Consumption Commitments and Risk Preferences*

Raj Chetty
UC-Berkeley and NBER

Adam Szeidl
UC-Berkeley

June 2006

Abstract

Many households devote a large fraction of their budgets to "consumption commitments" – goods that involve transaction costs and are infrequently adjusted. This paper characterizes risk preferences in an expected utility model with commitments. We show that commitments affect risk preferences in two ways: (1) they amplify risk aversion with respect to moderate-stake shocks and (2) they create a motive to take large-payoff gambles. The model thus helps resolve two basic puzzles in expected utility theory: the discrepancy between moderate-stake and large-stake risk aversion and lottery playing by insurance buyers. We discuss applications of the model such as the optimal design of social insurance and tax policies, added worker effects in labor supply, and portfolio choice. Using event studies of unemployment shocks, we document evidence consistent with the consumption adjustment patterns implied by the model.

*E-mails: chetty@econ.berkeley.edu, szeidl@econ.berkeley.edu. We thank Daron Acemoglu, George Akerlof, John Campbell, David Card, Gary Chamberlain, Martin Feldstein, Ed Glaeser, Caroline Hoxby, Larry Katz, Louis Kaplow, Matthew Rabin, Jonah Rockoff, Jesse Shapiro, two anonymous referees, and numerous seminar participants for helpful comments. James Sly provided outstanding research assistance. We are grateful for financial support from National Science Foundation grant SES 0522073.

Many households have "consumption commitments" that are costly to adjust when shocks such as job loss or illness occur. For example, most homeowners do not move during unemployment spells, and have a commitment to make mortgage payments. Consumption of many other durable goods (vehicles, furniture) and services (insurance, utilities) may also be difficult to adjust. Data on household consumption behavior show that more than 50% of the average household's budget is fixed over moderate income shocks (see section I).

This paper argues that incorporating consumption commitments into the analysis of risk preferences can help explain several stylized facts and yields a set of new normative implications. The canonical expected utility model of risk preferences does not allow for commitments because it assumes that agents consume a single composite commodity. This assumption requires that agents can substitute freely among goods at all times. When some goods cannot be costlessly adjusted, a composite commodity does not exist, and the standard expected utility model cannot be applied.

We analyze the effect of commitments on risk preferences in a model where agents consume two goods – one that requires payment of a transaction cost to change consumption (housing), and one that is freely adjustable at all times (food). Agents face wealth shocks after making a housing commitment. Under weak conditions, agents follow an (S,s) policy over housing: they move only if there is a large unexpected change in wealth. We focus on characterizing the value function over wealth for agents with a commitment, $v(W)$, which determines the welfare cost of shocks and risk preferences. In the conventional model without adjustment costs, $v(W)$ is a concave function. The introduction of commitments affects the shape of $v(W)$ relative to this benchmark in two ways.

First, commitments amplify risk aversion over small and moderate risks. Within the (S,s) band, the curvature of $v(W)$ is greater than it would be if housing were freely adjustable. Intuitively, for small or temporary shocks that do not induce moves, households are forced to concentrate all reductions (or increases) in wealth on changes in food consumption. For example, if an agent must reduce total expenditure by 10% and has pre-committed half of his income, he is forced to reduce spending on discretionary items by 20%. This concentrated reduction in consumption of adjustable goods raises the curvature of $v(W)$ within the (S,s) band. For a shock sufficiently large to warrant moving, additional reductions in wealth are accommodated by cutting both food and housing, restoring the curvature of $v(W)$ to its lower no-commitment level outside the (S,s) band.

Second, commitments generate a motive to take certain gambles. The gambling motive arises from non-concavities in $v(W)$ at the edges of the (S,s) band: the marginal utility of wealth at $S-\varepsilon$

is *lower* than the marginal utility of wealth at $S + \varepsilon$. As a result, committed agents may take bets that have large, move-inducing payoffs. Intuitively, an agent who earns an extra dollar can spend it only on food; but buying a lottery ticket gives him an opportunity to buy a better house or car, which can have higher expected utility than another dollar of food.

By changing $v(W)$ in these two ways, commitments can help resolve two basic puzzles that arise in expected utility theory. The first is that plausible degrees of risk aversion over small or moderate stakes imply implausibly high risk aversion over large-stake risks in the one-good expected utility model (Hansson 1988, Kandel and Stambaugh 1991, Rabin 2000). Commitments can potentially resolve this paradox because two distinct parameters control risk preferences over small and large risks. The second puzzle is the question of why individuals simultaneously buy insurance and lottery tickets. Friedman and Savage (1948) and Hartley and Farrell (2002) propose that non-concave utility over wealth can explain this behavior. Commitments provide microfoundations for such a non-concave utility over wealth and thus can help explain why agents repeatedly take some skewed gambles. The non-concave utility induced by commitments addresses many of Markowitz's (1952) criticisms of the Friedman and Savage utility specification because commitments create a reference point that makes the non-concavities shift with wealth levels.

Since commitments change preferences over wealth, which are a primitive in many economic problems, the model has a broad range of applications. We develop three applications in some detail in this paper, and describe several others qualitatively. In the first application, we evaluate the quantitative relevance of commitments for choice under uncertainty by performing calibrations similar to those of Rabin (2000). We find that the model can generate substantial risk aversion over "moderate stake" gambles – risks that involve stakes of roughly \$1,000-\$10,000 – while retaining plausible risk preferences over large gambles. Hence, commitments can explain high degrees of risk aversion in a lifecycle consumption-savings setting with respect to risks such as unemployment or health shocks. However, the model cannot explain risk aversion with respect to small gambles such as \$100 stakes (as documented e.g. by Cicchetti and Dubin 1994), because utility remains locally linear. Explaining such first-order risk aversion requires a kink in the utility function, as in models of loss aversion (e.g., Kahneman and Tversky 1979, Koszegi and Rabin 2005).[1]

The second application explores the normative implications of the model for social insurance policies. When agents have commitments, the marginal value of insurance can be larger for small

[1] The commitments model is not inconsistent with loss aversion, and quantifying the relative contribution of the two models over various stakes would be an interesting direction for future research.

