict drives the bargaining between selves. Both selves are able to project the outcome of conflict, and it is this commonly anticipated outcome that drives the bargaining.

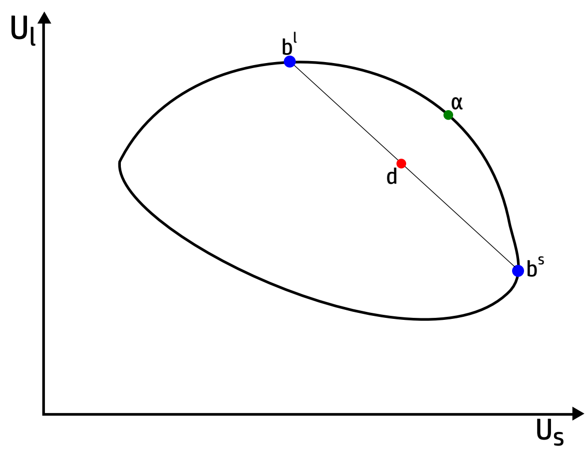

Figure 2: Utility vector set $\mathcal{U}$ with disagreement point d and outcome vector α.

[10] Proposition 5 has a variant equational form depending on which disagreement point is used, both with the same qualitative implications. The variant form is included in Appendix A.

[11] The Nash bargaining solution will give a unique utility vector, but it may be that multiple lotteries create the same utility vector, thus leading the *arg* function to be multi-valued. In this case, both selves are indifferent between the possible arguments, and I interpret this as the agent being indifferent between the options. In most applications this is a non-generic occurrence, and for the remainder of the paper I will assume that $\mathcal{D}(\cdot)$ is single valued.

Lemma 2. *The decision process employed at a given decision n is Pareto efficient with respect to the utilities of the selves. It is invariant to affine transformations of the shared payoff utility function $u(\cdot)$. If $\mathcal{U}_n$ is such that $(x,y) \in U_n$ if and only if $(y,x) \in U_n$, then $U_s(\mathcal{D}(\mathcal{A}_n(h_n))) = U_l(\mathcal{D}(\mathcal{A}_n(h_n)))$.*

An affine transformation could also be applied to both effort costs in the conflict game without changing the outcome, but the costs have been normalized so that the marginal cost of effort is 1. The last part of Lemma 2 is essentially the Symmetry axiom of Nash bargaining: the decision process does not favor one self over the other; the classic axiom is based on an exogenous, symmetric disagreement point, however.

Observe that if the Pareto frontier consists only of mixtures between the two bliss points, then the disagreement point (itself being a mixture between the bliss points) will lie on the Pareto frontier, and thus coincide with the outcome vector. So, for example, if there are only two actions, the outcome can be interpreted as the selves agreeing on the same mixing that would result from conflict, and by doing so bypassing the actual costs of conflict.[12]

2.4 Illustrative Example

To illuminate the workings of the model, we look at a simple example of savings-consumption. Consider an agent endowed with \$1 at time t_1, and nothing at time t_2, with $\Delta_1 = t_2 - t_1 = 1$. He must decide how much to consume at time t_1 and how much to save; assume there is no interest on savings. We take his consumption utility function to be $u(c) = \sqrt{c}$, and $\rho_l = 0$, $\gamma = 1$, $\rho_s = ln(2)$.

Denote the amount saved at time t_1 as a. Modeling the action as being the choice of a, we have $A_1 = [0,1]$, with each action granting a payoff utility as well as constraining the action set at time t_2. At t_2, both selves will agree to consume the full amount remaining, so decision 2 is trivial. Thus, each self in the first decision has a discounted utility given by $\sqrt{1-a} + e^{-\rho_i}\sqrt{a}$. To determine their bliss actions, we have:

$$b_1^i = arg \max_{a \in [0,1]} \sqrt{1-a} + e^{-\rho_i}\sqrt{a},$$

which gives

$$b_1^i = e^{-2\rho_i}/(1 + e^{-2\rho_i}); \;\; b_1^s = 0.2; \;\; b_1^l = 0.5.$$

This creates bliss points

[12]If we use the disagreement point without the effort costs added back in, then this is still the case, but it is not obvious; see Lemma A.3 in Appendix A for a proof.

$$(X_1^s,\ Y_1^s) = (U_s(0.2),\ U_l(0.2)) = \left(\frac{\sqrt{5}}{2},\ \frac{3\sqrt{5}}{5}\right);$$

$$\left(X_1^l,\ Y_1^l\right) = (U_s(0.5),\ U_l(0.5)) = \left(\frac{3\sqrt{2}}{4},\ \sqrt{2}\right).$$

Thus,

$$S_1 = \frac{2\sqrt{5} - 3\sqrt{2}}{4};\ L_1 = \frac{5\sqrt{2} - 3\sqrt{5}}{5};$$

$$d_1 = \left(X_1^l + \frac{S_1^2}{S_1 + L_1},\ Y_1^s + \frac{L_1^2}{S_1 + L_1}\right) \approx (1.086,\ 1.382).$$

Finally, the amount of savings undertaken by the agent is found by solving

$$\underset{a}{max}\ \left(\sqrt{1-a} + 0.5\sqrt{a} - d_{s,1}\right)\left(\sqrt{1-a} + \sqrt{a} - d_{l,1}\right),$$

which gives $a \approx 0.354$. This outcome is illustrated in figure 3.

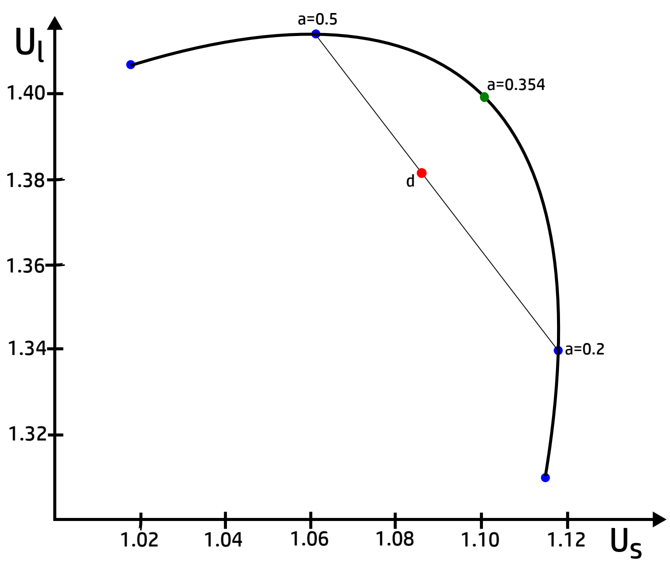

Figure 3: Savings decision of $a = 0.354$.

3 Diminishing Impatience

The source of diminishing impatience in the model is best introduced through a simple example. Consider an individual with payoff utility function given by $u(w) = w$ who, at time t_1 is given the choice between receiving \$100 at time t_2, or receiving \$120 at time t_3. The utility vectors created by the two options are, respectively,

$$\left(100e^{-\rho_s\Delta_1},\ 100e^{-\rho_l\Delta_1}\right);\ \left(120e^{-\rho_s(\Delta_1+\Delta_2)},\ 120e^{-\rho_l(\Delta_1+\Delta_2)}\right).$$

As this is a binary decision, the resulting $\mathcal{A}_t$ consists of all mixing lotteries between the two options, and $\mathcal{U}_t$ is as a result a line segment. Consider the interesting case where $120e^{-\rho_l\Delta_2} > 100 > 120e^{-\rho_s\Delta_2}$, so that the long-term self strictly prefers to wait for the second option, and the short-term self strictly prefers the sooner option. The time t_1 bliss points given by these options, then, are given by:

$$(X_1^s,\ Y_1^s) = \left(100e^{-\rho_s\Delta_1},\ 100e^{-\rho_l\Delta_1}\right);$$

$$\left(X_1^l,\ Y_1^l\right) = \left(120e^{-\rho_s(\Delta_1+\Delta_2)},\ 120e^{-\rho_l(\Delta_1+\Delta_2)}\right),$$

so that,

$$S_1 = 100e^{-\rho_s\Delta_1} - 120e^{-\rho_s(\Delta_1+\Delta_2)} = e^{-\rho_s\Delta_1}\left(100 - 120e^{-\rho_s\Delta_2}\right);$$

$$L_1 = 120e^{-\rho_l(\Delta_1+\Delta_2)} - 100e^{-\rho_l\Delta_1} = e^{-\rho_l\Delta_1}(120e^{-\rho_l\Delta_2} - 100).$$

We will use $\gamma = 1$, so we have

$$p = \frac{S}{S+L} = \frac{100 - 120e^{-\rho_s\Delta_2}}{(100 - 120e^{-\rho_s\Delta_2}) + e^{(\rho_s-\rho_l)\Delta_1}\,(120e^{-\rho_l\Delta_2} - 100)}.$$

As $\rho_s - \rho_l > 0$, it is straightforward to see that $\frac{\partial p}{\partial\Delta_1} < 0$. That is, as the payoffs are pushed into the future, the probability of the short-term self gaining control in the case of conflict shrinks: the disagreement point moves toward the bliss point of the long-term self. Thus, the model predicts that as rewards are delayed, the individual has a higher probability of waiting for the greater reward. The resulting $\mathcal{U}_n$ and outcomes are shown for multiple values of Δ_1 in figure 4, which uses $\rho_s = ln(2)$, $\rho_l = ln(\frac{20}{19})$, $\gamma = 1$, $\Delta_2 = 1$. For each $\mathcal{U}_n$ line segment, the higher point is the utility vector that results from waiting for the latter, larger reward, and the point in between the two endpoints is the outcome vector. One can see that as the delay (Δ_1) grows, the outcome moves to place more weight on waiting for the latter, greater reward.

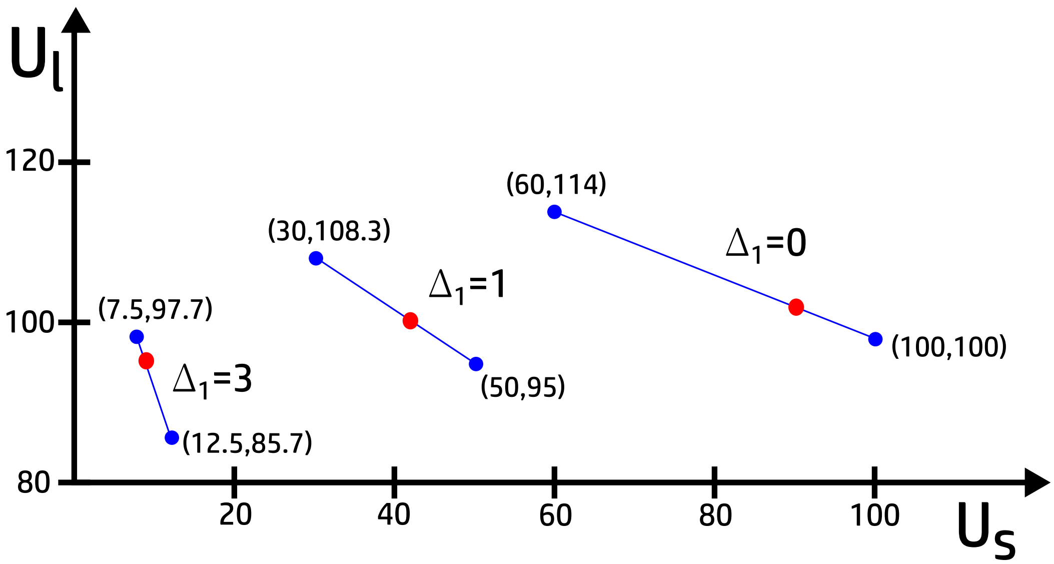

Figure 4: Utility vector set $\mathcal{U}_1$ for differing delays until payoffs.

The result on binary choices is generalized in the following.

Lemma 3. *Given a choice at time t_1 between a payout vector M_{Low} at time t_2 and a payout vector M_{High} at time t_3, with $u(M_{High}) > u(M_{Low})$, denote by p_L the probability an individual will choose M_L. Then,*

$$\tfrac{\partial}{\partial \Delta_1} p_L \leq 0,$$

with the derivative strict wherever $0 < p_L < 1$.

Intuitively, as rewards are pushed into the future, both selves "care" less about the difference between the two rewards; that is, the discounted value of the utility difference shrinks. However, the difference shrinks at a faster rate for the short-term self. Essentially, the long-term self sees a much bigger difference between far future rewards than the short-term self does, relatively. As a result, the short-term self has little incentive to exert effort in the conflict game, and is projected to win such a conflict with low probability, shifting the disagreement point in bargaining to favor the long-term self. Observationally, as Δ_1 increases, the behavior of the individual becomes "closer" to that predicted by standard geometric discounting, as though the long-term self was the entire individual. I will now present the general result for diminishing impatience.

Proposition 1 (Diminishing Impatience). *Consider an agent selecting an action from A_1, and suppose that all actions $a_1 \in A_1$ grant the same payoff utility at time t_1. Then,*

$$\frac{\partial}{\partial \Delta_1} \frac{Y_1^l - U_l(\mathcal{D}(\mathcal{A}_1))}{e^{-\rho_l \Delta_1}} \leq 0 \text{ and } \lim_{\Delta_1 \to \infty} \frac{Y_1^l - U_l(\mathcal{D}(\mathcal{A}_1))}{e^{-\rho_l \Delta_1}} = 0,$$

with the inequality strict if the Pareto frontier of $\mathcal{U}_1$ is smooth.

The term $Y_1^l - U_l(\mathcal{D}(\mathcal{A}_1))$ is the difference in time t_1 utilities the long-term self receives from his bliss point, Y_1^l, and from the decision made by the agent. Proposition 1 states that this difference in utilities is shrinking to 0 as Δ_1 grows. Note that, since all actions in A_1 grant the same payoff utility, Δ_1 is the amount of time between the decision at time t_1 and any possible payoff utility variation resulting from that decision; this is capturing the effect of shifting all payoffs into the future. Taking the future value of the utility difference, by dividing by $e^{-\rho_l \Delta_1}$, shows that Proposition 1 is not merely the result of both utilities shrinking to zero. Rather, the future payoff utility received draws close to the optimal future payoff utility of the long-term self.

The intuition of Proposition 1 is that as the consequences of a decision are pushed into the future, the outcome asymptotically approaches the bliss action of the long-term self. In the case of a discrete decision, this takes the form that the probability with which the agent will choose the bliss action of the long-term self approaches 1. For a continuous decision, the action chosen by the agent approaches the bliss action of the long-term self asymptotically. Putting Proposition 1 in terms of utilities allows the capture of both of these cases.

I'll now illustrate Proposition 1 for a continuous decision. Consider again a consumption-savings example, with no interest and $u(c) = \sqrt{c}$, $\gamma = 1$, $\rho_l = 0$, $\rho_s = ln(2)$. At time t_1, a consumer knows that they will receive endowment w at time t_2, and must now decide how much to consume at time t_2 and how much to save for time t_3, with $\Delta_2 = t_3 - t_2 = 1$. Denoting by a the savings decision made at time t_1.

The short-term self would like to choose a to maximize $e^{-ln(2)\Delta_1}\sqrt{w-a} + e^{-ln(2)(\Delta_1+1)}\sqrt{a}$, which gives $a = 0.2w$ as their their bliss action. Similarly, the bliss action of the long-term self is $a = 0.5w$. Note that the bliss actions of the selves are not individually affected by the delay between the decision and the first consumption, Δ_1, as both selves are individually geometric discounters. The bliss utility vectors at time t_1 are respectively given by:

$$\left(e^{-ln(2)\Delta_1}(\sqrt{0.8w} + 0.5\sqrt{0.2w}),\ \sqrt{0.8w} + \sqrt{0.2w}\right) = \left(2^{-\Delta_1}\sqrt{w}(\sqrt{0.8} + 0.5\sqrt{0.2}),\ \sqrt{w}(\sqrt{0.8} + \sqrt{0.2})\right);$$

$$\left(e^{-ln(2)\Delta_1}(\sqrt{0.5w} + 0.5\sqrt{0.5w}),\ \sqrt{0.5w} + \sqrt{0.5w}\right) = \left(2^{-\Delta_1}1.5\sqrt{0.5w},\ 2\sqrt{0.5w}\right).$$

Determining S and L, we find that

$$S_1 = 2^{-\Delta_1}\sqrt{w}\left(\sqrt{0.8} + 0.5\sqrt{0.2} - 1.5\sqrt{0.5}\right) = 2^{-\Delta_1}\sqrt{w}\left(\frac{2\sqrt{5} - 3\sqrt{2}}{4}\right) \equiv 2^{-\Delta_1}\sqrt{w}K_s;$$

$$L_1 = \sqrt{w}\left(2\sqrt{0.5} - \sqrt{0.8} - \sqrt{0.2}\right) = \sqrt{w}\left(\frac{5\sqrt{2} - 3\sqrt{5}}{5}\right) \equiv \sqrt{w}K_l.$$

The intuition becomes clear here again: S_1 is being discounted, which represents the short term seeing a smaller difference between the bliss points as Δ_1 increases. L_1 is discounted to a lesser degree (in the chosen example, it is not discounted at all), representing that the long-term self continues to see a difference between options for long delays. We can now calculate the probability of the short-term self winning the conflict game:

$$p_1 = \frac{2^{-\Delta_1}\sqrt{w}K_s}{2^{-\Delta_1}\sqrt{w}K_s + \sqrt{w}K_l} = \frac{K_s}{K_s + 2^{\Delta_1}K_l} \Rightarrow \frac{\partial}{\partial \Delta_1}p_1 < 0.$$

So, we see for the continuous case that the short-term self again has a lower probability of winning the conflict game as the payoffs are pushed into the future. Since this lower probability is anticipated by both selves, this will translate into a bargaining outcome more favorable to the long-term self: in this case, a higher savings rate. Figure 5 shows the result for $\Delta_1 = 1$; that is, the result when the agent is deciding on the amount of saving 1 unit of time in advance. The utility values for the short-term self contract and, as a result, the disagreement point moves closer to the bliss point of the long-term self; the savings rate increases from the previous 0.354 to 0.388.

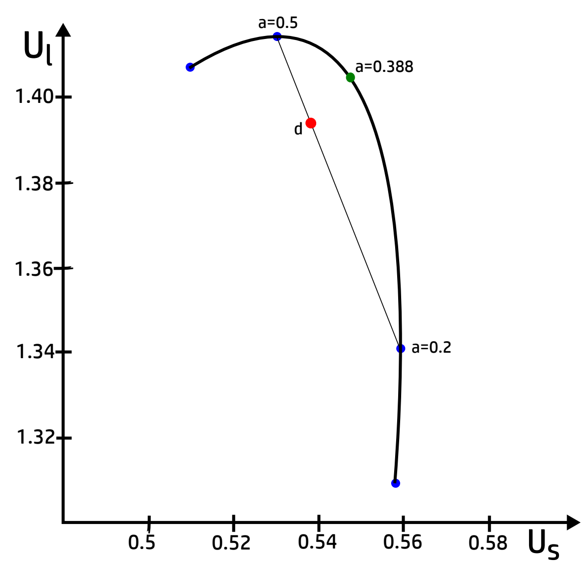

Figure 5: Savings decision of $a = 0.388$ made in advance.

The model thus predicts that when an individual is making a savings decision, the amount that they will choose to save increases the farther their decision is from the date of consumption. It is again this ability to account for "smooth" preference reversals that prevents the model from being reliant on period length specification, and aligns it more closely with hyperbolic discounting models, as opposed to the more popular β-δ specification.

4 Preferences for Commitment

To state the predictions of the model with regard to an agent's use of commitment devices, I will first introduce commitment sets, which are action sets that model committing to decisions in advance. It is then shown, in Lemma 4, that if the agent commits to a single action sufficiently far in advance, the action (or lottery) chosen will be different than the action chosen if he does not commit in advance. Building on this, Proposition 2 shows that, under the same time condition as Lemma 4, the agent will utilize costly, *voluntary* commitment devices to commit to single actions in advance, for sufficiently low commitment cost. Lemma 5 establishes that, intuitively, commitment devices are used to commit to actions favorable to the long-term self, as opposed to those favorable to the short-term self; one observes commitments to go to the gym, but not commitments to not go to the gym. Finally, Proposition 3 establishes conditions under which commitment devices for minimal actions, such as minimum amounts of savings, are equivalent to commitment devices for specific actions, such as exact amounts of savings. This result, in turn, implies that minimal action commitment devices have the same implications for behavior as single action commitment devices.

For this section, I exclusively consider two decision points, 1 and n, with $1 < n$; I do not exclude the possibility that there are decisions made in between the two. Define $\Delta_t = t_n - t_1$; note that this is potentially distinct from $\Delta_1 = t_2 - t_1$. Denote A_n^{nc} some fixed potential set of actions at decision n, with corresponding lottery set $\mathcal{A}_n^{nc}$, and utility vector set $\mathcal{U}_n^{nc}$. I will assume throughout that A_n^{nc} has at least two elements, and that the selves prefer different lotteries in $\mathcal{A}_n^{nc}$; this is equivalent to the assumption that the Pareto frontier of $\mathcal{U}_n^{nc}$ is not single valued.[13]

Definition 1. *A_1^r is a* required commitment set *for A_n^{nc} if the following conditions hold.*

1. *$\forall\ a_n \in A_n^{nc}\ \exists$ a unique $a_1^r \in A_1^r$ such that $a_1^r \in h_n \Rightarrow A_n(h_n) = \{a_n\}$. a_1^r is said to be a commitment action, and to induce action a_n.*

2. $u(a_i) = u(a_j)\ \forall\ a_i, a_j \in A_1^r$.

3. *$A_{n'}(h_{n'})$ is invariant to which action $a_1 \in A_1$ belongs to $h_{n'}$ if $n' \neq n$.*

[13] If this assumption does not hold, then the selves agree on what action to take, the resulting behavior is trivial, and nothing interesting results.

4. A_1^r *and* A_n^{nc} *have a one-to-one correspondence.*

A required commitment set A_1^r captures the concept of an agent having to decide now what action to take at a future date. The first condition specifies that any action in A_n^{nc} may be committed to by choosing a corresponding action in A_1^r. The second condition specifies that the immediate payoff utilities between the different commitment actions are the same, and the third condition specifies that the commitment actions do not affect action sets at any other decision points; these two together imply that the only difference between the commitment actions is through their effect on decision n actions. Finally, the fourth condition limits A_1^r to the commitment actions defined by condition 1; every action in A_n^{nc} has a unique corresponding commitment action in A_1^r, and every action in A_1^r is a commitment action for some unique action in A_n^{nc}.

A_1^r has an associated set of lotteries, $\mathcal{A}_1^r$, and utility vector set $\mathcal{U}_1^r$. Denote the lottery that the agent chooses from $\mathcal{A}_1^r$ as $\alpha_1^r = \mathcal{D}(\mathcal{A}_1^r)$; this is a lottery over commitment actions, which induces a realized lottery over actions at decision n. I refer to this resulting lottery over actions as α_n^r; thus, α_n^r is the realized lottery over actions at time t_n that results from the agent choosing from the required commitment set A_1^r at time t_1.

For expositional purposes, it is useful to note here that $\mathcal{U}_1^r$ is a time discounted version of $\mathcal{U}_n^{nc}$; the agent is essentially deciding from the actions of A_n^{nc}, he is just doing so earlier. Thus, the utility vectors in $\mathcal{U}_1^r$ are the same as those in $\mathcal{U}_n^{nc}$, discounted by the additional time between decisions 1 and n.[14] As the short-term self discounts more heavily than the long-term self, $\mathcal{U}_1^r$ is contracted more severely along the horizontal axis than the vertical. A contraction of this kind was illustrated previously in Figure 5.

Now, define α_n^{nc} to be the realized lottery over actions at time t_n that results if the agent does not previously commit to any action from A_n^{nc}. That is, $\alpha_n^{nc} = \mathcal{D}(\mathcal{A}_n^{nc})$. So, α_n^r is the realized lottery over actions at time t_n that results from commitment in advance, and α_n^{nc} is the realized lottery that results from no commitment in advance. The following lemma defines the condition under which these realized lotteries are different.

Lemma 4. *Let* A_1^r *be a required commitment set for* A_n^{nc}*, with* $n > 1$*. Then,* $\exists \underline{\Delta}$ *such that if* $t_n - t_1 = \Delta_t > \underline{\Delta}$*,* $\alpha_n^{nc} \neq \alpha_n^r$*, and* $U_{l,n}(\alpha_n^{nc}) < U_{l,n}(\alpha_n^r)$*. If the Pareto frontier of* $\mathcal{U}_n^{nc}$ *is smooth at* $(U_s(\alpha_n^{nc}),\ U_l(\alpha_n^{nc}))$*, then* $\underline{\Delta} = 0$.

Proof: Application of Proposition 1. Details in Appendix A.

Lemma 4, intuitively, says that decisions change when they are made earlier. Further, it says that if the earlier decision differs, it will differ in favor of the long-term self; these effects follow

[14] The payoff utility $u(\cdot)$ is not relevant at decision m; as it is the same for all actions, it can be normalized to zero.

from the diminishing impatience discussed in the previous section. Concerning the latter part of Lemma 4, the conditions under which the Pareto frontier is smooth at the point of decision are quite broad. If the action set is continuous and the utilities are continuously differentiable functions of the action set (such as in consumption-savings applications) it will generally be so. In specific, for consumption-savings decisions with strictly concave payoff utility, the Pareto frontier of $\mathcal{U}$ will be a strictly convex, and thus the second condition of Lemma 4 will always hold. It will also always hold for a binary action decision, as the outcome will be a strict mixing on the line between the two bliss points. If the Pareto frontier is not smooth at the outcome, which can occur generically for a discrete action set with more than two actions, then Lemma 4 says that a decision made sufficiently far in advance will still differ from one made at the time of the action.

This result only applies in situations where advanced commitment is required. In many applications it is more reasonable to think of commitments as something optional; the agent may commit, but they may choose not to as well.

Definition 2. *A_1^o is an* optional commitment set *for A_n^{nc} if $A_1^o = A_1^r \cup \{a_1^{nc}\}$ where A_1^r is a* required commitment set *for A_n^{nc}, and a_1^{nc} is such that $a_1^{nc} \in h_n \Rightarrow A_n(h_n) = A_n^{nc}$. a_1^{nc} is termed the no commitment action.*

An optional commitment set takes a required commitment set, and gives the agent the additional option to choose not to commit to any decision n action. Note that this definition does not require that $u(a_1^{nc})$ be equal to the payoff utility of the commitment actions; this will allow us to examine situations in which commitment carries with it some cost. If the realized action at decision 1 is a_1^{nc}, then this means that the agent will choose his period n lottery of actions from the full lottery set $\mathcal{A}_n^{nc}$, and the resulting lottery over actions will be α_n^{nc}, as previously defined.

A_1^o has an associated set of lotteries, $\mathcal{A}_1^o$, and utility vector set $\mathcal{U}_1^o$. Denote the lottery that the agent chooses from $\mathcal{A}_1^o$ as $\alpha_1^o = \mathcal{D}(\mathcal{A}_1^o)$. All actions in A_1^o induce a lottery over decision n actions, either α_n^{nc} for the no commitment action, or a degenerate lottery for the commitment actions. Thus, α_1^o induces a realized lottery over actions at time t_n, and I refer to this resulting lottery over actions as α_n^o. Denote as p_1^{nc} the probability that α_1^o places on the no commitment action, a_1^{nc}. For expositional purposes, note that $\mathcal{U}_1^r \subseteq \mathcal{U}_1^o$: $\mathcal{U}_1^o$ can be thought of as starting with $\mathcal{U}_1^r$, then adding the utility vector induced by the no commitment action, along with any mixings involving it that are Pareto improvements over existing vectors in $\mathcal{U}_1^r$.

Proposition 2 builds on Lemma 4 to examine the more interesting case of a decision made from the optional commitment lottery set $\mathcal{A}_1^o$.

Proposition 2 (Preference for Commitment). *Let A_1^r and A_1^o be a* required commitment set *and an* optional commitment set *for A_n^{nc}, respectively, with $n > 1$. Suppose that $u(a_1^{nc}) - u(a_1) = c \geq 0$ $\forall\ a_1 \in A_1^o/\{a_1^{nc}\}$, and that $\Delta_t > \underline{\Delta}$, where $\underline{\Delta}$ satisfies the condition of Lemma 4. Then, $\exists\ \bar{c} > 0$ such that if $c < \bar{c}$, $p_1^{nc} < 1$. Further, if $\mathcal{U}_n^{nc}$ contains Pareto improvements on all strict mixings between α_n^{nc} and α_n^r, then $\exists\ \widehat{c} > 0$ such that if $c < \hat{c}$, $p_1^{nc} = 0$, and $\alpha_1^o = \alpha_1^r$.*

Suppose that the agent can choose not to commit, and bears some utility cost c related to commitment actions. Proposition 2 says that as long as the agent would have chosen a different lottery if required to make the decision in advance (the broad condition for which was defined by Lemma 4), and the costs are sufficiently low, they will voluntarily use commitment options with a strictly positive probability.

To interpret the second part of the result, note that requiring $\mathcal{U}_n^{nc}$ to contain Pareto improvements on all strict mixings between α_n^{nc} and α_n^r is equivalent to the requirement that the line segment connecting the utility vectors created by α_n^{nc} and α_n^r lay on the interior of $\mathcal{U}_n^{nc}$, excepting the endpoints. Strict convexity of the Pareto frontier of $\mathcal{U}_n^{nc}$ is a sufficient condition for this, so the condition will always hold in savings-consumptions problems with concave utility, such as our illustrative example. It also holds generically in the case of discrete action sets when α_n^{nc} and α_n^r have different supports, since their induced utility vectors would then lie on different line segments of the Pareto frontier.[15] In the event that this condition holds, Proposition 2 says that for sufficiently low costs, the lottery over commitment actions chosen will be the same as the lottery they would have chosen if they were required to commit; the option not to commit will not be used. Two examples will greatly aid in the understanding of this result.

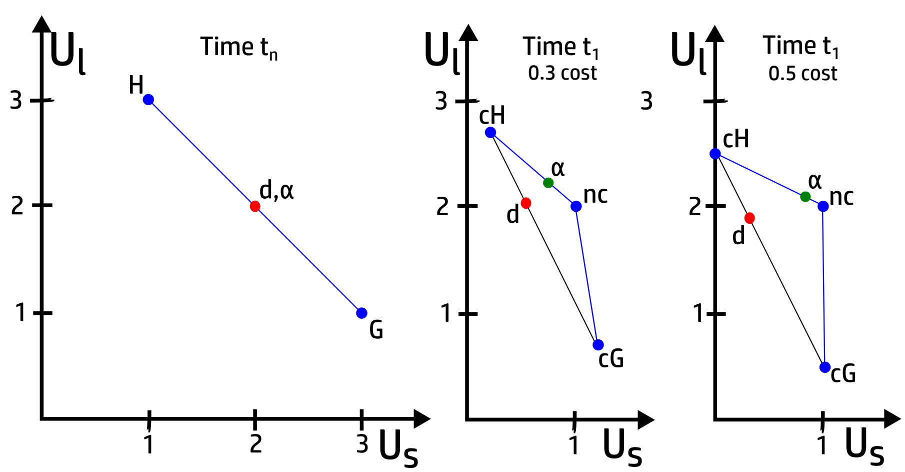

Figure 6: Utility vector sets at t_n and t_1 for optional commitment decision at t_1.

[15] Non-generically, one element of the discrete action set may induce the same utility vector as a mixing between two other elements of the action set, meaning that they would all lie on a line segment in $\mathcal{U}$.

Consider A_n^{nc} ={(H)ouse Salad, (G)rilled Steak}, $\Delta = 1$, $\gamma = 1$, $\rho_l = 0$, $\rho_s = ln(2)$. In this example, time t_n is the point at which the food is consumed. The time t_n discounted utilities of the two options are shown in the leftmost graph of Figure 6; the difference in payoffs shown can arise if the grilled steak gives higher payoff utility at decision n, but changes future action sets (perhaps through lower health) in a way that results in lower payoff utility at future decisions. The point which is labeled d, α indicates the utility vector induced by α_n^{nc}, the lottery that the agent would choose if making the decision at the time the food is consumed. In this case, it is a 50/50 mixture between the two options. The Pareto frontier is smooth at that point, so Lemma 4 implies that the mixing that the agent would choose in advance is different than 50/50.

The second graph illustrates the decision made from an optional commitment set at decision 1, where the payoff utility cost of commitment is 0.3. The three actions show are to commit to a house salad (cH), commit to a grilled steak (cG), and not to commit (nc). A_1^o ={cH, cG, nc} whereas A_1^r ={cH, cG}. The utility granted by cH, for example, is time discounted from the utility granted by H in the first graph, and further reduced by the cost of commitment.

If the agent were selecting from A_1^r, meaning if he had to select his food in advance, then the Pareto frontier would be the mixing between cH and cG on the graph, and the resulting mixing over actions is illustrated by point d; this is the utility vector induced by α_1^r. In this case, the probability that the agent will end up eating the house salad is higher than the 50/50 chance he had when selecting at the time of eating. The mixture over ultimate actions is α_n^r, so $\alpha_n^r \neq \alpha_n^{nc}$, in accordance with Lemma 4.

