5-2010

Commitment, Risk, and Consumption: Do Birds of a Feather Have Bigger Nests?

Stephen H. Shore

Todd Sinai
University of Pennsylvania

Recommended Citation

Shore, S. H., & Sinai, T. (2010). Commitment, Risk, and Consumption: Do Birds of a Feather Have Bigger Nests?. *The Review of Economics and Statistics, 92* (2), 408-424. http://dx.doi.org/10.1162/rest.2010.11380

Commitment, Risk, and Consumption: Do Birds of a Feather Have Bigger Nests?

Abstract

Consumption commitments—goods like housing for which adjustment is costly—change the relationship between risk and consumption. Commitment provides a motive to reduce consumption when possible future losses are too small to warrant adjustment but not when losses are large enough that adjustment would be worthwhile. This implies conditions under which mean-preserving increases in risk can increase housing consumption. Our empirical evidence exploits the interaction of these conditions with a novel proxy for unemployment risk: couples sharing an occupation. Consistent with our model, same-occupation couples consume more housing only when adjustment costs are high and potential losses are sufficiently large.

Disciplines

Economics | Real Estate

COMMITMENT, RISK, AND CONSUMPTION: DO BIRDS OF A FEATHER HAVE BIGGER NESTS?

Stephen H. Shore and Todd Sinai*

Abstract—Consumption commitments—goods like housing for which adjustment is costly—change the relationship between risk and consumption. Commitment provides a motive to reduce consumption when possible future losses are too small to warrant adjustment but not when losses are large enough that adjustment would be worthwhile. This implies conditions under which mean-preserving increases in risk can increase housing consumption. Our empirical evidence exploits the interaction of these conditions with a novel proxy for unemployment risk: couples sharing an occupation. Consistent with our model, same-occupation couples consume more housing only when adjustment costs are high and potential losses are sufficiently large.

I. Introduction

CONVENTIONAL economic wisdom suggests that households should respond to risk by saving more and consuming less. However, this precautionary saving intuition follows from the assumption that households can costlessly adjust their level and mix of consumption (Leland, 1968; Sandmo, 1970; Drèze & Modigliani, 1972; Kimball, 1990). This simplification is often at odds with reality, as adjusting the consumption of many goods carries some transaction cost. For example, to reduce housing consumption, home owners must incur the costs of selling a house, buying a new one, and moving. Goods with this feature are often referred to as consumption commitments. The empirical literature on precautionary saving has examined how saving, wealth, or consumption varies with risk but has paid remarkably little attention to adjustment costs.

We show that commitment introduces a motive for saving in anticipation of small losses in income that is not present for large losses; this can invert the usual negative precautionary saving relationship between risk and consumption. By way of intuition, when a household chooses the quantity of a consumption commitment such as housing, it recognizes that adjusting housing consumption (by moving) will be optimal only after a large loss. Following a small loss, the household will instead have to reduce nonhousing consumption substantially since adjusting housing consumption will not be warranted. The marginal utilities of housing and nonhousing consumption will diverge. By reducing housing consumption ex ante, a household that anticipates the possibility of such losses can mitigate the future divergence of housing and nonhousing consumption. By contrast, following a loss large enough to make adjustment optimal, moving will equate the marginal utilities of housing and nonhousing consumption. Foreseeing this rebalancing, a household that anticipates the possibility of large prospective losses lacks this motive for reducing housing consumption ex ante. As a result, an increase in risk that makes large losses more likely but small losses less likely can lead to greater housing consumption.

Using household-level microdata, we find that an increase in risk leads to more housing consumption when adjustment costs are relatively high but not when they are low. We bring two empirical innovations to bear on this question.

First, we exploit a novel source of variation to proxy for increasing risk. When couples share an occupation, the correlation of their unemployment events is higher. Couples with higher unemployment correlations face a higher probability that neither or both spouses will become unemployed but a lower probability that exactly one spouse will become unemployed. This approach, which follows much of the empirical precautionary saving literature in using occupation-based variation in unemployment risk as a proxy for income risk, enables us to control separately for each spouse's occupation and identify the risk solely from the pairing of couples.

Second, we interact our measure of risk with proxies for adjustment costs. Identification comes from comparing the housing consumption of same- and different-occupation couples when adjustment costs are high versus when they are low. This controls for the possibility that same-occupation couples may differ from other couples in dimensions besides risk and that these differences may affect housing consumption.[1] The model predicts that couples with higher unemployment correlations should consume more housing only when adjustment costs are large enough to deter moving. Therefore, we compare home owners (who

Received for publication October 18, 2006. Revision accepted for publication July 29, 2008.

* Shore: Johns Hopkins University; Sinai: Wharton School of the University of Pennsylvania and NBER.

We thank Andrew Au, Colleen Carey, Misha Dworsky, Alex Janovsky, and Daniel Simundza for outstanding research assistance. We received helpful comments from Raphael Bostic, Raj Chetty, Christopher Carroll, Tom Davidoff, Ed Glaeser, Joe Gyourko, Caroline Hoxby, Brigitte Madrian, Andrew Postlewaite, Julio Rotemberg, Albert Saiz, Nick Souleles, Joel Waldfogel, Matt White, Justin Wolfers, seminar participants at Columbia Business School, Harvard University, UC Berkeley, UC Davis, University of Illinois, University of Washington, and Wharton, and two anonymous referees. Special thanks to Joel Waldfogel for suggesting the second half of the title. S.H. gratefully acknowledges financial support from the National Institutes of Health–National Institute on Aging, grant number P30 AG12836, and the Boettner Center for Pensions and Retirement Security at the University of Pennsylvania as well as the Zell/Lurie Real Estate Center at the Wharton School. T.S. received support from the Research Scholars Program at the Zell/Lurie Real Estate Center.

[1] Lusardi (1997) has raised the concern that the individual-level, occupation-based proxies for risk used in the precautionary saving literature (Carroll & Samwick, 1997) are prone to omitted variable bias. People in high-risk occupations differ from other individuals in dimensions other than risk that affect saving and consumption. We control for this possible individual-level omitted variable bias by controlling for the occupation of each spouse directly and then exploit variation in within-household diversification. This provides an alternative to the natural experiment approach used in Fuchs-Schundeln and Schundeln (2005) to overcome omitted variable problems.

The Review of Economics and Statistics, May 2010, 92(2): 408–424

have high moving costs) to renters (who do not). We find that same-occupation home owners spend at least 2.1% more on their houses than do different-occupation couples. This result obtains even after controlling for income, each spouse's occupation, and a host of other demographic characteristics. By contrast, same-occupation renters spend no more on rent than do different-occupation renters. This pattern cannot be explained by same-occupation couples' selection into home ownership. The relationship between risk and housing consumption is also more positive when another measure of moving costs is high. When couples face effectively higher moving costs because they are unlikely to move for exogenous demographic reasons (age, education, presence of children), the difference between same- and different-occupation home owners' housing spending is greater.

The model also predicts that couples with higher unemployment correlations should consume more housing only when loss magnitudes are large enough to induce moving. We exploit cross-state variation in the nonlinearity of the unemployment insurance (UI) schedule. When an unemployed household faces a lower UI replacement rate, the potential loss to permanent income is greater, raising the odds of moving. We compare households with more and less generous unemployment insurance. We find that same-occupation home owners spend more on housing relative to different-occupation home owners only when unemployment insurance is less generous.

The remainder of this paper is arranged as follows. Section II sets up a simple model of consumption commitments that predicts that a mean-preserving increase in risk can increase a household's committed consumption in some settings but not others. We describe the data in section III and present results in section IV. Section V. concludes.

II. Model

In this section, we outline a stylized model of precautionary saving that incorporates consumption commitments. The model presented here contains standard features of precautionary saving and consumption commitment models. However, this is the first paper to focus on the implications for precautionary saving of a model with consumption adjustment costs. Chetty and Szeidl (2005) argue that commitment makes households more risk averse in the domain of small losses (too small to warrant adjusting committed consumption) than large ones (large enough to warrant adjusting committed consumption). This paper argues that this can be true for prudence as well as risk aversion. Commitment provides an incentive to reduce housing consumption (and therefore save) relatively more in anticipation of possible small losses than large ones. In a setting where adjustment is optimal only in response to large losses, an increase in risk that makes large losses more likely but small losses less likely (even one that is mean-preserving by construction) can increase committed consumption and reduce saving, thus inverting the standard precautionary saving result.

A. Setup

Following Chetty (2004), we present a model with two periods, $t = 1, 2$, and utility in each period, $u(h, f)$, a function of two goods, h (housing) and f (food).[2] The household's lifetime expected utility is the sum of expected utility from the two periods:

$$U = u(h_1, f_1) + E[u(h_2, f_2)]. \tag{1}$$

As in the precautionary saving literature, we treat household labor income risk as exogenous and endogenize consumption and saving. In the first period, the household receives income Y_1 and decides how much of each good, h_1 and f_1, to consume. The remainder, $Y_1 - h_1 - f_1$, is saved.[3] In the second period, the household receives an income $\tilde{Y}_2$ (which is not known in the first period). The household must then allocate second-period wealth, $Y_1 + \tilde{Y}_2 - h_1 - f_1$, between the two goods. If a household adjusts its consumption of good h, it pays a fixed transaction cost k.[4] It is this transaction cost that gives the h good its commitment feature. Therefore, a household's intertemporal budget constraint can be written as

$$\begin{aligned} Y_1 + \tilde{Y}_2 &= 2h_1 + f_1 + f_2 \text{ if } h_1 = h_2 \\ Y_1 + \tilde{Y}_2 &= h_1 + k + h_2 + f_1 + f_2 \text{ if } h_1 \neq h_2. \end{aligned} \tag{2}$$

To determine the optimal consumption in the first period, we determine optimal consumption and indirect utility in the second period and then work backward. In the second period, the household maximizes $u(h_2, f_2)$ subject to constraint (2), taking h_1 and f_1 as given. Assuming that u is

[2] Prices are normalized to 1, and there is no goods price risk. u is assumed to be symmetric, differentiable everywhere, strictly increasing, and strictly concave. Relaxing the symmetry assumption does not affect the substance of the results. We require that $u_{hh} < u_{hf}$, so that increasing the quantity of housing consumed reduces the marginal utility of housing more than the marginal utility of food.

[3] A standard user cost model of home ownership transforms the asset price of the house into the flow cost of the consumption of housing services, h (Hendershott & Slemrod, 1983; Poterba, 1984). The consensus of the literature is that the demand for housing is determined by the consumption motive rather than investment (Henderson & Ioannides, 1983; Goetzmann, 1993; Brueckner, 1997; Flavin & Yamashita, 2002). Consequently, the dual nature of housing as an asset and consumption good does not preclude its use as an indicator of consumption. Showing the conditions under which households might save precautionarily in a housing asset is beyond the scope of this model. We will come back to this issue later to make sure a savings motive is not driving our empirical results. In addition, because we abstract from the investment problems examined in other work, we make the simplifying assumptions that there are no risky assets, the riskless interest rate is 0, and the household cannot save in the housing asset.

[4] Fixed transaction costs provide greater analytic tractability than the case of proportional transaction costs ($k \propto h_1$) with no effect on the qualitative predictions of the model.

symmetric—$u(x, y) = u(y, x)$—for simplicity, optimal consumption is

$$\begin{aligned} h_2 = f_2 &= \frac{1}{2}(Y_1 + \tilde{Y}_2 - h_1 - f_1 - k) \text{ if } h_2 \neq h_1 \\ h_2 = h_1; f_2 &= Y_1 + \tilde{Y}_2 - 2h_1 - f_1 \text{ if } h_2 = h_1, \end{aligned} \tag{3}$$

and second-period indirect utility, $v(Y_1 + \tilde{Y}_2 - h_1 - f_1, h_1)$, a function of second-period wealth and first-period housing, is

$$v(Y_1 + \tilde{Y}_2 - h_1 - f_1, h_1) = \max \begin{bmatrix} u(h_1, Y_1 + \tilde{Y}_2 - 2h_1 - f_1), \\ u(\frac{1}{2}(Y_1 + \tilde{Y}_2 - h_1 - f_1 - k), \\ \frac{1}{2}(Y_1 + \tilde{Y}_2 - h_1 - f_1 - k)) \end{bmatrix}. \tag{4}$$

Adjusting housing consumption is optimal if and only if the indirect utility from moving (the second term) is greater than the indirect utility from not moving (the first term). The household will not adjust housing consumption, h, unless the shock to income is large enough that the benefit of rebalancing consumption exceeds the cost of moving.