(or temporary) shocks than large (or permanent) shocks. As a result, the optimal wage replacement rate for shocks such as unemployment may be higher than for larger shocks such as long-term disability. The mechanism underlying this result is that agents abandon commitments when large shocks occur, mitigating their welfare cost. We document evidence consistent with this mechanism by comparing households' consumption responses to small vs. large shocks using panel data. Following "small" unemployment shocks (wage income loss in year of unemployment between 0 and 33%), many households leave housing fixed while cutting food consumption significantly. However, households are more likely to reoptimize on both food and housing in response to "large" shocks (wage loss greater than 33%). Coupled with the consumption adjustment patterns documented in section I, these findings show that households deviate from ideal unconstrained consumption plans for moderate-scale shocks, supporting the basic channel through which commitments affect risk preferences.

In the third application, we illustrate implications of the model in an area that does not involve uncertainty. We introduce an endogenous labor supply decision and show that commitments make the wealth elasticity of labor supply larger in the short run relative to the long run. Commitments can therefore help explain the "added worker effect" in labor economics, as well as the small short-run labor supply responses to taxation estimated in the public finance literature.

Our model is related to several papers in the literature on consumption and durable goods, starting with the seminal work of Grossman and Laroque (1990), who analyze consumption and asset pricing in a model with a single durable good. Our analysis builds on and contributes to this literature in two ways. First, because Grossman and Laroque study a continuous time model with a smooth diffusion process for wealth, they do not explore how risk preferences vary with the size of the gamble as we do here. Second, by introducing an adjustable good, we show that commitments affect risk preferences by changing the sensitivity of adjustable consumption to shocks, an observation that is useful in calibrating the model and using it in applications.

Our analysis is also related to other recent studies that have explored two-good adjustment cost models. Flavin and Nakagawa (2003) analyze asset pricing in a two-good adjustment cost model, and find that it fits consumption data better than neoclassical models. Fratantoni (2001), Li (2003), Postlewaite, Samuelson, and Silverman (2004), Shore and Sinai (2005), and Gollier (2005) explore other implications of commitments models.

The remainder of the paper is organized as follows. The next section presents motivating evidence for the model by examining how consumption adjustment patterns vary across goods.

Section II develops the model and characterizes risk preferences. Section III presents applications, and Section IV concludes. All proofs are given in the appendix.

I Motivating Evidence

We begin by documenting some facts about consumption adjustment patterns at the household level to motivate the model developed below. Table 1 shows expenditure shares for broad consumption categories such as housing and food using data from the Consumer Expenditure Survey. Our goal is to classify each of these categories as "commitment" or "adjustable," in order to estimate the share of consumption that is fixed over moderate-scale shocks. This estimate is useful in calibrating the model and gauging the extent to which commitments amplify risk aversion.[2]

Frequency Classification. A simple classification method is to define categories whose consumption is infrequently reduced as commitments. Let g_{cit} denote the growth rate of consumption of category c by household i from year $t-1$ to year t. Figure 1 plots histograms for the distribution of g_{cit} for several consumption categories. We use food and housing data from the PSID (1968-1997), following the conventions established by Zeldes (1989) and Gruber (1997) in defining these variables. The data for the remaining categories is from the Consumer Expenditure Surveys (1990-1999). See the data appendix for details on the data construction and definitions.[3]

For nondurable goods and services, g_{cit} can be computed simply as the log change in reported expenditures from one year to the next. Figures 1a-1c show the resulting histograms for changes in nominal expenditure on food, entertainment, and health insurance. Measuring changes in consumption of housing and other durable goods requires a different approach because the data do not give measures of the flow consumption value of these goods. For housing, we first determine whether a household moved from year $t-1$ to year t. If it did not, we define the change in housing consumption as zero. For households that did move, we define the change in housing consumption as the log change in nominal rent for renters and the log change in the nominal market value of

[2] Calibrating the model requires a measure of the flow *consumption* share of commitments rather than the *expenditure* shares reported in Table 1. While data on flow consumption is unavailable for durables, aggregate expenditure shares in a representative sample are approximately equal to average consumption shares in a steady state economy with constant growth, depreciation, and interest rate. This fact allows us to interpret the shares in Table 1 as average consumption shares.

[3] We use housing data from the PSID because the CEX drops movers. We use food data from the PSID for consistency with the event-study analysis below (results from the CEX data are similar). Since the PSID gives data on only food and housing consumption, we use CEX data to classify the remaining categories.

the houses for homeowners.[4,5] We use an analogous approach in the CEX to define changes in car consumption (see data appendix for details).

For other durable goods, namely apparel and furniture/appliances, we do not have a measure of initial value of the durable stock in the data. For these categories, we analyze the *level* change in consumption instead of growth rate of consumption. We define the level change in consumption as net expenditures (purchases minus sales) over a one year period in real 2000 dollars. Note that households can actively reduce their consumption of these durables (beyond natural decline due to depreciation) only by selling these durables in the secondary market, i.e. by having negative net expenditures. Figure 1f shows the distribution of level changes in furniture consumption using this methodology.

Figure 1 indicates that food and entertainment are easily adjustable goods since their distributions are quite dispersed. In contrast, housing is infrequently adjusted, consistent with the large transaction costs inherent in changing housing consumption (e.g. broker fees, monetary and utility costs of moving). Similarly, consumption of durable goods such as cars and furniture is infrequently adjusted, particularly downward, perhaps because market resale values are significantly lower than actual values, creating an adjustment cost. Finally, the histogram for health insurance indicates that commitments extend beyond durables to services that involve contracts and penalties for early termination.[6]

Figure 1 suggests a natural statistic for classifying categories into the "commitment" and "adjustable" bins – the fraction of households that report an active reduction in consumption. A category whose consumption is difficult to cut is more of a commitment. Table 1 reports this statistic for the consumption categories in the CEX, following the methodology illustrated in the histograms. Using a cutoff of 33% on the frequency of downward adjustments, the first five consumption categories fall into the "commitment" category. The frequency of adjustment measure thus implies that commitments comprise approximately 50% of the average household's budget.