In the event that the agent is selecting from A_1^o, meaning he can commit to his food in advance, but does not have to, the Pareto frontier is given by the line segments connecting cH to nc, and nc to cG. The ultimate lottery selected is a mixing between cH and nc, illustrated by α on the graph. This is the utility vector induced by α_1^o. α_n^r and α_n^{nc} have the same support (and lie on the same segment of the Pareto frontier), so that the second part of Proposition 2 does not apply here. Thus, there is a positive probability of commitment, but not certain commitment. Finally, the third graph simply shows the resulting outcome for a higher cost of commitment; the probability of choosing not to commit becomes higher (the agent chooses a lottery with higher weight on the nc action).

An important intuition to take away from this graph is that the agent does not mix with committing to the option favored by the short-term self, the grilled steak. The agent either commits to the action preferred by the long-term self, or they do not commit at all. This is a desirable property of the model, because we do not observe individuals committing to unhealthy actions. They either commit to going to the gym, or they do not commit either way; they never commit to not go to the gym. The following lemma captures this intuition.

Lemma 5. *Let A_1^r and A_1^o be a required commitment set and optional commitment set for A_n^{nc}, respectively, with $n > 1$. Suppose that $u(a_1^{nc}) - u(a_1) = c \geq 0 \ \forall \ a_1 \in A_1^o/\{a_1^{nc}\}$. Then, $\exists \underline{\Delta}$ such that if $\Delta_t > \underline{\Delta}$, the lottery α_1^o places zero weight on commitment actions that are strictly worse for the long-term self than the no commitment action.*

Let's consider commitment as applied to the illustrative consumption-savings example. From Section 3, if the agent has to select the savings amount 1 time in advance, he selects $a = 0.388$; this is shown in the left side of Figure 7; this graph corresponds to $\mathcal{U}_1^r$, and the induced action lottery α_n^r is degenerate on $a = 0.388$. Consider now giving the agent the option not to commit to a savings rate in advance, and suppose that the commitment options carry with them some cost of commitment. Recall that the agent chooses $a = 0.354$ if the decision is made without prior commitment. That is, α_n^{nc} is a degenerate lottery with a probability of 1 on $a = 0.354$. This is illustrated in the right side of Figure 7; the no commitment point extends the Pareto frontier outward, since it carries no payoff utility cost in contrast to the commitment options; this graph corresponds to $\mathcal{U}_1^o$. The α shown on the graph is the utility vector induced by α_1^o, the decision made from $\mathcal{A}_1^o$.

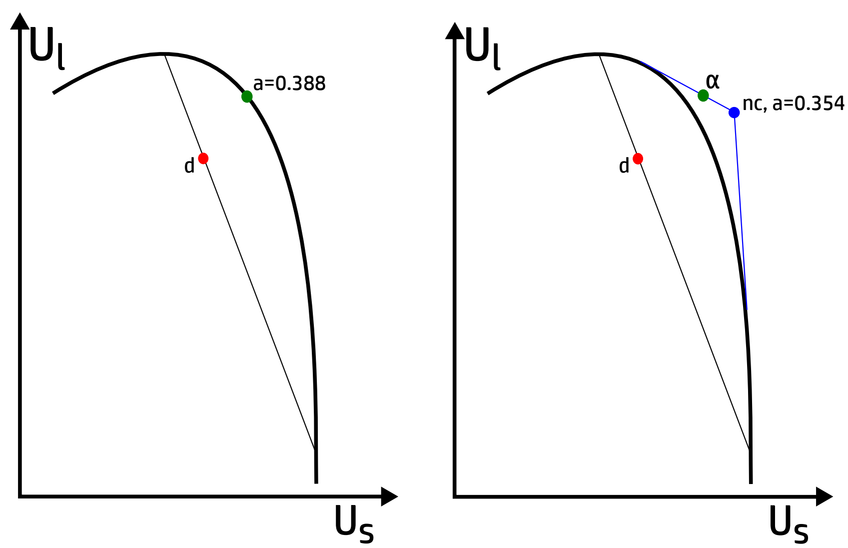

Figure 7: Required commitment versus optional commitment.

Note that since the Pareto frontier of $\mathcal{U}_n^{nc}$ is strictly convex, the second part of Proposition 2 applies, so that for sufficiently small costs of commitment, commitment should be guaranteed. This is illustrated in Figure 8, which shows what happens as the costs of commitment are reduced. The no commitment option draws closer to the curve created by the commitment options; as it does so, the portion of the Pareto frontier of $\mathcal{U}_1^o$ that includes mixings with the no commitment option shrinks and, if the cost is small enough, the option chosen when commitment was required, in this case $a = 0.388$, becomes part of the Pareto frontier, and the agent will choose that option.

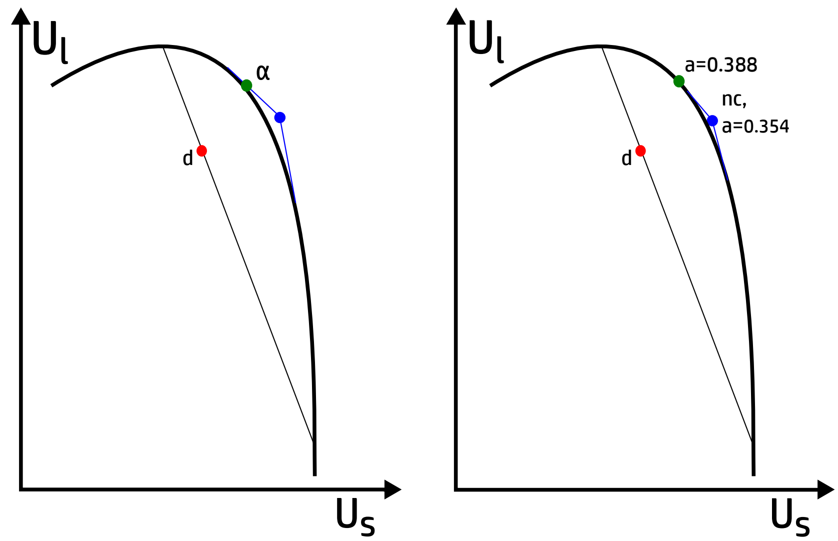

Figure 8: Optional commitment for differing commitment costs.

A reasonable objection to the applicability of this set of results is that savings commitments are more naturally thought of, and observed as, minimal savings commitments, as opposed to binding savings levels, so I now consider commitments to minimal actions.

Lemma 6. $\mathcal{D}(\mathcal{A}) = \mathcal{D}(\mathcal{D}(\mathcal{A}))$.

Proof: $\mathcal{D}(\mathcal{A})$ is the decision made from a given set of lotteries, $\mathcal{A}$. $\mathcal{D}(\mathcal{D}(\mathcal{A}))$ is the decision made when we limit the set of lotteries to the one lottery we would have picked from the full set. As there is only one lottery to pick from, the choice is the same.□

Lemma 6, while simple, is important to understand the application of Proposition 2 to savings problems, or other commitment decisions where we restrict options without committing to a single action. If an agent makes a commitment to a minimal amount of savings, then Lemma 6 implies that this is equivalent to committing to the exact amount of savings that the minimal amount of savings will induce. For example, if without any commitment an agent chooses savings $a = 0.35$, and committing to a minimal savings of 0.2 in advance induces him to ultimately choose savings of 0.4, then the commitment to a minimum of 0.2 is equivalent to committing to saving exactly 0.4 in the first place, both in terms of realized action of the amount saved, and in terms of utilities received by the selves.[16] Thus, a commitment decision to limit an action set can be reframed as a commitment decision to a specific action or lottery.

[16] Assuming that the costs of the commitment options are the same.

However, this does not mean that any set of commitment actions is equivalent to a required commitment set, or optional commitment set, that includes all possible actions. To see why this is, consider again our savings-consumption example. Suppose an agent has the ability to commit to a minimum savings rate in advance, but not a maximum one. Then, there is no minimum savings commitment which would be the equivalent of committing to saving 0, since no minimal commitment in advance would induce the agent to choose 0 as the ultimate savings decision. If a decision maker commits to a minimum savings rate of 0, this is equivalent to making no commitment at all.

Definition 3. *A_1^m is a* minimal action commitment set *for A_n^{nc} if the following conditions hold.*

1. A_n^{nc} is an interval in $\mathbb{R}$.
2. $\forall\ a_n \in A_n^{nc}\ \exists$ *a unique* $a_1^m \in A_1^m$ *such that* $a_1^m \in h_n \Rightarrow A_n(h_n) = \{a :\ a \geq a_n\}$. a_1^m *is said to be a minimal commitment action.*
3. $u(a_i) = u(a_j)\ \forall\ a_i, a_j \in A_1^m$.
4. $A_{n'}(h_{n'})$ *is invariant to which action* $a_1 \in A_1^m$ *belongs to* $h_{n'}$ *if* $n' \neq n$.
5. A_1^m *contains no actions other than those defined by 1.*

Proposition 3 (Minimal Commitment Equivalence). *Suppose that A_n is represented as an interval in $\mathbb{R}$, $a_n \in [a,\ b]$, and that $U_i(a_n)$ is continuous over A_n. Suppose further that for any interval $I \subseteq A_n$, the utility vector set, $\mathcal{U}_n(I)$ has a Pareto frontier that consists of utility vectors from only pure actions (degenerate lotteries). Then, there exists an interval $A_n^r = [a',\ b] \subseteq [a,\ b]$ such that if A_1^m is a minimal action commitment set for A_n, and A_1^r and a required commitment set for A_n^r, both with the same commitment cost c, then the same discounted utility vectors for the selves and decision n realized actions will result from the agent choosing from $\mathcal{A}_1^r$ or from $\mathcal{A}_1^m$.*

Proposition 3 says that allowing the agent to commit to a minimum action from within an interval has the same outcome as allowing them to commit to a single action from a (weakly) smaller interval; note that both intervals have the same maximum; this is because committing to a minimum of b is the same as committing to b. Crucially, this allows the application of Proposition 2 to minimal commitment actions.[17] Notably, consumption-savings decisions, with strictly concave consumption utility, will satisfy the suppositions of Proposition 3.

It follows that, under the suppositions of Proposition 3, the use of a minimal action commitment device has the same effect on agent decisions as exogenously forbidding low values in a single action commitment device. The intuition here is fairly straightforward: minimal action commitment sets

[17] As well as maximal commitment actions, as any maximum commitment set can be reframed as a minimal commitment set; e.g. amount consumed versus amount saved.

eliminate low options, but still allow the agent to revise upward when the period to make the decision arrives. Single action commitment sets prevent any revision upward. Thus, for example, the model predicts that higher savings rates will result from having the option of choosing minimal savings amount, followed by a later decision on additional saving, than from having the agent decide their full amount of savings in advance. The long-term self essentially uses the early commitment possibility to eliminate some particularly bad options (low savings rates), and then re-negotiates from a stronger position when the time for the ultimate decision on savings comes around.

5 Temptation

I now turn to temptation effects: the observation that agents' decisions do not always satisfy independence of irrelevant alternatives. To be clear, by independence of irrelevant alternatives, I am referring to Sen's (1971) condition α, which states that if $x \in B \subset A$, and $x \in C(A)$, then $x \in C(B)$, where $C(A)$ is the set of elements chosen from set A. If $C(\cdot)$ is a singleton, then this property says that the removal of unchosen options should not alter the chosen option. I interpret violations of this property in choice behavior as arising out of temptation. The presence of an option that one self finds very desirable (tempting) increases the anticipated effort that self would exert in conflict. This manifests as the option exerting a "pull" on the outcome decided upon by the agent, even when the option is not chosen.

In the multiself bargaining model, the action decided upon by the agent will depend on the bliss actions of both selves, even if these actions are "irrelevant" options. This is due to the fact that it is the bliss points determine the resolution of the projected conflict game, and therefore the disagreement point that the selves anticipate in their bargaining. "Irrelevant" alternatives that are not the bliss point of either self will not influence the outcome; in this sense the model is in accord with existing literature on temptation, which primarily takes the view that it is only the "most tempting" point that is relevant. For expositional purposes, this section focuses exclusively on the temptation created by the bliss point of the short-term self, but the results apply similarly for the bliss point of the long-term self.

Consider an agent at a restaurant choosing between a (H)ouse salad, (G)rilled steak, and a (B)acon cheeseburger. The individual believes that B is the most delicious and H the least, but that the opposite is the case regarding the health effects of the choices. Suppose further that, as a result of health effects being something in the future, the preferences of the selves are such that $H \succ_l G \succ_l B$ and $H \prec_s G \prec_s B$, where $\succ_i$ indicates the preference ordering of the self i. One possible set of utility pairs granted by the three choices, then, is illustrated by Figure 9.

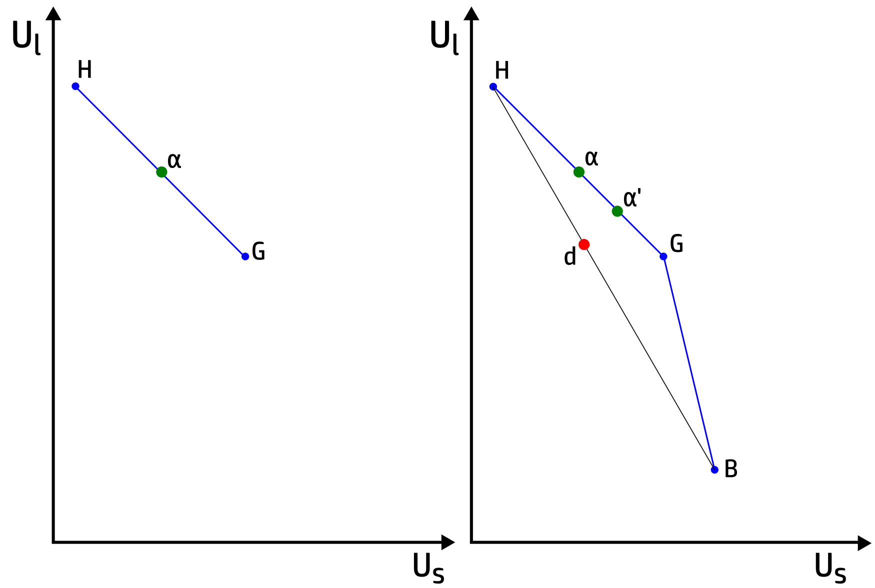

Figure 9: Menu choice with and without tempting option B.

The left graph of Figure 9 illustrates the decision made when the individual is choosing between only H and G. In the right graph we see how the outcome changes when we add in the third option, B. The point α shows the lottery chosen by the agent in the first case, while the addition of B moves the lottery to α', shown to the right. With the addition of B, though the individual still chooses B with zero probability (the outcome shown is a mixing between H and G), the weight placed on G grows. Intuitively, in the presence of an option more desirable to a self, that self is more willing to exert effort in the projected conflict game; this "pulls" the disagreement point, and thus the ultimate decision of the agent, closer to the bliss point of that self. It is important to note, though, that the long-term self is *also* more willing to exert effort with the addition of point B, since that self is much more opposed to B than they were to G. Temptation is always bi-directional in the model, and it is the differences in bliss utilities for *both* selves that determines the overall effect. It is therefore not always the case that adding an action to A_t which has a greater utility for the short-term self than their current bliss point moves the outcome in their favor.

To nail down the net temptation effect of a given option, I will break down the effect into component parts. First, Proposition 4 considers the effect of making the bliss point of the short-term self better or worse for the short-term self; in graphical terms this is moving B to the right or left in Figure 9.

Proposition 4 (Pure Temptation). *Consider a decision maker choosing α_n from $\mathcal{A}_n$, corresponding $\mathcal{U}_n$, and consider the effects of altering the utility granted to the short-term self by the short-term bliss point: X_n^s. As long as an open connected subset of the Pareto frontier of $\mathcal{U}_n$ containing $U_n = (U_s(\alpha_n),\ U_l(\alpha_n))$ remains a subset of the Pareto frontier of $\mathcal{U}_n$, then,*

$$\frac{\partial U_s(\mathcal{D}(\mathcal{A}_n))}{\partial X_n^s} \geq 0,\ \frac{\partial U_l(\mathcal{D}(\mathcal{A}_n))}{\partial X_n^s} \leq 0,$$

with the inequalities strict if the Pareto frontier is smooth at $(U_s(\alpha_n),\ U_l(\alpha_n))$.

The interpretation of Proposition 4 is straightforward: if we increase the utility granted to the short-term self by their bliss point, then the action chosen by the agent changes to one that grants a higher utility to the short-term self. This is the temptation effect in its purest and most intuitive form: as the desserts on the menu become more delicious, the agent is pulled more toward them. If the Pareto frontier is smooth, then the strictness of this change implies a continuous temptation effect, one which applies even if the bliss point is chosen with zero probability. However, if the slope is not defined at the outcome, then the temptation effect ceases to be strictly increasing.

For example, in Figure 9, if we move point B to the right, the decision will shift smoothly to the right as well; move B far enough, and the individual will eventually choose G for certain, and remain there for an interval as B is moved farther to the right. Once B is moved far enough, the individual would begin mixing between G and B, and the temptation effect would again be continuously increasing.[18]

To see why the result is true, note that in moving B in this way, the utility granted to the long-term self by B remains the same, and thus the difference to the long-term self between the bliss points remains the same (L_n constant). This means that the long-term self's incentive to exert effort in a potential conflict is unchanged, but the short-term self's incentive is increased. Thus, the projected probability of the short-term self winning a conflict is conclusively increasing, moving the disagreement point closer to B along the mixing line between H and B; this is compounded by the fact that the mixing line itself is shifting in favor of the short-term self (since the B endpoint is moving in favor of the short-term self). Finally, the fact that the disagreement point has shifted in favor of the short-term self means that the bargaining outcome will as well.

The relevance of requiring that a section of the Pareto frontier surrounding the decision be unchanged is to ensure that the changing point is an "irrelevant" alternative. Doing away with that condition, we obtain a more limited result:

[18] Though, once the agent was mixing with G and B, the Pareto frontier around the decision would also be changing, and it would no longer be a pure temptation effect; B would no longer be an "irrelevant" alternative.

Lemma 7. *Consider a decision maker choosing α_n from $\mathcal{A}_n$, with corresponding $\mathcal{U}_n$, and consider the effects of altering the utility granted to the short-term self by the short-term bliss point: X_n^s. Then,*

$$\frac{\partial U_s(\mathcal{D}(\mathcal{A}_n))}{\partial X_n^s} \geq 0.$$

Without the limitation on the Pareto frontier, an improvement in the bliss utility of the short-term self still moves the outcome in a direction favorable to the short-term self. However, it may also add Pareto improvements on the previous outcome, and so the net effect on the utility granted to the long-term self is ambiguous. This would occur if, in Figure 10, B was shifted far enough to the right that the Pareto frontier became the line segment between H and B (so that G was no longer on the frontier).

I now turn to the second component of temptation, which is the effect of making the short-term bliss point more or less desirable to the long-term self. In Figure 9, this would correspond to moving B up or down. If the bliss point of the short-term self becomes worse for the long-term self, then essentially the long-term self's (unchanged) bliss point becomes *relatively* better for the long-term self than it was before; intuitively, if the long-term self knows that losing a conflict would result in a very bad option, he will exert more effort in such a conflict. As it is the worsening of one bliss point that makes the other bliss point more appealing, I refer to this as indirect temptation. So, the projected probability that the short-term self wins a conflict conclusively decreases as a result of such a change.

However, if the short-term self wins, the outcome is worse for the long-term self than it was previous to the change. These two effects create an ambiguous net effect on the expected utility outcome of the projected conflict game, and thus on the disagreement point. Proposition 5 details the condition under which the the first effect dominates.[19] That is, the condition under which a worsening of the short-term self's bliss point for the long-term self leads to a better disagreement point, and thus a better outcome, for the long-term self.

[19]This is the only proposition in which the equational form is dependent on whether one uses the disagreement point with or without the effort costs added in, as discussed in section 2.3. Proposition 5 uses the disagreement point with effort costs added back in. Proposition 5', included in Appendix A, uses the disagreement point without effort costs added in; it has the same qualitative implications, and yields no additional intuition.

Proposition 5 (Indirect Temptation). *Consider a decision maker choosing* α_n *from* $\mathcal{A}_n$*, corresponding* $\mathcal{U}_n$*. Consider altering the utility granted to the long-term self by the short-term self bliss point:* Y_n^s*. As long as an open connected subset of the Pareto frontier of* $\mathcal{U}_n$ *containing the Nash bargaining outcome* $(U_s(\alpha_n),\ U_l(\alpha_n))$ *remains a subset of the Pareto frontier of* $\mathcal{U}_n$*, then*

$$\frac{S_n^\gamma + (1-\gamma)L_n^\gamma}{\gamma S_n L_n^{\gamma-1}} < \frac{U_l(\alpha_n) - d_l}{U_s(\alpha_n) - d_s} \Rightarrow \frac{\partial U_s(\mathcal{D}(\mathcal{A}_n))}{\partial Y_n^s} \geq 0, \frac{\partial U_l(\mathcal{D}(\mathcal{A}_n))}{\partial Y_n^s} \leq 0,$$

with the reverse strict inequality implying the reverse weak inequalities.

To interpret this proposition, I start by noting that the first term is a ratio of the changes in the coordinates of the disagreement point that result from an increase in Y_n^s: it is the slope along which the disagreement point moves as a result of such a change. The second term is the slope of the line connecting the disagreement point to the utility vector created by the current lottery choice, α_n. Note that if the Pareto frontier of $\mathcal{U}_n$ is smooth at $(U_s(\alpha_n),\ U_l(\alpha_n))$, then this slope is the negative of the slope of the Pareto frontier at that point; this follows from the nature of the Nash bargaining solution.

Thus, Proposition 5 says that if an increase in Y_n^s moves the disagreement point below the line connecting the original disagreement point to the original decision, then the lottery chosen by the agent changes to a lottery granting higher utility to the short-term self, and lower utility to the long-term self. Now, to get at some intuition for this, consider $\gamma = 1$; in which case

$$\frac{S_n^\gamma + (1-\gamma)L_n^\gamma}{\gamma S_n L_n^{\gamma-1}} = 1.$$

If additionally the Pareto frontier is smooth, then Proposition 5 can be more simply stated as: "If the slope of the Pareto frontier at the current decision is greater in magnitude than 1, then an increase in Y_n^s changes the outcome in favor of the short-term self (and similarly a decrease in Y_n^s changes the outcome in favor of the long-term self)." Tying this back into the menu-choice example, this condition would imply that shifting the B option downward would make the outcome worse for the short-term self; the additional incentive of the long-term self to exert effort in projected conflict would be the dominant effect of the shift. Intuitively, the relatively high slope of the Pareto frontier indicates that the long-term self has "more to lose" from shifts along the frontier; thus, said self will react more strongly to changes in options which have the potential to move the outcome along the frontier.

Now, putting together the implications of Propositions 4 and 5, consider the addition of a new option (such as menu choice B in the example) to an existing set of actions (such as $\{H,\ G\}$). Suppose that the new option is better for the short-term self than their existing bliss point, but worse for the long-term self. If the condition of Proposition 5 holds, then the net effect of such a new option is ambiguous: it increases X_n^s, which improves the outcome for the short-term self by Proposition 4, but decreases Y_n^s, which improves the outcome for the long-term self by Proposition 5. If the decrease in Y_n^s is sufficiently large relative to the increase in X_n^s, the outcome may actually improve in favor of the long-term self, in spite of the new point being a favorable option for the short-term self. I refer to this as a “backfire” effect, and such a situation is illustrated in Figure 10 below.

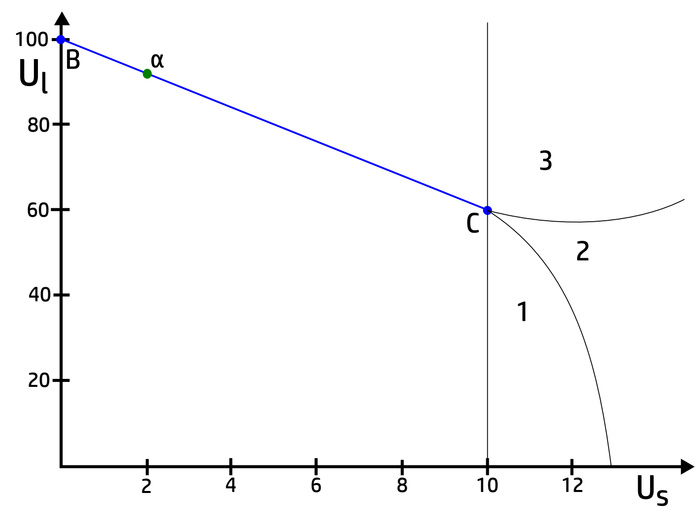

Figure 10: Regions dividing effect of adding a new bliss point; $\gamma = 1$

The Pareto frontier is Figure 10 is the line segment between B and C, with the outcome vector marked in between them as α. Consider adding another vector to the set of utility vectors. Region 1 in the graph is the “backfire” region; in this region the increase in bliss utility for the short-term self is relatively small, compared to the large potential decrease in utility for the long-term self. As a result, additions in this region cause the outcome to move to favor the long-term self; the indirect temptation effect dominates. Additions in region 2 cause the outcome to shift in favor of the short-term self while hurting the long-term self; the pure temptation effect dominates. For additions in region 3, the pure temptation effect still dominates, but the new option creates large enough Pareto improvements over the original outcome, due to shifting the Pareto frontier outward, that both selves benefit.

If the condition of Proposition 5 does not hold, however, there is no backfire effect, and region 1 is empty; is this case any new action which increases the bliss utility for the short-term self necessarily improves the outcome for the short-term self. Note now the distorted scale of Figure 10: the slope of this frontier is -4. The magnitude must be greater than 1 for region 1 to be non-empty, and a relatively high slope magnitude is required for region 1 to be significant. Connecting to the previous intuition, the steeper the slope, the more the long-term self has to lose from movement along the frontier, and thus the larger the potential "backfire" region will be. If the Pareto frontier is smooth, the curve separating regions 1 and 2 is defined by

$$(d_s^{'} - d_s)(U_l(\alpha) - d_l) = (d_l^{'} - d_l)(U_s(\alpha) - d_s),$$

where $(d_s,\ d_l)$ is the original disagreement point (defined by Lemma 1), $(d_s^{'},\ d_l^{'})$ is the disagreement point created by the original bliss point of the long-term self and the new option (a point on the curve), and α is the original outcome.[20] If the Pareto frontier is not smooth at the current outcome, then the boundary between regions 1 and 2 is itself a region as opposed to a curve; the addition of a new point anywhere inside this boundary region does not change the outcome.

I close this section by illustrating the temptation effect in a continuous action set. Consider again our illustrative savings example, and suppose that the agent has committed to saving at least 0.3. Then, when it is time to make the final savings decision, the bliss action of the short-term self is to save exactly 0.3, as 0.2 (their former bliss action) is no longer in A. As a result, the agent is less tempted toward low saving rates, and the action chosen by the agent moves from $a = 0.354$ to $a = 0.367$. This is shown in figure 11.

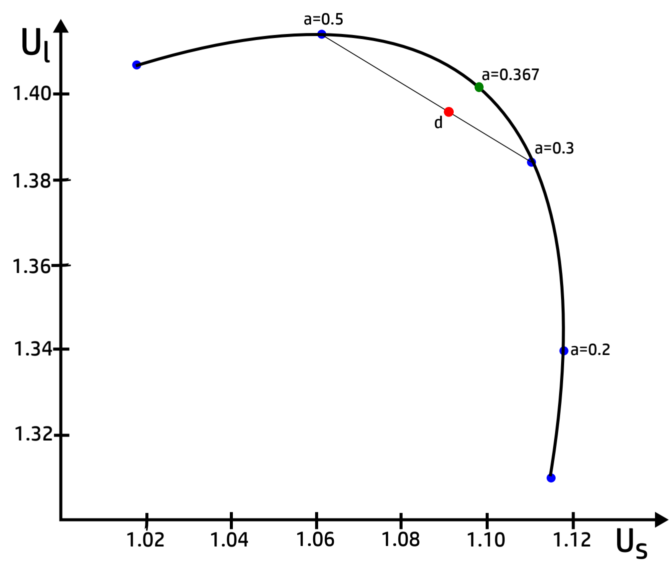

Figure 11: Savings decision with commitment to minimum of 0.3.

[20]This minor result is a direct implication of Lemma A.2, given in Appendix A.

6 Welfare Implications

While it presents more challenges to interpretation than the standard geometric discounting utility model, the dual-self bargaining model presented here need not be silent on questions of welfare. By choosing to view an agent as actually consisting of two individuals, there are several results that grant leverage for welfare evaluation. This gives an advantage over models in which there is a single individual whose preferences change in each period due to non-geometric time discounting. Rather than have one set of preferences for each period, here there are only two in total. There is also the observation that utilities are intrapersonal here, rather than interpersonal. Selves share a payoff utility from actions, and so utility comparisons arguably carry more weight here than they would when comparing across individuals. I first address the notion of Pareto improvements: actions that improved the utility for both selves over another action.

Lemma 8. *Given two actions a and b, a is never chosen from any action set A containing both a and b if and only if the utility vector created by b is a Pareto improvement over that created by a.*
Proof: If b is a Pareto improvement, then a will not be on the Pareto frontier, and so will never be chosen. If a is never chosen from $A = \{b, a\}$, then it must be that b is a Pareto improvement, as the agent would otherwise choose a strict mixing of the two actions.□

This has a close relation with the notion of the unambiguous choice relation developed by Bernheim and Rangel (2009). Essentially, if a is never chosen when b is available, then we say that b is *unambiguously preferred* to a, written as bP^*a, in the terminology defined by their work. Thus, in this model, b is unambiguously preferred to a if it represents a Pareto improvement over a in regard to the two selves. An immediate extension is that for a given action, a, if there exists another action, b, which is unambiguously preferred, a can definitively said to be Pareto inefficient from a welfare perspective. It also has the important implication that if an agent always makes the same decision from a binary choice (probability 1), then the choice made is a Pareto improvement over the other, and thus can be definitively said to be welfare improving.

I now turn to the more difficult question of welfare evaluation of options when neither is a Pareto improvement over the other. If we wish to make meaningful statements about welfare on such questions, it is necessary to consider some aggregation of the utilities received by the two selves. One method by which to do so is a weighted utility welfare function, $W(a) = wU_l(a) + (1-w)U_s(a)$, famously advanced by Harsanyi (1955, 1977) as a method of aggregating interpersonal social welfare. For what follows, I will confine attention to this form of welfare function.[21]

[21] A criticism of such a weighted utility function as a measurement of social welfare involves the difficulty in interpersonal utility comparisons, e.g. Sen (1977). In my model the utility comparisons are intrapersonal: the selves in fact share a payoff utility function, and I argue that this lessens the strength of the critique. Second, more practically, as the selves share their payoff utility function, it is desirable that any measure of welfare be invariant to

Lemma 9. *Consider a choice set given by $A = \{a,\ b\}$, suppose that b is the action preferred by the short-term self, and that the agent is observed to choose b with probability p. Then,*

$$p < \frac{1}{\left(\frac{1-w}{w}\right)^{\gamma} + 1} \iff wU_l(a) + (1-w)U_s(a) > wU_l(b) + (1-w)U_s(b).$$

Lemma 9 implies that, if welfare of an individual is evaluated by a weighting between options, observation of choice from a binary set of two actions is sufficient to say which is welfare superior.[22] If $w = 0.5$, so that welfare is considered as an equal weighting between the selves, then the condition simply becomes $p < 0.5$, so that γ need not be known, nor which action is preferred by which self. This is appealing for application: it would imply that in a series of random choices between two options the agent would be observed to choose the option one granting higher welfare more frequently.

We may also wish to consider welfare weightings that place greater weight on the utility of the long-term self.

Lemma 10. *Consider a choice set given by $A = \{a,\ b\}$, and suppose an agent is observed to choose b from A with probability p, and that b is the action preferred by the short-term self. Further, consider a utility weighting given by $wU_l(\cdot) + (1-w)U_s(\cdot)$, with $w \geq 0.5$. Then,*

$$p < 0.5 \Rightarrow wU_l(a) + (1-w)U_s(a) > wU_l(b) + (1-w)U_s(b).$$

Proof: Consider $w = 0.5$. From Lemma 9, $p < 0.5 \iff 0.5U_l(a) + 0.5U_s(a) > 0.5U_l(b) + 0.5U_s(b)$
$\Rightarrow U_l(a) - U_l(b) > U_s(b) - U_s(a) \Rightarrow w(U_l(a) - U_l(b)) > (1-w)(U_s(b) - U_s(a))$ for $w \geq 0.5$
$\Rightarrow wU_l(a) + (1-w)U_s(a) > wU_l(b) + (1-w)U_s(b)$ for $w \geq 0.5$.$\square$

Lemma 10 says that if a utility weighting welfare function places at least equal weight on the long-term self's utility, then observing the agent choosing the long-term self's preferred action more frequently than an alternate action implies that the more frequently chosen action grants higher welfare. This is of interest because if, for example, a bag of potato chips is resisted more often than not, then it implies that it is welfare improving to remove the bag of chips as an option. Additionally, it allows a degree of welfare evaluation to take place without taking a stance on whether an equal

affine transformations of this underlying payoff function, and a weighted utility welfare function meets this criterion.

[22] In cases where it is not apparent, which action is preferred by which self is easily established by time delay due to the diminishing impatience phenomenon

weighting between selves, or a higher weighting on the long-term self, is the correct choice in a welfare function.[23]

Finally, consider the welfare effects of the availability of commitment devices.