Once we have solved for the optimal consumption rule in the second period, we can solve for optimal consumption in the first period. The household's lifetime utility function (1) can be rewritten as

$$U(h_1, f_1) = u(h_1, f_1) + E \begin{bmatrix} \max(u(h_1, Y_1 + \tilde{Y}_2 - 2h_1 - f_1), \\ u(\frac{1}{2}(Y_1 + \tilde{Y}_2 - h_1 - f_1 - k), \\ \frac{1}{2}(Y_1 + \tilde{Y}_2 - h_1 - f_1 - k)) \end{bmatrix}. \tag{5}$$

In equation (5), $u(h_1, f_1)$ is the utility of first-period consumption. The expectation term is the expected utility of second-period consumption. Note that second-period utility is the maximum of the utility from not moving (the first term in the max operator) and the utility from moving (the second term).

To better understand optimal first-period consumption $\{h_1^*, f_1^*\}$, we must add structure by making assumptions about the distribution of $\tilde{Y}_2$. We assume that the household has two wage earners, a husband and wife, and that uncertainty comes from the possibility that one or both may receive a negative wage shock, which we refer to as becoming unemployed. Income for either husband or wife is Y_2 if employed and $Y_2 - L$ if unemployed. The husband's probability of unemployment is p, while the wife's is q. There is a correlation ρ between the employment status of the husband and wife.[5] Therefore, the distribution of household income in the second period, $\tilde{Y}_2$ can be written as:

[5] Since unemployment is binomial, the range of possible values of ρ lies in the interval

$$\rho \in \left[\frac{\max(p + q - 1,0) - pq}{\sqrt{pq(1-p)(1-q)}}, \frac{\min(p,q) - pq}{\sqrt{pq(1-p)(1-q)}} \right].$$

$$\begin{aligned} &\text{Both spouses employed:} \\ &\quad \tilde{Y}_2 = 2Y_2 \text{ with probability } 1 - p - q + \phi \\ &\text{Exactly one spouse unemployed:} \\ &\quad \tilde{Y}_2 = 2Y_2 - L \text{ with probability } p + q - 2\phi \\ &\text{Both spouses unemployed:} \\ &\quad \tilde{Y}_2 = 2Y_2 - 2L \text{ with probability } \phi, \\ &\quad \text{where } \phi \equiv pq + \rho\sqrt{pq(1-p)(1-q)}. \end{aligned} \tag{6}$$

Increasing the correlation of the couple's unemployment events ρ (or, equivalently, ϕ while holding p and q fixed) is equivalent to adding a mean-preserving spread in the distribution of household labor income, increasing the probability of the best and worst outcomes (neither or both unemployed) while decreasing the probability of the medium outcome (exactly one unemployed). However, expected household income, $E[\tilde{Y}_2] = 2Y_2 - (p + q)L$, is independent of ρ or ϕ.

In this setting, there are three possible states (none, one, or both unemployed) with two possible adjustment actions for each (move or do not move) so there are $2^3 = 8$ possible patterns of adjustment (for example, move if and only if one or both spouses become unemployed).

B. *Optimal Consumption*

In general, it is not possible to solve for the optimal h_1 and f_1 analytically. Instead, we consider the special case in which unemployment probabilities p and q become arbitrarily small. We also include numerical results with more realistic unemployment probabilities (10%) and proportional adjustment costs for the empirically plausible case of separable log utility. These numerical results are similar to the analytic results. In the appendix, we also present analytic results for quadratic utility without assuming that unemployment probabilities go to 0. Again, the results are similar. In all cases, we make the empirically realistic assumption that unemployment rates (p and q) are low in the sense that dual employment is the most common state. This implies that the distribution of income is negatively skewed ($E[(\tilde{Y}_2 - E[\tilde{Y}_2])^3] < 0$) so that the only relevant housing adjustments are those that reduce housing consumption.[6]

Lemma 1. *Assume utility $u(h, f)$ satisfies $u(x, y) = u(y, x)$, u_h, $u_f > 0$, u_{hh}, $u_{ff} < 0$, and u_{hh}, $u_{ff} < u_{hf}$ and*

[6] If we considered lottery risk (where the income distribution is positively skewed) instead of income risk (where it is negatively skewed), one possible increase in risk would increase the probability of winning a lottery large enough to induce moving to a larger house, while decreasing the probability of winning a lottery too small to induce a move. The model would then predict even greater reductions in housing consumption than would be predicted by a model of precautionary saving without consumption commitments. In that case, individuals would adjust consumption in good states but not bad ones.

u_{hhh}, u_{fff} are finite for all $h, f \in [\frac{1}{4}(Y_1 + 2Y_2) - 2L, \frac{1}{4}(Y_1 + 2Y_2)]$. Also assume Y_1, Y_2, and L are strictly positive. In the limit as p and q go to 0, any value of $k \in [0, \infty]$ lies in one of four nonempty, contiguous ranges in which adjustment is (a) always optimal, (b) optimal if and only if at least one spouse becomes unemployed, (c) optimal if and only if both spouses become unemployed, and (d) never optimal, where these ranges are listed in ascending order of k:

(*a*) *When k is such that adjustment is always optimal, then $\lim_{p,q\to 0} \frac{dh_1^*}{d\phi} \gtreqless 0$ if and only if $u_{hhh} \gtreqless 0$. As p, $q \to 0$, this pattern of adjustment is only optimal for $k = 0$.*

(*b*) *When k is such that adjustment is optimal if and only if at least one spouse becomes unemployed; then $\lim_{p,q\to 0} \frac{dh_1^*}{d\phi} = 0$.*

(*c*) *When k is such that adjustment is optimal if and only if both spouses become unemployed; then $\lim_{p,q\to 0} \frac{dh_1^*}{d\phi} > 0$.*

(*d*) *When k is such that adjustment is never optimal; then $\lim_{p,q\to 0} \frac{dh_1^*}{d\phi} \gtreqless 0$ if and only if $u_{hhh} \gtreqless 0$.*

Proof. See appendix A for the proof and for closed-form expressions for $\lim_{p,q\to 0} \frac{dh_1^*}{d\phi}$, $\lim_{p,q\to 0} \frac{df_1^*}{d\phi}$ and ranges of k.

Cases a and d: Adjustment is always optimal or never optimal. When adjustment costs are very low (case a) or very high (case d), the adjustment decision will be the same in all states. Increasing risk (here, ϕ, the covariance of spouses' unemployment spells) increases or decreases consumption depending on the third derivative of the utility function. When utility is quadratic ($u_{hhh} = 0$), increasing risk has no impact on consumption. These results are identical to those commonly found in the precautionary saving literature.

In the case of prudence ($u_{hhh} > 0$) considered in standard precautionary saving models, increasing risk reduces consumption. Marginal utility goes up more than twice as much when both husband and wife become unemployed than when only one becomes unemployed. Therefore, increasing risk (increasing ϕ) raises expected marginal utility and reduces the optimal level of consumption.

This is illustrated by figure 1, which presents numerical results for the case of log utility, 10% expected unemployment rates, and no moving costs.[7] (The results are qualitatively the same when moving costs are so high that adjustment is never optimal.) The figure shows the marginal lifetime utility of first-period housing consumption, $dU(h_1, f_1^*(h_1))/dh_1$ (y-axis), as a function of initial housing consumption, h_1 (x-axis) in different unemployment states. These lines represent the first-order conditions for first-period housing in various states if the second-period realization were known. In other words, how does a marginal increase in first-period housing consumption affect lifetime utility if both spouses (or one or none) are employed in the second period? If both the husband and wife are unemployed, represented by the △ plot, then the marginal utility of first-period housing consumption is strongly negative; the family could have increased lifetime utility had it bought a smaller house initially. By contrast, if both spouses are employed, the □ plot, then the marginal utility of first-period housing is positive; the family could have increased lifetime utility had it bought a bigger house initially. The *o* plot, representing the marginal utility when exactly one spouse becomes unemployed, is in between.

FIGURE 1.—IMPACT OF INCREASING H_1 ON LIFETIME EXPECTED UTILITY, NO MOVING COSTS

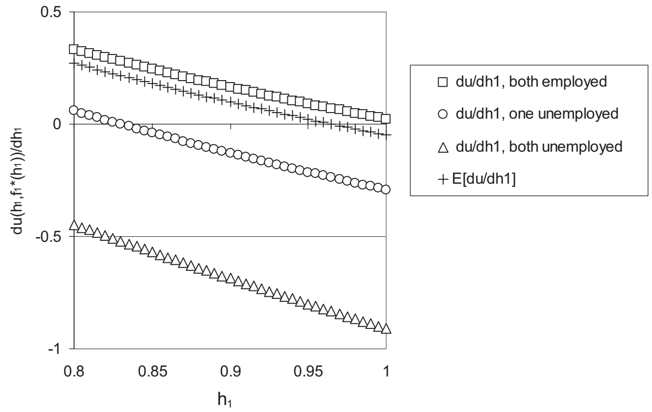

Note: This figure plots the marginal lifetime utility of first-period housing consumption against first-period housing consumption. There are no moving costs, so $k = 0$. First-period income, $Y_1 = 2$; second-period income for a given spouse is either $Y_2 = 1$ with probability $1 - p = 1 - q = 0.9$ or $Y_z - L = 0.5$ with probability $p = q = 0.1$. As a result, total household second-period income is 2, 1.5, or 1. The correlation of the household's unemployment shocks is $\rho = 0.2$. Lifetime utility is given as the sum of log food and log housing consumption in periods 1 and 2.

Plot + in this figure is merely an average of the △, *o*, and □ plots, weighted by the respective probabilities of the three employment outcomes. Since the first-order condition for h_1 is

$$E[dU(h_1, f_1^*(h_1))/dh_1] = 0, \tag{7}$$

the optimal level of consumption is simply the point where the expected marginal utility plot, +, crosses the y-axis. A mean-preserving spread increases the weight on both the unemployed and employed states (△ and □ plots) by reducing the weight on the one unemployed state (*o* plot). Because $u_{hhh} > 0$, and △ plot (both unemployed) is substantially lower than the *o* plot (one unemployed) and the □ plot (both employed). Therefore, a mean preserving spread will move the expected marginal utility (the + plot) down and reduce the optimal level of initial housing consumption. This is a graphical representation of precautionary saving.

[7] These plots assume that the level of food consumption in the first period is chosen optimally given first-period housing consumption and an unknown second-period employment realization. We use the following parameters: $Y_1 = 2$, $Y_2 = 1$, $L = 0.5$, $p = q = 0.1$, $\rho = 0.2$, $k = 0$.

FIGURE 2.—IMPACT OF INCREASING H_1 ON LIFETIME EXPECTED UTILITY, WITH 10% MOVING COST

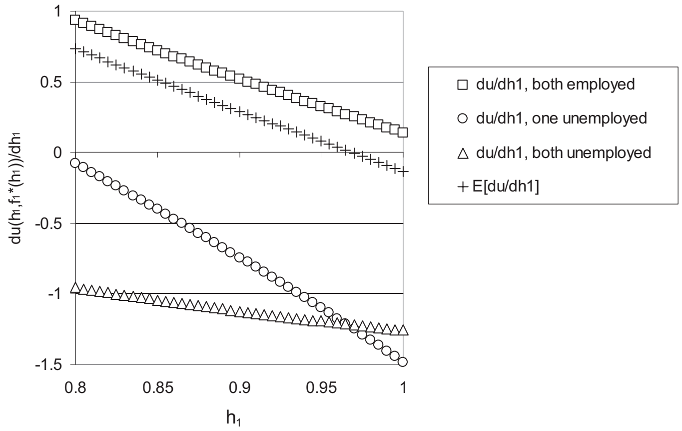

Note: This figure plots the marginal lifetime utility of first-period housing consumption against first-period housing consumption, h_1. The cost of adjusting housing consumption is 10% of h_1. First-period income, $Y_1 = 2$; second-period income for a given spouse is either $Y_2 = 1$ with probability $1 - p = 1 - q = 0.9$ or $Y_z - L = 0.5$ with probability $p = q = 0.1$. The correlation of the household's unemployment shocks is $\rho = 0.2$. As a result, total household second-period income is 2, 1.5, or 1. Lifetime utility is given as the sum of log food and log housing consumption in periods 1 and 2. Given these parameters, it is optimal to adjust housing consumption in the second period only if both spouses become unemployed within the range of values for h_1 shown.