Wealth Shock Classification. A drawback of the preceding classification method is that it does

[4] We use this methodology because fluctuations in rents or housing values due to asset price variation do not correspond to changes in housing consumption for non-movers. While this method misses changes in housing consumption among non-movers due to remodelling, reductions in housing consumption through remodelling are presumably infrequent.

[5] For simplicity, we do not compute housing consumption growth rates for individuals who switch from owning to renting or vice versa (5.5% of observations). Since omitting these observations understates the frequency of moves, we correct the fraction of zeros in the histogram (Figure 1b) to correct for this bias. See the data appendix for details.

[6] Late payments do not eliminate commitments. In the event of an income shock, the individual must still eventually pay the bill. The ability to make late payments is simply an additional credit channel.

not shed light on how consumption is adjusted in response to shocks. Examining the response of consumption to wealth shocks is important because the key concept underlying our model is that some portions of the household's consumption bundle remain fixed when moderate shocks occur.

To provide such evidence, we conduct an event study analysis of unemployment shocks using data from the PSID spanning 1968 to 2003. In particular, we compare the food and housing consumption responses of homeowners and renters to examine whether higher adjustment costs for housing lead to less adjustment on that margin. We focus on heads of household between the ages of 20 and 65 who report one unemployment spell during the panel. We construct event-study graphs by normalizing the year of unemployment as 0 for all individuals, and defining all other years relative to this base year (e.g., the year before the shock is -1, the year after the shock is +1). Observations with changes in the number of people in the household are excluded to avoid introducing noise in the consumption measures because of changes in household composition.

Figure 2a plots mean real annual growth rates of food and housing consumption, defined as above, for individuals who rented in year -1.[7] In the year of job loss, food consumption falls from its pre-unemployment level by 6.7%, while housing consumption falls by 4.3%. The fraction of renters who move rises from 39.4% in year -1 to 45.3% in year 0. Renters who do not move reduce food consumption by 5.7% on average, indicating that many renters reduce food but not housing consumption in response to the shock. Figure 2b shows analogous figures for individuals who were homeowners in year -1. Homeowners reduce food consumption by 9.1% in year 0, but reduce housing consumption by less than 0.1% on average. The fraction of homeowners who move rises from 12.8% in year -1 to only 14.2% in year 0, indicating that homeowners' consumption bundles become more distorted in terms of food-housing composition than renters' bundles during unemployment spells.[8] These findings support the claim that adjustment costs induce households to concentrate expenditure reductions on certain goods when moderate income shocks occur.

The event studies suggest a second, more general definition of "commitments" – categories that are insensitive to moderate income shocks. Unfortunately, since the PSID contains data on only food and housing, the effect of unemployment shocks on consumption of other goods cannot be estimated. However, we expect that more goods would be classified as commitments under this

[7] These figures omit individuals who switch from owning to renting or vice versa. In Chetty and Szeidl (2006), we show that including these switchers by computing the rental value of owning a house using Himmelberg et. al.'s (2005) user cost of housing series does not affect the results. We also show that the results are robust to the inclusion of controls and several other specification checks.

[8] Note that the magnitude of the food consumption drop is not much larger for homeowners than renters presumably because homeowners are wealthier on average, and thus experience a smaller effective shock.

wealth-shock definition than under the frequency of adjustment definition used above. For example, while utility bills fluctuate over time as temperatures and energy costs change, it may be difficult to voluntarily change heating expenditures sharply in response to a shock if one does not move. Similarly, consumption of education and gasoline for vehicles may also be relatively unresponsive to moderate shocks despite being volatile over time because of variation in input costs, lifecycle patterns, etc. While utilities, fuel costs, and education do not involve explicit adjustment costs, they have a commitment-like property in that they may be difficult to adjust and force additional adjustment on other margins. Hence, in our stylized two-good model, these categories affect risk aversion in the same way as housing and other goods with adjustment costs. Including these goods as "commitments" suggests an effective commitment share closer to 65%.

Irrespective of the details of the method used to define commitments, the key point is that a substantial portion of consumption is inflexible. The next section explores the implications of this observation for risk preferences.

II Theory

II.A Setup

Consider a household that lives for T periods and consumes two goods: adjustables, such as food (f), and commitments, such as housing (x). Adjusting commitment consumption that provides x units of services per period requires payment of an adjustment cost $k \cdot x$ where $k > 0$. The household chooses a consumption path to maximize the expected value of lifetime utility:

$$\mathrm{E}_0 \sum_{t=1}^{T} u(f_t, x_t)$$

where period utility $u(f, x)$ is strictly increasing in both arguments, strictly concave, and three times continuously differentiable. For simplicity, we assume that the household does not discount future utility and that the interest rate is zero.

The household has a risky income stream denoted by y_t. Let W_t denote wealth at the beginning of period t. The change in wealth, $W_{t+1} - W_t$, is determined by income minus expenditure on current food and housing consumption, less moving costs kx_{t-1} in the event that the household

decides to move.[9] The household's dynamic budget constraint is

$$W_{t+1} = W_t + y_{t+1} - f_t - x_t - kx_{t-1} \cdot 1\{x_{t-1} \neq x_t\} \qquad \text{for } 1 \leq t \leq T \tag{1}$$

together with the terminal condition $W_{T+1} = 0$. The timing of the household's problem is as follows. In period zero, the household selects a house x_0 to maximize expected utility given initial wealth W_0 and its knowledge of the distribution of the future path of income. The household begins consuming f and x in period 1, which it enters with wealth $W_1 = W_0 + y_1$ and commitment x_0. In period 1 and all subsequent periods, consumption of f_t and x_t is chosen to maximize expected utility, taking into account the cost of adjusting prior commitments.

Note when choosing x_0 in period 0, the household is aware of the commitment nature of the good and takes into account its effects on the welfare cost of shocks in subsequent periods. We focus here on characterizing consumption behavior and risk preferences in period 1, after the commitment choice has been made. From the perspective of period 1, x_0 is a state variable that is effectively treated as exogenous by the household. In this paper, we do not characterize the optimal policy for x_0. Shore and Sinai (2005) explore this decision, and show that an increase in risk can either increase or reduce x_0, depending on the distribution of the risk.