Lemma 11. *Consider action set A_n, and a welfare function $W(a) = wU_l(a) + (1-w)U_s(a)$, with $w > 0$. Consider decision 1, with $t_n - t_1 = \Delta_t \geq 0$. Then, $\exists \underline{\Delta}$ such that $\forall \Delta_t > \underline{\Delta}$, allowing the agent access to a zero-cost commitment action set at time t_1 increases welfare.*

Lemma 11 says that as long as you put positive weight on the utility of the long-term self, then a commitment option given sufficiently far in advance of the action being committed to is welfare improving. The implication is that if weighted utilities between selves is regarded as an acceptable method of welfare evaluation then, regardless of the weights used, commitment devices can create welfare improvements. In fact, the result will easily extend to any welfare function which is bounded, continuous, and strictly increasing in U_l. Thus, the model gives strong support to the notion that commitment devices create welfare improvements for agents.

7 Extensions, Future Work and Conclusion

One alteration to the model of interest is in the timing of the conflict game. Currently, selves compete for control, and then choose their action. An alternative method of modeling would be for the selves to first commit to an action, and then have their effort game, allowing for strategic choice of actions on the part of the selves. The general results for commitment and diminishing impatience are conjectured to remain unchanged, but the temptation results alter in that there would no longer be a "backfire" effect: neither self would strategically choose an action which would backfire in the way described, meaning that the decision process would become monotonic in the sense of Kalai and Smorodinsky (1975): adding an action which increased the maximal possible utility for a self would increase the outcome for that self. This is not necessarily an undesirable alteration, as the backfire effect is relatively limited in scope, and does not seem to correspond to any strong empirical regularities. However, at this time it is not known whether the equilibrium of a conflict game in which selves strategically choose their bliss points is generically unique.

Another alteration to consider is the bargaining procedure used; Nash bargaining was selected here for tractability, but the general properties of this model are dependent on the conflict game, not the bargaining procedure. Indeed, it is not difficult to show that the qualitative results of diminishing impatience, preference for commitment, and temptation, result from any bargaining procedure that satisfies two uncontroversial properties. First, invariance to affine transformations of utility. Second,

[23] Arguments for considering a short-term, or impatient self as *more* important than the long-term self, in contrast, do not exist in the literature on multiple selves, nor does the concept seem to carry any introspective weight.

the utility granted to each bargainer by the outcome should be monotonic in the utility granted to that self by the disagreement point. Other bargaining procedures which satisfy these properties, then, generate the same qualitative results.

Extending from this, an axiomatization of bargaining procedures which generate the desired behavior would be of keen interest. Obvious axioms of choice are Pareto efficiency and Symmetry, as well as invariance to affine transformations of the shared payoff utility function; this form of invariance is equivalent to the bargaining problem being invariant to shifts and non-distorting scaling of the bargaining set. Independence of Irrelevant Alternatives cannot be included, as it is the lack of this property which generates the temptation effects. Nor can the more general form of invariance to affine transformations that Nash bargaining builds upon be included; it is the lack of invariance when separate transformations are applied to the selves which generates time inconsistency in the model. Monotonicity, as Kalai-Smorodinsky bargaining, is certainly a plausible axiom. As mentioned, this removes the backfire effect. However, including Monotonicity is insufficient to pin down a unique solution, and it is not obvious if there is another natural axiom or axioms of choice. Other future work will certainly include symmetric multiself models in which the selves vary in dimensions other than time discounting. In particular, selves that vary in a risk aversion parameter is of interest in attempting to generate regularities related to risk.

To conclude, this work formalizes intuition about conflicting internal preferences into a model that provides a unified explanation for a number of behavioral regularities. A smooth time inconsistency arises from the relative difference in time preference between the selves. Temptation effects, or violations of independence of irrelevant alternatives, result from the differing incentives of the selves in conflict over control. Use of commitment devices derives from the long-term self's advantage in foresight over long time horizons. All of these come from a tightly parameterized difference in time preference between the selves, which additionally creates novel intuition about the nature of such behavior.

Part II

Temporal Reference Points

1 Introduction

Extensive evidence exists that agents are sensitive to the pattern and timing in which information is revealed to them, even when they cannot make use of that information. Gneezy and Potters (1997) and Haigh and List (2005) find that agents exhibit greater risk aversion the more frequently that they evaluate the state of financial outcomes. In particular, Bellmare et al. (2005) isolates information feedback, and not investment flexibility, as the variable of greater influence. Köszegi and Rabin (2009) develop a model in which agents prefer to have their information clumped together rather than spread out.

Concurrently, there is a body of work indicating that the way in which time is framed and presented influences how individuals evaluate situations put before them. For example, Chandran and Menon (2004) finds that agents judge risk to be greater when health risks are presented to them in daily terms rather than yearly, even though the overall risk level is objectively the same. Gourville (1998) examines the effect of the temporal framing of costs (e.g. "just pennies per day!") on the the decisions of individuals, and finds them to be significant.

These two groups of research suggest a common thread: there may be something special about the points in time to which an agent's attention is drawn. I leverage this idea to create a model which unites the concept of diminishing impatience with preferences for clustered information. Further, the model provides a novel explanation for some discordant results in studies on diminishing impatience.

I make the assumption that the value of a prospect is not only dependent on the distribution over ultimate outcomes, but also on a set of subjectively important times in the life of that prospect, which I term "temporal reference points". When evaluating a prospect the agent is assumed to have a discount function which is applied to each of the durations between these temporal reference points. This replaces standard time discounting, in order to model the sensitivity the agent may have with regard to these points in time.

I remain agnostic about how agents form these temporal reference points (though I believe it is a promising future avenue of research). Instead, I examine two natural cases. The first is the case where an agent forms the temporal reference points solely from the times he expects to receive new information about the outcome of a prospect. This agent is sensitive to updates of information on the ultimate value of a prospect, but not sensitive to other possible framing effects (such as emphasis on a particular point in time in the presentation of a set of prospects). I refer to such an agent as "information focused". The second is a case where shifting a prospect into the future causes an agent to treat the new starting point of the prospect as a temporal reference point.

Section 2 develops the notation and fundamentals of the model, which addresses the preferences of agents over prospects set in continuous time. In particular, agents are characterized by a utility function, a discounting function, and a framing mapping, the last of which represents the process by which the agent decides upon the temporal reference points of the prospect. If the framing maps to $\{0\}$, so that the only temporal reference point is *now*, then the model becomes standard expected utility, with the discount function reverting to standard time discounting. Thus, the model can be viewed as a generalization of standard models. Prospects are characterized by distributions over the set of possible outcomes, of which exactly one will occur in finite time. Further, each prospect has a set of possible histories of informational updates given to the agent (signals). The history at any given time affects both the distribution over future signals, as well as the expected distribution over outcomes.

Section 3 examines the case where agents are information focused. I develop the notion of preference for grouped information (PGI), which indicates a preference for informational updates to be folded into one another, reducing the total number of times at which information is gained. PGI is similar in spirit to Dillenberger (2010), which develops Preference for One-shot Resolution of Uncertainty (PORU), but here I place the notion in a time-sensitive setting, whereas PORU is defined in time neutrality. I show an equivalence between PGI and diminishing impatience in this environment. Diminishing impatience, an extensively documented property of agents' behavior (e.g. Thaler 1981) is the idea that an agent becomes less impatient between an immediate and a delayed alternative as both are delayed further.

A second, more refined, type of information preference is also developed. This is a preference for informational updates to be closer to each other temporally, which I call a preference for less dispersed information (PLDI). I show the equivalence between PLDI and strongly diminishing impatience, a refinement of diminishing impatience which requires an agent to become continually less impatient as rewards become more distant. As shown by the survey Frederick et al. (2002) there is evidence for strongly diminishing impatience, though it ceases to diminish after a period (they show that after a year there is zero evidence of further diminishing impatience). I am unaware of any empirical studies examining PLDI at the present time, but this work suggests such studies may be of interest. PLDI is further shown to be a generalization of PGI, with grouped information being a limiting case of less dispersed information. This limiting relationship is of importance in the continuous time setting, in which defining a single point in time at which information is received is potentially problematic.

Section 4 relaxes the assumption of information focused agents, which opens up many possibilities as to how agents may determine temporal reference points. I focus on a single case where the act of shifting a prospect into the future causes the agent to treat the new starting point of the prospect as a temporal reference point. I show that in such a case the agent's indifference curves over prospects are also shifted into the future by the same amount, retaining indifference between equally shifted prospects.

This particular case of temporal reference points is important because it suggests a novel resolution to apparently contradictory experimental evidence concerning diminishing impatience. A very large body of work attests that diminishing impatience becomes less strong as prospects are pushed further into the future, e.g. Frederick et al. (2002). Numerous studies show that subjects exhibit preference reversals when a set of options is made more distant, e.g. Keren and Roelofsma (1995). However, the thorough experimental survey Glimcher et al. (2007) examines the effect of shifting a set of prospects into the future, and finds disagreement with this previous work. The authors show that when the shifted prospects retain their temporal relation to one another, subjects' indifference curves remain statistically constant over the prospects. Specifically, the diminishing impatience across the set of prospects was just as strong, despite the set being more temporally distant. The model here offers explanation for this apparent conflict by supposing that the experimental design of the Glimcher study, where *all* of a very large number of delayed comparisons were made to a $20 payout in 60 days, caused the subjects to focus on 60 days as a subjectively important point in time. Thus, subjects plausibly form a temporal reference point in that study that would not be formed in other studies without a similar point to focus on.

Section 5 concerns the discounting function and how it can be interpreted in this model given the multiple steps of discounting occurring in the evaluation of prospects. I focus on two interpretations in particular. One in which the discounting is divided into geometric time discounting and an additional survival function, which can be thought of as encompassing implicit risk. In this interpretation, the agent sees the temporal reference points as problem spots, where something might go wrong with the explicitly stated outcomes of the prospect. This may be due to distrust of the explicitly stated prospect, mortality risk, or just general pessimism. The second interpretation treats the discounting function as time discounting, but has the agent place himself in the role of his future selves, evaluating his preferences "as if" he were those selves, before discounting backward. In essence, the temporal reference points become the times to which the agent projects himself.

Section 6 concludes and proposes further avenues for research, in particular research on the formation of temporal reference points.

2 The Model

2.1 Decision Problem and Notation

An agent is faced with a choice between multiple prospects at time 0, with the set of all possible prospects denoted $\mathcal{P}$. There is a closed and compact set of payouts, $\mathcal{A}$, and every prospect terminates with a payout in finite time with probability 1.[24] Further, there is a closed and compact set of pairs $\mathcal{O} = \{(\alpha, t_\alpha)\}$, where $\alpha \in \mathcal{A}$ and t_α indicates the time at which the payout occurs. Elements of $\mathcal{O}$ are referred to as outcomes.

Time is continuous, but each prospect, $A \in \mathcal{P}$, has a finite number of distinct, ordered, positive points in time,

$$I_A = \{t_1^A,\ t_2^A,\ ....t_N^A\},$$

at which new information, a signal, is received by the agent as to the outcome of the prospect.[25] The signal from prospect A received at time t_n^A is denoted s_n^A.

At any point in time in the life of a prospect, t, the agent will have seen a history of signals $h_t^A \in H_t^A = \{s_1^A,\ s_2^A, ...s_i^A\}$, where t_i^A is the largest element of I_A strictly less than t, of what has occurred at each of the elements of I_A (what information was revealed) up to, but not including, t. The distribution over possible s_t^A is a function of the signals received to that point, h_t^A. The history of signals and the prospect itself, then, define a mapping that gives the distribution over the final possible outcomes of the prospect, denoted for a given history h_t^A by

$$\triangle_A(h_t^A) :\ \mathcal{O} \longrightarrow [0, 1].$$

These distributions are assumed to be consistent with the probabilities of signals. So, the distribution over outcomes just before an informational update is equivalent to the distribution over signals expected at that update composed with the distributions induced by the receipt of those signals. Intuitively, this represents the agent rationally and recursively determining the distribution over outcomes based on the probability of each signal and the distribution over outcomes induced by those signals.

[24] Although the elements of $\mathcal{A}$ are not limited to monetary payments, nothing will be lost in the intuition of the model by considering them to be.

[25] Note that this excludes the possibility that uncertainty is resolved at time 0. While this assumption would be very limiting in a discrete time setting, here it is not. Given that time 0 is the time at which the agent is *evaluating* the prospect, this is just assuming that uncertainty is not being resolved *while* this evaluation is occurring. Uncertainty can still be resolved arbitrarily close to time 0, which allows the accounting of cases where, say, an agent is offered a choice between two gambles, and the dice are rolled moments after their choice.

At the beginning of the prospect looking forward to any time t generates an expected distribution over histories, each of which generates a distribution over outcomes. Thus each time t for a prospect has associated with it a distribution over distributions over outcomes, denoted Λ_A^t. At time 0, due to there being no history, this distribution over distributions is degenerate on a single distribution over outcomes. That is, Λ_A^0 is degenerate on a single distribution, denoted $\triangle_A^0$.

Appended to the set I_A are the times 0 and t_F^A, where $t_F^A > t_N^A$ is the time of the last outcome(s) generated by prospect A.[26] This defines the set:

$$T_A = \{0,\ t_1^A,\ t_2^A,\ ....t_N^A,\ t_F^A\}.$$

For ease of notation, t_0^A will sometimes be used in place of zero, and t_{N+1}^A in place of t_F^A, so that $T_A = \{t_0^A,\ t_1^A,\ t_2^A,\ ....t_N^A,\ t_{N+1}^A\}$. Payouts may be received at any of the elements of T_A, and once a payout is received, the prospect terminates with no further payouts possible.

Each agent further associates with each prospect a set of non-negative ordered points in time,

$$R_A = \{0,\ r_1^A,\ r_2^A,\ r_M^A\},$$

which will be referred to as the *temporal reference points* of that prospect. At minimum, this set includes 0. So, each agent has a mapping from the set of prospects to the set of subsets of the non-negative real numbers,

$$\mathcal{R}:\ \mathcal{P} \longrightarrow \mathbb{P}(\mathbb{R}^+).$$

This mapping can conceptually be thought of as how the agent frames the subjectively important points in time when evaluating the prospect. As will be shortly seen, the elements of R_A will determine how the agent applies discounting to their evaluation of the prospect.

The agent further has a utility function over outcomes, $u(\cdot) : \mathcal{A} \longrightarrow \mathbb{R}^+$, as well as a discounting function, $D(\cdot) : \mathbb{R}^+ \longrightarrow [0,\ 1]$, which is decreasing, non-negative, and satisfies $D(0) = 1$. The discount rate should be understood to not only include time discounting, but potentially other forms of discounting future payouts as well.[27] Further interpretation of this discount function is dealt with in Section 5.

[26] Note that $t_F^A > t_N^A$ implies the assumption that there is at least one possible payout/utility received after all uncertainty is resolved. Like the case of excluding 0 from I_A, this is not limiting in a continuous time setting. A payout can be modeled as being arbitrarily close to knowledge of the payout, such as when an agent receives a monetary payment mere moments after knowing that he will.

[27] However, in the special case where $\mathcal{R}$ maps all prospects A to $\{0\}$, so that the only important time is *now*, the model becomes standard expected utility with $D(t)$ representing normal time discounting (and thus encompassing both geometric and non-geometric time discounting models).

In evaluating the present discounted utility of prospects the agent discounts based on his temporal reference points, rather than discounting based on the time between now and the payouts as is standard in the literature. For each payout he discounts between now (the time at which he is evaluating) and the first such temporal reference point, between each pair of temporal reference points before the payout, and between the last temporal reference point before the payout and the time of the payout.

Formally, let $r^A_{\hat{n}(t)}$ be the highest element of R_A which is less than t. Then an agent at time 0 evaluates the present discounted value of a prospect A as

$$V_A = \int_{(\alpha,t)\in\mathcal{O}} u(\alpha) * \triangle^0_A(\alpha,t) * D(r^A_1 - 0) * D(r^A_2 - r^A_1) .... * D(r^A_{\hat{n}(t)} - r^A_{\hat{n}(t)-1}) * D\left(t - r^A_{\hat{n}(t)}\right),$$

and prefers the prospect with the highest such present discounted value.[28]

Definition. *Consider an agent whose mapping,* $\mathcal{R}$ *satisfies* $\mathcal{R}(A) = I_A \cup \{0\},\ \forall A \in \mathcal{P}$. *This agent is said to be* information focused.

An information focused agent regards as important all the times at which new information will be received, in addition to right now. Section 3 will examine this type of agent, whereas Section 4 will relax this assumption on the mapping $\mathcal{R}$.

Noting that for an information focused agent each outcome occurs at an element of T_A, we can write the present discounted value of a prospect A for an information focused agent as:

$$V_A = \int_{(\alpha,t^A_n)\in\mathcal{O}} u(\alpha) * \triangle^0_A(\alpha,t^A_n) * D(t^A_1 - 0) * D(t^A_2 - t^A_1) .... * D(t^A_{n-1} - t^A_{n-2}) * D\left(t^A_n - t^A_{n-1}\right).$$

Finally, in cases where $D(t)$ is continuously differentiable, I denote $h(t) = \frac{-D'(t)}{D(t)}$, the *hazard rate* of $D(t)$. The hazard rate gives us what is essentially the instantaneous rate of discounting. For example, for the geometric $D(t) = e^{-\rho t}$, $h(t) = \rho$. This will be of aid in intuition for several of the results.

[28] This representation makes clear that for $D(t) = e^{-\rho t}$ the model collapses to standard geometric time discounting.

2.2 Example With Visualization

Given a prospect with countable support over signals after each history of signals (which implies countable outcomes due to finite history length), we can visualize the prospect as a branching tree. The visualization will aid intuition even in the non-countable case. Aligning the tree from left to right, with the root at $t = 0$, and points further to the right as points farther in time, at each of the elements of I, when new information is received, the tree splits.

For example consider the following prospect, A. At time t_1 the agent will receive a with probability p. If a is not received then at time t_2 the agent learns what payout he will receive at time t_3. There is a q probability of receiving b, and a $1 - q$ probability of receiving c. The prospect is visually represented in Figure 1.

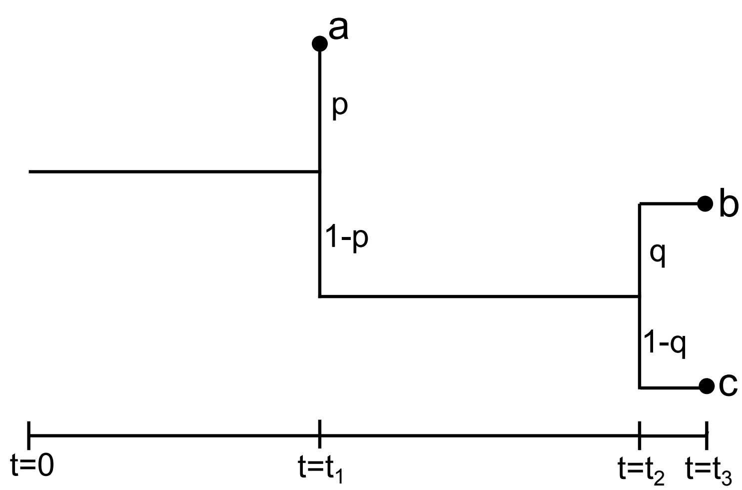

Figure 1: Prospect A with $T_A = \{0,\ t_1,\ t_2,\ t_3\}$

Each split in the tree is an informational update to the agent. At time t_1, they either receive a, which is information in itself, or do not, which is a signal that they will receive either b or c. For example, if h is the history containing the signal that a was not received at time t_1, then $\triangle_A(h)$ places q weight on an outcome of $(b,\ t_3)$ and $1 - q$ weight on an outcome of $(c,\ t_3)$.

If the agent is information focused for this example, then the value of the prospect would be

$$V_A = u(a)pD(t_1)+u(b)(1-p)qD(t_3-t_2)D(t_2-t_1)D(t_1)+u(c)(1-p)(1-q)D(t_3-t_2)D(t_2-t_1)D(t_1),$$

or more succinctly

$$V_A = D(t_1)\left[pa + (1-p)D(t_3 - t_2)D(t_2 - t_1)\left(qb + (1-q)c\right)\right].$$

3 Diminishing Impatience and Information Timing Preferences

In this section, I examine the case of an information focused agent. I begin by defining the notion of diminishing impatience in the context of this model. An agent who exhibits diminishing impatience becomes less impatient for a later outcome relative to an immediate outcome as both payouts become more distant. More formally,

Definition. *Denote by* $(a,\ t)$ *the deterministic prospect that provides payout* a *at time* t. *If an agent's preferences over prospects are such that*

$$(a,\ 0) \sim (b,\ t_1) \Rightarrow (a,\ t_2)\ \preceq (b,\ t_1 + t_2)\ \forall\ a, b \in \mathcal{A},\ \forall\ t_1, t_2 > 0,$$

then the agent exhibits diminishing impatience. *If the preference is everywhere strict, they exhibit* strictly diminishing impatience.

If the agent's preferences are further such that

$$(a,\ t_1) \sim (b,\ t_2) \Rightarrow (a,\ t_1 + t_3) \preceq (b,\ t_2 + t_3)\ \forall\ a, b \in \mathcal{A},\ \forall\ t_2 > t_1 \geq 0,\ t_3 > 0,$$

then the agent exhibits strongly diminishing impatience. *If the preference is everywhere strict, they exhibit* strictly strongly diminishing impatience.[29]

Diminishing impatience here says that if an agent is indifferent between an immediate reward and a later, better reward, he will prefer the later reward if both are equally delayed. Strongly diminishing impatience is a refinement of this which requires that preference change holds true for *any* sooner versus latter reward, not only when one is immediate.

Lemma 1. *An information focused agent exhibits (strictly) diminishing impatience if and only if*

$$D(t_1 + t_2)\ (>) \geq\ D(t_1)D(t_2)\ \forall\ t_1, t_2 > 0.$$

The agent exhibits (strictly) strongly diminishing impatience if and only if

$$D(t_2)D(t_1 + t_3)\ (<) \leq\ D(t_1)D(t_2 + t_3)\ \forall\ t_2 > t_1 \geq 0,\ t_3 > 0.$$

Further, if $D(t)$ *is continuously differentiable, then the agent exhibits (strictly) strongly diminishing impatience if and only if the associated hazard rate,* $h(t)$, *is (strictly) decreasing.*

Proof: All proofs not in the text are given in Appendix B.

[29] Similar notions are introduced by Chakraborty and Halevy (2015), with strictly diminishing impatience here corresponding to their definition of delay independent diminishing impatience (DIDI) and strictly strongly diminishing impatience here corresponding to their definition of delay independent strongly diminishing impatience (DISDI). In a discrete time setting without temporal reference points (or, equivalently, with the only temporal reference point being 0), the notions coincide exactly.

So, an agent exhibits diminishing impatience if and only if he discounts more heavily when a length of time is broken up into two intervals before discounting. This occurs in discounting functions with a steep "drop-off" at $t = 0$. For example, the well known $\beta - \delta$ discounting function,

$$D(t) \equiv \begin{cases} 1, & if\ t = 0 \\ \beta e^{-\delta t}, & otherwise. \end{cases}$$

Intuitively, breaking the length of time in two forces the overall discounting to have the "drop-off" twice.

As for strongly diminishing impatience in Lemma 1, the continuously differentiable case with a decreasing hazard rate lends us insight: an agent exhibits strongly diminishing impatience if he discounts steeply at first, with a high hazard rate, in the time immediately following a temporal reference point, and then less and less steeply as time continues. This matches, for one common example, the hyperbolic form of discounting.[30] The intuition matches that of the diminishing impatience case: if time is split in two, the overall discounting now has two periods of steep discounting rather than one.

Since we are dealing with deterministic prospects here in the referenced definition of diminishing impatience, there are no points at which new information is received. Thus, $T = \{0, t_F\}$, where t_F is the time the payout is occurring. The restriction to an information focused agent, then, ensures that there are no intermediate temporal reference points. Thus, discounting is based on the entire time between the present and the payout, becoming analogous to standard models. This should not be mistaken to mean that restricting attention to information focused agents in this model is no different than applying $\beta - \delta$ or hyperbolic discounting; the similarity only extends to evaluating deterministic prospects, when there is no new information to be had.

Similarly, Lemma 1 should not be misunderstood to be saying that the agent is exhibiting diminishing impatience *because* the time is broken up into two segments. Rather, the condition which implies diminishing impatience is the same as the condition while implies more discounting when time is broken up. It is this equivalence that will help lead to my first result, but first I must address the ranking of information levels of prospects, and preferences over them.

[30] Under generalized hyperbolic discounting, a reward in t time is discounted by $(1 + \alpha t)^{-(\gamma/\alpha)}$.

Definition. *Consider two prospects* $A, B \in \mathcal{P}$ *such that* $I_A \subset I_B$. *Then, prospect* A *is said to have* more grouped information *than prospect* B *if* $\|I_B\| - \|I_A\| = 1$ *and* $\exists\ i \geq 2$ *s.t.*

$$\Lambda_A^{t_n^A} = \Lambda_B^{t_n^B}\ \forall\ n\ s.t.\ 0 \leq n < i-1,\ and$$

$$\Lambda_A^{t_{n-1}^A} = \Lambda_B^{t_n^B}\ \forall\ n\ s.t.\ i \leq n \leq \|I_B\|,$$

or if $\exists\ C \in \mathcal{P}$ *s.t.* A *has* more grouped information *than* C, *and* C *has* more grouped information *than* B. *Further, an agent's preferences exhibit* preference for grouped information (PGI) *if they satisfy* $A \succeq B$ *for all such* A, B. *If the preference is everywhere strict, then denote it* SPGI.

This definition is slightly cumbersome notationally, but not hard to understand intuitively. In essence, A has more grouped information than B if the two prospects are the same except for an instance in which A gives the same information in one update that B gives in two updates (more than one instance of such is accounted for by the transitive part of the definition). I will illustrate the definition with two examples.

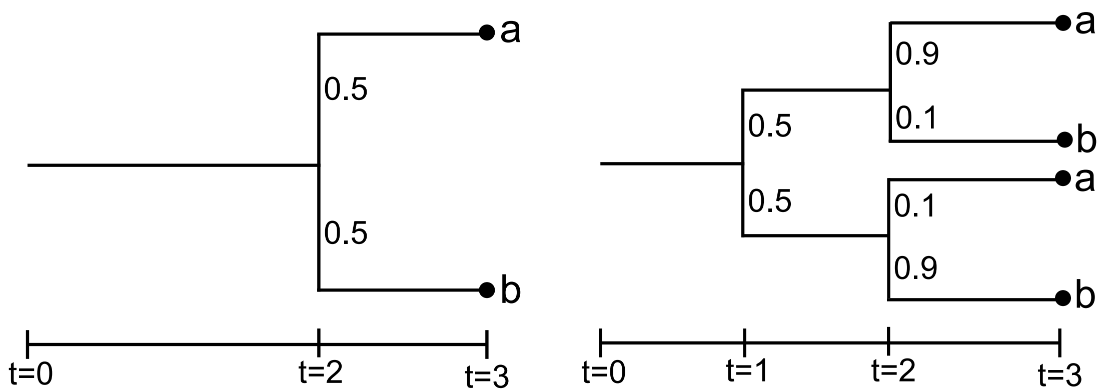

Figure 2: Prospect A, left, has more grouped information than prospect B, right.

Prospect A in Figure 2 has $I_A = \{2\}$, while prospect B has $I_B = \{1,\ 2\}$. In the scope of the definition, the relevant i is 2. So, to satisfy the definition, two things have to be true. First, it must be the case that $\Lambda_A^{t_0^A} = \Lambda_B^{t_0^B}$. This is just the requirement that they have the same initial distribution over outcomes, which they clearly do. Second, $\Lambda_A^{t_1^A} = \Lambda_B^{t_2^B}$. That is, the distribution over distributions expected at $t_2^B = 2$ under prospect B is the same as that at $t_1^A = 2$ under prospect A. Indeed, this is the case; at those times both prospects have 0.5 weight placed on the degenerate outcome of $(a,\ 3)$, and 0.5 weight placed on the degenerate outcome of $(b,\ 3)$. Intuitively, prospect B communicates the same information at times 1 and 2 that prospect A communicates at time 2 alone.

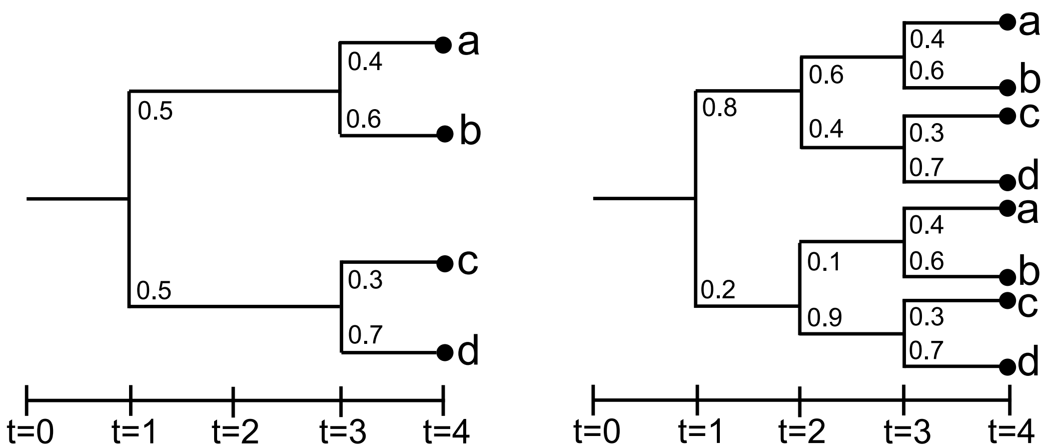

Figure 3: Prospect A, left, has more grouped information than prospect B, right.

In this slightly more complex example, the i that satisfies the definition is, again, 2. It is easy to confirm that the initial distribution over outcomes is the same. It must also be the case that $\Lambda_A^{t_1^A} = \Lambda_B^{t_2^B}$ and $\Lambda_A^{t_2^A} = \Lambda_B^{t_3^B}$. We see that at time $t_2^B = 2$, $\Lambda_B^{t_2^B}$ places weight $0.8 * 0.6 + 0.2 * 0.1 = 0.5$ on the distribution placing 0.4 weight on outcome a and 0.6 weight on outcome b, and weight $0.8 * 0.4 + 0.2 * 0.9 = 0.5$ weight on the distribution placing 0.3 weight on outcome c, and 0.7 weight on outcome d. $\Lambda_A^{t_1^A}$ is an identical distribution over distributions. $\Lambda_A^{t_2^A}$ and $\Lambda_B^{t_3^B}$ can be similarly verified to be the same distribution over degenerate outcomes. Again, in essence, Prospect A does in one step what prospect B does in two.

The agent who has preferences for grouped information prefers to eliminate points in time at which information is revealed by grouping information together.

Proposition 1. *An information focused agent exhibits (S)PGI if and only if he exhibits (strictly) diminishing impatience.*

Here the first result draws an equivalence between a preference for more grouped information and diminishing impatience. The intuition for the result is not difficult: a prospect with less grouped information implies that the times which are being discounted over are further broken up. As was previously established, the condition for diminishing impatience is the same as that which implies more discounting when lengths of time are subdivided.

The restriction on the nature of the agent's framing function, $\mathcal{R}$, is necessary here; if $\mathcal{R}$ were unrestricted it may be that a prospect A contained more temporal reference points in spite of having more grouped information. The restriction to information focused agents is not the minimum restriction needed for the equivalence, but it is the most natural and intuitive such restriction.[31]

[31] The minimum necessary restrictions for the equivalence would be first that all informational update times are temporal reference points, second that removing an informational update time removes that point as a temporal reference point, and third that removing an informational update time does not add new temporal reference points.

I now turn to a different notion of information preferences, and that is preferences over the relative dispersion of informative points in time.

Definition. *Consider* $A,\ B \in \mathcal{P}$ *such that* $\|I_A\| = \|I_B\|$. *Then* A *has* less dispersed information *than prospect* B *if* $\exists\ C \in \mathcal{P}$ *s.t.* A *has* less dispersed information *than* C *which has* less dispersed information *than* B*, or if all of the following is true. First,* $\|I_A \setminus I_B\| = \|I_B \setminus I_A\| = 1$*. Second, the single unique element in* I_A *and* I_B *have the same index in their respective sets, denoted* n'*. Third,*

$$\Lambda_A^{t_n^A} = \Lambda_B^{t_n^B}\ \forall\ n,\ 0 \leq n \leq \|I_A\|.$$

Fourth,

$$min\{t_{n'}^A - t_{n'-1}^A,\ t_{n'+1}^A - t_{n'}^A\} < min\{t_{n'}^B - t_{n'-1}^B,\ t_{n'+1}^B - t_{n'}^B\}.$$

Further, an agent is said to have preference for less dispersed information (PLDI) *if his preferences are such that* $A \succeq B$ *for all such* $A,\ B$*. If the preference is everywhere strict, then denote it* SPLDI*.*

The definition essentially says: take a prospect, B, take a point in time at which new information is revealed in that prospect (but not one where outcomes are received, since from $\Lambda_A^{t_0^A} = \Lambda_B^{t_0^B}$, A and B must have the same distribution over outcomes) and shift it so that its nearest neighbor in the set T is closer than its previous nearest neighbor. Then the resulting new prospect, A, has less dispersed information. Note that the conditions of the definition imply that the same information is revealed at the shifted point in both prospects; all that changes is *when* it is revealed.