Case c: Adjustment is optimal only in the worst state. When adjustment costs k are high enough to deter moving in the face of small losses (one unemployed) and low enough to make moving optimal in the face of large losses (both unemployed), a precautionary dissaving motive is introduced. Increasing risk now leads to more housing consumption, inverting the standard precautionary saving result.

This can be seen graphically in figure 2, which is identical to figure 1 except for a 10% proportional cost of adjusting housing consumption. Relative to the size of the loss, this adjustment cost is low enough to make moving optimal if both spouses become unemployed but high enough to deter moving when only one spouse becomes unemployed.

The noteworthy feature of this figure is the extremely low marginal utility of first-period housing consumption when exactly one spouse becomes unemployed—the *o* line. When exactly one spouse becomes unemployed, the household is stuck in a house that is more expensive than it would prefer. There is an imbalance between the high level (and low marginal utility) of housing consumption and the low level (and high marginal utility) of food consumption. Reducing first-period housing consumption by \$1 increases lifetime utility in this event because it allows the household to increase second-period food consumption, which has a high marginal utility, by \$2. The marginal utility of first-period housing, the *o* line, is highly negative for households with exactly one unemployed spouse. These households have to reduce food consumption dramatically in the second period because it is too costly to adjust housing consumption. They would have been much better off had initial housing consumption been lower. This provides a strong motive to reduce housing consumption in advance of possible single unemployment.

When both spouses remain employed, the household chooses not to move. Reducing first-period housing consumption to increase food consumption in the second period (which has a low marginal utility) would decrease lifetime utility. When both spouses become unemployed, the household chooses to pay a moving cost to rebalance housing and food consumption; because consumption is low, the marginal utility of wealth (which will be spent on food and housing equally) will be high, but lower (or not much higher) than the marginal utility in the one-unemployed state. Had initial housing consumption been \$1 lower, both housing and food consumption would be \$0.50 higher in this state.[8]

A mean-preserving increase in risk (increasing ϕ) raises the weight on both the unemployed and employed states ($\triangle$ and $\square$ plots) by reducing the weight on the one unemployed state (*o* plot). Since the *o* plot (one unemployed) is substantially below the weighted average of the $\triangle$ plot (both unemployed) and the $\square$ plot (both employed), a mean-preserving increase in risk will move the expected marginal utility (the + plot) up and therefore raise the optimal level of initial housing consumption.

This setup implies a substantial positive relationship between income correlation and housing consumption. In the numerical example given in figure 2, increasing the correlation of unemployment from no correlation to perfect correlation increases optimal spending on housing by 2.9% (and decreases optimal nonhousing consumption by 1.0%). The saving rate falls from 3.8% to 2.9% when the correlation of income increases. This effect is similar in size to—but the opposite sign from—what would be predicted by a standard model of precautionary saving without moving costs. Without moving costs, the same increase in income correlation leads to a 1.2% reduction in both housing and food consumption and an increase in the saving rate from 3.3% to 4.4%.

Lemma 1 proves that increasing risk (formally, increasing ρ or ϕ holding p and q fixed) will raise housing and lower food consumption. The net effect on aggregate consumption (housing plus food) depends on the curvature of the utility function (u_{hhh}), though it is strongly positive even for quadratic utility (lemma 2, $u_{hhh} = 0$) and log utility (figure 2, $u_{hhh} > 0$). Unlike a setting without adjustment costs, it is no longer sufficient (though still necessary) that $u_{hhh} > 0$ for increasing risk to lead to greater saving. In addition, in this case, increasing risk makes the household more likely to rent because the compensating differential needed for the household to accept the risk of paying the adjustment cost k (as opposed to facing no adjustment cost) is increasing with ϕ.

[8] The increase is actually 50 cents plus half the nominal reduction in moving costs. In this numerical example, moving costs are proportional to initial housing consumption.

Case b: Adjustment is optimal in all but the best state. In this case, increasing ϕ has no effect on housing consumption. It is optimal to make all adjustments in f_1 and not h_1, and the household responds to increased risk by adjusting food consumption. Since the household moves in all but the best state, its goal is to minimize the wedge between housing and food consumption in the event of the best outcome. Adjusting h_1 commits the household to a larger gap between h_2 and f_2 in the good state than adjusting f_1 would. By contrast, in the other states, utility is the same whether adjustment takes place in h_1 or f_1 because the household readjusts on moving.

Depending on the curvature of the utility function, $df_1^*/d\phi$ could be of either sign. However, there exists a $k >$ 0 such that $df_1^*/d\phi < 0$ for any utility function with $u_{hhh} >$ 0. Absent wealth effects that come from the impact of paying an adjustment cost on marginal utility, increasing risk leads to reduced food consumption and increased saving.

III. Data and Variable Construction

To estimate our model, we use two microdata sets. For data on changes in employment status, occupation and the probability of moving, we use the April 1996 panel of the Survey of Income and Program Participation (SIPP), which follows a panel of households for 48 months between April 1996 and March 2000. When we examine the effect of sharing an occupation on housing consumption and home ownership, we use a pooled cross-section of households from the 1980, 1990, and 2000 Integrated Public Use Microdata Series (IPUMS) of the U.S. Census. These data are a 1% random sample of responses to the U.S. Decennial Census and contain self-reported house values, incomes, and occupations, as well as employment status, a limited moving history, and a number of demographic variables and geographic identifiers.[9]

The SIPP initially contains 3,897,211 person $\times$ month observations, and the three waves of the IPUMS together initially contain 2,778,194 household-level observations. We impose several restrictions on our samples, which, taken together, reduce the number of usable observations to 307,154 household $\times$ month observations for the SIPP and 290,062 household observations for the IPUMS. These restrictions are detailed in table A1. In both the SIPP and IPUMS, we limit our attention to married couples in which both spouses are currently employed. In the IPUMS, we also impose the restrictions that both spouses work full time and live in a Metropolitan Statistical Area (MSA).[10] We restrict our attention to MSAs so we can control for local housing costs.

TABLE 1.—SUMMARY STATISTICS

A: IPUMS Summary Statistics

Variable	Owners Only Mean	Owners Only S.D.	Renters Only Mean	Renters Only S.D.
Husband and wife report same occupation (1950 definitions)	0.096	0.294	0.096	0.295
Husband and wife report same industry (1950 definitions)	0.141	0.348	0.127	0.333
House value; monthly rent	\$175,893	\$129,027	\$666	\$332
Family income	\$91,252	\$59,064	\$61,777	\$39,942
Husband's imputed unemployment rate (p)	0.065	0.022	0.071	0.025
Wife's imputed unemployment rate (q)	0.135	0.038	0.147	0.041
Husband's share of income	0.621	0.170	0.598	0.181
Imputed probability of moving	0.148	0.083	0.196	0.093
Sample average probability of moving	0.112	0.315	0.338	0.473
Number of observations	231,598		58,464	

B: SIPP Summary statistics

Variable	Mean	S.D.
Husband and wife report same occupation	0.031	0.172
Husband and wife report same industry	0.091	0.288
Family income	\$72,060	\$57,647
Husband's unemployment rate	0.070	0.256
Wife's unemployment rate	0.122	0.327

Note: Panel A: Data are from the 1980, 1990, and 2000 IPUMS. Sample construction is detailed in table A1. Dollar amounts are in real (2000) dollars. The number of observations for the "same industry" row is 240,680 for owners, and 59,987 for renters. The sample size differs because a larger fraction of the IPUMS sample reports their industry than do their occupation.

Panel B: Data are from the April 1996 panel of the Survey of Income and Program Participation, which covers 48 months between April 1996 and March 2000. Sample construction is detailed in table A1.

We make extensive use of occupation data in both data sets. The IPUMS reports one occupation variable with 227 categories that is consistently defined over all three waves, based on occupation definitions from 1950. As detailed in table 1A, the average rate of same-occupation couples across all occupations in the sample is 9.6%. In the SIPP data summarized in table 1B, the prevalence of same-occupation couples is somewhat lower, at 3.1 percent, since occupation definitions in the SIPP are more granular, with 463 three-digit codes. Table 2 lists the 20 occupations in the IPUMS with the highest fraction of same-occupation couples. The fraction of same-occupation couples varies widely by occupation: It ranges from 15% for physicians to 0 for many occupations. The third column of this table shows the distribution of occupations among same-occupation couples. These couples are drawn from a large number of common occupations such as managers (28% of same-occupation couples), teachers (13%), and clerical workers (9%). High-socioeconomic-status professions such as physicians, lawyers, and professors are found in less than 5% of these couples.

Consistent with the framework developed in section II, our proxy for income risk will be unemployment. In the SIPP, we define a person as employed when he or she has a

[9] Since house value and income are recorded as ranges, we assign the midpoint of the range, or 1.5 times the top code. All dollar values are converted to real (2000) dollars using the CPI.

[10] In the IPUMS, we discard households containing part-time or unemployed spouses because it is difficult to accurately measure their occupation or potential earnings capacity. In particular, one spouse might keep his or her labor supply in reserve as a buffer in case the other spouse becomes unemployed (Cullen & Gruber, 2000).

TABLE 2.—TOP 20 OCCUPATIONS BY PERCENTAGE OF COUPLES WHO SHARE THE SAME OCCUPATION

	Occupation, 1950 Basis	(1) Same Occupation Share of the Occupation	Occupation Share of Sample	Percentage of Same Occupation with Occupation
1	Physicians and surgeons	15.49%	0.52%	1.62%
2	Teachers	12.30	5.45	13.33
3	Operative and kindred workers	12.14	6.53	15.77
4	Managers, officials, and proprietors	11.55	11.97	27.52
5	Lawyers and judges	10.65	0.85	1.80
6	Professors (subject matter unspecified)	7.85	0.60	0.94
7	Professional, technical, and kindred workers	7.58	3.33	5.01
8	Managers and superintendents, building	7.40	0.38	0.56
9	Real estate agents and brokers	6.81	0.86	1.17
10	Members of the armed services	6.08	0.65	0.79
11	Salesmen and sales clerks	5.61	4.42	4.94
12	Clerical and kindred workers	5.11	8.98	9.12
13	Janitors and sextons	4.85	1.47	1.42
14	Editors and reporters	4.61	0.41	0.38
15	Cooks, except private household	4.46	0.99	0.88
16	Policemen and detectives	3.60	0.79	0.57
17	Insurance agents and brokers	3.39	1.15	0.77
18	Mail carriers	3.34	0.34	0.23
19	Service workers, except private household	3.13	0.58	0.36
20	Social and welfare workers, except group	3.05	0.68	0.41

Note: Only occupations comprising at least 0.25% of the sample are shown in this table. Column 1 presents the ratio of the number of same-occupation couples in an occupation to the number of couples where either (or both) spouse has that occupation. Column 2 is the ratio of the number of couples where either (or both) spouse has that occupation to the total number of couples. Column 3 is the fraction of same occupation couples with this occupation. Data are from the 1980–2000 IPUMS.

job all month or a job part of the month but spends no time as laid off or searching for a job; we define as unemployed a person who spends all month unemployed or has a job only part of the month and spends some of the month as laid off or searching for a job. In the IPUMS, we define a person as unemployed when his or her stated usual hours of work in that year are zero. Although we restrict our sample to dual-employed couples, we control for the probability of unemployment. We impute that probability for a husband (wife) as the average rate of unemployment for husbands (wives) in the same occupation and year, excluding the husband's (wife's) own observation, imposing the sample restrictions described in table A1 (except for the full-time worker restriction).[11] Table 1 shows that the annual unemployment rate for home-owning husbands (p) in the IPUMS averages 6.5% and the unemployment rate for wives (q) averages 13.5%. The unemployment rates in the SIPP, 7.0% and 12.2%, respectively, are similar.

IV. Empirical Evidence

In this section, we test the empirical implications of the theory developed in section II. Do couples with more highly correlated unemployment events spend more on housing? The model predicts that this will be the case when adjustment costs are high enough to deter moving in all but the worst states. Otherwise the standard precautionary saving result will be obtained: increasing the correlation of couples' unemployment events will reduce housing consumption.