We make two assumptions to simplify the analysis of the household's problem. First, we abstract from capital market imperfections by allowing the household to borrow against future income. Second, we assume that all uncertainty in the economy is resolved in period 1: in particular, the household learns the entire realization of $\{y_t\}_{t=1}^T$ at the beginning of period 1. We discuss the implications of relaxing these assumptions below.

II.B Consumption Behavior

The optimal consumption policy in period 1 is governed by two state variables: prior commitments x_0 and lifetime wealth $W = W_0 + \sum_{t=1}^T y_t$. Since the discount factor and interest rate are zero, and all uncertainty is resolved in period 1, the optimal consumption path is deterministic and flat: $f_t = f_1$ and $x_t = x_1$ for all $t = 1, 2..., T$. As a result, if the household ever moves, it moves in period 1.[10]

[9] While the consumption adjustment patterns documented above suggest that adjustment costs are larger for reductions than increases, we assume a symmetric adjustment cost for simplicity. Permitting asymmetric adjustment costs would not affect the main qualitative results.

[10] Since consumption is flat in periods $1, ..., T$, many of our results can also be presented in a two-period model. However, allowing for T periods is useful when we introduce borrowing constraints below.

The decision to move in period 1 depends on a trade-off between the benefit of having the optimal bundle of food and housing consumption and the transaction cost associated with reaching that bundle. To characterize this decision formally, let $v^c(W, x_0)$ denote the household's value function in period 1. Then

$$v^c(W, x_0) = \max\left(v^0(W, x_0), v^m(W, x_0)\right) \tag{2}$$

where $v^0(W, x_0)$ is maximized utility conditional on never moving, and $v^m(W, x_0)$ is the maximized utility of a household who moves in period 1. The optimal consumption choice of the household has a simple analytical solution if the utility function satisfies the following conditions.

A1 Limit properties of utility: $\lim_{f\to\infty} u_1(f, x) = \lim_{x\to\infty} u_2(f, x) = 0$; and $\lim_{f\to 0} u(f, x) = \inf_{f',x'} u(f', x')$ for all x.

A2 The marginal utility of food is nondecreasing in housing consumption: $u_{1,2}(f, x) \geq 0$.

Both of these conditions are satisfied for a wide class of utility functions, including (1) the constant elasticity of substitution (CES) specification when the elasticity of substitution is not greater than 1; and (2) separable power utility over the two goods as long as the coefficient of relative risk aversion for food is 1 or higher. Under these conditions, the agent's consumption policy in period 1 can be written as an (S,s) rule:

Lemma 1 *Under assumptions A1 and A2, for each $x_0 > 0$ there exist $s < S$ such that*

(i) when $W \in (s, S)$, the optimal policy is not to move:

$$x_t = x_0 \quad \textit{and} \quad f_t = W/T - x_0.$$

(ii) when $W \notin (s, S)$, the optimal policy is to move, and

$$(f_t, x_t) = \arg\max\left\{u(f, x) \mid f + x = (W - kx_0)/T\right\}.$$

(iii) when k increases, s falls and S increases.

The intuition underlying this result is straightforward. Suppose the agent experiences a large negative wealth shock in period 1. Then he will be forced to reduce food consumption drastically in order to maintain the housing commitment that he previously made. Since such a sharp reduction

in food consumption causes a large reduction in utility, it becomes optimal to pay the adjustment cost and move into a smaller home. Conversely, if wealth rises sharply, rather than allocating all of the extra wealth to food, whose marginal utility eventually diminishes to zero, it is preferable to pay the adjustment cost and upgrade to a large house. For smaller shocks, the utility gain from fully reoptimizing the consumption bundle is insufficient to offset the transaction cost, so there is an (S,s) band where the agent does not move.

Assumptions A1 and A2 are useful in obtaining the (S,s) result because they guarantee that the agent moves for large shocks. A1 ensures that it is preferable to move to a smaller house than to cut food consumption to zero. To understand the role of A2, consider the case where f and x are perfect substitutes. In that case, A2 is violated, and the household would never move, because housing and food are interchangeable. Note that empirical evidence confirms that households follow (S,s) policies for goods that involve adjustment costs (Eberly (1994), Attanasio (2000), Martin (2003)), suggesting that violations of assumptions A1 or A2, if any, are not large enough to cause deviations from the intuitive (S,s) policy in practice.

II.C Risk Preferences

The welfare cost of shocks, and thus the household's risk preferences, are determined by the shape of the value function $v^c(W, x_0)$ in period 1. We formally characterize this function in a series of steps. Before proceeding, it is helpful to introduce some notation. Let $v^n(W)$ denote the value function of a hypothetical agent who pays no adjustment costs for housing, and let $\gamma^n(W) = -v^n_{WW}W/v^n_W$ represent the period 1 coefficient of relative risk aversion (CRRA) over wealth for this agent. Let $\gamma^c(W)$ represent the analogous parameter for an agent with commitments.[11] Let $\gamma_f = \frac{-u_{11}f}{u_1}(f, x)$ represent the CRRA over food, and define $\varepsilon_{u_f,x} = \frac{u_{12}x}{u_1}$ as the elasticity of $u_1(f, x)$ with respect to x. Note that both γ_f and $\varepsilon_{u_f,x}$ are pure preference parameters: they depend on the level of f and x only, and not on the presence or absence of adjustment costs.

Let $f^n(W)$ and $x^n(W)$ represent the optimal consumption choices in period 1 as a function of lifetime wealth W for a consumer who faces no adjustment costs. Let $\varepsilon^n_{f,W}$ denote the elasticity of food consumption with respect to wealth W in period 1 for this consumer and $\varepsilon^c_{f,W}$ the corresponding elasticity in the presence of adjustment costs. Define the wealth elasticities of housing $\varepsilon^n_{x,W}$ and $\varepsilon^c_{x,W}$ analogously.

[11] Throughout, we use the convention that a superscript c refers to the presence of adjustment costs while a superscript n refers to the case of no adjustment costs.