More succinctly, if you push informative elements of T_A closer together, the resulting prospect has less dispersed information. The transitive part of the definition says that if you do this multiple times, the end result has less dispersed information than what you started with. A few examples will make this clear. Consider, $T_A = \{0,\ 1,\ 4,\ 6\}$ and $T_B = \{0,\ 2,\ 4,\ 6\}$.

Prospect A satisfies the fourth criterion for having less dispersed information than prospect B. If their informational update points give the same information, then A has less dispersed information than B. The relevant n' for the definition is 1, since it is element t_1 that is different. t_1^A is closer to its nearest neighbor than t_1^B. Specifically,

$$min\{t_1^A - t_0^A,\ t_2^A - t_1^A\} = 1 \text{ and } min\{t_1^B - t_0^B,\ t_2^B - t_1^B\} = 2.$$

Now consider $T_C = \{0,\ 1,\ 5,\ 6\}$. Assuming the first three criteria are satisfied for prospects C and A, Prospect C has less dispersed information than prospect B because it has less dispersed information than A. Conceptually, t_2^A, at 4, has been shifted to t_2^C, 5.

Proposition 2. *An information focused agent exhibits (S)PLDI if and only if he exhibits (strictly) strongly diminishing impatience.*

The intuition here is a bit less obvious. Less dispersed information, for example, means moving from discounting over two moderate lengths of time to discounting over one shorter and one longer length of time. If one pictures the tail of one of the moderate lengths being removed and added to the other moderate length, this creates one shorter and one longer period. If it is the case that the discounting becomes less steep as time goes on, then the tail removed is discounted more heavily then the tail added. This, then, connects back to Lemma 1. In that Lemma, one can see t_2 and $t_1 + t_3$ as being the moderate times, whereas t_1 and $t_2 + t_3$ are the shorter and longer times (because $t_1 < t_2$).

Lemma 2. *An information focused agent exhibits (S)PGI if they exhibit (S)PLDI.*
Proof: Direct from Propositions 1 and 2, and the fact that (strictly) strongly diminishing impatience implies (strictly) diminishing impatience.□

This relationship concerning different types of information preferences is not surprising. A prospect with more grouped information can be thought of as a limiting case of a prospect with less dispersed information, in which a shifted informative element of T is shifted close enough to its nearest neighbor to coincide with it in the limit. This limiting relationship is important conceptually, as it allows the sidestepping of tricky issues of how to define a single point in time at which an agent received information. Further, the combination of both types of preference allows us to fill in some gaps left by the definition of PGI.

To see what is meant by this, consider two prospects, A and B, with $T_A = \{0,\ 3,\ 6\}$ and $T_B = \{0,\ 2,\ 4,\ 6\}$, where prospect A gives the same information at time 3 that prospect B spreads over times 2 and 4. Intuition may lead toward regarding prospect A as having more grouped information than prospect B, but they do not satisfy the definition. However, prospect C, with $T_C = \{0,\ 3,\ 4,\ 6\}$ does have less dispersed information than B, and A has more grouped information than C. So, an agent exhibiting PLDI would satisfy $A \succeq C \succeq B$.

However, this does not necessarily extend to any such situation. Consider instead $T_A = \{0,\ 1,\ 2.5,\ 4\}$ and $T_B = \{0,\ 1,\ 1.1,\ 3.9,\ 4\}$ where, similar to the previous example, prospect A gives the same information at time 2.5 as B spreads over 1.1 and 3.9. In this case, there is no intermediate C with less dispersed information than B. Intuitively, this is because B's information is already clustered in tight temporal groups, and an agent exhibiting PLDI might prefer they stay clustered. Thus, knowledge of PLDI is not enough here to determine the agent's preference between the two prospects.

An agent whose $D(t)$ is represented by a $\beta - \delta$ function, and thus who exhibits SPGI as well as complete indifference to dispersion, would prefer prospect A; such an agent cares only about the number of groupings he receives the information in. However, a $\beta - \delta$ discounting function is arguably a poor choice for modeling in a continuous time setting, as it sees no distinction between a minute gap and a day. This is especially true if uncertainty resolution is to take place arbitrarily soon after the evaluation of prospects. For this reason, in this setting, a hyperbolic discounting function with very steep initial discounting that rapidly approaches geometric discounting is likely a superior choice to model agents who exhibit diminishing impatience, but show little evidence of strongly diminishing impatience.

4 Non-Informative Temporal Reference Points

In this section, the assumption that agents are information focused is relaxed, so the agents may regard non-informative points in time as important. This is done to accommodate the theoretical case that an agent may focus on a particular point in time based on how prospects are presented or framed, such as if an experiment repeatedly involves a particular point in time.

First, some notation. For a given set, S, denote:

$$S^{+t} \equiv \{s + t|\ s \in S\}$$

Definition. *A prospect X is said to be* τ-shifted from prospect Y, *if*

$$I_X = I_Y^{+\tau},$$

and if $\forall$ distributions over outcomes, $\triangle_X$ and $\triangle_Y$ satisfying $\triangle_X(\alpha,\ t+\tau) = \triangle_Y(\alpha,\ t),\ \forall\ (\alpha, t) \in \mathcal{O}$, it is the case that

$$\Lambda_X^{t_n^X}(\triangle_X) = \Lambda_Y^{t_n^Y}(\triangle_Y)\ \forall\ n,\ 1 \leq n \leq \|I_X\|.$$

Simply put, a prospect X is τ-shifted from Y if X is equivalent to taking everything about prospect Y (outcome mapping function, timing and distributions over signals) and shifting it forward in time by τ. Note that, in particular, this means that $\triangle_X^0(\alpha,\ t+\tau) = \triangle_Y^0(\alpha,\ t),\ \forall\ (\alpha, t) \in \mathcal{O}$, whenever X is τ-shifted from Y.

Proposition 3. *An agent's preferences satisfy* $A \sim B \Leftrightarrow C \sim D,\ \forall\ A, B, C, D\ \in \mathcal{P}$ *s.t.* C *is* τ*-shifted from* A*, with* $\mathcal{R}(C) = \mathcal{R}(A)^{+\tau} \cup \{0\}$*, and* D *is* τ*-shifted from* B*, with* $\mathcal{R}(D) = \mathcal{R}(B)^{+\tau} \cup \{0\}$*.*

Proposition 3 takes two competing prospects and shifts them into the future; the result says that if the point they are shifted *to* becomes a temporal reference point in addition to any that already existed, then an agent's indifference between the two prospects is maintained. The key here is that $\mathcal{R}(C)$ and $\mathcal{R}(D)$ both contain the element τ as their first non-zero element. As the agent still discounts back to that same point for both, before discounting to the present, the indifference between the prospects A and B is maintained. In other words, the added discounting of $D(\tau)$ is the same for both shifted prospects. The particular discount function, $D(t)$, used is irrelevant to this result.

The significance of this is as a resolution to the apparent conflict of differing experimental evidence on diminishing impatience. As discussed in the introduction, Glimcher et al. (2007) shows results of an experiment in which the indifference curves of agents were maintained when the prospects were shifted into the future by 60 days. The model here can account for this through the suggestion that, because *every* one of a large number of choice pairs in the delayed set of prospects involved a comparison to \$20 in 60 days, agents were induced to regard 60 days as a temporal reference point, a subjectively important date in time, in their evaluation and comparison of prospects.

Thus, the combination of Proposition 3, with a discount function implying diminishing impatience (in accordance with Lemma 1), can account for both the results of Glimcher, and the pervasive evidence showing that diminishing impatience weakens at the very least, as in Frederick et al. (2002), and arguably disappears, as rewards become more distant. This is done through the simple argument that the former experiment provided ample reason for agents to focus on a particular point of time in the future, whereas other experiments on diminishing impatience have not done so. Further contrast of this evidence is examined in Section 6.

5 Interpretation of Discounting in the Model

In this section, I will discuss possible interpretations of discounting in this model, and what it might mean, in an intuitive sense, to discount between temporal reference points.

One possible interpretation of the discount function, $D(t)$, is to split it into two components: geometric time discounting, along with a survival function representing implicit risk, $s(t)$: $D(t) \equiv e^{-\rho t} s(t)$.

In this case, the survival function, $s(t)$, can be interpreted as capturing the subjective view of the probability with which the explicitly stated outcomes of the prospect will endure over time t. $1 - s(t)$, then, is the subjective probability assigned to not explicitly stated (implicit) outcomes by the agent. This implicit risk may be a result of pessimism, distrust, or simply non-explicit death probability. It is straightforward to show that the same requirements on $D(t)$ which imply diminishing impatience, PGI, PLDI, etc., apply to $s(t)$ directly under this interpretation. Thus, for example, if $s(t)$ is continuously differentiable, then a time focused agent exhibits PLDI if and only if $s(t)$ has a decreasing hazard rate, $h(t) = \frac{-s'(t)}{s(t)}$.

A decreasing hazard rate lends itself well to an interpretation where focusing on particular points in time causes a rise in pessimism at those points which fades until brought back into focus by another such point. An agent, whether consciously or not, may view a point in time on which he is focused as an opportunity for something to go wrong with regard to the explicitly stated outcome of the prospect. While this may seem superficially similar to loss aversion, and the preference for grouped information that can arise out of it, such as in Köszegi and Rabin (2009), they are functionally quite distinct. Loss aversion implies greater weight being placed on negative adjustments in expectations than is placed on positive adjustments; no mechanism for the distinction of positive and negative adjustments exists in this model. This interpretation *can* however be extended to encompass the risk of some single external outcome taking place, provided that this outcome is weakly worse than any of the explicitly stated ones.[32] For example, prospect death risk as examined by Halevy (2008), or the implicit risk that an experimenter offering you a deferred payment will renege on his agreement at some point.

A second interpretation is to view the discount function as time discounting, with the addition that the agent focuses on temporal reference points "as if" he was his future self at that point. So, for example, if he regards one year from now as being a temporal reference point, he would first discount payouts occurring after that point in time "as if" he were himself in a year, and then discount back to the present from that point. Thus, the significance of the temporal reference points would be as points to which the agent "projects" himself in evaluating prospects. In this interpretation, $D(t)$ as hyperbolic time discounting, such as $D(t) = (1 + \alpha t)^{-(\gamma/\alpha)}$, would satisfy the conditions for diminishing and strongly diminishing impatience, and thus for PGI and PLDI.

[32]This could be fit in the model simply by normalizing the utility payout of such an outcome to zero.

6 Conclusion and Future Research

I conclude having shown that temporal reference points hold promise in giving novel explanation for a number of experimental results. Even with this promise, the question of how agents form temporal reference points remains to be explored in detail.

I dealt with, first, the case in which temporal reference points are formed by informational updates. This assumption immediately leads to the question of information thresholds; that is, how much information is needed to form a temporal reference point. For the information focused consumer, it was assumed here that *any* information creates a temporal reference point. However, in daily evaluation of prospects, informational updates are often, in a sense, ubiquitous. When considering the returns on prospects derived from stock investment for example, information is received on a daily basis by the simple fact that the agent knows that the stock market hasn't crashed today. It is not plausible, however, to suppose that such low-granularity information would cause the individual to focus on every instance at which he becomes freshly aware that, yes, the stock market is still intact. Thus, what the threshold for informational granularity is in order to trigger a temporal reference point is an important question for the viability of this model.

I also dealt with the case where uniform delay of two prospects caused the delay factor to become a temporal reference point. This particular deviation from information focused agents was chosen due to the results of Glimcher et al. (2007). Glimcher, as earlier discussed, proposes that the variable of interest explaining their results is the time of the soonest possible reward, in contrast to the explanation proposed here. The data from their experiment alone is not enough to disentangle the two explanations, but such disentanglement should be possible through choice experiments in which focus is made on a latter point in time (e.g. 60 days), but choice sets include options from that time, later times *and* earlier times.

Further, while the formation of a temporal reference point by the subjects in Glimcher is *plausible* due to the focus and repetition of a particular point in time, this does not go very far towards developing a actual theory of how temporal reference points are formed. Issues of attention and focus, repetition and ubiquity in temporal framing may all be factors in the formation of such. This is not an issue in applications where information focused consumers is a natural and obvious assumption to make, but if a theory of temporal reference points is to be extended beyond these cases, more investigation into the method of the formation of such reference points will be needed.

Part III

Related Literature

Time Inconsistency

> "Good resolutions are useless attempts to interfere with scientific laws. Their origin is pure vanity. Their result is absolutely nil." (Oscar Wilde, The Picture of Dorian Gray)

Samuelson (1937) gave us the discounted utility function that became ubiquitous in economic literature. In spite of Samuelson's concerns about the model ("It is completely arbitrary to assume that the individual behaves so as to maximize an integral of the form envisaged in [the DU model]"), it became the standard due to its elegance and tractability. As mounting evidence since that time has shown, Samuelson was right to disavow the descriptive accuracy of his model. Numerous inconsistencies in behavior have arisen inconsistent with the DU model, among them time inconsistency, the simple observation that what an agent wants for himself tomorrow is not the same as what he wants when tomorrow arrives. The empirical evidence for and theoretical work on time inconsistency and the subtopic of diminishing impatience is broad and spans disciplines, seeing particular interest in psychology and economics. I focus here on a few of the more influential and relevant studies. Frederick et al. (2002) provides a far more extensive and thorough review of the economic literature on time preference, with a special focus on time inconsistency, up to the date of that work. The literature in this section is relevant to both the multiself bargaining model and the temporal reference points model.

The classic work Strotz (1956) is the first to formally analyze time inconsistency. Strotz sets out a model in which an agent must choose a plan for lifetime consumption with the added complication of a non-geometric discount function. He identifies the inherent fluctuations in behavior and preference reversals that this would lead to and identifies two alternate strategies that a non-myopic agent can adopt in the face of such challenges. The first strategy is pre-commitment: noting that the ability to lock in or lock out certain future actions would always be desirable to a non-myopic agent, Strotz explains the observed use of costly commitment devices. The second strategy is consistency: in the face of a lack of commitment ability an agent would naturally wish to choose the best current plan given the foreknowledge of the differing preferences of the agent's self in the future and the ability of that future self to change the plan.

Strotz maintains that exponential discounting is the proper way of discounting, and concludes with the speculation that exponential discounting is a learned behavior which, with proper education and upbringing, can override the inherent non-exponential form. In the best case the agent becomes a true geometric discounter, having supplanted his natural inborn form of time inconsistency. In the less effective cases, we may see individuals being weak, and failing to uphold the correct exponential discounting that they have been taught. This is the cause of the "splurges, binges and extravagances" of human behavior.

Ainslie (1975) gives a wide overview of literature on impulsiveness in economics, sociology, psychology, psychiatry, and therapy. He identifies three proposed explanations for impulsive behavior. First, that the agents do not truly understand the consequences of their actions. Second, that agents are compelled by a "lower principle" to act against their own best interest. Third, that agents innately distort their valuation of consequences with imminent consequences having a greater weight. Through the aid of extensive animal research, Ainslie identifies the third as the most promising in explanatory power. He establishes that a hyperbolic discounting function and its associated smoothly diminishing impatience is the best and most elegant fit for the behaviors prescribed by the animal research.[33] He further notes that the same discounting function creates an elegant explanation for many observed human behaviors of impulsivity. Finally, Ainslie makes the important observation that Strotz was partially incorrect to identify pre-commitment and consistency as alternatives, noting simply that an agent would ideally employ both: committing against as many undesirable behaviors as possible, while retaining consistency with those that remained. In Ainslie (1992) he follows this earlier work with further evidence showing that both animal and human behavior fits hyperbolic discount functions.

Note here that while the multiself bargaining model I propose generates behavior in line with hyperbolic discounting, it *could* be interpreted to fit into Ainslie's second category. That is, with the impatient self representing a low impulse attempting to pull the individual away from their "true" preferences. While tempting, however, I do not see a good argument to be made for why the patient self is any more "true" than the impatient one. This leaves the model, by default, to fall into Ainslie's third category, if any of them.

Thaler (1981) provides an influential study testing three hypotheses concerning patterns of decision making. First, he finds that implicit discount rates decrease with length of time, indicating that the discount rate is higher over proximate time than over distant time. The explanation he promotes as most promising is that the agent inherently sees a greater difference between today and tomorrow than between a year from now and a year and a day. This is the explanation that has been most prominently adopted in the economic literature to explain preference reversal. Second, he finds that implicit discount rates decrease with the size of the rewards under consideration. He identifies as a plausible explanation a cost of self control that does not increase with the size of the reward. Third, he identifies that discount rates are smaller for losses than for gains. The explanation

[33] Specifically, under generalized hyperbolic discounting, a reward in t time is discounted by $(1+\alpha t)^{-(\gamma/\alpha)}$.

put forth for this is one of loss aversion: an agent does not equate the opportunity costs of forgone gain with out of pocket costs, over-weighing the latter compared to the former.

Loewenstein and Prelec (1992) argue that insufficient attention has been applied to violations of the canonical discounted utility (DU) of Samuelson (1937). They enumerate four identified anomalies in behavior not consistent with the DU model, and develop a model that accounts for them. Three of these anomalies match with the three results of Thaler above. The fourth, documented by Loewenstein (1988), finds that the amount required to compensate for delaying a reward by a given interval was two to four times greater than the amount subjects were willing to sacrifice to speed up consumption by the same interval. This is evidence of a framing effect because both choices were just different representations of the same underlying options. They encompass these anomalies in their model by including these framing effects, a value function which is steeper for losses than for gains and, most relevant to the present work, a hyperbolic time discounting function.

Keren and Roelofsma (1995) examine the connection between the certainty effect (the observation that certain outcomes are over-weighed relative to near-certain outcomes) and diminishing impatience. In their experiment, they develop six subject groups. In one group, subjects were offered a choice between $100 now and $110 in four weeks; 82% chose the $100. In another group, subjects were offered a choice between $100 in 26 weeks and $110 in 30 weeks; 37% chose the $100. This demonstrates the classic diminishing impatience and is in line with previous work. However, consider another two groups in their study. In one, subjects were offered a choice between a 50% chance of $100 now, and a 50% chance of $110 in four weeks; *39%* chose the first option. In the other, subjects were offered a choice between a 50% chance of $100 in 26 weeks, and a 50% chance of $110 in 30 weeks; 33% chose the first option. Similar results were shown for a 90% chance set. In other words, when payments are made less certain the present bias and diminishing impatience quickly fades. Their results are suggestive that at least part of the reason that subjects overweight the present is that the present is absolutely certain; that is, diminishing impatience is at least in part caused by the certainty effect. When the present is made very uncertain, present bias fades.

Halevy (2008), building on the results of Keren and Roelofsma, develops a discrete time model with a rank-dependent expected utility maximizer with constant stopping (death) probability. He develops his model to show the equivalence of diminishing impatience and the common ratio effect in such a setting. A correction by Saito (2011) argues that the equivalence is between diminishing impatience and the certainty effect, and between *strongly* diminishing impatience and the common ratio effect.

Chakraborty and Halevy (2015) further corrects both Halevy (2008) and Saito (2011). The authors show a problem in both works stemming from a too-weak definition of diminishing impatience. In the first two papers diminishing impatience and strongly diminishing impatience are defined based on a delay of only one period in two competing rewards. In the 2015 paper, Chakraborty and Halevy show that these time preference notions must be defined for all possible (discrete) delays in order for the desired equivalence to hold. With their strengthened notions of delay independent diminishing impatience (DIDI) and delay independent strongly diminishing impatience (DISDI), they show the equivalence in their setting between DIDI and the certainty effect (CE) and between DISDI and the common ratio effect (CRE). Notably, this is an equivalence between a time preference based in a discrete time setting, and risk preferences over probabilities in the unit interval. Both are united by the property of increasing elasticity of the probability distortion function inherent in rank-dependent utility. This difference in domains for time and risk preferences, one discrete and one continuous, is what led the original weaker definition to fail to create equivalence. With the original definition, the risk preference (CE, CRE) implied the corresponding time preference, but the converse did not hold, as Chakraborty and Halevy illustrate with an example. Intuitively, the strengthening of the time preference definitions the authors employ changes their notion of diminishing impatience to one that has a natural connection to a continuous time setting. To see this intuition, note that if the definition holds true for *any* delay, then it can approximate a notion of diminishing impatience in a continuous time setting (such as the definition employed by the temporal reference points model) by employing arbitrarily small period lengths. There is no such connection possible to continuous time with the original definition based on a delay of a single period length, which in a sense causes the original definition to be unable to bridge the gap between the discrete time preferences and continuous risk preferences.

Turning back to non-rank-dependent settings, Laibson (1997) popularized the use of the $\beta - \delta$ model of discounting in discrete time models. In such a discount function, consumption τ periods in the future is discounted using discount factor $\beta\delta^\tau$ when $\tau > 0$ (and by factor 1 for $\tau = 0$). When $0 < \beta < 1$, the author argues, this discount function retains the important qualitative properties of hyperbolic discounting while remaining relatively tractable. He terms this discounting "quasi-hyperbolic". Laibson's work led to widespread use of the $\beta - \delta$ form of discounting in applications. In Laibson (1998), he applies this model to known stylized facts about consumption choices, showing that it predicts many well-documented regularities. He further expounds on the welfare implication of the hyperbolic model, specifically how under-saving can lead to welfare losses in the Pareto sense: to all present and future selves. In Laibson et al. (1998), the authors calibrate a model to $\beta - \delta$ discounting based on data on saving for retirement. They show that patterns of asset accumulation are consistent with quasi-hyperbolic discounting, using observed regularities to distinguish between quasi-hyperbolic and exponential savers.

Frederick et al. (2002), as mentioned, provides a very thorough survey of the literature on non-exponential time discounting. One result from their work bears particular mention here. By analyzing the numerous studies measuring discount rates, they find that after a horizon of around a year there is no evidence at all of a further decline in discount rates. This is strong evidence that discounting is asymptotically exponential; that is, discount functions become increasingly close to exponential as time spans increase. This evidence fits in with the multiself bargaining model, provided that the long-term and short-term selves have reasonably distinct discount rates. In the language of the temporal reference points model, this would be consonant with a discount function exhibiting preference for grouped information, and also preference for less dispersed information within a limited time horizon. Such an agent would only care about the proximity of information to other information when the updates were relatively close.

Harris and Laibson (2013) provide an alternate approach to avoiding the discontinuity of predictions that come from the standard "beta-delta" model by introducing an element of uncertainty. In their model, there remains a discontinuous distinction between *now* and *later*, as in standard $\beta - \delta$, but the agent, in evaluating the discounted value of rewards, is uncertain about when *now* will end, and *later* will begin. The interpretation of this is challenging, as it implies the agent is internally uncertain about how they themselves are discounting future rewards. However, it is shown that this model of uncertainty generates preferences equivalent to a deterministic discounting function which is qualitatively similar to true hyperbolic discounting. Thus, their work presents an alternate form of a continuous time discounting function (as opposed to the hyperbolic discounting function).

Of special significance to the temporal reference points model is the work of Glimcher et al. (2007), whose well-documented results stand in direct opposition to several of the others discussed here. A study performed on individuals in MRIs given choices over pairs of prospects leads the authors to disagree with the fact that diminishing impatience swiftly disappears as contrasted prospects become more distant, and also disagree with the fact that preference reversals occur when options are pushed into the future. In the study one set of options, the "immediate set", gave choices between $20 immediately and delayed amounts greater than $20. An indifference curve was fit to individuals based on the choices made, and the implicit discounting indicated by this curve was found to be well fit by a hyperbolic function. This first results confirms previous work. However, another set of options in their study, the "delayed set", gave choices between $20 in two months, and an amount greater than $20 at a delay of more than two months. It is important to note that every choice in the delayed set had "$20 in two months" as one of the two options. Two important results of theirs stand out.

First, and most importantly, they find that the hyperbolic discount function stochastically fitting the indifference curve generated by choices over the immediate set is the same as the hyperbolic discount function for the indifference curve of the delayed set. In the words of the authors, "subjects were just as hyperbolic when making choices at delays of two months as they were when making choices of no delay". Second, they find that neural activity for the delayed set was scaled down by approximately as much as the 60 day option was for the immediate set. For instance, the 120 days option in the delayed option set was scaled by that same factor twice. Neither of these results are in line with previous results concerning time discounting. The authors propose that the additional variable of interest is the time of the soonest possible reward, supposing that the time of this reward becomes a default "starting point" for discounting. In other words, they argue that individuals are not "present biased" as much as they are "as soon as possible" biased. This explanation is not fully sufficient to rectify evidence, such as that of Keren and Roelofsma (1995) mentioned above, that agents *do* exhibit behavior consonant with diminishing impatience. The evidence presented there, and in many other surveys, cannot be explained by a bias for the soonest possible reward. Indeed, no such bias appears at all when the rewards become distant enough. The results of Glimcher et al. as well as the literature preceding it *can* be rectified in the temporal reference points model, as I discussed in that work. The fact that *every* comparison was made to 60 days is what makes it plausible that an agent would focus in on 60 days as an "important" point in time, whereas such a focus would be absent from other studies done on diminishing impatience.

Mullainathan and Banerjee (2010) take a different approach to time inconsistency by proposing a class of "temptation" goods. These goods are assumed to generate utility for the current self, but not for earlier selves that anticipate their consumption. The source of inconsistency is immediately apparent: an agent never wishes his future selves to consume temptation goods, yet always wishes to consume them immediately himself. By assuming the fraction of marginal earnings spent on temptation good is decreasing with overall consumption, the authors predict behavior consistent with regularities concerning the behavior of the poor relative to the rich, such as borrowing repeatedly at extremely high interest rates.

Going one step further in their notion of time inconsistency, Jamison and Wegener (2010) propose that people regard their future selves as truly separate persons. They draw upon neuroscientific and functional imaging studies that seem to indicate that mental systems associated with mentalizing other agents are the same as those associated with imagining oneself in the future. Thus, they argue that modeling intertemporal choice as a strategic game between present and futures selves is more than just a convenient modeling device, but has direct descriptive value.

Other evidence of diminishing impatience includes Angeletos et al. (2001), who simulate both hyperbolic and exponential households, finding that the former fit consumption data much better. Kirby and Herrenstein (1995) provide another study documenting preference reversal. Myerson and Green (1995) show further evidence of hyperbolic discounting in particular via a experimental study. Fang and Silverman (2009) estimate a structural model of labor supply and welfare program participation. They use their estimation to reject exponential discounting in favor of hyperbolic, and then use counterfactuals to quantify a measure of the utility loss stemming from a lack of commitment power.

As mentioned first by Strotz (1956), and followed up on by many papers mentioned, an immediate corollary of diminishing impatience is a desire for commitment power, unless the agent is completely naive about future preference reversals. The evidence for this preference for commitment is also substantial. Ariely and Wertenbroch (2002) show that students voluntarily choose binding intermediate deadlines for paper submissions in a college class. However, they do not set them optimally, and external deadlines are still superior. Thaler and Benartzi (2004) report on a proposed, and tested, savings plan which allows workers to voluntarily commit to automatically saving a fraction of future raise increases for retirement. Their evidence shows that the majority of subjects subscribe and stay in the program. Kaur et al. (2009) report on an experiment on workers in an Indian data entry firm. Workers had a wage per unit of work of w and could voluntarily choose to have that wage drop to $w/2$ unless they achieved a certain minimum target which they could also choose. Despite the fact that they stood to lose quite a lot, a significant number of subjects committed to non-trivial targets.

Temptation & Multiple Selves

The term "temptation" has been used to refer to two types of behavioral phenomenon observed in the literature. The first is related to the tendency of agents to be tempted to make decisions biased toward payoffs in the present or near future; this use is tied directly into time inconsistency. The second is the observation of violations of independence of irrelevant alternatives when multiple immediate choices are available, where agents may be "tempted" by an appealing option which influences their choice without actually choosing the appealing option. More precisely, violations of Sen's (1971) condition α, which states that if $x \in B \subset A$, and $x \in C(A)$, then $x \in C(B)$, where $C(A)$ is the set of elements chosen from set A.

Temptation in either form will lead to agents having a preference for commitment: the ability to constrain the effect of these tempting options either by removing options in advance, or by exercising costly self-control. Due to a large number of models addressing both forms of temptation simultaneously, and simple conflation of the word "temptation" to mean either of these effects, this section will necessarily have some spill-over to the previous section. I particularly focus on literature examining these concepts through the lens of models containing multiple sets of conflicting preferences.

First, there is copious research, largely in psychology, showing that self-regulation seems to be a limited resource. Baumeister and Heatherton (1996) examine patterns of "self-regulatory failure" in the literature. They conclude that the evidence supports a limited resource model of self-regulation, and that people often "voluntarily" lose control. Muraven et al. (1998) follows this by examining the limited resource model of self-control in more detail. Through a series of studies they are able to show that a limited resource model fits the data better than other models of self-control. Baumeister et al. (2007) provides a further short survey of evidence for the limited resource model to that date.

Vohs and Faber (2007) take these insights to economics by examining the hypotheses of the limited resource model as related to impulse buying. They find that participants who had participated in a willpower depleting activity felt stronger urges to buy and spent more money when confronted with unanticipated buying opportunities. Similarly, Ozdenoren et al. (2012) apply a limited willpower model to the domain of endowment consumption over time. Their model generates a number of qualitative predictions, showing that time preference may be domain specific, that previous actions affect preference for future ones (due to potential expenditures of willpower in the recent past), and that intertemporal smoothing will not, in general, appear.

A limited self-control resource implies violations of independence of irrelevant alternatives; the act of resisting one alternative will affect an agent's ability to resist another. While willpower as a limited resource is not a part of the multiself bargaining model, the same behavior is generated by both. In the multiself bargaining case, the introduction of more tempting options, that is more options desirable to the short-term self, increases the bargaining power of that self, and thus forces the long-term self to 'give ground'. This can thus be interpreted as a limited capacity to resist multiple tempting options.

Gul and Pesendorfer introduced the axiomatic treatment of temptation and self control. In Gul and Pesendorfer (2001), they develop a very important axiomatic representation result for "temptation preferences with self-control" in a two period setting. In the first period, the agent chooses a set of lotteries. In the second period, the agent chooses one lottery from the set chosen in the first period. An agent has both an *a priori* ranking over singleton sets of lotteries, a "commitment ranking", represented by $u(\cdot)$, and an instantaneous urge or temptation toward singleton sets of lotteries, a "temptation ranking", represented by $v(\cdot)$. An agent must compromise between what he would have chosen without temptation, represented by $arg\ \underset{x}{max}\ u(x)$, and the psychological cost of *not* choosing the most tempting option in each set, which would be $arg\ \underset{x}{max}\ v(x)$. The cost of this self-control is represented by the temptation utility difference between the choice made and the most tempting choice, so that the further from the temptation, the greater the cost. Thus, an individual chooses x to maximize $u(x) - (\underset{y}{max}\ v(y) - v(x))$, so that preferences over sets of options in the first period can be represented by:

$$U(A) = \underset{x \in A}{max}\ u(x) - (\underset{y \in A}{max}\ v(y) - v(x)).$$

Gul and Pesendorfer then define preferences in the second period over $\{(A, x) : x \in A\}$ where

A is the set of lotteries chosen in period 1, and $x \in A$ is the lottery chosen in period 2. In other words, the preference is over both the menu and the choice from the menu. This extended preference, denoted by $\succ^*$, allows them to formalize the notion of temptation, saying *y tempts x* if $(\{x\},\ x) \succ^* (\{x, y\},\ x)$. In other words, the agent would rather choose x in the absence of the temptation of y, then choose x in the presence of the temptation of y. This preference allows them to develop a notion of dynamically consistent preferences while accounting for self-control and temptation effects. More specifically, since the preferences are defined over the menu and the choice, and not the choice itself, the representation allows the accounting of apparent temptation effects while maintaining a notion of independence.

In Gul and Pesendorfer (2004), they extend their two period analysis to an infinite horizon to provide an alternative to non-exponential time discounting. This gives rise to a representation with a recursive definition of preferences, in which in each period the consumer need pay a self-control cost based on the difficulty of *not* consuming their entire wealth. Or, more precisely, based on the difference between the temptation utility granted by consuming everything, and the temptation utility of consuming the choice made. They interpret this as an individual whose temptation utility interferes with his long-run self-interest. Notably, they find that removing *non-binding* constraints changes equilibrium allocations, and that steady state consumption is independent of initial endowments and increasing in self-control (or decreasing in self-control costs).

As for economic literature on multiple selves, Thaler and Shefrin (1981) deal with the issue of self-control and intertemporal choice by modeling the individual as an organization in discrete time. This organization consists of a farsighted *planner* and a series of myopic *doers* who have no care beyond the current period. Their work stresses the similarities of this conflict within an individual to the principle-agent conflict of actual firms and is, to the best of my awareness, the first formal economic model to view an individual as having two sets of preferences at a single point in time that are in conflict. They argue that such a model is unavoidable to explain self-control, because "self-control" implies conflict within the self.