We implement this empirically by regressing log housing spending, $\ln(H_{i,t})$, on our indicator of increased risk (1_ρ), proxies for low moving costs (k) or small potential losses from unemployment (L) that we will denote by $X_{i,t}$, and the interaction of the two ($1_{\rho,i,t} \times X_{i,t}$), for household i in year t in the IPUMS data:

$$\begin{aligned}\ln(H_{i,t}) &= \alpha 1_{\rho,i,t} + \beta(1_{\rho,i,t} \times X_{i,t}) + \gamma_X X_{i,t} + \gamma_Z Z_{i,t} \\ &\quad + \delta_{m,t} + \gamma_{XZ}(X_{i,t} \times (Z_{i,t} + \delta_{m,t})) + \gamma_\Omega \Omega_{i,t} + \varepsilon_{i,t}.\end{aligned} \tag{8}$$

We include controls for other demographic characteristics ($Z_{i,t}$) and MSA $\times$ year ($\delta_{m,t}$), $Z_{i,t}$ and $\delta_{m,t}$ interacted with $X_{i,t}$, and also the occupations of the husband and wife ($\Omega_{i,t}$).

Standard tests of precautionary saving omit $X_{i,t}$ to measure $\alpha + \beta \times \bar{X}_{i,t}$, the relationship between risk and housing consumption (or more commonly, wealth or saving) for the average $\bar{X}_{i,t}$ in the sample.[12] A model of precautionary saving without adjustment predicts $\alpha + \beta + X_{i,t} < 0$;

[11] Since people who are unemployed may state that they have no occupation—even when they have worked and plan to work in a given occupation—this procedure likely underestimates the true unemployment rate by occupation.

[12] These tests (including Skinner, 1988; Dynan, 1993; Carroll and Samwick, 1997) regress individual-level occupation-based proxies for income risk ($1_{occ,i,t}$, in contrast to our use of the household-level same-occupation proxy for risk, $1_{\rho,i,t}$) on various measures of household wealth (in contrast to our use of housing consumption). Other papers take a similar approach in the housing context. Those empirical studies find a negative relationship between income risk and home ownership (Diaz-Serrano, 2005; Haurin, 1991; Robst, Deitz, & McGoldrick, 1999). Prior evidence on housing spending is ambiguous. Haurin and Gill (1987) find that military husbands' incomes positively affect their housing spending, but their wives' incomes (which they argue are more uncertain) do not. Haurin (1991) examines the effect of income risk on house spending and fails to find a statistically significant effect.

increased risk leads to reduced housing consumption for any $X_{i,t}$. By contrast, our model predicts that $\alpha + \beta \times X_{i,t} < 0$ for high values of $X_{i,t}$ (for example, when moving costs are low) but $\alpha + \beta \times X_{i,t} > 0$ for low values of $X_{i,t}$. Equation (8) identifies α and β—and therefore $\alpha + \beta \times X_{i,t}$, providing a test of predictions from our model and from standard precautionary saving models—so long as high- and low-risk couples do not differ in their unobserved taste for housing ($E[\varepsilon_{i,t} \times 1_{\rho,i,t}] = 0$). We argue that this assumption is not only valid, but also much more likely to be valid than the analogous assumption required for most precautionary saving tests.[13]

Since our model predicts that $\alpha + \beta \times X_{i,t} < 0$ for high values of $X_{i,t}$ but $\alpha + \beta \times X_{i,t} > 0$ for low values of $X_{i,t}$, it also predicts that $\beta < 0$. For estimates of β to be unbiased, we need only assume that any difference in the unobserved taste for housing between low- and high-risk households is uncorrelated with $X_{i,t} : E[\varepsilon_{i,t} \times 1_{\rho,i,t} \times X_{i,t}] = 0$. We can identify β correctly even if the unobserved taste for housing differs between high- and low-risk couples ($E[\varepsilon_{i,t} \times 1_{\rho,i,t}] \neq 0$) or between couples with high and low $X_{i,t}(E[\varepsilon_{i,t} \times X_{i,t}] \neq 0)$. This weaker identifying assumption provides a second test of predictions from our model.

To proxy for risk, 1_ρ, we use whether a couple shares the same occupation. Same-occupation couples, as we show in the next section, have a lower probability of just one spouse becoming unemployed and higher probability of neither spouse or both spouses becoming unemployed. We include separate dummies to control for the direct effects of the husband's and wife's occupations, $\Omega_{i,t}$, to exploit variation in within-household diversification. Once we remove the effect of either spouse's occupation on housing consumption, does sharing an occupation further increase that consumption? In other words, do a dual-teacher couple and a dual-realtor couple together spend more on owned housing than two mixed teacher-realtor couples? Unlike papers on precautionary saving that use individual-level occupation-based proxies for risk ($1_{occ,i,t}$), including same-occupation dummies allows us to identify α even when risky occupations are correlated with unobserved taste for consumption. In other words, our couple-level assumption that $E[\varepsilon_{i,t} \times 1_{\rho,i,t}] = 0$, which we need to identify α, may hold after controlling for individual occupation, even when its individual-level analog $E[\varepsilon_{i,t} \times 1_{occ,i,t}] = 0$, needed to identify α in other papers, does not.

We use two different proxies for low adjustment costs, $X_{i,t}$, which will range from 0 (the highest adjustment costs, where moving is optimal only in the worst state) to 1 (the lowest adjustment costs, where moving is optimal given smaller shocks). As we will show, this mapping of adjustment costs to the propensity to move matches what we observe in the data.

First, we compare home owners ($X_{i,t} = 0$) to renters ($X_{i,t} = 1$), since home ownership involves greater costs of moving. Second, we use an estimate of the expected length of stay in the house, based on demographics ($X_{i,t}$ = expected probability of moving). This too is a proxy for the moving cost, k. A household that anticipates more frequent moving has a lower effective cost of a forced move. For a household that is likely to move anyway, the cost of a forced move is a minor shift in the timing of the move. If a household was unlikely to move in the absence of joint unemployment, the effective moving cost is the full moving cost.

As a proxy for a low magnitude of the income loss due to unemployment, L, we use the relative generosity of unemployment insurance ($X_{i,t} = 0$ if UI generosity is in the bottom decile, 1 otherwise). Access to more generous UI, whether due to differences in the UI rules across states and over time or nonlinearity in the reimbursement schedule, makes it less likely that even dual-unemployed couples would choose to move.

In the vector of controls, Z, we include the imputed probability of each of the husband and wife becoming unemployed since that affects expected income. We include the product of these probabilities to control for the probability that both spouses would be jointly unemployed if their risks were independent. The Z vector also includes the squared unemployment rates for the husband and wife in case the relationship between the risk of unemployment and housing demand is nonlinear. We control for family income, the share of the income earned by the husband, and dummies for the number of people in the household, the number of children, the educational attainment of the husband and wife, and the age brackets of the husband and wife.

A. *Sharing an Occupation, Unemployment Correlation, and the Probability of Moving*

Before estimating equation (8), we present empirical evidence that (a) sharing an occupation increases the correlation of couples' unemployment events, so that $1_{\rho,i,t}$ is a good proxy for ϕ in the model, and (b) dual unemployment dramatically increases the odds that a household moves relative to single unemployment or dual employment. These results indicate that our proxies for risk and moving costs line up with case c of the model, in which moving is optimal if and only if both spouses become unemployed, and therefore increased risk leads to higher housing consumption.

First, using the data from the SIPP, we estimate the within-household unemployment correlation, controlling for the observable characteristics, $Z_{i,t}$, δ_t, and $\Omega_{i,t}$, for same- and different-occupation couples. The procedure is described in appendix B. We find that same-occupation couples face an increase in risk stemming from more highly correlated unemployment. Over the course of a year, such

[13] Such papers assume that individual-level occupation-based proxies for risk ($1_{occ,i,t}$) are uncorrelated with unobservable tastes ($E[\varepsilon_{i,t} \times 1_{occ,i,t}] = 0$), an assumption that has been criticized (Browning & Lusardi, 1996; Lusardi, 1997).

TABLE 3.—PROBABILITY OF ONE OR BOTH SPOUSES BECOMING UNEMPLOYED AT SOME POINT DURING A TWELVE-MONTH WINDOW CONDITIONAL ON BOTH INITIALLY EMPLOYED, BY WHETHER THE COUPLE SHARES AN OCCUPATION

	(1) Unemployment Correlation	(2) Probability No Spouses Unemployed During Subsequent 12 Months	(3) Probability at Most One Spouse Unemployed During Subsequent 12 Months	(4) Probability Both Spouses Unemployed at Some Point During Subsequent 12 Months	(5) Probability Husband Becomes Unemployed During Subsequent 12 Months	(6) Probability Wife Becomes Unemployed During Subsequent 12 Months	(7) Number of Observations
Different occupation	0.051	0.813	0.173	0.014	0.075	0.126	172,348
	(0.002)	(0.001)	(0.001)	(0.0003)	(0.001)	(0.001)	
Same occupation	0.215	0.857	0.122	0.021	0.054	0.110	5,481
	(0.023)	(0.005)	(0.004)	(0.002)	(0.003)	(0.005)	
Difference	0.163	0.044	−0.051	0.007	−0.021	−0.016	
	(0.024)	(0.005)	(0.005)	(0.002)	(0.003)	(0.005)	

Note: The unit of observation is a couple × month. The sample consists of married couples who both report being employed in any given month. Same-occupation couples have the same three-digit occupation codes. The columns are computed using the formula in equation EEE and the estimated unemployment correlations and probabilities from footnote FFF. All estimates condition on household income, the husband's share of the income, the number of people in the household, the number of children, age, education, and occupation of the husband and the wife, and year. Data are from the April 1996 panel of the Survey of Income and Program Participation. Standard errors (in parentheses) are bootstrapped using sampling with replacement and 200 replications.

couples are more likely to be either both employed or both unemployed and are less likely to have just one spouse unemployed, even controlling for each spouse's occupation. The estimates are reported in table 3. In the first column, different-occupation households have low unemployment correlations of 5.1% (with a bootstrapped standard error of 0.2%), reflecting their within-household diversification. Same-occupation households, by contrast, have a 21.5% correlation in their unemployment risks, yielding a difference between same- and different-occupation couples of 16.3% (2.4% standard error).[14] Same-occupation couples' spread in risk can be seen in columns 2 through 4. Same-occupation couples have higher rates of both spouses becoming unemployed (2.1% versus 1.4%) and both spouses remaining employed (85.7% versus 81.3%), and lower rates of just one spouse becoming unemployed (12.2% versus 17.3%). The differences between each of these are statistically significant. The overall rate of becoming unemployed is somewhat lower for same-occupation husbands and wives: 2.1 percentage points lower for husbands (column 5) and 1.6 percentage points lower for wives (column 6).

Second, dual-unemployment appears to proxy for a "large shock" because it dramatically increases the likelihood of moving relative to single unemployment of dual employment. For the sample of couples in which both spouses are employed in the current month, we estimate a probit to predict moving over the following twelve months. As covariates, we include dummy variables for whether no spouse, one spouse, or two spouses are unemployed in the next month and also $Z_{i,t}$, δ_t, $\Omega_{i,t}$, and a dummy variable for home ownership.

Table 4 reports the marginal effect on the probability of moving from going from zero unemployed to one unemployed and from one to two unemployed. In the first column, just one spouse becoming unemployed raises the probability of moving over the next twelve months by 2.7 percentage points. If both spouses become unemployed, the likelihood of a move rises by an additional 6.3 percentage points over single-unemployment couples. Since the average annual moving rate is just 4.3% for dual-employed home owners, dual unemployment yields an enormous jump in the likelihood of moving.[15]

TABLE 4.—MARGINAL PROBABILITY OF MOVING OVER THE NEXT TWELVE MONTHS, BY UNEMPLOYMENT STATUS

LHS Variable	(1) Moves in the Next 12 Months	(2)
Difference between one newly unemployed and none unemployed	0.027	0.025
	[7.18]	[6.73]
Difference between two newly unemployed and one newly unemployed	0.063	0.062
	[2.35]	[2.07]
Occupation dummies?	No	Yes
Adjusted R^2	0.1399	0.1861
Number of observations	201,413	183,430

Note: z-statistics are in square brackets. The unit of observation is a couple × month. The sample consists of married couples who both report being employed in the current month. The columns report the estimated marginal effects of the number of spouses jointly unemployed on the probability of moving, estimated using a probit. The probability of moving measures whether there will be at least one change of address during subsequent twelve months. All estimates control for home ownership status, household income, the husband's share of the income, the number of people in the household, the number of children, age, education, and year. Data are from the April 1996 panel of the Survey of Income and Program Participation.