Effect of commitments on CRRA. The CRRA over wealth in the benchmark case of a consumer who faces no adjustment costs is

$$\begin{aligned}\gamma^n(W) &= \frac{-v^n_{WW}W}{v^n_W} = -W\frac{u_{11}\frac{\partial f}{\partial W} + u_{12}\frac{\partial x}{\partial W}}{u_1} \\ &= \gamma_f \varepsilon^n_{f,W} - \varepsilon_{u_f,x}\varepsilon^n_{x,W}. \end{aligned} \tag{3}$$

The intuition for this expression is as follows. By the envelope theorem, marginal utility over wealth (v_W) equals the marginal utility of food (u_1) at the optimum. Hence the coefficient of relative risk aversion, defined as the elasticity of v^n_W with respect to W, also equals the elasticity of u_1 with respect to W. The marginal utility of food itself changes with W for two reasons, which correspond to the two terms in (3): first, because increased food consumption reduces u_1 and second, because increased housing consumption affects u_1 through complementarity between the two goods.

For a consumer with adjustment costs, a similar calculation gives

$$\gamma^c(W, x_0) = \gamma_f \varepsilon^c_{f,W} \tag{4}$$

as long as $W \in (s, S)$. Note that $\varepsilon_{x,w}$ does not appear in (4) because within the (S, s) band, housing does not change with W.

In general, the terms in (3) and (4) are not directly comparable, because γ_f is evaluated at different values of (f, x) in the two cases. However, at the wealth level $\overline{W}$ implicitly defined by $x^n(\overline{W}) = x_0$, the optimal choice of x and f is the same with and without adjustment costs.[12] At this wealth level, we can compute an exact value for the change in risk aversion due to commitments:

$$\gamma^c(\overline{W}, x_0) - \gamma^n(\overline{W}) = \gamma^f \cdot \left[\varepsilon^c_{f,W} - \varepsilon^n_{f,W}\right] + \varepsilon_{u_f,x}\varepsilon^n_{x,W} > 0. \tag{5}$$

Equation (5) captures the key intuition of our model. It shows that commitments magnify risk aversion for two reasons, corresponding to the two non-negative terms on the right hand side. The first term arises from the fact that wealth shocks are borne solely on the food margin when housing consumption is fixed. This makes the elasticity of food consumption with respect to wealth larger in the case with commitments. Consequently, the marginal utility of food, and hence the marginal utility of wealth, rises more quickly as wealth falls when the agent has commitments. More

[12] Such a $\overline{W}$ exists because A2 ensures that $x^n(W)$ is increasing (Chipman, 1977), $x^n(W)$ is continuous, $x^n(0) = 0$ and by A1 $x^n(W) \to \infty$ as $W \to \infty$.

concretely, consider an individual who loses his job and cuts back sharply on food or entertainment expenditures to pay the mortgage and other bills. The welfare cost of unemployment is larger for this individual because he cannot reduce commitment consumption easily, forcing him to make larger reductions on adjustable goods.

The second term in (5) shows that commitments magnify risk aversion further in the presence of complementarity ($u_{12} > 0$). Without adjustment costs, the optimal response to a reduction in wealth is to reduce both x and f. When $u_{12} > 0$, the reduction in x reduces the marginal utility of f, cushioning the effect of the cut in f on marginal utility. When x cannot be adjusted, this cushioning effect is shut down, and a given drop in f has a larger impact on marginal utility. More concretely, suppose the adjustable and commitment goods are electricity and housing. A given reduction in electricity (used for heating) has a larger cost in terms of marginal utility when the agent cannot simultaneously reduce housing consumption, magnifying risk aversion.

A Special Case. The following specification of utility yields a simple expression that is useful in calibrating the effect of commitments on risk aversion:

$$u\left(f,x\right)=\frac{f^{1-\gamma_f}}{1-\gamma_f}+\mu\cdot\frac{x^{1-\gamma_x}}{1-\gamma_x}. \tag{6}$$

From (3), the ratio of curvatures over wealth with and without adjustment costs at $\overline{W}$ is

$$\frac{\gamma^c}{\gamma^n}=1+\frac{x}{f}\frac{\gamma_f}{\gamma_x}. \tag{7}$$

Equation (7) shows that the commitment share of the budget is a key factor in determining how much commitments magnify risk aversion over small shocks. When commitments constitute a higher share of expenditures, shocks are concentrated on a smaller set of goods, and risk aversion is higher. When $\gamma_f > \gamma_x$, the consumer is particularly risk averse over adjustable goods, increasing the amplification effect.

Curvature of the value function. The amplified risk aversion result in (5), which was derived for wealth level $\overline{W}$, can be extended to all wealth levels in the (S,s) band by imposing restrictions on how the curvature of $u(f,x)$ varies over a range of consumption bundles. In particular, we assume that the utility function satisfies one of the following conditions.

A3 $u(f,x)$ is homogenous of some degree $1-\gamma$.

A4 $u(f,x)$ is separable, γ_f is constant and $\sup_x \gamma_x(x) < \gamma_f$.

These assumptions allow for most common specifications of preferences over two goods, including CES utility and separable power utility (6) as long as $\gamma_f \geq \gamma_x$. The following proposition characterizes the curvature of $v(W, x_0)$ at all wealth levels.

Proposition 1 *Assume that A1 and A2 hold. Then,*

(i) when either A3 or A4 holds, commitments magnify risk aversion in the (S,s) band: $\gamma^c(W, x_0) > \gamma^n(W)$ *for all* $W \in (s, S)$.

(ii) when either A3 or A4 holds, risk aversion is higher inside the (s, S) *band than outside:* $\gamma^c(W) > \gamma^c(W')$ *for all* $W \in (s, S)$ *and* $W' \notin (s, S)$ *such that* $W' > kS/T$.

(iii) agents have a gambling motive: $v^c(W, x_0)$ *is locally convex at* $W = s$ *and* $W = S$, *and hence there exist wealth levels* W *and random variables* $\widetilde{Z}$ *with* $E\widetilde{Z} = 0$ *such that*

$$Ev^c(W + \widetilde{Z}, x_0) > v^c(W, x_0).$$

The results of this proposition are illustrated in Figure 1a, which plots the value functions $v^c(W, x_0)$ and $v^n(W)$. Part (i) extends the result in (5) by establishing that commitments increase the curvature of the value function at all $W \in (s, S)$. Part (ii) of the proposition shows that a consumer with commitments is more risk averse inside the (S,s) band than outside. This result follows closely from (i), because outside the (S,s) band the consumer has abandoned prior commitments and thus has preferences similar to those of a consumer who faces no adjustment costs. One difference between the two consumers is that the one who abandoned prior commitments has effective wealth reduced by the transaction cost kx_0 relative to the consumer with no adjustment costs. At small wealth levels, this reduction in wealth can significantly increase the CRRA. Hence, (ii) holds only at wealth levels that are not too small relative to the transaction cost of moving, explaining the condition $W' > kS/T$.