To expand a bit on the details of their model, a series of single-period lived *doers* derive utility from consumption, and have direct control over the consumption in that period. The single *planner* derives utility from the consumption of all doers, rather than getting consumption directly. Each doer would consume the maximum amount possible if allowed, but the planner is able to restrict their choices through "discretion" and "rules". Discretion takes the form of a parameter choice which is costly in utility terms but limits the maximum consumption that can be chosen. This is the essence of costly self control. Rules refers to the use of either pre-commitment, or the imposition of a "rule of thumb" that the doer has to follow, such as a self-imposed ban on borrowing (perhaps with exceptions to be made for houses and automobiles). They apply their model first to data from pension plans, which showed that pension savings and non-pension savings were not perfect substitutes. More specifically, the introduction of forced savings was *not* offset by an equal decrease in voluntary savings. This is explained by the involuntary savings being essentially willpower-free, allowing the planner to focus self-control on voluntary savings.

It is shown in Benabou and Pycia (2002) that Gul and Pesendorfer's (2001) representation can also be interpreted through the dual-self view of Thaler and Shefrin, in which there is an endogenous probability of "losing control" to your more myopic urge. Their re-interpretation bears an interesting conceptual connection to the multiself bargaining model. Specifically, Benabou and Pycia postulate two selves who lobby the brain for control, each expending a resource cost, and receiving probability of control proportional to the expenditure of the effort. This bears a close resemblance to the conflict that I use to determine the disagreement point. There are important distinctions, most saliently of which is that the multiself bargaining here allows bargaining between the selves to a mutually preferred point, rather than stopping at strict randomization between their most favored points. This allows the accommodation of randomization in the case of discrete decisions, as well as deterministic decisions for actions such as savings-consumption choice. In contrast, Benabou and Pycia's interpretation implies randomization in all cases where the selves prefer different outcomes. Further, I distinguish selves by a single parameter rather than distinct evaluations. This more restrictive view is limiting, but allows the generation of novel insights.

Chatterjee and Krishna (2006, 2009), develop a model of conflicting preferences closely related to the representation result of Gul and Pesendorfer (2001). In the first period an individual has to choose a set of feasible choices from which he will select one in the second period. With a probability dependent on the set chosen, the individual in essence "loses control" to an "alter ego" in the second period, who makes a choice based on their own preferences, represented by $v(\cdot)$. model, in which an "alter-ego" has a probability of appearing and overriding the decision of the far-sighted decision maker. The preferences over sets of lotteries, then, is given by:

$$U(A) = (1 - p_A) \underset{x \in A}{max} u(x) + p_A \underset{y \in B_v(A)}{max} u(y),$$

where p_A is the probability of "losing control" when faced with set A, and $B_v(A)$ is the set of v maximizers in A. The authors show that this dual-self representation is a relaxation of the axioms of Gul and Pesendorfer. That is, Gul and Pesendorfer's representation implies that a decision maker behaves "as if" he has an alternate ego in the second period when he is making choices in the first period. The authors argue that this reinterpretation allows unambiguous welfare statements in the face of dynamic inconsistency.

Building on Thaler and Shefrin (1981), work of Fudenberg and Levine (2006, 2011, 2012) bears the closest resemblance in the literature to the multiself bargaining model. In Fudenberg and Levine

(2006), they develop a discrete time model consisting of a single long-lived patient self, and a series of one period lived short-term selves, in which the short-term self has full control over the action in each period, and the long-term self exerts costly effort to constrain the actions available to the short-term self, where the cost is increasing in the utility denied to the short-term self. Their model is more general than that of Thaler and Shefrin, applying to a wide range of situations, as opposed to simply consumption-savings. Further, they specify the self-control aspect more precisely, allowing the model to be more tractable and make more precise predictions.

Expanding on the details of their model, each period is played in two stages. First, the long-run self chooses a self-control action that influences the utility function of the myopic self. The choice of this self-control action reduces utility for both selves, and changes the preferences of the short-term self. Second, the short-term self chooses the optimal action (for himself). They use their model, first, to explain diminishing impatience. Further, by arguing that self-control costs are convex rather than linear, they explain violations of independence of irrelevant alternatives as well as experimental evidence showing decreased self-control in the face of increased cognitive load. Rabin's paradox of risk aversion in the large and small is explained by a case in which cash-on-hand is used as a commitment device, leading the agent to consume all small but unexpected winnings. When winning large amounts, self-control allows the agent to save some of it. Intertemporal smoothing makes agents less risk averse in this latter case, so that they are less risk averse to large gambles than to small ones. Further, they argue that the axioms of Gul & Pesendorfer (2001, 2004) need be relaxed in order to account for the experimental evidence on the effects of cognitive load.

In Fudenberg and Levine (2011), they extend their model, with the assumption a convex cost of self control, in order to account for a wider range of puzzling experimental phenomena. One important qualitative prediction derived in their work is that preference reversal is less likely when the probability of rewards is smaller. This is in line with the data of Keren and Roelsofsma (1995) expounded on earlier. Second, they are able to show how convex costs of self-control can explain the Allais paradox as well as the common ratio effect. The essential force behind both of these qualitative additions is that in cases where payoffs are less likely (such as when two options are mixed with a chance of getting nothing), less self-control is needed. The convex costs of self-control, then, allow the long-term self to exert more influence over situations where decisions are more likely to be payoff irrelevant. Similarly, they predict that the Allais paradox should disappear as payoffs become more distant (as the long-term self is able to exert more control). Further, by estimating a version of the model, they are able to demonstrate that a wide range of phenomena can be explained by a stable set of parameters.

In Fudenberg and Levine (2012), they further extend the model, relaxing the assumption of completely myopic short-term selves in order to remove discontinuities in such effects as time discounting and cognitive load. As the short-term selves now care somewhat about the future, the authors modify self-control costs to depend on the present value denied to the short-term self, rather than the period utility. This specification requires, as the authors point out, the assumption that the short-term selves, in spite of caring about the future, are still strategically naive, unable to conceive of future self-control. Finally, in order for temptation and limited self-control to have influence in more than just the present period, they make the crucial addition of a cognitive resource variable that tracks self-control over time.

The Multiself Bargaining model developed in this work departs from this strand of work by Fudenberg and Levine conceptually primarily by treating selves symmetrically. The key innovation is that many behaviors of interest can be generated solely by a difference in geometric time preference between two selves, without the introduction of self-control cost or cognitive resources as modeling elements.

Ainslie (1986) speculates on the nature of intra-personal conflict, and its value to psychology and economics, in particular touching on the idea of inner bargaining. His form of bargaining is not symmetric, and focuses on selves that act at different times. However the work is significant for being the first, to my awareness, to mention the concept of two inner selves deliberately coming to a mutually beneficial agreement rather than conflict.

Green and Hojman (2009) in a working paper develop an important set of results for welfare analysis related to multiple selves and to cases of IIA violations in general. Starting with the statement that "revealed preference theory cannot be used as a basis for welfare analysis because rationality cannot be reasonably assumed," they explain choices are arising from compromise between conflicting preference relations.[34] They develop a method to determine a set of "explanatory preferences" that can give rise to the observed data, and find that, in general, their method does not yield a unique set of explanatory preferences. Thus, they compute bounds on welfare based on the set of all possible sets of explanatory preferences. The authors do not assume multiple sets of preferences, but rather show how such sets are generated from a set of choice data. They also develop a cardinal welfare theory based on the weighting the different possible sets of explanatory preferences place on each preference within the explanatory set.

Finally, Ambrus and Rozen (2013) show in a working paper that a limitation on the number of selves is necessary for any multiple self model to have predictive value. They establish a measure of the number of independence of irrelevant alternatives violations that a choice rule generates. They then show that there is a linear relationship between the number of selves aggregated into a given choice rules and the number of such IIA violations. They conclude, then, having shown that without

[34] The authors by "rationality" mean that there exists a preference relation defined over singleton sets which accounts for the choice data; they do not allow for the Gul & Pesendorfer view of rationality, which broadens the definition of the preference to encompass IIA violations.

an *a priori* limitation on the number of selves, any number of IIA violations can be rationalized.

Neurological studies on the descriptive accuracy of multiple concurrent selves are mixed. McClure et al. (2004) shows MRI evidence consistent with the notion that decisions made with different time horizons engage very different areas of the brain, and argues that this is evidence for multiple structures within the brain. In contrast, the aforementioned Glimcher et al. (2007) uses separate MRI evidence to contradict this assertion, claiming that there is no evidence that multiple selves in conflict is any more than a convenient modeling assumption. I am at present unaware of a resolution to this conflict among neurological evidence.

Information Preferences & Temporal Framing

In this section, I give an overview of the literature on the preferences of agents over the timing of information, as well as studies on the effects of temporal framing. The literature in this section is relevant primarily to the temporal reference points model.

I begin with an overview of both disappointment aversion and loss aversion with narrow framing, neither of which forms part of the temporal reference points model, but which are both popular sources of alternate hypotheses as to why information timing matters to agents. Both involve agents being more sensitive to bad outcomes than to good ones. The central intuition of the literature employing these to explain information preference is that additional information exposure provides more opportunities for both good and bad news, but the negative effect of the chance of bad news outweighs the positive effect of the change of good news.

First, Gul (1991) sets forth the theory of disappointment aversion, in which agents evaluate a lottery based on an endogenous reference point. Outcomes worse than the endogenous reference point are weighted at $1 + \beta$, while those better than the endogenous reference point are weighted at 1. He proposes that this model is the most restrictive one which both includes expected utility theory as a special case and also accounts for the Allais paradox. The reference point is the unique value, x, such that if the expected value of the lottery is calculated with the additional weight put on values less than x, that expected value will be equal to x. Disappointment aversion fares well empirically for many aspects of choice behavior. For example, Camerer and Ho (1994) use a survey of violations of independence and betweenness to show that disappointment aversion fits the data far better than expected utility.

Loss aversion, introduced by Kahneman and Tversky (1979) in their seminal paper on prospect theory, is the tendency of agents to be more sensitive to losses than to gains. Narrow framing suggests that individuals examine each outcome or decision in isolation. In combination, these two have been used to explain a number of economic regularities, being jointly referred to as "myopic loss aversion".

Turning now to evidence evidence of the effects of information feedback on agents, Gneezy and Potters (1997) allowed subjects to sequentially bet $2 on a repeated known lottery. In one case, subjects learned the outcome and made a new decision each period. In another case, subjects made a decision for the next three periods, and learned only the aggregate outcome. Subjects were found to bet much less in the first case than in the second. The authors note that this increased risk aversion in the presence of greater feedback is in line with myopic loss aversion. Haigh and List (2005) did a similar study to Gneezy and Potters with undergraduates and professional traders. They not only replicated the result of Gneezy and Potters, but interestingly found that the informational feedback effect was greater for the traders than the undergraduates.

Bellemare et al. (2005) were able to separate the effect of information feedback from that of investment flexibility. Building from the Gneezy and Potters experiment, they added a third case where subjects could only change their decision every third period, but still received informational updates every period. They significantly found that even when investment flexibility was fixed, informational feedback alone explained most of the difference in risk aversion. Thus, they argue, "myopic loss aversion is driven by information feedback".

Köszegi and Rabin (2009) develop a dynamic consumption model in which utility is based on current consumption and changes in beliefs about consumption (both present and future). In their model, agents are assumed to be loss averse with respect to changes in these beliefs; that is, negative news is worse than positive news is good. The informational updates received in their model are the changes in expected future consumption. One result is that agents prefer to get information clumped together. This is due to the loss aversion: information more spread out gives more opportunities for fluctuations in the expectations of consumption, and thus more opportunities to experience loss. This preference is in line with the PGI developed in the temporal reference points model, though the source is different. (As their model is in discrete time, there is no equivalence to PLDI). Their model further shows a preference for receiving the same information earlier. The agent prefers to get bad news about a future event sooner so that the loss is felt to be more distant.

Palacios-Huerta (1999), building on Gul's disappointment aversion, is the first in the economic literature to raise the idea that a disappointment averse individual would prefer all uncertainty to be resolved at once, rather than sequentially. This is shown by working out an example. The fundamental insight is that a weakening of the independence axiom will induce a preference for the way in which uncertainty is resolved. The author supports the use of disappointment aversion over prospect theory (i.e. myopic loss aversion) based explanations on the basis of the fact that disappointment aversion is only one parameter richer than expected utility, and retains as much as possible of expected utility theory while also being consistent with the Allais paradox and other behavioral anomalies.

Dillenberger (2010) deals with the nature of what he calls Preferences for One-Shot Resolution of Uncertainty (PORU), a formalization of the idea raised by Palacios-Huerta. An agent exhibiting PORU would prefer that any two stage lottery be resolved in a single stage. Dillenberger further introduces Negative Certainty Independence (NCI), which says that an agent who prefers a non-degenerate lottery to a certain outcome will not reverse preferences when both options are mixed with a common third option. Essentially, NCI is a formalization of the certainty effect, as it says that the certainty of the certain outcome must (weakly) give it some additional appeal over the lottery.[35] Dillenberger then establishes that these properties, PORU and NCI, are equivalent. He further quantifies PORU by defining the "gradual resolution premium", which is the amount the agent would pay to replace a two stage lottery with its single stage equivalent. That is, to avoid the gradual resolution of uncertainty.

It is important to note that, while there are certainly similarities between PORU and a preference for grouped information (PGI) in the temporal reference points model, they are not the same property because they occupy different settings. Dillenberger (2010) uses time neutrality (the idea that an agent does not care about timing), as one of the conditions for the equivalence he develops in order to isolate the effects of gradual resolution of uncertainty apart from the timing. Thus, PGI cannot be seen as a direct application of PORU. The closest connection to PORU would be via an information focused agent with discount function described by a continuous time analog of $\beta - \delta$ discounting, so that

$$D(t) \equiv \begin{cases} 1, & if\ t = 0 \\ \beta e^{-\delta t}, & otherwise. \end{cases}$$

In the temporal reference points model, such an agent would only have preferences with regard to the number of stages in which information is revealed in this model, without caring at all about the relative timing of those informational updates.

Artstein-Avidan and Dillenberger (2010) show directly that disappointment aversion implies PORU, and then, as an application, extend disappointment aversion to lotteries with an arbitrary number of stages, noting that a disappointment averse agent prefers to replace each two-stage sub-lottery with its single-stage counterpart. This mirrors closely the property of PGI in the temporal reference points model.

Turning to the effects of direct temporal framing, Chandran and Menon (2004) examine the effects on temporal framing on the decisions of individuals. Through survey data, they examine

[35] The certainty effect is an empirically supported observation that agents over-weigh certain outcomes in comparison to outcomes with even the slightest bit of uncertainty.

the differences in framing health risks by per day risk versus per year risk. They find that risks presented in per day terms appear more proximal and concrete to subjects, leading to increased self-risk perceptions. They suggest, through an adoption of the Construal Level Theory of Liberman and Trope (1998, 2000) that this is due to the idea of "a day" triggering the notion of a proximal event, and "a year" a distant one. In the language of the temporal reference points model, an alternate hypothesis is that framing health risks in terms of daily events creates additional temporal reference points, which decreases the value of the health-based lottery, making the subjects appear more risk averse.

Gourville (1998) looks at the "pennies-per-day" phenomenon where a consumer cost is presented as a drawn out, minimal daily (or weekly, or monthly) cost in contrast to a single larger payment. He finds that the alternate types of cost framing trigger different associations, and lead to significantly different consumer behavior. In contrast to the work of Chandran and Menon, Gourville argues that the associations of the two types of payments are not temporal. Rather, the minimal daily cost is shown to be compared with other small ongoing costs by the agent (cup of coffee per day!), while the single larger payment is associated with other large expenses. Thus it is not clear here that the effect is due to the temporal framing or other types of framing. In other words, it may be due to framing based on the size of the payment, rather than the timing of it.

Bargaining

In the realm of bargaining, a small existing strand of literature explores methods of generating endogenous disagreement points in bargaining problems, which is relevant to the multiself bargaining model.

Vartiainen (2007) studies a bargaining problem without a disagreement outcome. He develops a solution that determines the outcome and the disagreement point simultaneously. Further, he shows that there is a unique solution of this kind which satisfies Pareto-optimality, independence of irrelevant alternatives, symmetry and scale invariance, and that this solution maximizes the Nash product with respect to the solution and the reference point. Essentially the way this is defined is to find two points in the set of outcomes, x and y , such that x is the Nash bargaining solution when y is the disagreement point, and that when the entire set is inverted arithmetically, the inverted y becomes the Nash bargaining solution with the inverted x as the disagreement point. This solution concept requires strict convexity of the utility set in order to generate a unique solution.

Bozbay et al. (2012) develop a similar extension for Kalai-Smorodinsky (1975) bargaining, where both a disagreement point and a compromise point are determined simultaneously. The compromise point is the classic Kalai-Smorodinsky outcome for the disagreement point. The disagreement point is determined by following the line joining the compromise point and the "anti-utopia point" (consisting of the worst possible utility for each player), and then finding the worst point on this line within the set of outcomes. Once again, this solution concept requires strict convexity of the utility set in order to generate a unique solution.

Unfortunately, neither of these methods of endogenous disagreement points can be applied to the multiself bargaining model. This is primarily because both require *strict* convexity of the bargaining set, which would exclude from consideration discrete action sets (which create utility vector sets which are not strictly convex). Secondarily, the methods of disagreement point generation, while mathematically elegant, do not seem to be any more conceptually fitting for the application than the one used.

Part IV

Appendices

Appendix A: Proofs for Multiself Bargaining Model

I begin the proofs by formalizing a few results, not included in the main body of the paper, that derive from the nature of the Nash bargaining solution.

Lemma A.1 formalizes the notion that, taking the line drawn between the disagreement point and the Nash bargaining outcome, if the disagreement point is moved to the right side of that line, the outcome will also move to the right, and visa versa. In Figure A.1, the shaded area represents where a moved disagreement point would move the outcome to the right. This result will be used in several of the latter proofs.

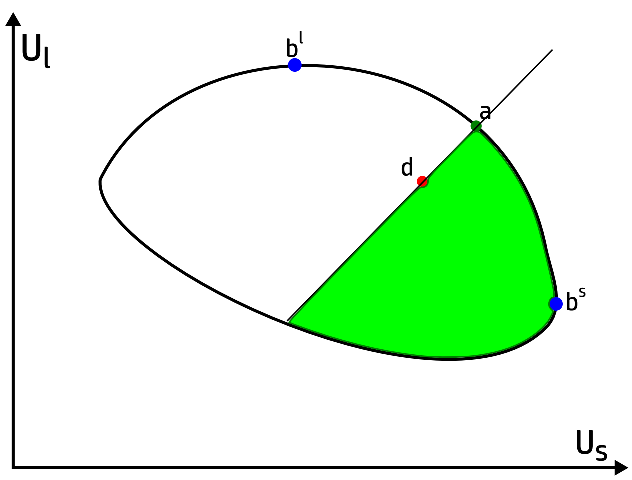

Figure A.1: Line dividing effect of moving disagreement point.

Lemma A.1. *Consider a bargaining set, $\mathcal{U}$ with disagreement point $(d_s^1,\ d_l^1)$, and denote the Nash bargaining outcome by $(U_s^1,\ U_l^1)$. Consider the same bargaining set with a different disagreement point $(d_s^2,\ d_l^2)$, and denote Nash bargaining outcome by $(U_s^2,\ U_l^2)$. Denote $\triangle d_s = d_s^2 - d_s^1,\ \triangle d_l = d_l^2 - d_l^1$, and assume that neither disagreement point is on the Pareto frontier of $\mathcal{U}$. Then,*

$$d_s^2 \geq U_s^1 \Longrightarrow U_s^2 > U_s^1;\ d_l^2 \geq U_l^1 \Longrightarrow U_l^2 > U_l^1.$$

If $d_s^2 < U_s^1$, $d_l^2 < U_l^1$, then

$$\triangle d_s(U_l^1 - d_l^1) - \triangle d_l(U_s^1 - d_s^1) < 0 \Rightarrow U_l^2 \geq U_l^1;$$

$$\triangle d_s(U_l^1 - d_l^1) - \triangle d_l(U_s^1 - d_s^1) > 0 \Rightarrow U_s^2 \geq U_s^1;$$

$$\triangle d_s(U_l^1 - d_l^1) - \triangle d_l(U_s^1 - d_s^1) = 0 \Rightarrow U_s^2 = U_s^1,$$

with the right inequalities strict if the Pareto frontier of $\mathcal{U}$ is smooth at (U_s^1, U_l^1).

Proof: The first line follows directly from fact that if the disagreement point is not on the Pareto frontier, then the Nash bargaining outcome will necessarily be a Pareto improvement on it. Meaning, $U_s^2 > d_s^2$.

Now, consider the second case, where the original outcome remains a Pareto improvement on the new disagreement point. Note that the Pareto frontier of $\mathcal{U}$ is a continuous, not necessarily differentiable, curve with endpoints at the bliss points of the two selves. Take any continuous bijective mapping f that maps $[0,1]$ onto this curve. Without loss of generality, assume that it maps 0 to the bliss point of the long-term self, and 1 to the bliss point of the short-term self. f can then be divided into two mappings, one for each of the two coordinates of the points on the Pareto frontier, $f_s : [0,1] \to [X^l, X^s]$ and $f_l : [0,1] \to [Y^l, Y^s]$. For $x \in [0,1]$ denote by $(U_s(x),\ U_l(x)) \equiv (f_s(x),\ f_l(x))$ the point on the Pareto frontier which x is mapped to. Note that $U_s(x)$ is strictly increasing in x, while $U_l(x)$ is strictly decreasing in x. Then, we can rewrite the outcome of the Nash bargaining procedure to be the solution to:

$$\underset{x \in [0,1]}{max} \left(U_s(x) - d_s^1\right)\left(U_l(x) - d_l^1\right).$$

Denote x^1 the argument of the solution to this maximization (Nash bargaining gives us a unique solution). Then, the Nash bargaining outcome can be written as $\left(U_l(x^1),\ U_s(x^1)\right)$. Where $U_s(x)$ and $U_l(x)$ are differentiable, which is wherever the Pareto frontier is smooth, the first derivative of the objective is:

$$\frac{\partial U_s(x)}{\partial x}(U_l(x) - d_l^1) + \frac{\partial U_l(x)}{\partial x}(U_s(x) - d_s^1),$$

so that the first order condition is:

$$\frac{U_l(x) - d_l^1}{U_s(x) - d_s^1} = -\frac{\partial U_l(x)/\partial x}{\partial U_s(x)/\partial x}.$$

If the Pareto frontier is smooth at the Nash bargaining outcome, we then have

$$\frac{U_l(x^1) - d_l^1}{U_s(x^1) - d_s^1} = -\frac{\partial U_l(x^1)/\partial x}{\partial U_s(x^1)/\partial x}.$$

Now consider the new disagreement point $(d_s^2,\ d_l^2)$. Denote

$$x^2 = arg \max_{x\in[0,1]} \left(U_s(x) - d_s^2\right)\left(U_l(x) - d_l^2\right).$$

Assume that the Pareto frontier is smooth at $\left(U_l(x^1),\ U_s(x^1)\right)$, then,

$$\triangle d_s(U_l^1 - d_l^1) - \triangle d_l(U_s^1 - d_s^1) < 0 \Leftrightarrow \triangle d_s U_l(x^1) + \triangle d_l d_s^1 < \triangle d_l U_s(x^1) + d_l^1 \triangle d_s$$

$$\Leftrightarrow \triangle d_s U_l(x^1) + (d_l^1 + \triangle d_l) d_s^1 < \triangle d_l U_s(x^1) + d_l^1 (d_s^1 + \triangle d_s)$$

$$\Leftrightarrow (d_s^2 - d_s^1) U_l(x^1) + d_l^2 d_s^1 < (d_l^2 - d_l^1) U_s(x^1) + d_l^1 d_s^2$$

$$\Leftrightarrow U_s(x^1)U_l(x^1) - d_s^1 U_l(x^1) - d_l^2 U_s(x^1) + d_l^2 d_s^1 < U_s(x^1)U_l(x^1) - d_s^2 U_l(x^1) - d_l^1 U_s(x^1) + d_l^1 d_s^2$$

$$\Leftrightarrow \left(U_l(x^1) - d_l^2\right)\left(U_s(x^1) - d_s^1\right) < \left(U_l(x^1) - d_l^1\right)\left(U_s(x^1) - d_s^2\right)$$

$$\Leftrightarrow \frac{U_l(x^1) - d_l^2}{U_s(x^1) - d_s^2} < \frac{U_l(x^1) - d_l^1}{U_s(x^1) - d_s^1} \Leftrightarrow \frac{U_l(x^1) - d_l^2}{U_s(x^1) - d_s^2} < -\frac{\partial U_l(x^1)/\partial x}{\partial U_s(x^1)/\partial x}$$

$$\Leftrightarrow \frac{\partial U_s(x^1)}{\partial x}(U_l(x^1) - d_l^2) + \frac{\partial U_l(x^1)}{\partial x}(U_s(x^1) - d_s^2) < 0.$$

This last line shows that the Nash product, with the new disagreement point, is now decreasing in x at x^1. This, implies $x^2 < x^1$, since the Nash product is single-peaked along the Pareto frontier. Thus,

$$\Rightarrow x^2 < x^1 \Rightarrow U_s(x^2) < U_s(x^1),\ U_l(x^2) > U_l(x^1).$$

The rest of the cases for a smooth Pareto frontier follow similarly.

Now, we need to consider the case where the Pareto frontier is not smooth at $U^1 = \left(U_l(x^1),\ U_s(x^1)\right)$, and proceed by contradiction. Suppose that $\triangle d_s(U_l^1 - d_l^1) < \triangle d_l(U_s^1 - d_s^1)$ but that $U_l^2 < U_l^1$ (the other case will follow identically). This immediately implies $U_s^2 > U_s^1$ (U^2 is to the lower right of

U^1). Denote the line connecting d^1 to U^1 as $\overleftrightarrow{d^1U^1}$, with slope m^1. Similarly $\overleftrightarrow{d^2U^2}$, with slope m^2 for the line between d^2 and U^2, and $\overleftrightarrow{d^1d^2}$, with slope m^d, for the line between the disagreement points.

We begin by showing that d^2 is above $\overleftrightarrow{d^1U^1}$. Consider three cases. Case 1, $\triangle d_s > 0$. Then,

$$\triangle d_s(U_l^1 - d_l^1) < \triangle d_l(U_s^1 - d_s^1) \Rightarrow m^1 = \frac{(U_l^1 - d_l^1)}{(U_s^1 - d_s^1)} < \frac{\triangle d_l}{\triangle d_s} = m^d.$$

Thus, $m^d > m^1$. As $d_s^2 > d_s^1$, d^2 lies above $\overleftrightarrow{d^1U^1}$. Case 2, $\triangle d_s < 0$,

$$\triangle d_s(U_l^1 - d_l^1) < \triangle d_l(U_s^1 - d_s^1) \Rightarrow \frac{(U_l^1 - d_l^1)}{(U_s^1 - d_s^1)} > \frac{\triangle d_l}{\triangle d_s}.$$

Two subcases. If $\triangle d_l > 0$, then $m^d < 0$, (d^2 is to the left and above), and so d^2 is definitely above $\overleftrightarrow{d^1U^1}$. If $\triangle d_l < 0$, then $m^1 > m^d > 0$, and as $d_s^2 < d_s^1$, d^2 lies above $\overleftrightarrow{d^1U^1}$. Case 3, $\triangle d_s < 0$. $\triangle d_s(U_l^1 - d_l^1) < \triangle d_l(U_s^1 - d_s^1) \Rightarrow \triangle d_l > 0$, so that d^2 is directly above d^1.

Thus, d^2 lies above $\overleftrightarrow{d^1U^1}$. Now, U^2 is below $\overleftrightarrow{d^1U^1}$, as it lies to the right along the Pareto frontier of $\mathcal{U}$ from U^1. So, as d^2 lies above $\overleftrightarrow{d^1U^1}$ and U^2 lies below it, $\overleftrightarrow{d^1U^1}$ and $\overleftrightarrow{d^2U^2}$ necessarily cross, so that $m^2 < m^1$.

Now, consider a some set of utility vectors, $\mathcal{V}$, that has a Pareto frontier consisting of the line connecting U^1 and U^2. Note that all points on the line between U^1 and U^2 are necessarily elements of $\mathcal{U}$, as $\mathcal{U}$ is convex. Then, consider the bargaining problem given by $< \mathcal{V}, d^1 >$, and denote its solution vector as V^1. $V_s^1 \leq U_s^1$. This is because the Nash product is single-peaked along the Pareto frontier. Thus, if $V_s^1 > U_s^1$, then *any* point to the right of U^1 along the Pareto frontier of $\mathcal{V}$ induces a greater Nash bargaining product than that of U^1. But, this would imply that U^1 was not the Nash product maximizing choice from $< \mathcal{U}, d^1 >$, since there are points to the right of U that are elements of $\mathcal{U}$. Therefore, $V_s^1 \leq U_s^1$. This means that that the slope between d^1 and V^1, denote m_v^1, is higher than m^1. $m_v^1 \geq m^1$.

Similarly, for the bargaining problem $< \mathcal{V}, d^2 >$, with solution V^2, $V_s^2 \geq U_s^2 \Rightarrow m_v^2 \leq m^2$.

$$m_v^2 \leq m^2 < m^1 \leq m_v^1 \Rightarrow m_v^2 < m_v^1$$

However, since the slope of the Pareto frontier of $\mathcal{V}$ is smooth, from earlier in the proof we know that the line connecting a disagreement point to the outcome vector should be the negative of the slope of the frontier. This implies that $m_v^1 = m_v^2$, which gives us the contradiction. □

Lemma A.2 uses the fact that if we have two different utility vector sets whose Pareto frontiers coincide in an open interval around the Nash bargaining outcome, then small variations in the disagreement point will have the same effect on the outcome chosen from both sets. This allow us to apply Lemma 1 for different sets when the Pareto frontier around the outcome is the same for both sets.

Lemma A.2. *Consider a utility vector set set* $\mathcal{U}^1$*, with disagreement point* $(d_s^1,\ d_l^1)$*, and denote the Nash bargaining outcome by* $(U_s^1,\ U_l^1)$*. Consider another utility vector set* $\mathcal{U}^2$ *with disagreement point* $(d_s^2,\ d_l^2)$*, and denote the Nash bargaining outcome by* $(U_s^2,\ U_l^2)$*. Denote* $\triangle d_s = d_s^2 - d_s^1$*,* $\triangle d_l = d_l^2 - d_l^1$*, and assume that neither disagreement point is on the Pareto frontier of their respective utility vectors sets. Further, assume that there is an open connected subset of the Pareto frontier of* $\mathcal{U}^1$ *containing* $(U_s^1,\ U_l^1)$*, and that this subset is also a subset of the Pareto frontier of* $\mathcal{U}^2$*. and that* $(U_s^1,\ U_l^1)$ *is on the Pareto frontier of* $\mathcal{U}^1$*. Then,*

$$d_s^2 \geq U_s^1 \Longrightarrow U_s^2 > U_s^1;\ d_l^2 \geq U_l^1 \Longrightarrow U_l^2 > U_l^1.$$

Otherwise, if $d_s^2 < U_s^1$*,* $d_l^2 < U_l^1$*, then,* $\exists \bar{\triangle} > 0$ *such that if* $\triangle d_l < \bar{\triangle}$ *and* $\triangle d_s < \bar{\triangle}$*,*

$$\triangle d_s(U_l^1 - d_l^1) - \triangle d_l(U_s^1 - d_s^1) < 0 \Rightarrow U_l^2 \geq U_l^1;$$

$$\triangle d_s(U_l^1 - d_l^1) - \triangle d_l(U_s^1 - d_s^1) > 0 \Rightarrow U_s^2 \geq U_s^1;$$

$$\triangle d_s(U_l^1 - d_l^1) - \triangle d_l(U_s^1 - d_s^1) = 0 \Rightarrow U_s^2 = U_s^1,$$

with the right inequalities strict if the Pareto frontier of $\mathcal{U}^1$ *is smooth at* (U_s^1, U_l^1)*.*

Proof: The first line follows directly from fact that if the disagreement point is not on the Pareto frontier, then the Nash bargaining outcome will be a Pareto improvement on it. Meaning, $U_s^2 > d_s^2$.

Now consider the second line. First, the Nash bargaining outcome chosen from $\mathcal{U}^2$ with disagreement point $(d_s^1,\ d_l^1)$ is $(U_s^1,\ U_l^1)$. This follows from the fact that the Pareto frontier of $\mathcal{U}^2$ is the same as that of $\mathcal{U}^1$, and from the fact that the Nash bargaining product is continuous and has a unique local maximum. Thus, if the Nash product from $\mathcal{U}^2$ could be increased by moving away from $(U_s^1,\ U_l^1)$ in one direction, then so could the Nash product from $\mathcal{U}^1$.

Second, consider the Nash bargaining outcome chosen from $\mathcal{U}^1$ with disagreement point $(d_s^2,\ d_l^2)$. As the Nash bargaining outcome varies continuously with the disagreement point, for a sufficiently small change in the disagreement point, this outcome will lie within the open connected subset of the Pareto frontier containing $(U_s^1,\ U_l^1)$. Then, by the same reasoning as for the original disagreement point, this new outcome will also be the result of the Nash bargaining solution applied to $\mathcal{U}^2$ with disagreement point $(d_s^2,\ d_l^2)$. Thus, for a sufficiently small variation in the disagreement point, the effect on the outcome chosen from $\mathcal{U}^1$ is the same as the effect on the outcome chosen from $\mathcal{U}^2$. The result then follows directly from Lemma A.1.□

Lemma A.3. *Suppose that the Pareto frontier of $\mathcal{U}_n$ consists only of mixtures of the two bliss points. Then, the lottery selected via bargaining has the same mixing weights between actions as the equilibrium of the conflict game.*

Proof of Lemma A.3: First, if using the disagreement point with the effort costs added back in, the result follows directly from the fact that the disagreement point is on the Pareto frontier.