[14] We obtain similar results if we restrict the sample to home owners. In that case, the difference in correlation between same- and different-occupation couples is 0.134 (0.025). We also obtain similar results if we use the SIPP's layoff variable rather than our measure of unemployment or "same industry" rather than "same occupation."

[15] Since we do not observe unemployment severity, the realization of moves after an unemployment shock is only a rough indicator of the probability of crossing an S-s bound. Since some unemployment events in the data have a larger effect on permanent income than others, they do not correspond perfectly to the unemployment states in the model. In particular, a substantial fraction of unemployment shocks are temporary. As such, the estimates in tables 3 and 4 will not map perfectly to model parameters. Even if all households that suffer dual unemployment (as defined in our model) move, it is not surprising that many household listed as dual unemployed (but whose unemployment is very transitory) in the data do not move.

TABLE 5.—EFFECT OF HIGHER CORRELATION IN UNEMPLOYMENT RISK ON LOG HOUSE VALUE, FOR HOME OWNERS

LHS Variable: log(House Value)	(1)	(2)	(3)
Same occupation [1_ρ]	**0.043**	**0.027**	**0.021**
	(0.004)	**(0.003)**	**(0.004)**
Husband's unemployment rate	−6.026	−2.290	1.516
[p]	(0.271)	(0.227)	(0.375)
Husband's unemployment	23.872	6.804	−8.089
rate^2 [p^2]	(1.580)	(1.304)	(2.086)
Wife's unemployment rate	−0.098	0.455	0.213
[q]	(0.205)	(0.172)	(0.293)
Wife's unemployment rate^2	−1.234	−1.922	0.460
[q^2]	(0.707)	(0.586)	(1.027)
Husband's unemployment	−11.313	−6.554	−2.644
rate × Wife's rate [$p \times q$]	(1.496)	(1.239)	(1.282)
Z controls?	No	Yes	Yes
MSA × year dummies?	No	Yes	Yes
Occupation dummies?	No	No	Yes
Adjusted R^2	0.2976	0.3180	0.3370

Note: N = 231,598. Left-hand-side variable is log(house value). See table A1 for the sample construction. The **Z** vector of demographic controls in columns 2–4 include the log family income, husband's percentage share of the couple's income, a vector of dummy variables for the number of persons in the household, number of children, educational attainment of the husband and wife, and age of the husband and wife. Data are from the 1980–2000 IPUMS.

B. Relationship between Income Risk and House Value

We examine the two predictions of our model: first, $\alpha + \beta \times X_{i,t} > 0$ for low $X_{i,t}$ (high moving costs) but not otherwise, and second, $\beta < 0$. To test the first prediction, we must assume that same-occupation couples do not differ from other couples in their unobserved taste for housing, after including a host of controls ($E[\varepsilon_{i,t} \times 1_{\rho,i,t}] = 0$). To test the second, weaker prediction, we must assume only that any difference between same- and different-occupation couples in their unobserved taste for housing is uncorrelated with $X_{i,t}(E[\varepsilon_{i,t} \times 1_{\rho,i,t} \times X_{i,t}] = 0)$.

First, does greater risk increase housing consumption when moving costs are high enough to deter moving in all but the worst state? We begin by restricting our attention to households with high moving costs: home owners. We predict $\alpha > 0$; same-occupation home owners consume more housing than otherwise identical different-occupation home owners. These results are reported in table 5, where we find that same-occupation home-owning couples spend more on housing on average than do different-occupation home-owning couples. Column 1 of table 5 reports the results when the only additional covariates are the unemployment rate controls, log of family income, and the income share of the husband. Husbands and wives with the same occupation spend 4.3% more on housing ($\hat{\alpha} = 0.043$, with a 0.4% standard error) than couples with different occupations.[16] This positive estimated relationship between "same occupation" and housing consumption, $\hat{\alpha} > 0$, is consistent with our model but not consistent with a standard precautionary saving model without adjustment costs.

By including additional covariates in this regression, we can rule out many alternative reasons that same-occupation couples may have a greater preference for housing. In column 2, we control for MSA of residence in each year and a host of household demographic characteristics. We find that same-occupation couples buy houses that are on average 2.7% more expensive when compared to different-occupation couples in the same MSA and same year. In column 3, to control for the possibility that same-occupation couples are more prevalent in occupations that have a strong unobservable preference for housing, we add dummy variables for each spouse's occupation. The new coefficient on "same occupation" implies that after controlling for covariates, same-occupation couples spend 2.1% (0.4% standard error) more on their houses than do different-occupation couples. Since same-occupation couples have a 16 percentage point higher correlation in unemployment risk (from table 3), simple extrapolation gives an elasticity of house spending with respect to the unemployment correlation of 0.13 (0.021/0.16).[17]

Second, does greater risk raise housing consumption more when moving costs are higher ($\beta < 0$)? In table 6, we compare high- and low-moving cost households by interacting the same-occupation dummy variable with a proxy for moving costs. In columns 1 and 2, we pool renters and home owners together. The proxy for low moving costs is whether a household rents, with $X_{i,t} = 1$ for renters and $X_{i,t} = 0$ for home owners. Column 1 includes the full set of $Z_{i,t}$ covariates and the $\delta_{k,t}$ MSA × year dummies, all interacted with $X_{i,t}$, the dummy for renter status. Column 2 adds occupation dummies. The dependent variable is log housing spending, defined as log annual rent for renters and log annualized house price for owners. We annualize the house price to standardize renters' and owners' housing costs and calculate it by multiplying the house price by the sample average rent to house price ratio. Errors in this transformation will be absorbed by the renter dummy.

The first row of table 6 shows that same-occupation home owners spend about 2.7% more on housing than do different-occupation home owners ($\hat{\alpha} = 0.027$, with a standard error of 0.5%). The second row of table 6 reports that the

[16] While we generally treat the other estimated coefficients as nuisance parameters, some persistent results bear highlighting. Consistent with empirical precautionary saving papers that exploit occupation-level variation in unemployment risk (Carroll and Samwick, 1997, 1998), husbands in occupations with high risks of unemployment spend less on housing (the coefficient on the husband's imputed unemployment rate < 0). We also find that households that face greater occupation-based unemployment risk spend less on housing (the coefficient on the husband's unemployment rate × wife's rate < 0). But our key finding is that these results are attenuated or even reversed when the correlation of couples' unemployment increases (the coefficient on same occupation > 0).

[17] These results do not seem sensitive to our particular empirical choices. We obtain essentially the same results when we use "same industry" rather than "same occupation" as our proxy for couples' unemployment correlation; control for other dimensions of similarity, such as age (in ranges) and education; or split the sample by decade, running separate regressions for 1980, 1990, and 2000. However, we find a significantly larger effect of a higher correlation in unemployment risk on housing spending for same-occupation home owners when the husband is under age 45, where neither of the spouses has had any postsecondary schooling, or when the household income is below the sample median.

TABLE 6.—HOW THE EFFECT OF RISK VARIES WITH MOVING COSTS AND THE MAGNITUDE OF POTENTIAL INCOME LOSSES

LHS Variable:	(1) Log(Housing Spending)	(2)	(3) Log(House Price)	(4)	(5) Log(Rent)	(6)	(7) Log(House Price)	(8)
Sample:	Everyone		Owners		Renters		Owners	
X Variable:	X = 1 if renter, 0 if owner		X = P(Moving)		X = P(Moving)		X = 1 if in Top 90% of UI Replacement Rate, 0 Otherwise	
Same occupation	0.027	0.022	0.046	0.033	−0.057	−0.038	0.073	0.027
	(0.005)	(0.005)	(0.007)	(0.007)	(0.036)	(0.035)	(0.018)	(0.017)
Same occupation × **X**	−0.046	−0.039	−0.133	−0.079	0.157	0.175	−0.055	−0.016
	(0.010)	(0.010)	(0.041)	(0.041)	(0.165)	(0.160)	(0.019)	(0.018)
Z × **X** controls?	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
MSA × year × **X** controls?	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Occupation dummies?	No	Yes	No	Yes	No	Yes	No	Yes
Adjusted R^2	0.0863	0.0982	0.5353	0.5498	0.1477	0.2127	0.5575	0.5694
Number of observations	290,062	290,062	231,598	231,598	58,464	58,464	156,285	156,285

Note: Housing spending is defined as log annual rent for renters and log(house price × sample average(rent)/sample average(house price)) for owners. The **Z** vector of demographic controls include the log family income, husband's percentage share of the couples's income, a vector of dummy variables for the number of persons in the household, number of children, educational attainment of the husband and wife, and age of the husband and wife. The regressions in columns 7 and 8 also include state × year × **X** controls and report robust standard errors, corrected for correlation by state × year × being in the top 9% of UI replacement rate. Data are from the 1980 through 2000 IPUMS except for columns 7 and 8, which use the 1990 and 2000 IPUMS.

difference in housing spending between same- and different-occupation renters is 4.6 percentage points less than the difference for owners ($\hat{\beta} = -0.046$ with a standard error of 1.0%), yielding a negative relationship on net between risk and housing spending for renters. This difference between owners' and renters' response to risk is the pattern predicted by a model with consumption commitments. When we include controls for the occupations of the husband and wife in column 2, the results are similar.

Columns 3 and 4 of table 6 present the results from estimating equation (8) on a sample of home owners and using the imputed probability of moving as the proxy for low moving costs, $X_{i,t}$. We impute the likelihood of moving as the rate of recent moving by similar families. In the IPUMS, we construct the average rate of having moved in the previous year by husband's age × husband's education × presence of children cells. We define the bins using ten-year age brackets, nine education categories, and an indicator for whether the family has any children, and take the average for all of the households in that bin excluding the household in question. The covariates $Z_{i,t}$ control separately for each of the household attributes we use to impute the probability of moving (age, education, and presence of children), so $X_{i,t}$ can be identified separately from $Z_{i,t}$ by the fact that the age profiles of moving vary with education and offspring.

The first row of column 3 corresponds to home owners with a zero imputed probability of moving—that is, they are highly unlikely to move. Such households face the highest moving costs, and same-occupation households in that category spend 4.6% (with a standard error of 0.7%) more on housing than do otherwise identical different-occupation households. For a household that was planning to move anyway ($X_{i,t} \approx 1$), the effective cost of a forced move is very low; these couples should display the standard negative relationship between risk and consumption. This is what we find. As the likelihood of moving rises, reflecting a more mobile household, the differential between same- and different-occupation households is reduced (row 2). At a probability of moving of about 0.35—a three-year expected stay (1/0.35)—the estimated difference between same- and different-occupation households disappears.[18] And then, as the probability of moving increases further, same-occupation households are estimated to consume less housing than otherwise-equivalent different-occupation households. These highly mobile households have low moving costs and thus behave like precautionary savers when faced with an increase in risk. In column 4, adding occupation dummies as controls does little except reduce the magnitude of the estimated coefficients on the "same occupation" variables by one-third.

Columns 5 and 6 of table 6 present results of regressions identical to columns 3 and 4 for the subsample of renters. Given the lower moving costs renters face, there is no reason to believe that the probability of moving for demographic reasons would have any additional effect on the relationship between "same occupation" and housing consumption. Consistent with this prediction, we find that the estimated coefficient on the "same occupation" term is insignificantly different from 0, so that the relationship between "same occupation" and housing consumption for renters is independent of their exogenous move probability.

Finally, columns 7 and 8 present results comparing how the response to risk varies with unemployment insurance (UI) generosity for the subsample of home owners.[19] A

[18] When $X = 0.35$, $-0.133 \times X = -0.046$, exactly offsetting the estimated coefficient on same occupation.