Part (iii) shows that commitments create an incentive to take certain zero expected-value gambles in period 1. The gambling motive arises because the value function is locally convex at s and S: as shown in Figure 1a, the one-sided derivatives of v^c satisfy $v_1^c(W-, x_0) < v_1^c(W+, x_0)$ at $W = s$ and $W = S$. Since the value function is strictly concave within the (S,s) band, fair gambles that have payoffs within the (S,s) band are always rejected. Agents only pursue gambles which have payoffs that would make it optimal to drop prior commitments. To see why such gambles can be attractive, consider an individual who is deciding whether to buy a candy bar that costs \$1 or a

fair lottery ticket for \$1 that will pay \$1 million if he wins. A one good (no commitments) model assumes that the agent will buy one million candy bars if he wins the lottery (or one million units of the composite commodity). In this case, buying the lottery ticket is not optimal because the marginal utility of candy is diminishing, and the agent would be better off getting one candy bar with certainty. However, with commitments, the agent will buy more than just candy if he wins the lottery. While the \$1 in hand cannot be spent to buy a better house or car, the \$1 million can. Consequently, the expected utility of the skewed lottery may exceed the utility of the candy bar.

Borrowing constraints and persistent uncertainty. We now explore the implications of relaxing the two main assumptions made above. First consider the effect of borrowing constraints. Suppose that the consumer has access to a savings technology but can never borrow. To simplify further, assume that the income profile is flat starting in period 2, i.e., $y_2 = y_3 = ... = y_T = y$. This implies that the borrowing constraint will not bind in periods 2 through T; however, it may bind in period 1 if the realization of y_1 is low enough.

Figure 1b illustrates the effect of the borrowing constraint on $v(W, x_0)$. For high realizations of y_1, the borrowing constraint has no effect on the value function, because the consumer does not want to borrow. However, once y_1 is low enough, wealth falls below the point W^{BC} at which the constraint starts to bind. At this point, the inability to borrow forces the consumer to concentrate additional reductions in expenditure solely on period 1 food consumption, instead of splitting the reduction over T periods. This argument suggests that for wealth realizations $W < W^{BC}$, the difference in risk aversion between the commitment and no-commitment cases ($\gamma^c - \gamma^n$) is increased by a factor of at least T when the consumer does not move.[13] In the appendix, we establish this result formally under some regularity conditions on $u(f, x)$.

To see more concretely why commitments have a larger effect on risk aversion when agents are borrowing constrained, consider the example of a short unemployment spell. Individuals without borrowing constraints can smooth the income loss from this shock over several years. Hence, their risk aversion with respect to such a temporary shock is small regardless of whether they have commitments or not. In contrast, individuals who cannot borrow are forced to cut expenditure sharply while unemployed. This makes the potential welfare gain from relaxing commitments much greater for them.

[13] The optimal moving policy with borrowing constraints is more complex when $W < W^{BC}$ than in the basic model. Following a large negative shock in period 1, the consumer may decide to move into a small house permanently, or to set a small x_1 for one period only, and then move again in period 2. In the interest of space, we do not characterize the dynamics of housing consumption with borrowing constraints in more detail here.

We now briefly discuss what happens when the consumer faces additional uncertainty in future periods. In this case, (3) and (4) remain valid expressions for the CRRA (without borrowing constraints). However, the key elasticity $\varepsilon^c_{f,W}$ responsible for magnified risk aversion may be reduced. This is because future risk provides additional margins, namely future commitment consumption, that the consumer can adjust in response to a shock. For example, suppose there is a high probability of a move-inducing shock in period 2. In this case, commitments have a smaller impact on risk aversion in period 1, because future commitment consumption is effectively adjustable, and only current commitment consumption is not adjusted in response to the shock. While the commitment effect may be mitigated to some extent with future uncertainty, the option of future adjustment becomes less relevant with borrowing constraints, and risk aversion could still be substantially amplified.

II.D Welfare costs of non-negligible risks

Our analysis thus far has focused on characterizing the coefficient of relative risk aversion, which measures preferences over infinitesimal risks. Since most risks of interest are not infinitesimal in practice, we now analyze the effect of commitments on preferences over non-negligible risks.

To model non-negligible risks, let $\widetilde{W}$ denote a random variable which represents the realization of lifetime wealth in period 1. Assume that $\widetilde{W} > kS/T$ with probability one and that $\widetilde{W}$ has a finite expected value $\mathrm{E}\widetilde{W}$. Let $W^{CE}(\widetilde{W}, x_0, k)$ denote the certainty equivalent, defined such that an agent with adjustment cost k and prior commitment x_0 is indifferent between $\widetilde{W}$ and a sure payment of $W^{CE}(\widetilde{W}, x_0, k)$ at the beginning of period 1. The welfare cost of the risk $\widetilde{W}$ can be measured by the proportional risk premium $\pi(\widetilde{W}, x_0, k)$, which is defined by the equation $\mathrm{E}\widetilde{W} \cdot \left(1 - \pi(\widetilde{W}, x_0, k)\right) = W^{CE}(\widetilde{W}, x_0, k)$.