Consider the other case (without the effort costs added in). The bliss points are given by $(X_n^s,\ Y_n^s)$ and $(X_n^l,\ Y_n^l)$, so that a mixing between them can be expressed as $\left(wX_n^s + (1-w)X_n^l,\ wY_n^s + (1-w)Y_n^l\right)$, where w is the weight placed on the bliss point of the short-term self. From Lemma 1, the disagreement point is given by:

$$\left(X_n^l + \frac{S_n^{\gamma+1}\left(S_n^\gamma + (1-\gamma)L_n^\gamma\right)}{(S_n^\gamma + L_n^\gamma)^2},\ Y_t^s + \frac{L_n^{\gamma+1}\left(L_n^\gamma + (1-\gamma)S_n^\gamma\right)}{(S_n^\gamma + L_n^\gamma)^2}\right).$$

To find the Nash bargaining solution, then, we choose w to maximize the product of gains over the disagreement point utilities. This product is given by:

$$\left(wX_n^s + (1-w)X_n^l - X_n^l - \frac{S_n^{\gamma+1}\left(S_n^\gamma + (1-\gamma)L_n^\gamma\right)}{(S_n^\gamma + L_n^\gamma)^2}\right)\left(wY_n^s + (1-w)Y_n^l - Y_t^s - \frac{L_n^{\gamma+1}\left(L_n^\gamma + (1-\gamma)S_n^\gamma\right)}{(S_n^\gamma + L_n^\gamma)^2}\right)$$

$$= \left(wS_n - \frac{S_n^{\gamma+1}\left(S_n^\gamma + (1-\gamma)L_n^\gamma\right)}{(S_n^\gamma + L_n^\gamma)^2}\right)\left((1-w)L_n - \frac{L_n^{\gamma+1}\left(L_n^\gamma + (1-\gamma)S_n^\gamma\right)}{(S_n^\gamma + L_n^\gamma)^2}\right).$$

The first order condition is

$$S_n\left((1-w)L_n - \frac{L_n^{\gamma+1}\left(L_n^\gamma + (1-\gamma)S_n^\gamma\right)}{(S_n^\gamma + L_n^\gamma)^2}\right) - L_n\left(wS_n - \frac{S_n^{\gamma+1}\left(S_n^\gamma + (1-\gamma)L_n^\gamma\right)}{(S_n^\gamma + L_n^\gamma)^2}\right) = 0$$

$$\Rightarrow S_nL_n(S_n^\gamma + L_n^\gamma)^2 - S_nL_n^{\gamma+1}\left(L_n^\gamma + (1-\gamma)S_n^\gamma\right) + L_nS_n^{\gamma+1}\left(S_n^\gamma + (1-\gamma)L_n^\gamma\right) = 2wS_nL_n(S_n^\gamma + L_n^\gamma)^2$$

$$\Rightarrow S_nL_n(S_n^\gamma + L_n^\gamma)^2 - S_nL_n^{2\gamma+1} + L_nS_n^{2\gamma+1} = 2wS_nL_n(S_n^\gamma + L_n^\gamma)^2$$

$$\Rightarrow S_n^{2\gamma} + 2S_n^\gamma L_n^\gamma + L_n^{2\gamma} - L_n^{2\gamma} + S_n^{2\gamma} = 2w(S_n^\gamma + L_n^\gamma)^2$$

$$\Rightarrow S_n^\gamma(S_n^\gamma + L_n^\gamma) = w(S_n^\gamma + L_n^\gamma)^2 \Rightarrow w = \frac{S_n^\gamma}{S_n^\gamma + L_n^\gamma}.$$

This is exactly the same weight placed on the short-term self gaining control, and thus on the bliss action of the short-term self, by the equilibrium of the conflict game.□

Proof of Lemma 1: First, $e_s = e_l = 0$ is not an equilibrium; a marginal increase in effort by either self leads to a guarantee of control, and thus a strictly positive increase in utility. Similarly,$e_i > 0$, $e_{-i} = 0$ is not an equilibrium; self i would strictly prefer to lower their effort, as doing so does not reduce their probability of control, but does reduce their cost. Now, consider an interior equilibrium, $e_s > 0$, $e_l > 0$. The short-term self chooses e_s to maximize expected utility, given by

$$U_s(e_s, e_l) = \frac{e_s^\gamma}{e_s^\gamma + e_l^\gamma} X_n^s + \frac{e_l^\gamma}{e_s^\gamma + e_l^\gamma} X_n^l - e_s.$$

The first and second derivatives of this objective are

$$\frac{\partial}{\partial e_s} U_s(e_s, e_l) = \gamma e_s^{\gamma-1} e_l^\gamma (e_s^\gamma + e_l^\gamma)^{-2} S_n - 1;$$

$$\frac{\partial^2}{\partial^2 e_s} U_s(e_s, e_l) = \gamma e_s^{\gamma-2} e_l^\gamma (e_s^\gamma + e_l^\gamma)^{-3} \left[(\gamma - 1)(e_s^\gamma + e_l^\gamma) - 2\gamma e_s^\gamma\right] S_n,$$

similarly for the long-term self. Note that the second derivative is negative if $0 \le \gamma \le 1$, satisfying the second order condition. We can write the first order condition of the selves as

$$\gamma e_s^{\gamma-1} e_l^\gamma S_n = (e_s^\gamma + e_l^\gamma)^2; \quad \gamma e_l^{\gamma-1} e_s^\gamma L_n = (e_s^\gamma + e_l^\gamma)^2.$$

Defining $R_n = \frac{L_n}{S_n}$, and equating the two left hand sides we have

$$\gamma e_s^{\gamma-1} e_l^\gamma S_n = \gamma e_l^{\gamma-1} e_s^\gamma L_n \Rightarrow e_l = e_s \frac{L_n}{S_n} = e_s R_n.$$

Substituting into the short-term self's first order condition,

$$\gamma e_s^{\gamma-1} e_s^\gamma R_n^\gamma S_n = (e_s^\gamma + e_s^\gamma R_n^\gamma)^2 \Rightarrow \gamma e_s^{-1} R_n^\gamma S_n = (1 + R_n^\gamma)^2$$

$$\Rightarrow e_s = \frac{\gamma R_n^\gamma S_n}{(1 + R_n^\gamma)^2} = \frac{\gamma L_n^\gamma S_n^{\gamma+1}}{(S_n^\gamma + L_n^\gamma)^2}.$$

We get e_l similarly, and the probability p of the short-term self winning the contest.

$$e_l = \frac{\gamma L_n^{\gamma+1} S_n^\gamma}{(S_n^\gamma + L_n^\gamma)^2};\ p = \frac{e_s^\gamma}{e_s^\gamma + e_l^\gamma} = \frac{S_n^\gamma}{S_n^\gamma + L_n^\gamma} = \frac{1}{1+R_t^\gamma}.$$

This gives us a unique interior local maximum for $0 \leq \gamma \leq 1$; we must now verify that neither self prefers to deviate to $e_i = 0$. Given the above e_s and e_l, the expected utility for the short-term self is

$$\frac{S_n^\gamma}{S_n^\gamma + L_n^\gamma} X_n^s + \frac{L_n^\gamma}{S_n^\gamma + L_n^\gamma} X_n^l - \frac{\gamma L_n^\gamma S_n^{\gamma+1}}{(S_n^\gamma + L_n^\gamma)^2}$$

$$= \frac{S_n^\gamma}{S_n^\gamma + L_n^\gamma} X_n^s - \frac{S_n^\gamma}{S_n^\gamma + L_n^\gamma} X_n^l + \frac{S_n^\gamma + L_n^\gamma}{S_n^\gamma + L_n^\gamma} X_n^l - \frac{\gamma L_n^\gamma S_n^{\gamma+1}}{(S_n^\gamma + L_n^\gamma)^2}$$

$$= \frac{S_n^\gamma}{S_n^\gamma + L_n^\gamma}(X_n^s - X_n^l) + X_n^l - \frac{\gamma L_n^\gamma S_n^{\gamma+1}}{(S_n^\gamma + L_n^\gamma)^2} = \frac{S_n^{\gamma+1}}{S_n^\gamma + L_n^\gamma} + X_n^l - \frac{\gamma L_n^\gamma S_n^{\gamma+1}}{(S_n^\gamma + L_n^\gamma)^2}$$

$$= X_n^l + \frac{S_n^{\gamma+1}\left(S_n^\gamma + (1-\gamma)L_n^\gamma\right)}{(S_n^\gamma + L_n^\gamma)^2}.$$

If the short-term self exerts zero effort, their expected utility is X_n^l. So, we need

$$\frac{S_n^{\gamma+1}(S_n^\gamma + L_n^\gamma) - \gamma L_n^\gamma S_n^{\gamma+1}}{(S_n^\gamma + L_n^\gamma)^2} > 0 \Rightarrow S_n^{\gamma+1}(S_n^\gamma + L_n^\gamma) > \gamma L_n^\gamma S_n^{\gamma+1}$$

$$\Rightarrow (S_n^\gamma + L_n^\gamma) > \gamma L_n^\gamma \Rightarrow S_n^\gamma > (\gamma - 1)L_n^\gamma.$$

Which is true for $\gamma \leq 1$. Similarly for the long-term self. So, neither player wishes to deviate to $e_i = 0$; note that this implies that neither bliss point is a Pareto improvement over the disagreement point. Thus, we have a unique Nash equilibrium given by:

$$e_{s,n}^* = \frac{\gamma L_n^\gamma S_n^{\gamma+1}}{(S_n^\gamma + L_n^\gamma)^2},\ e_{l,n}^* = \frac{\gamma L_n^{\gamma+1} S_n^\gamma}{(S_n^\gamma + L_n^\gamma)^2},\ p_n^* = \frac{S_n^\gamma}{S_n^\gamma + L_n^\gamma} = \frac{1}{1+R_n^\gamma},$$

which creates an expected utility vector of:

$$\left(X_n^l + \frac{S_n^{\gamma+1}\left(S_n^\gamma + (1-\gamma)L_n^\gamma\right)}{(S_n^\gamma + L_n^\gamma)^2},\ Y_n^s + \frac{L_n^{\gamma+1}\left(L_n^\gamma + (1-\gamma)S_n^\gamma\right)}{(S_n^\gamma + L_n^\gamma)^2}\right).\square$$

Proof of Lemma 2: Pareto efficiency follows directly from the fact that we are applying the Nash bargaining solution, and so will only choose a point on the Pareto frontier of $\mathcal{U}$.

Consider replacing the shared payoff utility function $u(\cdot)$ with $v(\cdot) = \lambda u(\cdot) + \mu$. Denote $\mathcal{U}_n^u$ and

$\mathcal{U}_n^v$ the respective sets of utility vectors that result from using the two utility functions, and similarly for u and v superscripts on other terms. Every vector $(U_s, U_l) \in \mathcal{U}_n^u$ consists of weighted sums of $u(\cdot)$, and so becomes $(\lambda U_s + \mu,\ \lambda U_l + \mu) \in \mathcal{U}_n^v$.[36] As the bliss points are transformed similarly, $S_n^v = \lambda S_n^u$ and $L_n^v = \lambda L_n^u$. Putting these into the formula for d_n^v (with or without the effort costs added in), we find that $d_n^v = \lambda d_n^u + \mu$. Thus, each point in $\mathcal{U}_n^u$, as well as the disagreement point, have been subject to the same affine transformation. It follows from the properties of Nash bargaining that the outcome is subject to the same affine transformation, and so the resulting lottery (corresponding to that utility outcome) is unchanged.

If $\mathcal{U}_n$ is such that $(x, y) \in U_n$ if and only if $(y, x) \in U_n$, then the coordinates of the two bliss points are necessarily mirrors of each other. $X_n^s = Y_n^l$, $X_n^s = Y_n^l$, so that $S_n = L_n$. This immediately implies that the disagreement point is such that $d_{n,s} = d_{n,l}$. Given that the disagreement point is symmetric, the result follows from the Symmetry property of Nash bargaining.

Proof of Lemma 3: The time t_1 utility vectors given from choosing M_L or M_H are, respectively

$$\left(u(M_L)e^{-\rho_s \Delta_1},\ u(M_L)e^{-\rho_l \Delta_1}\right);\ \left(u(M_H)e^{-\rho_s(\Delta_1+\Delta_2)},\ u(M_H)e^{-\rho_l(\Delta_1+\Delta_2)}\right).$$

If $u(M_H)e^{-\rho_l \Delta_2} < u(M_L)$, then both selves prefer the sooner payoff, and $p_L = 1$. If $u(M_H)e^{-\rho_s \Delta_2} > u(M_L)$, then both selves prefer the latter payoff, and $p_L = 0$. In either case, the result follows directly. Consider now the case where:

$$u(M_H)e^{-\rho_l \Delta_2} > u(M_L) > u(M_H)e^{-\rho_s \Delta_2},$$

so that the long-term self strictly prefers the latter payoff, and the short-term self strictly prefers the sooner payoff. Then,

$$S_1 = u(M_L)e^{-\rho_s \Delta_1} - u(M_H)e^{-\rho_s(\Delta_1+\Delta_2)} = e^{-\rho_s \Delta_1}\left(u(M_L) - u(M_H)e^{-\rho_s \Delta_2}\right)$$

$$L_1 = u(M_H)e^{-\rho_l(\Delta_1+\Delta_2)} - u(M_L)e^{-\rho_l \Delta_1} = e^{-\rho_l \Delta_1}\left(u(M_H)e^{-\rho_l \Delta_2} - u(M_L)\right)$$

$$p_1 = \frac{S^\gamma}{S^\gamma + L^\gamma} = \frac{e^{-\rho_s \Delta_1 \gamma}\left(u(M_L) - u(M_H)e^{-\rho_s \Delta_2}\right)^\gamma}{e^{-\rho_s \Delta_1 \gamma}\left(u(M_L) - u(M_H)e^{-\rho_s \Delta_2}\right)^\gamma + e^{-\rho_l \Delta_1 \gamma}\left(u(M_H)e^{-\rho_l \Delta_2} - u(M_L)\right)^\gamma}$$

[36]The weighted sums are discounted sums of expectations of utilities induced by lotteries over actions.

$$= \frac{\left(u(M_L) - u(M_H)e^{-\rho_s \Delta_2}\right)^\gamma}{\left(u(M_L) - u(M_H)e^{-\rho_s \Delta_2}\right)^\gamma + e^{(\rho_2 - \rho_l)\Delta_1 \gamma}\left(u(M_H)e^{-\rho_l \Delta_2} - u(M_L)\right)^\gamma} \Rightarrow \frac{\partial}{\partial \Delta_1} p_1 < 0.$$

Thus, the probability the short-term self wins the conflict game is decreasing in Δ_1. From Lemma A.3, the decision of the agent will have the same weightings on the actions as the conflict game. Thus, the probability that the bliss point of the short-term self (the smaller, sooner payout) is chosen is decreasing in Δ_1. $\square$

Proof of Proposition 1 (Diminishing Impatience): Without loss of generality, we can consider the shared payoff utility to be zero. Note that this implies that the discounted utilities of the actions only differ through their impact on future action sets, and therefore future payoff utilities.

Consider the effects of an ε increase in Δ_1. Denote with superscript ε relevant terms after this increase. So, $\Delta_1^\varepsilon = \Delta_1 + \varepsilon$, $\mathcal{U}_1^\varepsilon$ is the set of utility vectors after the increase, S_1^ε and L_1^ε are the respective differences between bliss points utilities for the two selves after the increase. The time between the decision at time t_1 and all payoff utilities increases by ε, so that both selves discount the value of all lotteries at time t_1 by an additional $e^{-\varepsilon \rho_i}$ as a result of the change. The proof will proceed by considering three bargaining problems given by:

$$< \mathcal{U}_1,\ d_1 >; < \mathcal{U}_1^\varepsilon,\ d_1^a >; < \mathcal{U}_1^\varepsilon,\ d_1^\varepsilon >,$$

where

$$d_1^a = (e^{-\varepsilon \rho_s} d_{1,s},\ e^{-\varepsilon \rho_l} d_{1,l}).$$

We will show first that the first two bargaining problems result in the same chosen lottery, and second that the third bargaining problem improves the outcome for the long-term self relative to the second bargaining problem. This will then immediately imply that the ε problem gives higher utility to the long-term self.

Consider $< \mathcal{U}_1^\varepsilon,\ d_1^a >$. The first coordinate (short-term self utility) of all utility vectors in $\mathcal{U}_1^\varepsilon$, as well as the disagreement point, are multiplied by $e^{-\varepsilon \rho_s}$ relative to $< \mathcal{U}_1,\ d_1 >$. Similarly with the long-term self utility and $e^{-\varepsilon \rho_l}$. This means that $< \mathcal{U}_1^\varepsilon,\ d_1^a >$ is an affine transformation of $< \mathcal{U}_1,\ d_1 >$ and, by the invariance to affine transformations property of the Nash bargaining solution, both bargaining problems must result in the same action.

Now consider $< \mathcal{U}_1^\varepsilon,\ d_1^\varepsilon >$. Note that:

$$S_1^\varepsilon = e^{-\varepsilon \rho_s} X_1^s - e^{-\varepsilon \rho_s} X_1^l = e^{-\varepsilon \rho_s} S_1;\ L_1^\varepsilon = e^{-\varepsilon \rho_l} L_1^s - e^{-\varepsilon \rho_l} L_1^l = e^{-\varepsilon \rho_l} L_1.$$

We must consider the disagreement point both with and without the effort cost added in. Recall that without the cost added in:

$$d_1 = \left(X_1^l + \frac{S_1^{\gamma+1}}{S_1^\gamma + L_1^\gamma},\ Y_1^s + \frac{L_1^{\gamma+1}}{S_1^\gamma + L_1^\gamma} \right),$$

which implies:

$$d_1^\varepsilon = \left(e^{-\varepsilon\rho_s} X_n^l + \frac{(e^{-\varepsilon\rho_s} S_1)^{\gamma+1}}{(e^{-\varepsilon\rho_s} S_n)^\gamma + (e^{-\varepsilon\rho_l} L_n)^\gamma},\ e^{-\varepsilon\rho_l} Y_n^s + \frac{(e^{-\varepsilon\rho_l} L_1)^{\gamma+1}}{(e^{-\varepsilon\rho_s} S_n)^\gamma + (e^{-\varepsilon\rho_l} L_n)^\gamma} \right)$$

$$= \left(e^{-\varepsilon\rho_s} \left[X_n^l + \frac{{S_1}^{\gamma+1}}{{S_n}^\gamma + e^{(\rho_2-\rho_1)\gamma\varepsilon} {L_n}^\gamma} \right],\ e^{-\varepsilon\rho_l} \left[Y_n^s + \frac{{L_1}^{\gamma+1}}{e^{(\rho_1-\rho_2)\gamma\varepsilon} {S_n}^\gamma + {L_n}^\gamma} \right] \right).$$

Now, consider $< \mathcal{U}_1^\varepsilon,\ d_1^a >$ again. The disagreement point from this bargaining problem is:

$$d_1^a = \left(e^{-\varepsilon\rho_s} \left[X_n^l + \frac{{S_1}^{\gamma+1}}{{S_n}^\gamma + {L_n}^\gamma} \right],\ e^{-\varepsilon\rho_l} \left[Y_n^s + \frac{{L_1}^{\gamma+1}}{{S_n}^\gamma + {L_n}^\gamma} \right] \right).$$

Now comparing $< \mathcal{U}_1^\varepsilon,\ d_1^\varepsilon >$ to $< \mathcal{U}_1^\varepsilon,\ d_1^a >$, note that the two share the same set of utility vectors, but differ in their disagreement point. Specifically, since $\rho_2 > \rho_1$, $d_{1,s}^\varepsilon < d_{1,s}^a$, and $d_{1,l}^\varepsilon > d_{1,l}^a$. Applying Lemma A.1, with d_1^a as the first disagreement point, and d_1^ε as the second, we see that $\triangle d_s < 0$ and $\triangle d_l > 0$. Denoting in parenthesis the case where the Pareto frontier is smooth, the utility the long-term self receives under $< \mathcal{U}_1^\varepsilon,\ d_1^\varepsilon >$ is (strictly) greater than that received under $< \mathcal{U}_1^\varepsilon,\ d_1^a >$, meaning that it results in a lottery (strictly) more favorable for the long-term self. As $< \mathcal{U}_1^\varepsilon,\ d_1^a >$ and $< \mathcal{U}_1,\ d_1 >$ result in the same action, $< \mathcal{U}_1^\varepsilon,\ d_1^\varepsilon >$ likewise results in an action (strictly) preferred by the long-term self as compared to $< \mathcal{U}_1,\ d_1 >$. This gives the first part of the result.

We now show that this also holds for the disagreement point without the effort added in, which is given by:

$$d_1 = \left(X_n^l + \frac{S_n^{\gamma+1}\,(S_n^\gamma + (1-\gamma)L_n^\gamma)}{(S_n^\gamma + L_n^\gamma)^2},\ Y_t^s + \frac{L_n^{\gamma+1}\,(L_n^\gamma + (1-\gamma)S_n^\gamma)}{(S_n^\gamma + L_n^\gamma)^2} \right).$$

The disagreement points for the ε case and the affine transformation case are given by:

$$\begin{aligned} d_1^\varepsilon \quad = \quad & \Bigg(e^{-\varepsilon\rho_s} X_n^l + \frac{(e^{-\varepsilon\rho_s} S_n)^{\gamma+1}\,((e^{-\varepsilon\rho_s} S_n)^\gamma + (1-\gamma)(e^{-\varepsilon\rho_l} L_n)^\gamma)}{((e^{-\varepsilon\rho_s} S_n)^\gamma + (e^{-\varepsilon\rho_l} L_n)^\gamma)^2}, \\ & e^{-\varepsilon\rho_l} Y_t^s + \frac{(e^{-\varepsilon\rho_l} L_n)^{\gamma+1}\,((e^{-\varepsilon\rho_l} L_n)^\gamma + (1-\gamma)(e^{-\varepsilon\rho_s} S_n)^\gamma)}{((e^{-\varepsilon\rho_s} S_n)^\gamma + (e^{-\varepsilon\rho_l} L_n)^\gamma)^2} \Bigg) = \end{aligned}$$

$$\left(e^{-\varepsilon\rho_s}\left[X_n^l + \frac{{S_n}^{\gamma+1}\left(S_n^\gamma + (1-\gamma)e^{(\rho_2-\rho_1)\gamma\varepsilon}L_n^\gamma\right)}{(S_n^\gamma + e^{(\rho_2-\rho_1)\gamma\varepsilon}L_n^\gamma)^2}\right],\ e^{-\varepsilon\rho_l}\left[Y_t^s + \frac{L_n^{\gamma+1}\left(L_n^\gamma + (1-\gamma)e^{(\rho_1-\rho_2)\gamma\varepsilon}S_n^\gamma\right)}{(e^{(\rho_1-\rho_2)\gamma\varepsilon}S_n^\gamma + L_n^\gamma)^2}\right]\right);$$

$$d_1^a = \left(e^{-\varepsilon\rho_s}\left[X_n^l + \frac{{S_n}^{\gamma+1}(S_n^\gamma + (1-\gamma)L_n^\gamma)}{(S_n^\gamma + L_n^\gamma)^2}\right],\ e^{-\varepsilon\rho_l}\left[Y_t^s + \frac{L_n^{\gamma+1}(L_n^\gamma + (1-\gamma)S_n^\gamma)}{(S_n^\gamma + L_n^\gamma)^2}\right]\right).$$

If $d_{1,s}^\varepsilon < d_{1,s}^a$, and similarly, $d_{1,l}^\varepsilon > d_{1,l}^a$, then the application of Lemma 1 goes through as before.

$$d_{1,s}^\varepsilon < d_{1,s}^a \Leftrightarrow \frac{\left(S_n^\gamma + (1-\gamma)e^{(\rho_2-\rho_1)\gamma\varepsilon}L_n^\gamma\right)}{(S_n^\gamma + e^{(\rho_2-\rho_1)\gamma\varepsilon}L_n^\gamma)^2} < \frac{(S_n^\gamma + (1-\gamma)L_n^\gamma)}{(S_n^\gamma + L_n^\gamma)^2}$$

$$\Leftrightarrow \left(S_n^\gamma + (1-\gamma)e^{(\rho_2-\rho_1)\gamma\varepsilon}L_n^\gamma\right)(S_n^\gamma + L_n^\gamma)^2 < (S_n^\gamma + (1-\gamma)L_n^\gamma)(S_n^\gamma + e^{(\rho_2-\rho_1)\gamma\varepsilon}L_n^\gamma)^2$$

$$\Leftrightarrow \left(S_n^\gamma + (1-\gamma)e^{(\rho_2-\rho_1)\gamma\varepsilon}L_n^\gamma\right)\left(S_n^{2\gamma} + 2S_n^\gamma L_n^\gamma + L_n^{2\gamma}\right) <$$

$$(S_n^\gamma + (1-\gamma)L_n^\gamma)\left(S_n^{2\gamma} + 2e^{(\rho_2-\rho_1)\gamma\varepsilon}S_n^\gamma L_n^\gamma + e^{(\rho_2-\rho_1)\gamma\varepsilon}L_n^{2\gamma}\right)$$

$$\Leftrightarrow 2S_n^{2\gamma}L_n^\gamma + S_n^\gamma L_n^{2\gamma} + (1-\gamma)e^{(\rho_2-\rho_1)\gamma\varepsilon}L_n^\gamma S_n^{2\gamma} <$$

$$2e^{(\rho_2-\rho_1)\gamma\varepsilon}S_n^{2\gamma}L_n^\gamma + e^{(\rho_2-\rho_1)\gamma\varepsilon}S_n^\gamma L_n^{2\gamma} + (1-\gamma)L_n^\gamma S_n^{2\gamma}$$

$$\Leftrightarrow (1-\gamma)(e^{(\rho_2-\rho_1)\gamma\varepsilon} - 1)L_n^\gamma S_n^{2\gamma} < 2(e^{(\rho_2-\rho_1)\gamma\varepsilon} - 1)S_n^{2\gamma}L_n^\gamma + (e^{(\rho_2-\rho_1)\gamma\varepsilon} - 1)S_n^\gamma L_n^{2\gamma}$$

$e^{(\rho_2-\rho_1)\gamma\varepsilon} - 1$ is positive, so,

$$\Leftrightarrow (1-\gamma)L_n^\gamma S_n^{2\gamma} < 2S_n^{2\gamma}L_n^\gamma + S_n^\gamma L_n^{2\gamma} \Leftrightarrow 0 < (1+\gamma)S_n^\gamma + L_n^\gamma,$$

with the last true since $\gamma \geq 0$. Similarly for $d_{1,l}^\varepsilon > d_{1,l}^a$. Thus, the first part of Proposition 1 holds whether we use the disagreement point with or without the effort costs added in.

Now consider the second part of Proposition 1. The utility granted to the long-term self is bounded below by the utility granted to them by the disagreement point. Thus, the difference between the long-term self's bliss point utility and the utility granted to them by the action chosen by the agent is bounded above by the difference between their bliss point utility and the utility of the disagreement point. In other words,

$$\frac{Y_1^{l,\varepsilon} - U_l(\mathcal{D}(\mathcal{A}_1^\varepsilon))}{e^{-\rho_l(\triangle_1+\varepsilon)}} \leq \frac{Y_1^{l,\varepsilon} - d_{1,l}^\varepsilon}{e^{-\rho_l(\triangle_1+\varepsilon)}}.$$

So, let us consider the limit of the right hand side.

$$\lim_{\varepsilon\to\infty} \frac{Y_1^{l,\varepsilon} - d_{1,l}^\varepsilon}{e^{-\rho_l(\triangle_1+\varepsilon)}} = \lim_{\varepsilon\to\infty} e^{\rho_l(\triangle_1+\varepsilon)} \left(e^{-\varepsilon\rho_l} Y_1^l - e^{-\varepsilon\rho_l} \left[Y_1^s + \frac{{L_1}^{\gamma+1}}{e^{(\rho_1-\rho_2)\gamma\varepsilon} {S_n}^\gamma + {L_n}^\gamma} \right] \right) =$$

$$\lim_{\varepsilon\to\infty} e^{\rho_l\triangle_1} \left(Y_1^l - Y_1^s - \frac{{L_1}^{\gamma+1}}{e^{(\rho_1-\rho_2)\gamma\varepsilon} {S_n}^\gamma + {L_n}^\gamma} \right) = \lim_{\varepsilon\to\infty} e^{\rho_l\triangle_1} \left(L_1 - \frac{{L_1}^{\gamma+1}}{e^{(\rho_1-\rho_2)\gamma\varepsilon} {S_n}^\gamma + {L_n}^\gamma} \right)$$

$$= \lim_{\varepsilon\to\infty} e^{\rho_l\triangle_1} L_1 \left(1 - \frac{{L_1}^{\gamma}}{e^{(\rho_1-\rho_2)\gamma\varepsilon} {S_n}^\gamma + {L_n}^\gamma} \right) = 0,$$

the last following from the fact that $\rho_1 - \rho_2 < 0$. The result also follows with the disagreement point without effort costs added in.

$$\lim_{\varepsilon\to\infty} e^{\rho_l(\triangle_1+\varepsilon)} \left(e^{-\varepsilon\rho_l} Y_1^l - e^{-\varepsilon\rho_l} \left[Y_1^s + \frac{L_1^{\gamma+1}\left(L_1^\gamma + (1-\gamma)e^{(\rho_1-\rho_2)\gamma\varepsilon} S_1^\gamma\right)}{(e^{(\rho_1-\rho_2)\gamma\varepsilon} S_1^\gamma + L_1^\gamma)^2} \right] \right)$$

$$= \lim_{\varepsilon\to\infty} e^{\rho_l\triangle_1} \left(Y_1^l - Y_1^s - \frac{L_1^{\gamma+1}\left(L_1^\gamma + (1-\gamma)e^{(\rho_1-\rho_2)\gamma\varepsilon} S_1^\gamma\right)}{(e^{(\rho_1-\rho_2)\gamma\varepsilon} S_1^\gamma + L_1^\gamma)^2} \right)$$

$$= \lim_{\varepsilon\to\infty} e^{\rho_l\triangle_1} L_1 \left(1 - \frac{L_1^{\gamma}\left(L_1^\gamma + (1-\gamma)e^{(\rho_1-\rho_2)\gamma\varepsilon} S_1^\gamma\right)}{(e^{(\rho_1-\rho_2)\gamma\varepsilon} S_1^\gamma + L_1^\gamma)^2} \right) = 0.$$

So, for either disagreement point,

$$\lim_{\varepsilon\to\infty} \frac{Y_1^{l,\varepsilon} - U_l(\mathcal{D}(\mathcal{A}_1^\varepsilon))}{e^{-\rho_l(\triangle_1+\varepsilon)}} \leq \lim_{\varepsilon\to\infty} \frac{Y_1^{l,\varepsilon} - d_{1,l}^\varepsilon}{e^{-\rho_l(\triangle_1+\varepsilon)}} = 0.\square$$

Proof of Lemma 4: Consider a given commitment action $a_1 \in A_1^r$ which constrains A_n to a_n. The discounted utility a self receives from a_1 is given by:

$$U_i(a_1) = u(a_1) + K + e^{-\rho_i \Delta_t} U_i(a_n),$$

where K is the total of discounted payoff utilities resulting from all decisions strictly between 1 and n, of which there may be none. Following from the definition of required commitment set, and the fact that commitment actions a_1 cannot affect the action sets in periods other than n, $u(a_1)$ and K are the same $\forall\ a_1$, so without loss of generality we can normalize $u(a_1) + K$ to zero.[37] Similarly, a commitment lottery $\alpha_1 \in \mathcal{A}_1^r$ induces a realized lottery over actions at decision n, α_n. As the utilities of lotteries are equal to the lottery over utilities, we have:

$$U_i(\alpha_1) = e^{-\rho_i \Delta_t} U_i(\alpha_n) \Rightarrow e^{\rho_i \Delta_t} U_l(\alpha_1) = U_l(\alpha_n).$$

Now, note that $\alpha_n^{nc} \neq b_n^l$. That is, the agent will not choose a degenerate lottery on the bliss action of the long-term self at decision n; this follows from the fact that the bliss points are not Pareto improvements over the disagreement point (shown in the proof of Lemma 1), and thus cannot be the Nash bargaining solution. So,

$$U_l(\alpha_n^{nc}) < U_l(b_n^l) \Rightarrow 0 < U_l(b_n^l) - U_l(\alpha_n^{nc}).$$

Consider now the case where $\Delta_t = 0$. Because $\mathcal{U}_1^r$ is not discounted relative to $\mathcal{U}_n^{nc}$, $\mathcal{U}_n^{nc} = \mathcal{U}_1^r$, and the lottery over actions induced by the commitment lottery choice will be the same as the lottery chosen without commitment, so that $\alpha_n^{nc} = \alpha_n^r$. Now consider increasing Δ_t. By Proposition 1,

$$\lim_{\Delta_t \to \infty} \frac{Y_1^l - U_l(\mathcal{D}(\mathcal{A}_1^r))}{e^{-\rho_l \Delta_t}} = 0 \Rightarrow \lim_{\Delta_t \to \infty} e^{\rho_l \Delta_t} U_l(b_1^l) - e^{\rho_l \Delta_t} U_l(\alpha_1^r) = 0$$

$$\Rightarrow \lim_{\Delta_t \to \infty} U_l(b_n^l) = U_l(\alpha_n^r).$$

So, for $\Delta_t = 0$, $U_l(\alpha_n^r) = U_l(\alpha_n^{nc}) < U_l(b_n^l)$, while in the limit of Δ_t, $U_l(\alpha_n^{nc}) < U_l(\alpha_n^r) = U_l(b_n^l)$. Thus, there $\exists$ some $\underline{\Delta}$ for which $U_l(\alpha_n^{nc}) < U_l(\alpha_n^r)$. By the first part of Proposition 1, $U_l(\alpha_n^r)$ is monotonically increasing in Δ_t, so $U_l(\alpha_n^{nc}) < U_l(\alpha_n^r)\ \forall\ \Delta_t \geq \underline{\Delta}$. The difference in the utilities for the long-term self trivially implies that $\alpha_n^{nc} \neq \alpha_n^r$.