[19] There are several sources of variation in the generosity of UI, which we define as the replacement rate of the couple's wages. While state UI programs typically compensate the unemployed for up to 50% of lost wages up to a cap, the level of the cap and the replacement rate schedule vary across states, over time, and according to the number of dependent children. In addition, the generosity of UI is a nonlinear function of income: the replacement rate remains constant until income reaches the cap, at which point it declines with income. Once income exceeds the cap,

higher UI replacement rate reduces the effective size of shocks to permanent income since it affords the unemployed the ability to set a higher reservation wage in their job search (Feldstein & Poterba, 1984). A theory incorporating consumption commitments predicts that there would be a positive relationship between "same occupation" and housing spending only when the household couple experience a loss large enough to induce moving. As the UI replacement rate increases, the odds of such a sizable loss fall, and the difference in housing spending between same- and different-occupation couples should decrease (become less positive). By contrast, a model of precautionary saving without commitment would predict that same-occupation couples would spend less on housing than different-occupation couples and that gap should decrease (become less negative) as UI becomes more generous and replaces the precautionary function of the household's own savings. (This precautionary saving mechanism is described in Engen & Gruber, 2001.)

To test this theory, we interact an indicator variable for the couple being above the tenth percentile of the UI replacement rate with the same-occupation indicator.[20] $X_{i,t} = 1$ for those with relatively more generous unemployment insurance (low L) and $X_{i,t} = 0$ for those with less generous unemployment insurance (high L). We also control for the UI replacement rate dummy interacted with the demographic covariates, MSA × year and state × year. The standard errors are corrected for correlation by state × year × X.

The results are reported in column 7. Same-occupation households that face low UI replacement rates spend 7.3% (1.8% standard error) more on housing relative to different-occupation households. This difference is attenuated when UI becomes more generous. A same-occupation household facing the higher replacement rate spends just 1.8% (7.3 − 5.5) more on their house than an otherwise identical different-occupation household. In column 8, we add separate occupation dummies for the husband and wife. While the signs on the "same occupation" coefficients remain the same, they are smaller in magnitude than in column 7 and no longer are statistically significant.

TABLE 7.—EFFECT OF SAME OCCUPATION ON THE DEMAND FOR HOME OWNERSHIP AND THE IMPACT OF EFFECTIVE MOVING COSTS

	(1)	(2)
LHS Variable	Own = 1	
Sample	Everyone	
X Variable	P(moving)	
Same occupation	−0.018	−0.017
	(0.005)	(0.005)
Same occupation × **X**	0.019	0.019
	(0.027)	(0.026)
Z × **X** controls?	Yes	Yes
MSA × year × **X** controls?	Yes	Yes
Occupation dummies?	No	Yes
Adjusted R^2	0.1942	0.2024
Number of observations	290,062	290,062

Note: Linear probability model where the left-hand-side variable takes the value of 1 if the household owns its own home. See table A1 for the sample construction. The **Z** vector of demographic controls includes the log family income, husband's percentage share of the couples's income, a vector of dummy variables for the number of persons in the household, number of children, educational attainment of the husband and wife, and age of the husband and wife. Data are from the 1980–2000 IPUMS.

C. Self-Selection

Because high-moving-cost households increase housing consumption when income risk rises, while low-moving-cost households do not, we can reject both the standard precautionary saving model and the possibility that our results can be explained by an unobserved taste for housing among high-risk households. Suppose that same-occupation couples had a higher mean unobserved taste for housing than different-occupation couples, so their preference distribution shifted to the right. That form of heterogeneity would cause same-occupation home owners to spend more on their homes. But it would also suggest, counterfactually, that same-occupation renters would spend more on rent.

We can also reject other alternative explanations for the results presented in section IVC. Suppose instead that same-occupation couples merely had a preference for spending a lot on home ownership but not on rental housing. If that were the case, they should also be more likely to own their houses, which is rejected by the data. In table 7, we regress an indicator variable for being a home owner on the "same occupation" dummy variable (1_ρ) and also interact it with our estimated probability of moving. Same-occupation couples actually are less likely than different-occupation couples to be home owners, a result that is strongest for couples who are unlikely to move.

the absolute benefit amount remains level at 50% of the cap amount but is declining as a percentage of income. This nonlinearity in the replacement rate implies that the share of income earned by each spouse influences the couple's total replacement rate in the event that both spouses lose their jobs. For example, if both spouses earn exactly the cap amount and both become unemployed, they collectively will receive 50% of their former wages. But if one spouse earns twice the cap amount and the other earns almost nothing—total family income is the same but its allocation across spouses is not—the family can at best replace 25% of their former income.

To implement this proxy, we calculate UI replacement rates for each spouse using the unemployment insurance calculator developed by Cullen and Gruber (2000) and extended by Chetty. We are grateful to Raj Chetty for letting us use his UI calculator and benefit data. Since the calculator contains information about UI since 1984, we restrict our sample to 1990 and 2000. We then calculate the household's replacement rate as the average replacement rate for each spouse, weighted by their respective income shares.

[20] The tenth percentile household average replacement rate is about 15%. The replacement rate increases rapidly above that point to a maximum of 0.5: the mean replacement rate in the bottom decile is 0.076 and in the top 90% of it is 0.352. In columns 7 and 8 we weight the unemployment risk variables by the husband's and wife's income shares. (The husband's unemployment rate is multiplied by s and the wife's by $(1 - s)$. The same occupation indicator is multiplied by $s(1 - s)$.) We also control for s and $s(1 - s)$ separately. This weighting is necessary to separately identify the UI replacement rate effect, which is a function of the program rules and the within-household distribution of income, from the distribution of income alone. In table 6, we report the coefficients on the same occupation variables evaluated at the average $s(1 - s)$, which is 0.208.

If same-occupation couples had the same mean but higher variance in their unobserved taste for housing than different-occupation couples, it could explain more of the empirical regularities we find. Since same-occupation households would have thicker tails in the distribution of their preference for housing—they would either love housing or hate it—same-occupation households who loved housing would own and also spend more than the more neutral different-occupation households. Those same-occupation households that disliked housing would rent, and not spend much on rent, relative to different-occupation households. In addition, depending on the clearing price of owned housing, it is possible that more different-occupation than same-occupation households prefer owning. In that case, same-occupation households would have a lower rate of home ownership. If this explanation were true, it would imply a straightforward and testable prediction. The residuals for the same-occupation, home-owning couples in the housing demand regression should be more right skewed than those for the different-occupation couples. Similarly, the residuals for the same-occupation renting couples in the rent regression should be more left skewed. In our data, there is no distinguishable difference in skewness in residuals between same- and different-occupation couples, so variation in the second moment of unobserved taste for housing cannot explain our results.

Furthermore, it seems unlikely that any difference between same- and different-occupation couples in the taste for housing is present only for households with a low exogenous probability of moving or less generous unemployment insurance. UI generosity is determined by state interacted with the within-household wage distribution. In our regressions, we control for state and the wage distribution separately, identifying the UI effect using the interaction of the two sources of variation. Similarly, the probability of moving is determined by age $\times$ education $\times$ presence of children. Again, we control for each of these covariates separately. It would be highly unlikely for house-loving households to be concentrated within one of these subgroups. Furthermore, there is no reason that any such effect would be limited to home owners and not renters, which would be needed to reconcile columns 3 and 4 of table 6 with columns 5 and 6.

V. Conclusion

This paper shows that adjustment costs can invert the usual negative relationship between risk and consumption. The result is driven by the strong desire to reduce committed consumption in advance of possible shocks too small to warrant adjustment relative to shocks large enough to make adjustment worthwhile. An increase in risk that makes small losses less likely but large losses more likely will then lead to increased committed consumption.

We illustrated this idea in the context of a dual-career household that faces unemployment risk and consumes housing. This result requires adjustment costs to be high enough to deter moving in all but the worst states. Therefore, it should not apply when moving costs are low, as they are for renters or those who expect to move soon. We exploit this feature of the model in our empirical tests, which compare the effect of increased risk in settings where moving costs are high to ones where they are low. This comparison provides a test of the model that differentiates it from leading alternative hypotheses.

When we proxy for increased risk—more precisely, a higher correlation in unemployment risk—with whether a married couple shares the same occupation, we find the predicted pattern pervasive in the data. Controlling for each spouse's characteristics, including their individual occupations and probabilities of unemployment, we find that same-occupation households spend relatively more on housing. As expected, this result is confined to home owners and is strongest for owners who face effectively higher moving costs due to a lower exogenous probability of moving. Finally, same-occupation couples spend relatively more on housing consumption compared to different-occupation couples when unemployment insurance is less generous.

In this paper, we focus on the role of adjustment dynamics in affecting the relationship between risk and the level of consumption. The precautionary saving literature is concerned with the impact of prudence on the level of saving or consumption but not on adjustment dynamics. We show that adjustment costs can have a large effect on the relationship between risk and consumption in both theory and data.[21] Of course, our finding that commitments affect the relationship between risk and consumption does not deny the importance of prudence in generating precautionary saving. It merely follows from a difference in the precautionary motive in the domains of large and small losses.

Our findings may also explain why precautionary saving results are quite sensitive to the sample in which they are measured (Hurst et al., 2005). Some papers find little evidence that those with more risk save more and spend less. These include Dynan (1993) and Guiso, Jappelli, and Terlizzese (1992), who use household consumption variability and expectations of household risk, respectively, as measures for risk. Other papers find strong evidence of this relationship, including Carroll and Samwick (1997, 1998) and Fuchs-Schundeln and Schundeln (2005), who exploit individual-level, occupation-based variation in income risk. Our results help to reconcile these varied findings. We show

[21] This result is at the nexus of the precautionary saving, S-s adjustment, and consumption commitments literatures. The consumption commitments literature examines the impact of adjustment dynamics on risk aversion and not on prudence and the level of consumption (Chetty & Szeidl, 2007; Browning & Crossley, 2004; Postlewaite, Samuelson, & Silverman, 2006). The empirical literature on durable goods and S-s bands considers the impact of shocks on adjustment dynamics but not on the target level of consumption (Attanasio, 2000; Eberly, 1994; Bertola, Guiso, & Pistaferri, 2005). Following Glossman and Laroque (1990), portfolio choice has also been studied at length (Fratantoni, 1998, 2001; Flavin, 2001; Flavin & Yamashita, 2002; Chetty & Szeidl, 2005).

that results in the domain of individual-level risks to single unemployment may be different from household-level risks to dual unemployment. In the range of individual-level losses, our findings mirror the standard precautionary saving result: couples in occupations with higher unemployment probabilities consume less housing. When the range of losses is expanded to include household-level dual unemployment, holding total unemployment risk fixed, this result is reversed.

REFERENCES

Attanasio, O., "Consumer Durables and Inertial Behaviour: Estimation and Aggregation of (S,s) Rules for Automobile Purchases," *Review of Economic Studies* 67:4 (2000), 659–677.

Bertola, G., L. Guiso, and L. Pistaferri, "Uncertainty and Consumer Durables Adjustment," *Review of Economic Studies* 72:4 (2005), 973–1007.

Browning, M., and T. F. Crossley, "Shocks, Stocks and Socks: Smoothing Consumption over a Temporary Income Loss," CAM working paper (2004).

Browning, M., and A. Lusardi, "Household Saving: Micro Theory and Micro Facts," *Journal of Economic Literature* 34:4 (1996), 1797–1855.

Brueckner, J., "Consumption and Investment Motives and the Portfolio Choices of Homeowners," *Journal of Real Estate Finance and Economics* 15:2 (1997), 159–180.

Carroll, C. D., and A. A. Samwick, "The Nature of Precautionary Wealth," *Journal of Monetary Economics* 40:1 (1997), 41–71.

——— "How Important Is Precautionary Saving?" this REVIEW 80:3 (1998), 410–419.

Chetty, R., "Consumption Commitments, Unemployment Durations, and Local Risk Aversion," NBER working paper no. 10211 (2004).

Chetty, R., and A. Szeidl, "Marriage, Housing, and Portfolio Choice: A Test of Grossman-Laroque," UC Berkeley working paper (2005).

——— "Consumption Commitments and Risk Preferences," *Quarterly Journal of Economics* 122:2 (2007), 831–877.

Cullen, J. B., and J. Gruber, "Does Unemployment Insurance Crowd Out Spousal Labor Supply?" *Journal of Labor Economics* 18:3 (2000) 546–572.

Diaz-Serrano, L., "Labor Income Uncertainty, Skewness and Homeownership: A Panel Data Study for Germany and Spain," *Journal of Urban Economics* 58:1 (2005), 156–176.

Drèze, J. H., and F. Modigliani, "Consumption Decisions under Uncertainty," *Journal of Economic Theory* 5:3 (1972), 308–335.

Dynan, K. E., "How Prudent Are Consumers?" *Journal of Political Economy* 101 (1993), 1104–1113.