To compare risk aversion over gambles of different sizes, we introduce the concept of "equivalent relative risk aversion." Consider a gamble $\widetilde{W}$ for which a consumer with commitments has certainty equivalent $W^{CE}(\widetilde{W}, x_0, k) < \mathrm{E}\widetilde{W}$. Let $\gamma\left(\widetilde{W}, x_0, k\right)$ denote equivalent relative risk aversion, defined such that an agent with power utility preferences and CRRA equal to $\gamma\left(\widetilde{W}, x_0, k\right)$ would be indifferent between $\widetilde{W}$ and a sure payment of W^{CE}. This definition implies that for gambles similar to $\widetilde{W}$, the agent with commitments has essentially the same risk preferences as one with no commitments and a constant relative risk aversion utility with curvature $\gamma\left(\widetilde{W}, x_0, k\right)$.[14] Equivalent

[14] When $W^{CE}(\widetilde{W}, x_0, k) > \mathrm{E}\widetilde{W}$ so that the consumer is risk loving over $\widetilde{W}$, we define $\gamma\left(\widetilde{W}, x_0, k\right) = 0$. The appendix shows that equivalent relative risk aversion is well-defined for all risks considered here.

relative risk aversion can be thought of as an approximate measure of the risk premium per unit of risk:

$$\gamma(\widetilde{W}, x_0, k) \approx 2\frac{-\log\left(1-\pi(\widetilde{W}, x_0, k)\right)}{\text{var}[\log \widetilde{W}]}. \tag{8}$$

For small risks, the validity of (8) follows from the Arrow-Pratt approximation (Arrow 1965, Pratt 1964). For large gambles, (8) is an approximation, which is exact for the class of lognormally distributed risks.

To study the effect of commitments on the welfare costs of moderate vs. large-stake risks, we characterize how equivalent relative risk aversion varies over a sequence of risks of increasing size. First, define $\widetilde{W}$ to be a "moderate risk" with respect to k and x_0 if $W^{CE}(\widetilde{W}, x_0, k) \in (s, S)$. That is, $\widetilde{W}$ is a moderate risk if its certainty equivalent does not induce the household to move.[15] Now consider a sequence of risks $\widetilde{W}_\sigma$ with a common expected value $\text{E}\widetilde{W}$, indexed by their standard deviation $\sigma \geq 0$.[16] We label this a sequence of increasing risks with respect to x_0 and k if (1) $\lim_{\sigma\to\infty}\text{Prob}[W_\sigma \in (s, S)] = 0$ so that for σ large, the risks assign close to full probability to the region outside the (S,s) band; (2) $\lim_{\sigma\to\infty} \sup W^{CE}\left(\widetilde{W}_\sigma, x_0, k\right) < s$ so that for σ large the risks are not moderate in the sense defined above; and (3) as $\sigma \to 0$, the third absolute central moment of $\widetilde{W}_\sigma$ is of smaller order than σ^2. The last condition (Pratt 1964) ensures that for σ small, the tails of the distribution do not dominate the distribution's shape. This condition holds for most commonly used families of distributions, including the normal, the lognormal, the exponential and the uniform.

Proposition 2 *Assume that A1 and A2 hold.*

(i) Adjustment costs magnify moderate-stake risk aversion: For any moderate risk $\widetilde{W}$ with respect to k_1 and x_0, and any $k_2 > k_1$, $\pi(\widetilde{W}, x_0, k_1) \leq \pi(\widetilde{W}, x_0, k_2)$.

(ii) When either A3 or A4 holds, for any sequence of increasing risks, equivalent relative risk aversion is greater for small risks than large risks: $\lim_{\sigma\to 0}\gamma\left(\widetilde{W}_\sigma, x_0, k\right) = \gamma^c(E\widetilde{W}, x_0)$ *while* $\lim_{\sigma\to\infty}\sup\gamma\left(\widetilde{W}_\sigma, x_0, k\right) < \gamma^c(E\widetilde{W}, x_0)$.

This proposition formalizes the main result of the paper: Consumption commitments magnify

[15] Note that whether a shock of given size is "moderate" depends on the consumer's expected location in the (S,s) band. If the consumer is very close to the edge, even small shocks may fail to qualify as "moderate" as they may induce moves.

[16] In this exercise, we vary the distribution of $\widetilde{W}$ while holding x_0 fixed. This is consistent with x_0 being endogenously chosen in period zero if the consumer learns about the true distribution of $\widetilde{W}$ after x_0 is set. In that case, the consumer has a prior over possible distributions $\widetilde{W}$ at date zero.

the welfare cost of moderate risks (risks that usually do not induce households to abandon commitments) relative to large risks. Part (i) shows that the risk premium demanded for moderate shocks is greater for households that are more committed in the sense of having higher adjustment costs for the commitment good. Part (ii) shows that households with commitments require greater compensation per unit of risk for small shocks than they do for large shocks. Both results in the proposition arise from the intuition that households with commitments are forced to deviate from ideal consumption plans when faced with moderate shocks.

The results are illustrated in Figure 4, which plots $\gamma(.)$ as a function of the standard deviation (σ) for a family of truncated lognormal risks for two agents with different adjustment costs ("renters" and "homeowners"). Consistent with the proposition, the figure shows that equivalent relative risk aversion is high for small risks, but asymptotes to a low level as the size of risk increases. In contrast, for a consumer with power utility and no commitments, equivalent relative risk aversion is flat by definition. The figure also shows that the consumer who faces higher adjustment costs is more risk averse, consistent with part (i) of the proposition.

III Applications

III.A Moderate-Stake Risk Aversion

As noted above, an important puzzle in expected utility theory is that plausible degrees of risk aversion over small or moderate stakes imply implausibly high risk aversion over large-stake risks. In this section, we assess the potential of the commitments model to resolve this puzzle using calibrations similar to Rabin's (2000). In particular, we compute numerically a range of "moderate stake" shocks over which the model can explain high degrees of risk aversion.

The puzzle is illustrated in Table 2, which has a structure analogous to Table 1 in Rabin (2000). While Rabin gives calibration results that do not rely on functional form assumptions about utility, we illustrate the puzzle here by assuming that agents have constant relative risk aversion preferences for simplicity. Consider such an agent who, at an initial wealth level of \$300,000, is indifferent between accepting and rejecting a 50-50, lose \$1000, gain \$$g$ gamble, where g = \$1025, \$1050, \$1100, or \$1250. In columns 1, 3, 5 and 7 of the table, we report G values such that the same agent is indifferent between accepting a 50-50, lose \$$L$, gain \$$G$ gamble for various levels of L. For example, column 3 shows that if the individual is indifferent between a 50-50 lose \$1000/gain \$1050 gamble, he will also be indifferent between an 50-50 lose \$15,000/gain \$98,027 gamble.