If the Pareto frontier of $\mathcal{U}_n^{nc}$ is smooth at $(U_s(\alpha_n^{nc}),\ U_l(\alpha_n^{nc}))$ then again consider the limiting case with $\Delta_t = 0$, so that $\alpha_n^{nc} = \alpha_n^r$, implying

$$(U_s(\alpha_n^{nc}),\ U_l(\alpha_n^{nc})) = (U_s(\alpha_n^r),\ U_l(\alpha_n^r)).$$

Since the utility vectors in $\mathcal{U}_1^r$ are the discounted utility vectors of $\mathcal{U}_n^{nc}$, the Pareto frontier of $\mathcal{U}_1^r$ is smooth at a vector if the Pareto frontier of $\mathcal{U}_n^{nc}$ is smooth at the corresponding vector. As $\mathcal{U}_n^{nc}$ is smooth at $(U_s(\alpha_n^r),\ U_l(\alpha_n^r))$, $\mathcal{U}_1^r$ is smooth at $(U_s(\alpha_1^r),\ U_l(\alpha_1^r))$. Finally, we can apply the first part of Proposition 1 to see that a marginal increase in Δ_t causes a strict increase in $U_l(\alpha_n^r)$, which

[37] This amounts to a non-distorting shifting of the entire bargaining problem, which has no impact on the lotteries chosen.

indicates for $\Delta_t > 0$ the utilities, and thus the actions, will differ.□

Proof of Proposition 2 (Preference for Commitment): First, denote $\alpha_1^{cnc} \in \mathcal{A}_1^o$ the lottery over commitment actions at decision 1 that induces the same lottery over decision n actions as α_n^{nc}. That is, α_1^{cnc} is committing to doing what you would have done if you hadn't committed; cnc indicates committing to the non-committing action lottery. It differs from not committing in that it carries with it some commitment cost, c, and so is Pareto dominated by the no commitment action, a_1^{nc} in $\mathcal{A}_1^o$. Now note that $\mathcal{U}_1^o$ consists of $\mathcal{U}_1^r$ with the addition of $(U_s(a_1^{nc}),\ U_l(a_1^{nc})) = (U_s(\alpha_1^{cnc}) + c,\ U_l(\alpha_1^{cnc}) + c)$, as well as mixings between that vector and the Pareto frontier of $\mathcal{U}_m^r$; see the right graph of Figure 7 for a visualization.

Start with the case where $\mathcal{U}_n^{nc}$ contains Pareto improvements on all strict mixings between α_n^{nc} and α_n^r, both of which lie on the Pareto frontier. Since $\mathcal{U}_1^r$ consists of discounted utility vectors from $\mathcal{U}_n^{nc}$, the utility vectors induced by α_1^{cnc} and α_1^r, $(U_s(a_1^{cnc}),\ U_l(a_1^{cnc}))$ and $(U_s(a_1^r),\ U_l(a_1^r))$, lie on the Pareto frontier of $\mathcal{U}_1^r$, and the line segment between their induced utility vectors lies inside of $\mathcal{U}_1^r$. However, $(U_s(a_1^{cnc}),\ U_l(a_1^{cnc}))$ is not on the Pareto frontier of $\mathcal{U}_1^o$, since it is dominated by the no commitment action in $a_1^{nc} \in A_1^o$.

We will proceed in two steps. Step 1 is to show that for sufficiently small cost of commitment, α_1^r is on the Pareto frontier of $\mathcal{U}_1^o$. Step 2 is to show that for sufficiently small cost of commitment, if α_1^r is on the Pareto frontier of $\mathcal{U}_1^o$, it will be the lottery chosen.

Step 1. Consider the lottery over decision 1 actions defined by $0.5a_1^{cnc} + 0.5\alpha_1^r$. As $\mathcal{U}_1^r$ contains Pareto improvements on all strict mixings between a_1^{cnc} and α_1^r, there is some point α_1^P with induced utility vector $\left(U_s(a_1^P),\ U_l(a_1^P)\right)$ on the Pareto frontier of $\mathcal{U}_1^r$ which Pareto dominates this 50/50 mixture. Specifically,

$$U_s(a_1^P) > 0.5U_s(a_1^{cnc}) + 0.5U_s(\alpha_1^r) \Rightarrow U_s(a_1^P) - 0.5U_s(a_1^{cnc}) - 0.5U_s(\alpha_1^r) \equiv \varepsilon_s > 0;$$

$$U_l(a_1^P) > 0.5U_l(a_1^{cnc}) + 0.5U_l(\alpha_1^r) \Rightarrow U_l(a_1^P) - 0.5U_l(a_1^{cnc}) - 0.5U_l(\alpha_1^r) \equiv \varepsilon_l > 0.$$

Now consider the lottery over decision 1 actions defined by $0.5a_1^{nc} + 0.5\alpha_1^r$. This lottery is Pareto dominated by a_1^P if

$$U_s(a_1^P) - 0.5U_s(a_1^{nc}) - 0.5U_s(\alpha_1^r) > 0 \Leftrightarrow U_s(a_1^P) - 0.5\left[U_s(\alpha_1^{cnc}) + c\right] - 0.5U_s(\alpha_1^r) > 0$$

$$\Leftrightarrow U_s(a_1^P) - 0.5U_s(\alpha_1^{cnc}) - 0.5c - 0.5U_s(\alpha_1^r) > 0 \Leftrightarrow \varepsilon_s - 0.5c > 0 \Leftrightarrow 2\varepsilon_s > c.$$

So, for sufficiently small c, $0.5a_1^{nc} + 0.5\alpha_1^r$ is Pareto dominated by a_1^P.

Now, we proceed by contradiction. Suppose that $0.5a_1^{nc} + 0.5\alpha_1^r$ is Pareto dominated by a_1^P, and α_1^r is not on the Pareto frontier of $\mathcal{U}_1^o$.

Note that $U_l(\alpha_1^{cnc}) < U_l(\alpha_1^r) \Rightarrow U_l(\alpha_1^r) - U_l(\alpha_1^{cnc}) \equiv \delta_l > 0$, by Lemma 4. As $U_l(\alpha_1^{cnc}) + c = U_l(a_1^{nc})$, it follows that for $c < \delta_l$, $U_l(a_1^{nc}) < U_l(\alpha_1^r)$, meaning that α_1^r is not Pareto dominated by a_1^{nc}. Then, it must be the case that α_1^r is dominated by a mixing between a_1^{nc} and some other lottery $\alpha_1^D \in \mathcal{A}_1^r$ that induces a utility vector on the Pareto frontier of $\mathcal{U}_1^r$. So,

$$wU_s(\alpha_1^D) + (1-w)U_s(a_1^{nc}) > U_s(\alpha_1^r);\ wU_l(\alpha_1^D) + (1-w)U_l(a_1^{nc}) > U_l(\alpha_l^r).$$

Then, it follows that since the mixture between a_1^{nc} and α_1^D dominates α_1^r, it must be that $U_l(\alpha_1^r) < U_l(\alpha_1^D)$. Similarly, as α_1^P is a mixture of α_1^r and a_1^{nc}, $U_l(a_1^{nc}) < U_l(\alpha_1^P) < U_l(\alpha_1^r)$. Altogether,

$$U_l(a_1^{nc}) < U_l(\alpha_1^P) < U_l(\alpha_1^r) < U_l(\alpha_1^D).$$

For ease of notation, I will refer to the utility vectors induced by these four as U^{nc}, U^P, U^r and U^D. I will use $m(U^{nc}, U^P)$ to refer to the magnitude of the slope between U^{nc} and U^P.

As U^P was established as a Pareto improvement on a mixing between U^{nc} and U^r, it must be the case that U^P lies above the mixing line segment connecting U^{nc} and U^r. This implies

$$m(U^P, U^{nc}) > m(U^r, U^{nc}) > m(U^r, U^P).$$

Next, U^P, U^r and U^D are all vectors on the Pareto frontier of $\mathcal{U}_1^r$. Since $\mathcal{U}_1^r$ is convex,

$$U_l(\alpha_1^P) < U_l(\alpha_1^r) < U_l(\alpha_1^D) \Rightarrow m(U^r, U^P) \geq m(U^r, U^D).$$

Together,

$$\Rightarrow m(U^r, U^{nc}) > m(U^r, U^D) \Rightarrow m(U^r, U^{nc}) > m(U^{nc}, U^D) > m(U^r, U^D).$$

But, if the last is true, then U^r lies above the mixing line between U^D and U^{nc}, which contradicts the assertion that U^r is Pareto dominated by a mixture of the two. Thus, we conclude that for

$$c < min\{\varepsilon_s, \varepsilon_l, \delta_l\},$$

α_1^r lay on the Pareto frontier of $\mathcal{U}_1^o$. Figure A.2 gives a visualization of the relevant slopes.

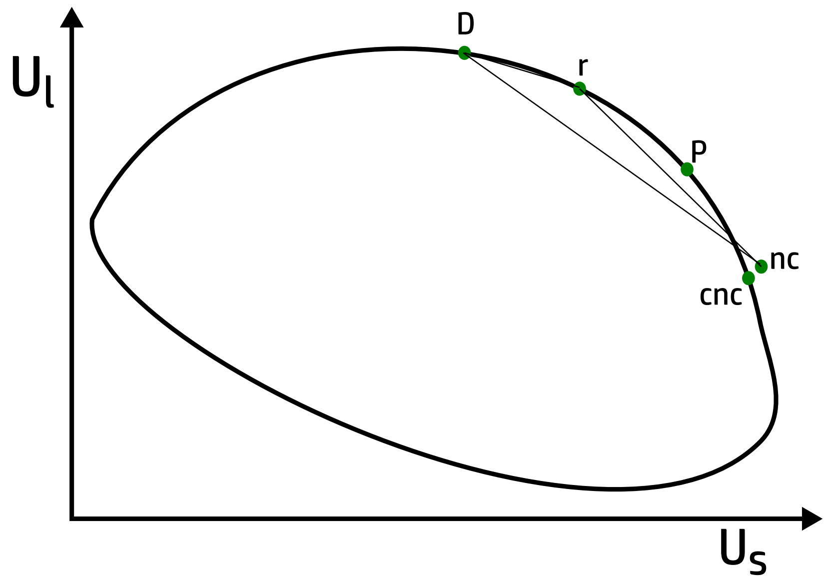

Figure A.2: Illustration of slopes.

Illustrated by the thick curve is $\mathcal{U}_1^r$. $\mathcal{U}_1^o$ consists of this plus the addition of U^{nc}, along with mixings with it. Finally, as $\mathcal{U}_1^r$ is convex, if the line segment connecting U^r to U^{cnc} is strictly inside the Pareto frontier (excepting the endpoints), as it is by supposition, then there exists an open subset of the Pareto frontier of $\mathcal{U}_1^r$ around U^r such that the line segment connecting any vector in this subset to U^{cnc} is also strictly inside the Pareto frontier. The argument that was used to show that α_1^r is on the Pareto frontier of $\mathcal{U}_1^o$, then, applies to the points in this open subset as well, which establishes that there is an open subset of the Pareto frontier of $\mathcal{U}_1^r$ containing α_1^r which is a subset of the Pareto frontier of $\mathcal{U}_1^o$.

Step 2. We now want to show that if α_1^r lay on the Pareto frontier of $\mathcal{U}_1^o$, then for sufficiently small c, $\alpha_1^r = \mathcal{D}(A_1^o)$. We begin by showing that for small c the bliss points of $\mathcal{U}_1^r$ are the same as for $\mathcal{U}_1^o$.

$$U_l(\alpha_1^{cnc}) < U_l(b_l^r) \Rightarrow U_l(b_l^r) - U_l(\alpha_1^{cnc}) = \zeta_l > 0 \Rightarrow U_l(b_l^r) - U_l(\alpha_1^{nc}) = \zeta_l - c;$$

$$U_s(\alpha_1^{cnc}) < U_s(b_s^r) \Rightarrow U_s(b_s^r) - U_s(\alpha_1^{cnc}) = \zeta_s > 0 \Rightarrow U_s(b_l^r) - U_s(\alpha_1^{nc}) = \zeta_s - c.$$

So, for $c < min\{\zeta_s, \zeta_l\}$, the bliss points of $\mathcal{U}_1^r$ grant higher utility than a_1^{nc}, which means that $\mathcal{U}_1^o$ has the same bliss points as $\mathcal{U}_1^r$. As the disagreement point is solely dependent on the bliss point, this means that the decision from $\mathcal{A}_1^r$ will utilize the same disagreement point as the decision from $\mathcal{A}_1^o$. Finally, the application of Lemma A.2, with $\triangle d_s = \triangle d_l = 0$, $(U_s(\alpha_1^r), U_l(\alpha_1^r)) = (U_s(\alpha_1^o), U_l(\alpha_1^o)) \Rightarrow \alpha_1^r = \alpha_1^o$.

Thus, for $c < min\{\varepsilon_s, \varepsilon_l, \delta_l, \zeta_l, \zeta_s\}$, $\alpha_1^r = \alpha_1^o$. This concludes the case for when $\mathcal{U}_n^{nc}$ contains Pareto improvements on all strict mixings between α_n^{nc} and α_n^r.

Now we consider the case where $\mathcal{U}_n^{nc}$ does not contain Pareto improvements on strict mixings between α_n^{nc} and α_n^r. In this case, the utility vectors induced by the two lotteries must lie on a line segment which is a subset of the Pareto frontier of $\mathcal{U}_n^{nc}$. Similarly, then, α_1^{cnc} and α_1^r must lie on a line segment on the Pareto frontier of $\mathcal{U}_1^r$. As for the first part of the proof, this is because $\mathcal{U}_1^r$ is a distortion of $\mathcal{U}_n^{nc}$. Consider the endpoints of this line segment, and the corresponding lotteries. Denote these lotteries g and h. We can then write:

$$\alpha_1^r = \theta^r g + (1-\theta^r)h;\ \alpha_1^{cnc} = \theta^{nc} g + (1-\theta^{nc})h,$$

for some values of θ^r and θ^{nc}. Denote $g_l = U_l(g)$, $g_s = U_s(g)$, $h_l = U_l(h)$, $h_s = U_s(h)$. As α_1^r is the result of the Nash bargaining solution for $\mathcal{U}_1^r$, and thus maximizes the Nash product, we know that

$$(\theta^r g_l + (1-\theta^r)h_l - d_l^r)(\theta^r g_s + (1-\theta^r)h_s - d_s^r)$$

$$-(\theta^{nc} g_l + (1-\theta^{nc})h_l - d_l^r)(\theta^{nc} g_s + (1-\theta^{nc})h_s - d_s^r) = \varepsilon_N > 0.$$

Now we look at the no commitment lottery in $\mathcal{A}_1^o$,

$$(U_s(a_1^{nc}),\ U_l(a_1^{nc})) = (U_s(\alpha_1^{cnc}) + c,\ U_l(\alpha_1^{cnc}) + c)$$

$$= (\theta^{nc} g_s + (1-\theta^{nc})h_s + c,\ \theta^{nc} g_l + (1-\theta^{nc})h_l + c)\,.$$

Consider the lottery given by

$$\tilde{\alpha}_1^r = \frac{\theta^r - \theta^{nc}}{1-\theta^{nc}} g + \frac{1-\theta^r}{1-\theta^{nc}} a_1^{nc}$$

Intuitively, $\tilde{\alpha}_1^r$ maintains the same relative position as α_1^r, but is shifted up to the Pareto frontier that has been expanded by the addition of a_1^{nc} to the choice set. Now,

$$(U_s(\tilde{\alpha}_1^r),\ U_l(\tilde{\alpha}_1^r)) =$$

$$\left(\frac{\theta^r-\theta^{nc}}{1-\theta^{nc}}g_s+\frac{1-\theta^r}{1-\theta^{nc}}\left[\theta^{nc}g_s+(1-\theta^{nc})h_s+c\right],\ \frac{\theta^r-\theta^{nc}}{1-\theta^{nc}}g_l+\frac{1-\theta^r}{1-\theta^{nc}}\left[\theta^{nc}g_l+(1-\theta^{nc})h_l+c\right]\right)$$

$$=\left(\theta^r g_s+(1-\theta^r)h_s+\frac{1-\theta^r}{1-\theta^{nc}}c,\ \theta^r g_l+(1-\theta^r)h_l+\frac{1-\theta^r}{1-\theta^{nc}}c\right).$$

As previously established, if $c < min\{\zeta_s, \zeta_l\}$ then the disagreement points of $\mathcal{U}_1^o$ and $\mathcal{U}_1^r$ are the same. Now we look at the difference in the Nash products that result from selecting $\tilde{\alpha}_1^r$ or selecting a_1^{nc} from $\mathcal{U}_1^o$.

$$\left(\theta^r g_s+(1-\theta^r)h_s+\frac{1-\theta^r}{1-\theta^{nc}}c-d_s\right)\left(\theta^r g_l+(1-\theta^r)h_l+\frac{1-\theta^r}{1-\theta^{nc}}c-d_l\right)-$$

$$\left(\theta^{nc}g_s+(1-\theta^{nc})h_s+c-d_s\right)\left(\theta^{nc}g_l+(1-\theta^{nc})h_l+c-d_l\right)=$$

$$\frac{1-\theta^r}{1-\theta^{nc}}c\left(\theta^r(g_l+g_s)+(1-\theta^r)(h_l+h_s)+2\frac{1-\theta^r}{1-\theta^{nc}}c-(d_l+d_s)\right)-$$

$$c\left(\theta^{nc}(g_l+g_s)+(1-\theta^{nc})(h_l+h_s)+2c-(d_l+d_s)\right)+\varepsilon_N.$$

For sufficiently low c, this difference is positive, as all terms except ε_N approach zero. This means that for sufficiently low c, $\tilde{\alpha}_1^r$ results in a higher Nash product than a_1^{nc}, so that a degenerate lottery on a_1^{nc} cannot be the choice of the agent.□

Proof of Lemma 5: Two cases. If $c > 0$, then because the discounted utility received from actions committed to vanishes as $\Delta_t \to \infty$, and the no commitment action a_1^{nc} has strictly higher payoff utility than all other actions in $\mathcal{A}_1^o$, for sufficiently high t the utility vector induced by the no commitment action Pareto dominates all other utility vectors, and the result follows.

If $c = 0$, then $U(a_1^{nc}) = U(\alpha_1^{cnc})$. For $\Delta_t = 0$, $\mathcal{U}_1^o$ and $\mathcal{U}_n^{nc}$ are all the same utility vector set. Thus, as we know the agent chooses α_n^{nc} at decision n, he must choose some mixture of the equivalent vectors from $\mathcal{U}_1^o$, which are a_1^{nc} and α_1^{cnc}. By Proposition 1, increasing the time between commitment and payoff (Δ_t) cannot decrease the future payoff to the long-term self, meaning that α_n^o (the lottery of decision n actions induced by α_1^o) must weakly improve for the long-term self (move up on the Pareto frontier of $\mathcal{U}_n^{nc}$). But this means that α_1^o must also move up the Pareto frontier of U_1^o from a_1^{nc}, and the result follows.□

Proof of Proposition 3 (Minimal Commitment Equivalence): First, note that A_1^m may be represented as $[a,\ b]$, with $a_1^m \in [a,\ b]$ being the minimal action committed to, so that choosing $a_1^m \in A_1^m$ causes the action set at decision n to be limited to $[a_1^m,\ b]$. By supposition, $\mathcal{U}_n([a_1^m,\ b])$ will have a Pareto frontier of pure actions, and so the decision made by the agent in decision n is a pure action in $[a_1^m,\ b]$. By Lemma 6, then, the commitment to a minimal a_1^m is equivalent to the commitment to some $a_n \in [a_1^m,\ b]$.

Now, as $U_i(a_n)$ is continuous by supposition, $\mathcal{U}_n$ varies continuously with the size of the interval of actions available at decision n. Likewise, as the disagreement point varies continuously with the bliss points, which vary continuously with $\mathcal{U}_n$, the disagreement point varies continuously with the size of the interval of actions available. Then, as the Nash bargaining solution is continuous with respect to the disagreement point values, and points on the Pareto frontier, the decision made by the agent at decision n varies continuously with the size of the interval of actions available.

Consider $a_1^m = b$. This limits the action set at decision n to $[b,\ b] = \{b\}$. Now consider decreasing a_1^m continuously, so that the action set at decision n, $[a_1^m,\ b]$ is increasing. By the previous continuity argument, this continuous decrease of a_1^m results in decision actions n that vary continuously, starting from b; they cannot go above b since that is the defined maximum action. Varying a_1^m from b to a, the set of decision n actions that result vary from b to some minimum of a'. Every $a_n \in [a',\ b]$ is induced by some $a_1^m \in [a,\ b]$, and a_1^m induces no actions outside of $[a',\ b]$. Thus, $\mathcal{A}_1^r$ and $\mathcal{A}_1^m$ create the same set of decision n actions, and therefore the same decision 1 discounted utilities. As they have the same commitment cost c by supposition, the result follows.□

Proof of Proposition 4 (Pure Temptation): I will drop the n subscript here for ease of notation. First, note that the disagreement point, with or without effort cost added in, can be expressed as a function of the four bliss point utility values $X^s,\ Y^s,\ X^l,\ Y^l$. We can then consider the change in the disagreement point utilities with respect to a change in the bliss point utility values. Start with the disagreement point with effort costs added in.

$$d = (d_s,\ d_l) = \left(X^l + \frac{S^{\gamma+1}}{S^\gamma + L^\gamma},\ Y^s + \frac{L^{\gamma+1}}{S^\gamma + L^\gamma}\right).$$

$S = X^s - X^l$ and $L = Y^l - Y^s$, so that

$$\frac{\partial d_s}{\partial X^s} = \frac{(\gamma+1)S^\gamma}{S^\gamma + L^\gamma} - \frac{S^{\gamma+1}\gamma S^{\gamma-1}}{(S^\gamma + L^\gamma)^2} = \frac{S^{2\gamma} + (\gamma+1)S^\gamma L^\gamma}{(S^\gamma + L^\gamma)^2} > 0;$$

$$\frac{\partial d_l}{\partial X^s} = -\frac{L^{\gamma+1}S^{\gamma-1}}{(S^\gamma + L^\gamma)^2} < 0;$$

$$\Rightarrow \frac{\partial d_s}{\partial X^s}(U_l^1 - d_l^1) - \frac{\partial d_l}{\partial X^s}(U_s^1 - d_s^1) > 0 \Rightarrow \frac{\partial U_s(\mathcal{D}(\mathcal{A}))}{\partial X^s} \geq 0,$$

with the last from application of Lemma A.2, with the inequality strict if the Pareto frontier is smooth. Since by supposition there are no Pareto improvements on the original choice, we also have

$$\frac{\partial U_l(\mathcal{D}(\mathcal{A}))}{\partial X^s} \leq 0,$$

also with the inequality strict if the Pareto frontier is smooth.

Now consider the disagreement point with effort costs added in,

$$d = (d_s, d_l) = \left(X^l + \frac{S^{2\gamma+1} + (1-\gamma)S^{\gamma+1}L^\gamma}{(S^\gamma + L^\gamma)^2},\ Y^s + \frac{L^{2\gamma+1} + (1-\gamma)L^{\gamma+1}S^\gamma}{(S^\gamma + L^\gamma)^2} \right).$$

$$\frac{\partial d_s}{\partial X^s} = \frac{(2\gamma+1)S^{2\gamma} + (1-\gamma)(1+\gamma)S^\gamma L^\gamma}{(S^\gamma + L^\gamma)^2} - 2\gamma S^{\gamma-1}\frac{S^{2\gamma+1} + (1-\gamma)S^{\gamma+1}L^\gamma}{(S^\gamma + L^\gamma)^3}$$

$$= \frac{S^{3\gamma} + (\gamma^2+2)S^{2\gamma}L^\gamma + (1-\gamma^2)S^\gamma L^{2\gamma}}{(S^\gamma + L^\gamma)^3} > 0;$$

$$\frac{\partial d_l}{\partial X^s} = \frac{\gamma(1-\gamma)S^{\gamma-1}L^{\gamma+1}}{(S^\gamma + L^\gamma)^2} - 2\gamma S^{\gamma-1}\frac{L^{2\gamma+1} + (1-\gamma)L^{\gamma+1}S^\gamma}{(S^\gamma + L^\gamma)^3}$$

$$= \frac{\gamma(\gamma-1)S^{2\gamma-1}L^{\gamma+1} - \gamma(\gamma+1)S^{\gamma-1}L^{2\gamma+1}}{(S^\gamma + L^\gamma)^3} < 0;$$

So, the same reasoning applies.□

Proof of Lemma 7: An increase in X^s has two potential effects to consider. One is the alteration of the disagreement point, and one is the alteration of the Pareto frontier. The first increases $U_s(\mathcal{D}(\mathcal{A}))$, as seen in the proof of Proposition 4. For the second, an increase in X^s can expand the Pareto frontier, but not contract it. The utilities granted by the Nash bargaining solution both strictly increase if Pareto improvements upon a bargaining outcome are added. Thus, both effects increase $U_s(\mathcal{D}(\mathcal{A}))$.□

Proof of Proposition 5: We drop the n subscripts for ease of notation. As in the proof of Proposition 4, we look at the effect on the disagreement point of altering the bliss values. Note that $\frac{\partial L}{\partial Y^s} = -1$, and consider first the disagreement point with effort costs added in.

$$\frac{\partial d_l}{\partial Y^s} = 1 + \frac{-(\gamma+1)L^\gamma}{S^\gamma + L^\gamma} + \frac{\gamma L^{\gamma-1}L^{\gamma+1}}{(S^\gamma + L^\gamma)^2} = \frac{(S^\gamma + L^\gamma)^2 - (\gamma+1)L^\gamma(S^\gamma + L^\gamma) + \gamma L^{\gamma-1}L^{\gamma+1}}{(S^\gamma + L^\gamma)^2}$$

$$= \frac{(1-\gamma)S^\gamma L^\gamma + S^{2\gamma}}{(S^\gamma + L^\gamma)^2} > 0; \ \frac{\partial d_s}{\partial Y^s} = \frac{\gamma S^{\gamma+1} L^{\gamma-1}}{(S^\gamma + L^\gamma)^2} > 0.$$

As the disagreement point improves for both selves, the application of Lemma A.2 is less straightforward than for the proof of Proposition 4.

$$\triangle d_s(U_l(\alpha) - d_l) - \triangle d_l(U_s(\alpha) - d_s) < 0 \Leftrightarrow$$

$$\frac{\gamma S^{\gamma+1} L^{\gamma-1}}{(S^\gamma + L^\gamma)^2}(U_l(\alpha) - d_l) - \frac{(1-\gamma)S^\gamma L^\gamma + S^{2\gamma}}{(S^\gamma + L^\gamma)^2}(U_s(\alpha) - d_s) < 0$$

$$\Leftrightarrow (\gamma S L^{\gamma-1})(U_l(\alpha) - d_l) < ((1-\gamma)L^\gamma + S^\gamma)(U_s(\alpha) - d_s) \Leftrightarrow \frac{U_l(\alpha) - d_l}{U_s(\alpha) - d_s} < \frac{(1-\gamma)L^\gamma + S^\gamma}{\gamma S L^{\gamma-1}}.$$

So,

$$\frac{U_l(\alpha) - d_l}{U_s(\alpha) - d_s} < \frac{(1-\gamma)L^\gamma + S^\gamma}{\gamma S L^{\gamma-1}} \Rightarrow \triangle d_s(U_l(\alpha) - d_l) - \triangle d_l(U_s(\alpha) - d_s) < 0 \Rightarrow$$

$$\frac{\partial U_s(\mathcal{D}(\mathcal{A}))}{\partial Y^s} \leq 0, \frac{\partial U_l(\mathcal{D}(\mathcal{A}))}{\partial Y^s} \geq 0. \square$$

Proposition 5' (Indirect Temptation with original disagreement point): *Consider a decision maker choosing α_n from $\mathcal{A}_n$, corresponding $\mathcal{U}_n$, and consider the effects of altering the utility granted to the long-term self by the short-term self bliss point: Y_n^s. As long as an open connected subset of the Pareto frontier of $\mathcal{U}_n$ containing the Nash bargaining outcome $(U_s(\alpha_n),\ U_l(\alpha_n))$ remains a subset of the Pareto frontier of $\mathcal{U}_n$, then*

$$\frac{L}{S} * \frac{(S_n^\gamma + L_n^\gamma)^2 + \gamma^2 L_n^\gamma (S_n^\gamma - L_n^\gamma)}{\gamma L_n^\gamma (S_n^\gamma + L_n^\gamma) + \gamma^2 L_n^\gamma (S_n^\gamma - L_n^\gamma)} < \frac{U_l(\alpha_n) - d_l}{U_s(\alpha_n) - d_s} \Rightarrow \frac{\partial U_s(\mathcal{D}(\mathcal{A}_n))}{\partial Y_n^s} \geq 0, \frac{\partial U_l(\mathcal{D}(\mathcal{A}_n))}{\partial Y_n^s} \leq 0,$$

with the reverse strict inequality implying the reverse weak inequalities.