Eberly, J. C., "Adjustment of Consumers' Durables Stocks: Evidence from Automobile Purchases," *Journal of Political Economy* 102:3 (1994), 403–436.

Engen, E. M., and J. Gruber, "Unemployment Insurance and Precautionary Saving," *Journal of Monetary Economics* 47:3 (2001), 545–579.

Feldstein, M., and J. Poterba, "Unemployment Insurance and Reservation Wages," *Journal of Public Economics* 23:1–2 (1984), 141–167.

Flavin, M., "Owner-Occupied Housing in the Presence of Adjustment Costs," UCSD working paper (2001).

Flavin, M., and T. Yamashita, "Owner-Occupied Housing and the Composition of the Household Portfolio," *American Economic Review* 91:1 (2002), 345–362.

Fratantoni, M. C., "Homeownership and Investment in Risky Assets," *Journal of Urban Economics* 44:1 (1998), 27–42.

——— "Homeownership, Committed Expenditure Risk, and the Stockholding Puzzle," *Oxford Economic Papers* 53:2 (2001), 241–259.

Fuchs-Schundeln, N., and M. Schundeln, "Precautionary Savings and Self-Selection: Evidence from the German Reunification 'Experiment,' " *Quarterly Journal of Economics* 120:3 (2005), 1085–1120.

Goetzmann, W., "The Single Family Home in the Investment Portfolio," *Journal of Real Estate Finance and Economics* 6:3 (1993), 201–222.

Grossman, S., and G. Laroque, "Asset Pricing and Optimal Portfolio Choice in the Presence of Illiquid Durable Consumption Goods," *Econometrica* 58:1 (1990), 25–51.

Guiso, L., T. Jappelli, and D. Terlizzese, "Earnings Uncertainty and Precautionary Saving," *Journal of Monetary Economics* 30:2 (1992), 307–337.

Haurin, D. R., "Income Variability, Homeownership, and Housing Demand," *Journal of Housing Economics* 1:1 (1991), 60–74.

Haurin, D. R., and H. L. Gill, "Effects of Income Variability on the Demand for Owner-Occupied Housing," *Journal of Urban Economics* 22:2 (1987), 136–151.

Hendershott, P., and J. Slemrod, "Taxes and the User Cost of Capital for Owner-Occupied Housing," *AREUEA Journal* 10:4 (1983), 375–393.

Henderson, J. V., and Y. M. Ioannides, "A Model of Housing Tenure Choice," *American Economic Review* 73:1 (1983), 98–113.

Hurst, E., A. Kennickel, A. Lusardi, and F. Torralba, "Precautionary Savings and the Importance of Business Owners," NBER working paper no. 11731 (2005).

Kimball, M. S., "Precautionary Saving in the Small and in the Large," *Econometrica* 58:1 (1990), 53–73.

Leland, H. E., "Saving and Uncertainty: The Precautionary Demand for Saving," *Quarterly Journal of Economics* 82:3 (1968), 465–473.

Lusardi, A., "Precautionary Saving and Subjective Earnings Variance," *Economics Letters* 57 (1997), 319–326.

Postlewaite, A., L. Samuelson, and D. Silverman, "Consumption Commitments and Employment Contracts," University of Michigan working paper (2006).

Poterba, J., "Tax Subsidies to Owner-Occupied Housing: An Asset Market Approach," *Quarterly Journal of Economics* 99:4 (1984), 729–752.

Robst, J., R. Deitz, and K. McGoldrick, "Income Variability, Uncertainty, and Housing Tenure Choice," *Regional Science and Urban Economics* 29:2 (1999), 219–229.

Sandmo, A., "The Effect of Uncertainty on Saving Decisions," *Review of Economic Studies* 37:3 (1970), 353–360.

Skinner, J., "Risky Income, Life Cycle Consumption, and Precautionary Savings," *Journal of Monetary Economics* 22:2 (1988), 237–255.

APPENDIX A

Proofs

Plugging the distribution of unemployment shocks from equation (6) into the objective function from equation (5) yields the following objective function:

$$
\begin{aligned}
U(h_1, f_1) = u(h_1, f_1) & \\
&+ (1 - p - q + \phi)\max\begin{pmatrix} (u(h_1, Y_1 + 2Y_2 - 2h_1 - f_1), \\ u(\tfrac{1}{2}(Y_1 + 2Y_2 - h_1 - f_1 - k), \\ \tfrac{1}{2}(Y_1 + 2Y_2 - h_1 - f_1 - k)) \end{pmatrix} \\
&+ (p + q - 2\phi)\max\begin{pmatrix} u(h_1, Y_1 + 2Y_2 - 2h_1 - f_1 - L), \\ u(\tfrac{1}{2}(Y_1 + 2Y_2 - h_1 - f_1 - k - L), \\ \tfrac{1}{2}(Y_1 + 2Y_2 - h_1 - f_1 - k - L)) \end{pmatrix} \\
&+ \phi\max\begin{pmatrix} u(h_1, Y_1 + 2Y_2 - 2h_1 - f_1 - 2L), \\ u(\tfrac{1}{2}(Y_1 + 2Y_2 - h_1 - f_1 - k - 2L), \\ \tfrac{1}{2}(Y_1 + 2Y_2 - h_1 - f_1 - k - 2L)) \end{pmatrix}.
\end{aligned}
\tag{A1}
$$

The first maximum refers to second-period utility in the dual-employment state, the second to the single-unemployment state, and the third to the dual-unemployment state. In case a (always move), the second argument in each maximum is assumed to be greater; in case b (move only under single or dual unemployment), the first argument is assumed to be greater only in the first maximum; in case c (move only under dual unemployment), the second argument is assumed to be greater only in the third maximum; in case d (never move), the first argument is assumed to be greater in each maxima.

Proof of Lemma 1: $p, q \rightarrow 0$ Optimal consumption. Relabel $U(h_1, f_1)$ as $U(h_1, f_1, p, q, \phi, k)$. The existence of u_{hh}, u_{ff}, u_{hhh}, and u_{fff} over the relevant ranges of h and f ensures that u_h, u_f, u_{ff}, and u_{hh} are continuous over this range. We ignore the values Y, L, and k that form knife-edge cases in which the two arguments in the second or third maximum are exactly equal (where the conditions for this are present below). These cases occur with probability 0 if model parameters are chosen at random within an arbitrarily small window. This implies that $\partial U(h_1, f_1)/\partial f_1$, $\partial U(h_1, f_1)/\partial h_1$, $\partial U^2(h_1, f_1)/\partial f_1 \partial \phi$, and $\partial U^2(h_1, f_1)/\partial h_1 \partial \phi$ are continuous in the neighborhood around $h_1, f_1 = \frac{1}{4}(Y_1 + 2Y_2)$; $p = q = \phi = 0$. Consider the solution to the system of four equations (which are straightforward to calculate in terms of equation (A1)):

$$\partial U(h_1, f_1)/\partial f_1 = 0$$

$$\partial U(h_1, f_1)/\partial h_1 = 0$$

$$\partial U^2(h_1, f_1)/\partial f_1 \partial \phi = 0$$

$$\partial U^2(h_1, f_1)/\partial h_1 \partial \phi = 0$$

in four unknowns $\{f_1^*, h_1^*, \partial h_1^*/\partial \phi, \partial f_1^*/\partial \phi\}$. When $p = q = \phi = 0$, $h_1^*, f_1^* = \frac{1}{4}(Y_1 + 2Y_2)$ will be part of one such solution (as they solve the first-order conditions described for an interior optimum described in the first two equations); it is trivial to show that this is the global optimum in the maximization problem. The implicit function theorem ensures that there exists a ball for p, q, and ϕ around $p = q = \phi = 0$ such that all solutions $\{f_1^*, h_1^*, \partial h_1^*/\partial \phi, \partial f_1^*/\partial \phi\}$ for these values of p, q, and ϕ are arbitrarily close to the solutions when $p = q = \phi = 0$. Therefore, $\lim_{p,q\rightarrow 0} \{h_1^*, f_1^*, h_1^*/\partial \phi, \partial f_1^*/\partial \phi\}$ can be found by evaluating this system of equations at $p = q = \phi = 0$. $\lim_{p,q\rightarrow 0} h_1^*, f_1^* = \frac{1}{4}(Y_1 + 2Y_2)$. The solutions for $\partial h_1^*/\partial \phi$ and $\partial f_1^*/\partial \phi$ in each of the optimal adjustment cases are shown below. The ranges of k that rationalize these patterns of optimal adjustment are given below.

Case a: Always move:

$$\frac{dh_1^*}{d\phi} = \frac{df_1^*}{d\phi} = \frac{\left[\begin{array}{r} u_h(\frac{1}{4}(Y_1 + 2Y_2) - \frac{1}{2}k, \frac{1}{4}(Y_1 + 2Y_2) - \frac{1}{2}k) \\ - 2u_h(\frac{1}{4}(Y_1 + 2Y_2) - \frac{1}{2}k - \frac{1}{2}L, \frac{1}{4}(Y_1 + 2Y_2) - \frac{1}{2}k - \frac{1}{2}L) \\ + u_h(\frac{1}{4}(Y_1 + 2Y_2) - \frac{1}{2}k - L, \frac{1}{4}(Y_1 + 2Y_2) - \frac{1}{2}k - L) \end{array}\right]}{u_{hh}(\frac{1}{4}(Y_1 + 2Y_2)) + u_{hh}(\frac{1}{4}(Y_1 + 2Y_2) - \frac{1}{2}k) + u_{hf}(\frac{1}{4}(Y_1 + 2Y_2) - \frac{1}{2}k)}.$$

Note that the denominator is negative since by assumption, $u_{hh} < 0$ and $u_{hh} < u_{hf}$. The numerator will be positive if $u_{hhh} > 0$ everywhere, negative if $u_{hhh} < 0$, and 0 if $u_{hhh} = 0$. Therefore, $dh_1^*/d\phi < 0$ if $u_{hhh} > 0$ everywhere, $dh_1^*/d\phi > 0$ if $u_{hhh} < 0$ everywhere, and $dh_1^*/d\phi = 0$ if $u_{hhh} = 0$ everywhere.

Case b: Move in all but best state:

$$\frac{dh_1^*}{d\phi} = 0;$$

$$\frac{df_1^*}{d\phi} = \frac{1}{2}\frac{\left[\begin{array}{r} u_h(\frac{1}{4}(Y_1 + 2Y_2), \frac{1}{4}(Y_1 + 2Y_2)) - 2u_h(\frac{1}{4}(Y_1 + 2Y_2) - \frac{1}{2}L \\ - \frac{1}{2}k, \frac{1}{4}(Y_1 + 2Y_2) - \frac{1}{2}L - \frac{1}{2}k + u_h(\frac{1}{4}(Y_1 + 2Y_2) \\ - L - \frac{1}{2}k, \frac{1}{4}(Y_1 + 2Y_2) - L - \frac{1}{2}k \end{array}\right]}{u_{hh}(\frac{1}{4}(Y_1 + 2Y_2), \frac{1}{4}(Y_1 + 2Y_2))}.$$

Housing consumption does not change in response to increased risk, ϕ. Food consumption does. Again, the denominator of $df_1^*/d\phi$ is negative. By the same continuity argument using the implicit function theorem as above, as $k \rightarrow 0$, the numerator converges to the numerator in case a. Therefore, there exists a $k > 0$ such that $df_1^*/d\phi < 0$ for any utility function with $u_{hhh} > 0$ everywhere, since the numerator converges to a value that must be positive if $u_{hhh} > 0$. For larger values of k, depending on the curvature of the utility function, $df_1^*/d\phi$ could be of either sign. $df_1^*/d\phi = \frac{1}{4}k$ for $u_{hhh} = 0$ (quadratic utility), though again this converges to zero as k does. Note that $df_1^*/d\phi$ in the quadratic case merely represents a wealth effect, as $\frac{1}{4}dE[f_1^* + h_1^* + \tilde{h}_2 + \tilde{f}_2]/d\phi = \frac{1}{4}k$.

Case c: Move only in worst state:

$$\frac{dh_1^*}{d\phi} = \frac{1}{2}\frac{\left[\begin{array}{r} u_h(\frac{1}{4}(Y_1 + 2Y_2), \frac{1}{4}(Y_1 + 2Y_2) - L) \\ - u_f(\frac{1}{4}(Y_1 + 2Y_2), \frac{1}{4}(Y_1 + 2Y_2) - L) \end{array}\right]}{u_{hh}(\frac{1}{4}(Y_1 + 2Y_2), \frac{1}{4}(Y_1 + 2Y_2)) - u_{hf}(\frac{1}{4}(Y_1 + 2Y_2), \frac{1}{4}(Y_1 + 2Y_2))}.$$

Again, the denominator is negative. The numerator is positive because marginal utility is falling in consumption and $\frac{1}{4}(Y_1 + 2Y_2) - L < \frac{1}{4}(Y_1 + 2Y_2)$. Therefore $dh_1^*/d\phi > 0$. The general closed-form expression for $df_1^*/d\phi$ is complex and uninformative, but under separability when $u_{hf} = 0$, and when utility is of the form $u(h, f) = g(h) + g(f)$, it reduces to

$$\frac{df_1^*}{d\phi} = \frac{1}{2}\frac{g'(\frac{1}{4}(Y_1 + 2Y_2) - L - \frac{1}{2}k) - g'(\frac{1}{4}(Y_1 + 2Y_2) - L)}{g''(\frac{1}{4}(Y_1 + 2Y_2))}.$$

As a result, $d(f_1^* + h_1^*)/d\phi$ can be expressed under separability as

$$\frac{dh_1^* + df_1^*}{d\phi} = \frac{1}{2}\frac{\left[\begin{array}{r} g'(\frac{1}{4}(Y_1 + 2Y_2) - L - \frac{1}{2}k) - 2g'(\frac{1}{4}(Y_1 + 2Y_2) - L) \\ + g'(\frac{1}{4}(Y_1 + 2Y_2)) \end{array}\right]}{g''(\frac{1}{4}(Y_1 + 2Y_2))}.$$

While the sign of this expression will depend on the concavity of g, note that it is positive when $g''' = 0$ because $L > \frac{1}{2}k$. We know that $L > \frac{1}{2}k$ because if not, then adjustment would reduce both housing and food consumption in the dual-unemployment state.

Case d: Never move:

$$\frac{dh_1^*}{d\phi} = -\frac{1}{4}\frac{\left[\begin{array}{r} u_h(\frac{1}{4}(Y_1 + 2Y_2), \frac{1}{4}(Y_1 + 2Y_2) - 2L) - 2u_h(\frac{1}{4}(Y_1 + 2Y_2), \\ \frac{1}{4}(Y_1 + 2Y_2) - L + 2u_f(\frac{1}{4}(Y_1 + 2Y_2), \frac{1}{4}(Y_1 + 2Y_2) - L \\ - u_f(\frac{1}{4}(Y_1 + 2Y_2), \frac{1}{4}(Y_1 + 2Y_2) - 2L) \end{array}\right]}{4u_{ff}(\frac{1}{4}(Y_1 + 2Y_2), \frac{1}{4}(Y_1 + 2Y_2)) - u_{hf}(\frac{1}{4}(Y_1 + 2Y_2), \frac{1}{4}(Y_1 + 2Y_2))}.$$

Note that the denominator is negative. The numerator is negative when the third derivatives are positive. To put this formally, when utility is separable and of the form $u(h, f) = g(h) + g(f)$, then

$$\frac{dh_1^*}{d\phi} = \frac{df_1^*}{d\phi} = \frac{1}{4}\frac{\left[\begin{array}{r} g'(\frac{1}{4}(Y_1 + 2Y_2)) - 2g'(\frac{1}{4}(Y_1 + 2Y_2) \\ - L + g'(\frac{1}{4}(Y_1 + 2Y_2) - 2L) \end{array}\right]}{g''(\frac{1}{4}(Y_1 + 2Y_2))},$$

which is positive if and only if $g''' > 0$. As in case c, the general closed-form expression for $df_1/d\phi$ is complex and uninformative without separability.

Ranges of k. Define $v(x) \equiv u(x, x)$. Note that since utility is strictly increasing it is straightforward to define $v^{-1}(u) = x$ such that $v(x) = u$. In the limit as $p, q \to 0$, the following ranges for k characterize the optimal pattern of moving as a function of k given Y and L.

Case a: Always move:

$$k = 0.$$

Case b: Move in all but best state:

$$0 < k \leq \tfrac{1}{2}(Y_1 + 2Y_2) - L - 2v^{-1}(u(\tfrac{1}{4}(Y_1 + 2Y_2), \tfrac{1}{4}(Y_1 + 2Y_2) - L)).$$

Case c: Move only in worst state:

$$k > \tfrac{1}{2}(Y_1 + 2Y_2) - L - 2v^{-1}(u(\tfrac{1}{4}(Y_1 + 2Y_2), \tfrac{1}{4}(Y_1 + 2Y_2) - L))$$

$$k \leq \tfrac{1}{2}(Y_1 + 2Y_2) - 2L - 2v^{-1}(u(\tfrac{1}{4}(Y_1 + 2Y_2), \tfrac{1}{4}(Y_1 + 2Y_2) - 2L)).$$

Case d: Never move:

$$k > \tfrac{1}{2}(Y_1 + 2Y_2) - 2L - 2v^{-1}(u(\tfrac{1}{4}(Y_1 + 2Y_2), \tfrac{1}{4}(Y_1 + 2Y_2) - 2L)).$$

The three cut-off points are obtained by setting the indirect utility from moving equal to the indirect utility from not moving for $\tilde{Y} = 2Y_2$, $2Y_2 - L$, $2Y_2 - 2L$, respectively, and setting $h_1^* = f_1^* = \frac{1}{4}(Y_1 + 2Y_2)$.

Proof of Lemma 2: Quadratic Utility Lemma 2. *Assume utility is quadratic, so that* $u(h, f) = h - \frac{1}{2}\alpha h^2 + f - \frac{1}{2}\alpha f^2$, $\alpha > 0$:

(*a*) When k is such that adjustment is always optimal, then $\frac{dh_1^*}{d\phi} = 0$.
(*b*) When k is such that adjustment is optimal if and only if at least one spouse becomes unemployed, then $\frac{dh_1^*}{d\phi} > 0$, but approaches zero as p, q $\to 0$.
(*c*) When k is such that adjustment is optimal if and only if both spouses become unemployed, then $\frac{dh_1^*}{d\phi} > 0$.
(*d*) When k is such that adjustment is never optimal, then $\frac{dh_1^*}{d\phi} = 0$.

Proof. Calculating $\partial U(h_1, f_1)/\partial h_1 = 0$ and $\partial U(h_1, f_1)/\partial f_1 = 0$, two equations in two unknowns (h_1, f_1), and solving this pair of equations in each case yields:

Case a: Always move:

$$f_1^* = h_1^* = (Y_1 + 2Y_2 - (p + q)L - k)/4, \text{ so that } df_1^*/d\phi = dh_1/d\phi = 0.$$

Case b: Move in all but best state:

$$h_1^* = (Y_1 + 2Y_2)/4 - ((p + q)L + (p + q - \phi)^2 k)/(16 - 12p - 12q + 12\phi);$$

$$f_1^* = (Y_1 + 2Y_2 - (p + q)L - (p + q - \phi k)/4; \text{ so that}$$

$$dh_1^*/d\phi = (3(p + q)L + (p + q - \phi)(8 - 3p - 3q + 3\phi)k)/(4 - 3p - 3q + 3\phi)^2 > 0;$$

$$df_1^*/d\phi = k/4 > 0.$$

Case c: Move only in worst state:

$$h_1^* = (Y_1 + 2Y_2 - (p + q)L - k\phi)/4 + (\phi(2 - p - q)L + (1 - \phi)k)/(4 - 3\phi);$$

$$f_1^* = (Y_1 + 2Y_2 - (p + q)L - k\phi)/4, \text{ so that}$$

$$dh_1^*/d\phi = (4(2 - p - q)L + (-8\phi + 3\phi^2)k/4)/(4 - 3\phi)^2;$$

$$df_1/d\phi = -k/4 < 0.$$

$2L > k$ is sufficient (but not necessary) for $dh_1/d\phi > 0$. $2L > k$ can be shown most simply by noting that the decision in the dual unemployment state to move necessarily means having a lower level of housing consumption relative to not moving. So as not to be dominated, moving must allow for a higher level of food consumption than not moving:

$$\tfrac{1}{2}(Y_1 + 2Y_2 - 2L - k - f_1 - h_1) > Y_1 + 2Y_2 - 2L - f_1 - 2h_1. \tag{A2}$$

If $k > 2L$, then equation (A2) implies that $f_1 + 3h_1 > Y_1 + 2Y_2$, which is inconsistent with the optimal values f_1^* and h_1^*. Therefore, $2L > k$ so that $dh_1^*/d\phi > 0$.

TABLE A1.—SAMPLE CONSTRUCTION

Restriction	Number Lost	Total Remaining
Data source: IPUMS		
Original sample		2,778,194
Live in an MSA	1,016,455	1,761,767
Married	779,536	982,231
Husband and wife both age 25 or over	63,992	918,239
Listed occupations	20,499	897,740
Husband and wife both work full time	572,470	325,270
8 or fewer people in household	1,513	323,757
Not a farm household	2,318	321,439
Family income above 0 and not missing	113	321,326
Both husband and wife have income ≥ 0	1,160	320,166
Occupation not rare (contains > 200 persons/year)	17,806	302,360
Cell size for imputing probability of moving ≥ 30	185	302,175
House value or rent nonmissing and > 0	12,113	290,062
Data source: SIPP		
Original sample (person $\times$ month)		3,897,211
Married couple households $\times$ month	3,117,752	779,459
Drop extended families	160,775	618,684
Husband and wife both age 25 or over	137,147	481,537
Observe monthly change in employment status	173,786	307,751
Family income above zero and not missing	597	307,154
Observe continuous 12-month moving history		201,413

Sources: 1980, 1990, and 2000 IPUMS; April 1996 panel of the SIPP.

Case d: Never move:

$$h_1^* = f_1^* = (Y_1 + 2Y_2 - (p+q)L)/4, \text{ so that}$$
$$dh_1^*/d\phi = df_1^*/d\phi = 0.$$

APPENDIX B: Calculating the Difference in Correlations in Table 3

We use the following procedure to calculate the difference in correlation:

1. Using probit, estimate $I_{UE,i,t} = \lambda 1_{\rho,i,t} + \varsigma_Z Z_{i,t} + \varsigma_\delta \delta_t + \varsigma_\Omega \Omega_{i,t} + \nu_{i,t}$ separately for husbands and wives, where $I_{UE,i,t}$ is an indicator variable for whether a currently employed husband (wife) becomes unemployed at some point over the subsequent twelve months.
2. Estimate $\nu_{i,t}$, and predict $I_{UE,i,t}|1_{\rho,i,t} = 0$, and $I_{UE,i,t}|1_{\rho,i,t} = 1$ for each spouse (H, W).
3. Regress $\nu_{H,i,t} \times \nu_{W,i,t} = \tau 1_{\rho,i,t} + \upsilon_Z Z_{i,t} + \nu_\delta \delta_t + \upsilon_\Omega \Omega_{i,t} + \bar{\omega}_{i,t}$ using OLS.
4. Predict $\nu_{H,i,t} \times \nu_{W,i,t}|1_{\rho,i,t} = 0$, and $\nu_{H,i,t} \times \nu_{W,i,t}|1_{\rho,i,t} = 1$.
5. Compute

$$\frac{\overline{\nu_{H,i,t} \times \nu_{W,i,t}|1_{\rho,i,t} = 1}}{\sqrt{\overline{I_{UE}|1_{H,\rho,i,t} = 1}(1 - \overline{I_{UE}|1_{H,\rho,i,t} = 1})\,\overline{I_{UE}|1_{W,\rho,i,t} = 1}\,(1 - \overline{I_{UE}|1_{W,\rho,i,t} = 1})}}$$
$$- \frac{\overline{\nu_{H,i,t} \times \nu_{W,i,t}|1_{\rho,i,t} = 0}}{\sqrt{\overline{I_{UE}|1_{H,\rho,i,t} = 0}(1 - \overline{I_{UE}|1_{H,\rho,i,t} = 0})\overline{I_{UE}|1_{W,\rho,i,t} = 0}(1 - \overline{I_{UE}|1_{W,\rho,i,t} = 0})}}.$$

6. Estimate the standard errors by bootstrapping the previous steps using sampling with replacement and 200 replications.