Proof: The proof is the same as that for Proposition 5, other than the derivatives of the disagreement point values.

$$\frac{\partial d_l}{\partial Y^s} = 1 - \frac{(2\gamma+1)L^{2\gamma} + (1-\gamma)(1+\gamma)S^\gamma L^\gamma}{(S^\gamma + L^\gamma)^2} + 2\gamma L^{\gamma-1} \frac{L^{2\gamma+1} + (1-\gamma)L^{\gamma+1} S^\gamma}{(S^\gamma + L^\gamma)^3}$$

$$= \frac{(S^\gamma + L^\gamma)^3 - (S^\gamma + L^\gamma)\left((2\gamma+1)L^{2\gamma} + (1-\gamma)(1+\gamma)S^\gamma L^\gamma\right) + 2\gamma L^{3\gamma} + 2\gamma(1-\gamma)L^{2\gamma}S^\gamma}{(S^\gamma + L^\gamma)^3}$$

$$= \frac{S^{2\gamma}L^\gamma\left(2+\gamma^2\right) + S^{3\gamma} + S^\gamma L^{2\gamma}\left(1-\gamma^2\right)}{(S^\gamma + L^\gamma)^3}.$$

$$\frac{\partial d_s}{\partial Y^s} = -\frac{\gamma(1-\gamma)S^{\gamma+1}L^{\gamma-1}}{(S^\gamma + L^\gamma)^2} + 2\gamma L^{\gamma-1}\frac{S^{2\gamma+1} + (1-\gamma)S^{\gamma+1}L^\gamma}{(S^\gamma + L^\gamma)^3}$$

$$= \frac{-(S^\gamma + L^\gamma)\gamma(1-\gamma)S^{\gamma+1}L^{\gamma-1} + 2\gamma L^{\gamma-1}S^{2\gamma+1} + 2\gamma(1-\gamma)S^{\gamma+1}L^{2\gamma-1}}{(S^\gamma + L^\gamma)^3}$$

$$= \frac{(\gamma-\gamma^2)S^{\gamma+1}L^{2\gamma-1} + (\gamma+\gamma^2)S^{2\gamma+1}L^{\gamma-1}}{(S^\gamma + L^\gamma)^3}.$$

So that,

$$\frac{\partial d_l/\partial Y^s}{\partial d_s/\partial Y^s} = \frac{S^{2\gamma}L^\gamma\left(2+\gamma^2\right) + S^{3\gamma} + S^\gamma L^{2\gamma}\left(1-\gamma^2\right)}{(\gamma-\gamma^2)S^{\gamma+1}L^{2\gamma-1} + (\gamma+\gamma^2)S^{2\gamma+1}L^{\gamma-1}}$$

$$= \frac{L}{S} * \frac{S^\gamma L^\gamma\left(2+\gamma^2\right) + S^{2\gamma} + L^{2\gamma}\left(1-\gamma^2\right)}{(\gamma-\gamma^2)L^{2\gamma} + (\gamma+\gamma^2)S^\gamma L^\gamma}$$

$$= \frac{L}{S} * \frac{(S^\gamma + L^\gamma)^2 + \gamma^2 L^\gamma(S^\gamma - L^\gamma)}{\gamma L^\gamma(S^\gamma + L^\gamma) + \gamma^2 L^\gamma(S^\gamma - L^\gamma)}$$

The rest of the proof follows as for Proposition 5.□

Proof of Lemma 9: Define $R = \frac{L}{S}$. Then,

$$wU_l(a) + (1-w)U_s(a) > wU_l(b) + (1-w)U_s(b) \Leftrightarrow w(U_l(a) - U_l(b)) > (1-w)(U_s(b) - U_s(a))$$

$$\Leftrightarrow wL > (1-w)S \Leftrightarrow R > \frac{1-w}{w} \Leftrightarrow p = \frac{1}{R^\gamma + 1} < \frac{1}{\left(\frac{1-w}{w}\right)^\gamma + 1},$$

the last connection from Lemma 1 and A.3.□

Proof of Lemma 11: By lemma 4, the ultimate action will shift in favor of the long-term self for sufficiently large Δ_t. This shifted action, in turn, will increase the discounted utility for the long-term self and decrease it for the short-term self. For sufficiently large Δ_t, the ratio between this increase and decrease can be made arbitrarily large, as the short-term self discounts at a faster rate. Thus, for any weighting of the long-term self's utility, we can choose Δ_t high enough that the weighted sum will increase.$\square$

Appendix B: Proofs for Temporal Reference Points Model

Proof of Lemma 1: First, diminishing impatience. As the agent is information focused, and there are no new informational updates in a deterministic prospect, there are no temporal reference points that break up the length of time which the agent discounts over. Thus,

$$(a,\ 0) \sim (b,\ t_1) \Leftrightarrow u(a) = u(b)D(t_1), \text{ and}$$

$$(a,\ t_2) \preceq (b,\ t_1 + t_2) \Leftrightarrow u(a)D(t_2) \leq u(b)D(t_1 + t_2).$$

$$\text{So, } (a,\ 0) \sim (b,\ t_1) \Rightarrow (a,\ t_2) \preceq (b,\ t_1 + t_2)\ \forall\ a, b \in \mathcal{O},\ u(a),\ u(b) > 0,\ \forall\ t_1, t_2 > 0\ \Leftrightarrow$$

$$u(a) = u(b)D(t_1) \Rightarrow u(a)D(t_2) \leq u(b)D(t_1 + t_2)\ \forall\ a, b \in \mathcal{O},\ u(a),\ u(b) > 0,\ \forall\ t_1, t_2 > 0\ \Leftrightarrow$$

$$u(a)/u(b) = D(t_1) \Rightarrow u(a)/u(b) \leq D(t_1 + t_2)/D(t_2)\ \forall\ a, b \in \mathcal{O},\ u(a),\ u(b) > 0,\ \forall\ t_1, t_2 > 0\ \Leftrightarrow$$

$$D(t_1) \leq D(t_1 + t_2)/D(t_2)\ \forall\ t_1, t_2 > 0\ \Leftrightarrow D(t_1)D(t_2) \leq D(t_1 + t_2)\ \forall\ t_1, t_2 > 0.$$

For strongly diminishing impatience, similarly,

$$(a,\ t_1) \sim (b,\ t_2) \Leftrightarrow u(a)D(t_1) = u(b)D(t_2), \text{ and}$$

$$(a,\ t_1 + t_3) \preceq (b,\ t_2 + t_3) \Leftrightarrow u(a)D(t_1 + t_3) \leq u(b)D(t_2 + t_3).$$

$$\text{So, } (a,\ t_1) \sim (b,\ t_2) \Rightarrow (a,\ t_1 + t_3) \preceq (b,\ t_2 + t_3)\ \forall\ a, b \in \mathcal{O},\ u(a),\ u(b) > 0,$$

$$\forall\ t_2 > t_1 \geq 0,\ t_3 > 0,\ \Leftrightarrow$$

$$u(a)D(t_1) = u(b)D(t_2) \Rightarrow u(a)D(t_1 + t_3) \leq u(b)D(t_2 + t_3)\ \forall\ a, b \in \mathcal{O},\ u(a),\ u(b) > 0,$$

$$\forall\ t_2 > t_1 \geq 0,\ t_3 > 0,\ \Leftrightarrow$$

$$u(a)/u(b) = D(t_2)/D(t_1) \Rightarrow u(a)/u(b) \leq D(t_2 + t_3)/D(t_1 + t_3)\ \forall\ a, b \in \mathcal{O},\ u(a),\ u(b) > 0,$$

$$\forall\ t_2 > t_1 \geq 0,\ t_3 > 0,\ \Leftrightarrow$$

$$D(t_2)/D(t_1) \leq D(t_2 + t_3)/D(t_1 + t_3)\ \forall\ t_2 > t_1 \geq 0,\ t_3 > 0,\ \Leftrightarrow$$

$$D(t_2)D(t_1 + t_3) \leq D(t_1)D(t_2 + t_3)\ \forall\ t_2 > t_1 \geq 0,\ t_3 > 0,$$

with the strict cases following by simply making the inequalities strict at each step. As far as the relation to the hazard rate when the discounting function is continuously differentiable, it will suffice to show, that when the hazard rate is decreasing,

$$D(t_2)D(t_1 + t_3) \leq D(t_1)D(t_2 + t_3)\ \forall\ t_2 > t_1 \geq 0,\ t_3 > 0\ \Leftrightarrow h'(t) \leq 0\ \forall\ t \geq 0.$$

$$h'(t) \leq 0 \; \forall \; t \geq 0 \Leftrightarrow \frac{-D'(t+\alpha)}{D(t+\alpha)} \leq \frac{-D'(t)}{D(t)} \; \forall \; t \geq 0, \;\; \alpha > 0$$

$$\Leftrightarrow \int_t^{t+\gamma} \frac{D'(\tau+\alpha)}{D(\tau+\alpha)} d\tau > \int \frac{D'(\tau)}{D(\tau)} d\tau \; \forall \; t \geq 0, \;\; \alpha, \gamma > 0$$

$$\Leftrightarrow [ln(D(\tau+\alpha))]_t^{t+\gamma} \geq [ln(D(\tau))]_t^{t+\gamma} \;\; \forall \; t \geq 0, \;\; \alpha, \gamma > 0$$

$$\Leftrightarrow ln(D(t+\alpha+\gamma)) - ln(D(t+\alpha)) \geq ln(D(t+\gamma)) - ln(D(t)) \; \forall \; t \geq 0, \;\; \alpha, \gamma > 0$$

$$\Leftrightarrow ln\left(\frac{D(t+\alpha+\gamma)}{D(t+\alpha)}\right) \geq ln\left(\frac{D(t+\gamma)}{D(t)}\right) \;\; \forall \; t \geq 0, \;\; \alpha, \gamma > 0$$

$$\Leftrightarrow \frac{D(t+\alpha+\gamma)}{D(t+\alpha)} \geq \frac{D(t+\gamma)}{D(t)} \; \forall \; t \geq 0, \;\; \alpha, \gamma > 0$$

$$\Leftrightarrow D(t)D(t+\alpha+\gamma) \geq D(t+\gamma)D(t+\alpha) \; \forall \; t \geq 0, \;\; \alpha, \gamma > 0$$

Denote $t_1 = t$, $t_2 = t + \gamma$ and $t_3 = \alpha$,

$$\Leftrightarrow D(t_1)D(t_2+t_3) \geq D(t_2)D(t_1+t_3) \; \forall \; t_2 > t_1 \geq 0, \;\; t_3 > 0. \square$$

Proof of Proposition 1: Consider two prospects A and B, such that A has more grouped information than B, with $\|I_B\| - \|I_A\| = 1$. Recall that by definition the prospects have the same time 0 distribution over outcomes, and let the unique point in time that B gives an informational update at but A does not be denoted by t_b.

Let's first see that $\|I_B\| > 1$. If not, $I_A = \{\}$, which implies that prospect A has no informational updates, which means it is a deterministic prospect; any non-deterministic prospect necessarily gives information at some point in the resolution, whether it is before the payout is received, or at the time the payout is received. Since A and B have the same time 0 distribution over outcomes, this means that prospect B must also be deterministic. But this contradicts that prospect B has an informational update, and so by contradiction $\|I_B\| > 1$.

Recall that we can write the value of a prospect A as

$$V_A = \int_{(\alpha,t_n)\in\mathcal{O}} u(\alpha) * \triangle_A^0(\alpha,t) * D(t_1^A - 0) * D(t_2^A - t_1^A) .... * D(t_{n-1}^A - t_{n-2}^A) * D\left(t_n - t_{n-1}^A\right).$$

Denote by t_{b-1} and t_{b+1} as the elements of T_A and T_B before and after t_b respectively. Because $0 < t_b < t_F^B$, these elements both exist. Further, both T_A and T_B share these two elements, while only

T_B contains t_b. Note that in this case no outcomes can occur at time t_b, because both prospects have the same distribution over outcomes, and thus all times at which outcomes occur must be included in both T_A and T_B. Now consider two types of outcomes.

The first group are those occurring at or before time t_{b-1}; this group may be empty. If not, then all such outcomes are discounted by the same sequence of discounting multipliers in both prospects, as both T_A and T_B have the same series of elements at and before t_{b-1}. Thus, such outcomes give the same present value in the calculation of the value of the prospects, and can be ignored in a comparison of which prospect has the higher value.

The second group are those occurring at or after t_{b+1}. This group is non-empty, because at minimum some outcomes occur at $t_F^A = t_F^B$. The discounting given to these elements for prospect A will include:

$$.... * D(t_{b+1} - t_{b-1}) * ....,$$

while for prospect B it will include:

$$.... * D(t_b - t_{b-1}) * D(t_{b+1} - t_b) * ....,$$

where all discount factors before and after those shown are the same for both prospects, since T_A and T_B are the same except for t_b. Thus, these elements are discounted more heavily for prospect B, and thus B has a lower value than A, if and only if

$$D(t_b - t_{b-1}) * D(t_{b+1} - t_b) \leq D(t_{b+1} - t_{b-1}).$$

And B has a strictly lower value if and only if the inequality is strict.

Now, if $A\ (\succ) \succeq B$ for any two prospects such that A has more grouped information than B and $\|I_B\| - \|I_A\| = 1$, then the present value of prospect A is (strictly) higher than for prospect B. Since this is true for any such A and B, for any positive value of τ_1 and τ_2 we can find a pair of prospects such that $\tau_1 = t_b - t_{b-1}$ and $\tau_2 = t_{b+1} - t_b$, so it must be the case that

$$D(\tau_1 + \tau_2)\ (>) \geq\ D(\tau_1)D(\tau_2)\ \forall\ \tau_1, \tau_2 > 0.$$

In the other direction, if $D(\tau_1 + \tau_2)\ (>) \geq\ D(\tau_1)D(\tau_2)\ \forall\ \tau_1, \tau_2 > 0$, then for any such pair of prospects, we have B with some elements discounted more heavily, so that $A\ (\succ) \succeq\ B$.

Now turning attention to the transitive part of the definition of more grouped information. Note that for any two prospects, A and Z for which A has more grouped information than Z, it must be the case that there is a finite sequence of pairs of prospects, whose magnitudes of their informational update times differ by only 1, each with more grouped information than the last. So, A is more grouped than B, which is more grouped than C, which is more grouped than D..., which is more grouped than Z, with $\|T_B\| - \|T_A\| = \|T_C\| - \|T_B\| = \|T_D\| - \|T_C\| = ... = 1$.

Then, if $D(\tau_1 + \tau_2)\ (>) \geq\ D(\tau_1)D(\tau_2)\ \forall\ \tau_1, \tau_2 > 0$, by the reasoning above, each prospect in the sequence will have a (strictly) lower value than the one before it, so that prospect A will have a (strictly) higher value than prospect Z.

If A has a (strictly) higher present value than prospect Z whenever A has more grouped information than Z, then just consider the cases where $\|T_Z\| - \|T_A\| = 1$. Then by the reasoning above, $D(\tau_1 + \tau_2)\ (>) \geq\ D(\tau_1)D(\tau_2)\ \forall\ \tau_1, \tau_2 > 0$.

Lemma 1, with the observation that $D(\tau_1 + \tau_2)\ (>) \geq\ D(\tau_1)D(\tau_2)\ \forall\ \tau_1, \tau_2 > 0$ if and only if the agent exhibits (strictly) diminishing impatience completes the proof.□

Proof of Proposition 2: Due to Lemma 1, it will suffice to show that an information focused agent exhibits (S)PLDI if and only if $D(\tau_2)D(\tau_1 + \tau_3)\ (<) \leq\ D(\tau_1)D(\tau_2 + \tau_3)\ \forall\ \tau_2 > \tau_1 \geq 0,\ \tau_3 > 0$.

Consider $A,\ B \in \mathcal{P}$ with $\|I_A \setminus I_B\| = \|I_B \setminus I_A\| = 1$ that satisfy the four direct conditions for less dispersed information. Denote by n' the shared index of the unique elements $T_A \setminus T_B$ and $T_B \setminus T_A$. Note the times *not* shared by the two prospects ($t^A_{n'}$ and $t^B_{n'}$) cannot support outcomes, because if they did, $\triangle^0_A \neq \triangle^0_B$, in violation of the third condition of less dispersed information.

For any time index but n', I will drop superscript to refer to, for example, $t_{n'-1} = t^A_{n'-1} = t^B_{n'-1}$, since that time is shared by both T_A and T_B. Now, look at two groups of outcomes.

First, outcomes occurring at or before $t_{n'-1}$; this group may be empty. If not, then all such outcomes are discounted by the same sequence of discounting multipliers in both prospects, as both T_A and T_B have the same series of elements at and before $t_{n'-1}$. Thus, such outcomes give the same present value in the calculation of the value of the prospects, and can be ignored in a comparison of which prospect has the higher value.

Second, outcomes occurring at or after $t_{n'+1}$. This set is non-empty, because $t_{n'+1} \leq t^A_F$. These outcomes will be discounted differently by the two prospects.

In prospect A, they will be discounted by

$$... * D(t^A_{n'} - t_{n'-1}) * D(t_{n'+1} - t^A_{n'}) * ....,$$

in prospect B, by

$$... * D(t^B_{n'} - t_{n'-1}) * D(t_{n'+1} - t^B_{n'}) * ....$$

These outcomes will be discounted more heavily by prospect B if any only if

$$D(t^B_{n'} - t_{n'-1}) * D(t_{n'+1} - t^B_{n'}) \leq D(t^A_{n'} - t_{n'-1}) * D(t_{n'+1} - t^A_{n'}).$$

From the definition of less dispersed information,

$$min\{t^A_{n'} - t_{n'-1},\ t_{n'+1} - t^A_{n'}\} < min\{t^B_{n'} - t_{n'-1},\ t_{n'+1} - t^B_{n'}\}$$

Let $\tau_1 = min\{t^A_{n'} - t_{n'-1},\ t_{n'+1} - t^A_{n'}\}$, $\tau_2 = min\{t^B_{n'} - t_{n'-1},\ t_{n'+1} - t^B_{n'}\}$, and $\tau_3 = t_{n'+1} - t_{n'-1} - \tau_1 - \tau_2$. Then,

$$D(t^B_{n'} - t_{n'-1}) * D(t_{n'+1} - t^B_{n'}) \leq D(t^A_{n'} - t_{n'-1}) * D(t_{n'+1} - t^A_{n'}) \iff$$

$$D(\tau_2) * D(\tau_1 + \tau_3) \leq D(\tau_1) * D(\tau_2 + \tau_3).$$

By construction, $\tau_2 > \tau_1$, and $\tau_3 > 0$. To see why the latter is true, note that each of τ_1 and τ_2, being minimums of a split between $t_{n'+1}$ and $t_{n'-1}$, can be at most half of that distance, and τ_1 must be less than τ_2, so strictly less than half. Thus, their sum, $\tau_1 + \tau_2 < t_{n'+1} - t_{n'-1}$, making τ_3 strictly positive. Thus, B's outcomes, and thus its present value, will be discounted more heavily for all such A and B if and only if

$$D(\tau_2)D(\tau_1 + \tau_3) \leq D(\tau_1)D(\tau_2 + \tau_3)\ \forall\ \tau_2 > \tau_1 \geq 0,\ \tau_3 > 0,$$

with $V_B < V_A$ if and only if the inequality is strict.

Turning attention to all possible cases where A has less dispersed information than Z, note that for any two such prospects, it must be the case that there is a countable sequence of pairs of prospects starting with A and ending with Z, each pair only disjoint by a single element of I, and each with less dispersed information than the last. So, A has less dispersed information than B, which has less dispersed information than C, which has less dispersed information than D..., which has less dispersed information than Z.

Then, if $D(\tau_2)D(\tau_1+\tau_3)\ (<)\leq\ D(\tau_1)D(\tau_2+\tau_3)\ \forall\ \tau_2>\tau_1\geq 0,\ \tau_3>0$, by the reasoning above, each prospect in the sequence will have a (strictly) lower value than one before it, so that prospect A will have a (strictly) higher present value than prospect Z.

If A has a (strictly) higher present value than prospect Z whenever A has less dispersed information than Z, then only consider the cases where A and Z are disjoint by one element, and by the reasoning above, $D(\tau_2)D(\tau_1+\tau_3)\ (<)\leq\ D(\tau_1)D(\tau_2+\tau_3)\ \forall\ \tau_2>\tau_1\geq 0,\ \tau_3>0$.

Finally, since having a higher present value is the condition for preference, it is concluded that for an information focused consumer, $A \succeq Z$ whenever A has less dispersed information than Z if and only if $D(\tau_2)D(\tau_1+\tau_3)\ (<)\leq\ D(\tau_1)D(\tau_2+\tau_3)\ \forall\ \tau_2>\tau_1\geq 0,\ \tau_3>0$.

Lemma 1, with the observation that $D(\tau_2)D(\tau_1+\tau_3)\ (<)\leq\ D(\tau_1)D(\tau_2+\tau_3)\ \forall\ \tau_2>\tau_1\geq 0,\ \tau_3>0$ if and only if the agent exhibits (strictly) strongly diminishing impatience completes the proof.□

Proof of Proposition 3: Directly,

$$A \sim B \Leftrightarrow V_A = V_B \Leftrightarrow$$

$$\int_{(\alpha,t)\in\mathcal{O}} u(\alpha) * \triangle_A^0(\alpha,t) * D(r_1^A - 0) * D(r_2^A - r_1^A) .... * D(r_{\hat{n}(t)}^A - r_{\hat{n}(t)-1}^A) * D\left(t - r_{\hat{n}(t)}^A\right) =$$

$$\int_{(\alpha,t)\in\mathcal{O}} u(\alpha) * \triangle_B^0(\alpha,t) * D(r_1^B - 0) * D(r_2^B - r_1^B) .... * D(r_{\hat{n}(t)}^B - r_{\hat{n}(t)-1}^B) * D\left(t - r_{\hat{n}(t)}^B\right) \Leftrightarrow$$

$$D(\tau-0)\int_{(\alpha,t)\in\mathcal{O}} u(\alpha)*\triangle_A^0(\alpha,t)*D((r_1^A+\tau)-\tau)*...D((r_{\hat{n}(t)}^A+\tau)-(r_{\hat{n}(t)-1}^A+\tau))*D\left((t+\tau)-(r_{\hat{n}(t)}^A+\tau)\right) =$$

$$D(\tau-0)\int_{(\alpha,t)\in\mathcal{O}} u(\alpha)*\triangle_B^0(\alpha,t)*D((r_1^B+\tau)-\tau)*...D((r_{\hat{n}(t)}^B+\tau)-(r_{\hat{n}(t)-1}^B+\tau))*D\left((t+\tau)-(r_{\hat{n}(t)}^B+\tau)\right) \Leftrightarrow$$

$$D(r_1^C - 0)\int_{(\alpha,t)\in\mathcal{O}} u(\alpha) * \triangle_A^0(\alpha,t) * D(r_2^C - r_1^C) * ...D(r_{\hat{n}(t+\tau)}^C - r_{\hat{n}(t+\tau)-1}^C) * D\left((t+\tau) - r_{\hat{n}(t+\tau)}^C\right) =$$

$$D(r_1^D - 0)\int_{(\alpha,t)\in\mathcal{O}} u(\alpha) * \triangle_B^0(\alpha,t) * D(r_2^D - r_1^D) * ...D(r_{\hat{n}(t+\tau)}^D - r_{\hat{n}(t+\tau)-1}^D) * D\left((t+\tau) - r_{\hat{n}(t+\tau)}^D\right) \Leftrightarrow,$$

this step from the fact that $\mathcal{R}(C) = \mathcal{R}(A)^{+\tau} \cup \{0\}$ and $\mathcal{R}(D) = \mathcal{R}(B)^{+\tau} \cup \{0\}$,

$$\int_{(\alpha,t+\tau)\in\mathcal{O}} u(\alpha)*\triangle^0_C(\alpha,t+\tau)*D(r^C_1-0)*D(r^C_2-r^C_1)*...D(r^C_{\hat{n}(t+\tau)}-r^C_{\hat{n}(t+\tau)-1})*D\left((t+\tau)-r^C_{\hat{n}(t+\tau)}\right) =$$

$$\int_{(\alpha,t+\tau)\in\mathcal{O}} u(\alpha)*\triangle^0_B(\alpha,t+\tau)*D(r^D_1-0)*D(r^D_2-r^D_1)*...D(r^D_{\hat{n}(t+\tau)}-r^D_{\hat{n}(t+\tau)-1})*D\left((t+\tau)-r^D_{\hat{n}(t+\tau)}\right) \Leftrightarrow,$$

this from the fact that $\triangle^0_X(\alpha, t+\tau) = \triangle^0_Y(\alpha, t)$, whenever X is τ-shifted from Y,

$$V_C = V_D \Leftrightarrow C \sim D.\square$$

Part V

Bibliography

Ainslie, George, "Specious Reward: A Behavioral Theory of Impulsiveness and Impulsive Control," 1975. *Psychological Bulletin*, 82: 463-96.

Ainslie, George, "Beyond Microeconomics. Conflict Among Interests in a Multiple Self as a Determinant of Value," 1986. in Jan Elster, ed., *The Multiple Self,* Cambridge: Cambridge University Press, 133-175.

Ainslie, George, *Picoeconomics,* 1992. Cambridge: Cambridge University Press.

Ambrus, Attila and Kareen Rozen, "Rationalizing Choice with Multi-Self Models," 2013. Working Paper.

Angeletos, George Marios, David Laibson, Andrea Repetto, Jeremy Tobacman and Stephen Weinberg, "The Hyperbolic Consumption Model: Calibration, Simulation, and Empirical Evaluation," 2001. *The Journal of Economic Perspectives*, 15, 3: 47-68.

Ariely, Dan and Klaus Wertenbroch, "Procrastination, Deadlines, and Performance: Self-Control by Precommitment," 2002. *Psychological Science*, 13(3): 219-224.

Artstein-Avidan, S. and D. Dillenberger, "Dynamic Disappointment Aversion: Don't Tell Me Anything Until You Know For Sure," 2010. Working Paper.

Ashraf, Nava, Dean Karlan and Wesley Yin, "Tying Odysseus to the Mast: Evidence From a Commitment Savings Product in the Philippines," 2006. *The Quarterly Journal of Economics,* 121 (2): 635-672.

Baumeister, Roy F. and Todd F. Heatherton, "Self-Regulation Failure: An Overview," 1996. *Psychological Inquiry,* 7(1):1-15.

Baumeister, Roy F., Kathleen D. Vohs, and Dianne M. Tice, "The Strength Model of Self-Control," *Current Directions in Psychological Science*, 2007. 16(6): 351-355.

Bellemare C., M. Krause, S. Kröger, and C. Zhang, "Myopic Loss Aversion: Information Feedback vs. Investment Flexibility," 2005. *Economics Letters*, 87: 319–324.

Benabou, Roland and Marek Pycia, "Dynamic inconsistency and self-control: a planner–doer interpretation," *Economics Letters,* 2002, 77 (3): 419–424.

Bernheim, Douglas and Antonio Rangel, "Beyond Revealed Preference: Choice-Theoretic Foundations for Behavioral Welfare Economics," 2009. *The Quarterly Journal of Economics*, 124 (1): 51-104.

Bozbay, Irem, Franz Dietrich and Hans Peters, "Bargaining with endogenous disagreement: the extended Kalai-Smorodinsky solution", 2012. *Games and Economic Behavior,* 74: 407-17.

Camerer C. and T. Ho, "Violations of the Betweenness Axiom and Nonlinearity in Probability," 1994. *Journal of Risk and Uncertainty*, 8: 167-196

Chakraborty, Anujit and Yoram Halevy, "Strotz Meets Allais: Diminishing Impatience and the Certainty Effect: Reply and Corrigendum," 2015. Working Paper.

Chandran, Sucharita and Geeta Menon, "When a Day Means More than a Year: Effects of Temporal Framing on Judgments of Health Risk," 2004. *Journal of Consumer Research*, 31(2): 375-389.

Chatterjee, Kalyan and R. Vijay Krishna, "Menu Choice, Environmental Cues and Temptation: A "Dual Self" Approach to Self-Control," 2006. Working Paper.

Chatterjee, Kalyan and R. Vijay Krishna, "A "Dual Self" Representation for Stochastic Temptation," 2009. *American Economic Journal: Microeconomics*, 1(2): 148–167.

Dillenberger, David, "Preferences for One-Shot Resolution of Uncertainty and Allais-Type Behavior," 2010. *Econometrica*, 78(6): 1973-2004.

Eagleman, David, *Incognito: The Secrets Lives of the Brain*, 2011. New York City: Pantheon.

Fang, Hanming and Dan Silverman, "Time-Inconsistency and Welfare Program Participation: Evidence from the NLSY," 2009. *International Economic Review*, 50 (4): 2009

Frederick, Shane, George Loewenstein, and Ted O'Donoghue, "Time Discounting and Time Preference: A Critical Review," 2002. *Journal of Economic Literature*, 40(2): 351-401.

Fudenberg, Drew, and David Levine, "A Dual-Self Model of Impulse Control," 2006. *American Economic Review*, 96: 1449-1476.

Fudenberg, Drew, and David Levine, "Risk, Delay, and Convex Self-Control Costs," 2011. *American Economic Journal: Microeconomics* 3(3): 34–68

Fudenberg, Drew, and David Levine, "Timing and Self-Control," 2012. *Econometrica*, 80(1): 1-42.

Glimcher, Paul, Joseph Kable and Kenway Louie, "Neuroeconomic Studies of Impulsivity: Now or Just as Soon as Possible?," 2007. *American Economic Review*, 97(2): 142-147.

Gneezy, Uri and Jan Potters, "An experiment on risk taking and evaluation periods," 1997. *Quarterly Journal of Economics*, 112: 632-645.

Gourville, John T., "Pennies-a-Day: The Effect of Temporal Reframing on Transaction Evaluation," 1998. *Journal of Consumer Research,* 24 (March): 395-408.

Gul, Faruk, "A Theory of Disappointment Aversion," 1991. *Econometrica*, 59 667-686.

Gul, Faruk, and Wolfgang Pesendorfer, "Temptation and Self-Control," 2001. *Econometrica*, 69(6): 1403–35.

Gul, Faruk, and Wolfgang Pesendorfer, "Self-Control and the Theory of Consumption." 2004. *Econometrica*, 72(1): 119–158.

Haigh, Michael S. and John A. List, "Do Professional Traders Exhibit Myopic Loss Aversion? An Experimental Analysis," 2005. *Journal of Finance,* 60(1): 523-534.

Hanks, Andrew, David Just, and Brian Wansink, "Trigger Foods: The Influence of 'Irrelevant Alternatives' in School Lunchrooms," 2012. *Agricultural and Resource Economics Review*, 41:114 - 123.

Harris, Christopher, and David Laibson, "Instantaneous Gratification," 2013. *Quarterly Journal of Economics*, 128(1): 205-248.

Harsanyi, John, "Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons of Utility," 1955. *Journal of Political Economy*, 63(4): 309-321.

Harsanyi, John, "Morality and the Theory of Rational Behavior," 1977. *Social Research*, 44(4): 623-656.

Halevy, Yoram, "Strotz Meets Allais: Diminishing Impatience and the Certainty Effect," 2008. *American Economic Review*, 98(3): 1145-62.

Huang, Kevin, Zheng Liu and John Zhu, "Temptation and Self-Control: Some Evidence and Applications," 2013. Working Paper.

Jamison, Julian and Jon Wegener, "Multiple Selves in Intertemporal Choice," 2010. *Journal of Economic Psychology*, 2010. 31(5): 832-839.

Kahneman, Daniel and Amos Tversky, "Prospect Theory: An Analysis of Decision Under Risk," 1979. *Econometrica*, 47: 263-291.

Kalai, Ehud and Meir Smorodinsky, "Other Solutions to Nash's Bargaining Problem," 1975. *Econometrica*, 43 (3): 513-518.

Kaur, Supreet, Michael Kremer and Sendhil Mullainathan, “Self-Control and the Development of Work Arrangements,” 2010. *American Economic Review: Papers and Proceedings* 100(2): 624-628

Keren, Gideon and Peter Roelofsma, “Immediacy and Certainty in Intertemporal Choice,” 1995. *Organizational Behavior and Human Decision Making*, 63(3): 287-97

Kirby, Kris and R. J. Herrenstein, “Preference reversals due to myopic discounting of delayed reward,” 1995. *Psychological Science*, 6(2): 83-89.

Köszegi, Botond and Matthew Rabin, “Reference-Dependent Consumption Plans,” 2009. *American Economic Review*, 99: 909-936.

Laibson, David, “Golden Eggs and Hyperbolic Discounting,” 1997. *Quarterly Journal of Economics*, 112(2): 443–77.

Laibson, David, D, Repetto A, Tobacman J., “Self-Control and Saving for Retirement,” 1998. *Brookings Papers on Economic Activity*, 1: 91-196.

Laibson, David, “Life-cycle Consumption and Hyperbolic Discount Functions,” 1998. *European Economic Review,* 42(3-5): 861-871.

Liberman, Nira and Yaacov Trope, “The Role of Feasibility and Desirability Considerations in Near and Distant Future Decisions: A Test of Temporal Construal Theory,” 1998. *Journal of Personality and Social Psychology*, 75 (July): 5-18.

Liberman, Nira and Yaacov Trope, “Temporal Construal and Time-Dependent Changes in Preference,” 2000. *Journal of Personality and Social Psychology*, 79 (December): 876-89.

Loewenstein, George, "Frames of Mind in Intertemporal Choice," 1988. *Management Science*, 34: 200-214.

Loewenstein, George and Drazen Prelec, “Anomalies in Intertemporal Choice: Evidence and an Interpretation,” 1992. *Quarterly Journal of Economics* 107(2): 573-97.

McIntosh, Donald, *The Foundations of Human Society*, 1969. Chicago: University of Chicago Press.

McClure, Samuel M., David I. Laibson, George Loewenstein, and Jonathan D. Cohen, “Separate Neural Systems Value Immediate and Delayed Monetary Rewards,” 2004. *Science*, 306(5695): 503–07.

Mullainathan, S. and Abhijit Banerjee, “The shape of temptation: Implications for the economic lives of the poor,” 2010. *NBER Working Papers.*

Muraven, Mark, **Dianne M. Tice, and Roy F. Baumeister**, "Self-Control as a Limited Resource: Regulatory Depletion Patterns," 1998. *Journal of Personality and Social Psychology,* 74(3): 774-789.

Myerson, Joel and Leonard Green, "Discounting of delayed rewards: Models of individual choice," 1995. *Journal of the Experimental Analysis of Behavior*, 64(3): 263–276.

Ozdenoren, Emre, Stephen Salant and Dan Silverman, "Willpower and the Optimal Control of Visceral Urges," 2012. *Journal of the European Economic Association*, 10(2): 342-368

Palacios-Huerta, Ignacio, "The Aversion to the Sequential Resolution of Uncertanity," 1999. *Journal of Risk and Uncertainty*, 18: 249-269.

Samuelson, Paul, "A Note on Measurement of Utility," 1937. *The Review of Economic Studies*, 4(2): 155-161.

Segal, Uzi, "The Ellsberg Paradox and Risk Aversion: An Anticipated Utility Approach," 1987. *International Economic Review*, 28(1): 175-202.

Sen, Amartya, "Choice Functions and Revealed Preference," 1971. *Review of Economic Studies*, 38(3): 307-317.

Sen, Amartya, "On Weights and Measures: Informational Constraints in Social Welfare Analysis," 1977. *Econometrica*, 45, 1539-1572.

Strotz, Robert H., "Myopia and Inconsistency in Dynamic Utility Maximization," 1956. *Review of Economic Studies*, 23(3): 165–80.

Thaler, Richard H., "Some Empirical Evidence on Dynamic Inconsistency," 1981. *Economic Letters*, 8: 201-207.

Thaler, Richard H., and Hersh M. Shefrin, "An Economic Theory of Self-Control," 1981. *Journal of Political Economy*, 89(2): 392–406.

Thaler, Richard H., and Shlomo Benartzi, "Save More Tomorrow: Using Behavioral Economics to Increase Employee Saving," 2004. *Journal of Political Economy,* 112 (1).

Tullock, Gordon, "Efficient rent-seeking," 1980. In J.M. Buchanan, R.D. Tollison and G. Tullock (Eds.), *Toward a theory of the rent-seeking society*, 97-112. College Station: Texas A&M University Press.

Vartiainen, Hannu, "Collective choice with endogenous reference outcome," 2007. *Games and Economic Behavior*, 58: 172–180

Vohs, Kathleen D. and Ronald J. Faber, "Spent Resources: Self-Regulatory Resource Availability Affects Impulse Buying," 2007. *Journal of Consumer Research*, 4: 537-547.