As in Rabin's table, an entry of ∞ means that there is no finite gain that would make the individual indifferent between accepting and rejecting the particular gamble. The fact that $G = \infty$ for many large values of L in columns 1,3, 5, and 7 illustrates that plausible risk preferences for gambles involving small or moderate stakes ($L < \$10,000$) generate implausibly high levels of risk aversion for gambles over higher stakes in the conventional no-commitments model.

The remainder of the table demonstrates that commitments can help resolve this puzzle by breaking the link between risk preferences over moderate and large stakes. Columns 2,4, 6 and 8 show the risk preferences of a consumer with commitments who has period utility given by the special case in (6). The utility function is calibrated such that the agent (1) has a commitment budget share parameter of 50%, consistent with the evidence in section I and (2) is indifferent between accepting and rejecting an 50-50, lose \$1000, gain g gamble, where g = \$1025, \$1050, \$1100, \$1250 varies across the columns. The parameters of the utility function used in the calibrations are described in the notes to Table 2.

The effect of commitments on risk preferences can be seen in two ways. First, compare adjacent pairs of odd and even columns, for example columns 3 and 4. For gambles where $L < \$10,000$, the agents with and without commitments have virtually identical risk preferences. However, for gambles where $L = \$15,000$ or higher, the consumer with commitments exhibits manifestly lower risk aversion. For instance, the agent with commitments is indifferent between a 50-50 lose \$40,000/ gain \$68,897 gamble, whereas no finite gain is sufficient to compensate for a \$40,000 loss for the agent with standard power utility.

Now compare the four commitment consumers, across columns 2, 4, 6 and 8. This comparison shows that higher risk aversion over moderate stake risks need not translate into higher risk aversion for large gambles for agents with commitments. For example, the moderate loss of \$5,000 necessitates an offsetting gain of only \$5,695 for the consumer in column 2, while it requires a much higher gain of \$14,879 for the consumer in column 8. Yet for large gambles the risk preferences of these two consumers are almost identical: a loss of \$25,000 requires a gain of \$43,003 in column 2 and a gain of \$44,746 in column 8.

These two results arise from the fact that risk preferences over small and large risks are controlled by distinct parameters when agents have commitments. Risk aversion over small risks is essentially determined by the curvature over adjustable consumption (food) and the commitment share, while risk aversion over large risks is determined by curvature over commitment consumption (housing). In response to a large enough loss or gain, the agent drops his commitments and re-optimizes,

absorbing much of the shock by changing commitment consumption. As a result, agents with commitments can have high degrees of risk aversion over moderate stakes because they are forced to concentrate adjustments on food, while having lower large-stake risk aversion because they can adjust on other margins over which utility is less curved. In contrast, agents without commitments can only be risk averse over moderate stakes if they have very rapidly diminishing marginal utility over all goods. This rapidly diminishing marginal utility necessarily translates into implausible risk aversion over large stakes.

For a consumer with commitments, high levels of risk aversion over moderate stakes (as in the columns on the right side of Table 2) require high levels of curvature over food, γ_f. The table shows the calibrated γ_f parameter implied by the risk preferences in each column. One may be concerned that the γ_f values reported are high relative to familiar CRRA values between 1-5. It is important to note, however, that the criterion typically used to evaluate whether a particular value for the CRRA is plausible is to look at implied behavior over risky choices (as in Mankiw and Zeldes, 1991, Barsky et. al. 1997). As the table demonstrates, a CRRA over food in excess of 10 can indeed lead to plausible risk preferences in the commitments framework. The reason that such values of CRRA over food generate risk preferences consistent with introspection is that they apply only for moderate shocks. For larger shocks, the model implies much lower CRRA: equivalent relative risk aversion for the gamble in the last row is between 2.4 and 2.8 in all columns for the commitment consumer. This point mirrors Kandel and Stambaugh (1991) and Rabin and Thaler's (1998) point that observed risk preferences over moderate stakes are consistent with high values of the CRRA, while observed risk preferences over larger stakes imply much lower values of the CRRA. The commitments model provides a unified framework that endogenously generates higher CRRA values for moderate shocks than large shocks.

Table 2 points to an important limitation of the model in explaining risk preferences over small stakes. The first few rows in the table show that for small gambles (e.g. \$100-\$500 stakes), a consumer with commitments behaves in essentially the same way as a standard CRRA expected utility maximizer. Even in the most risk-averse specification reported (columns 7-8), the consumer is willing to accept an 50-50 gamble of lose \$100/ gain \$102. Hence, commitments cannot generate substantial risk aversion with respect to gambles with stakes below \$500. The reason is that the commitments model exhibits approximate risk neutrality for small stakes, since the induced utility $v^c(W)$ remains locally linear.

These calibrations suggest that commitments can explain a high level of risk aversion with

respect to gambles that involve stakes of approximately \$1,000-\$10,000. This range includes many risks of interest in the economics literature, including unemployment, temporary illness or injury, and vehicle or property damage.

III.B Social Insurance

The optimal structure of social insurance programs and redistributive tax policies has attracted considerable attention in public finance. Since risk preferences are a key factor in this analysis, commitments have implications for the design of such programs, which we briefly explore below.

Consider an example where an agent faces the risk of losing $\$Z$ of wealth with probability p. Let W_g denote the agent's lifetime wealth in the good state, where the shock does not occur. An actuarially fair insurance program that raises wealth in the bad state by \$1 must lower wealth in the good state by $\frac{p}{1-p}$. The change in expected utility from such a program is

$$\widetilde{MWG}(Z) = pv'(W_g - Z) - (1-p)\frac{p}{1-p}v'(W_g) = p(v'(W_g - Z) - v'(W_g))$$

To convert this expression into a money metric, we normalize $\widetilde{MWG}$ by the welfare gain from a \$1 increase in consumption in the good state. The marginal welfare gain from an extra \$1 of insurance relative to a \$1 increase in wealth in the good state equals:

$$MWG(Z) = \frac{p}{1-p}\frac{v'(W_g - Z) - v'(W_g)}{v'(W_g)}$$

This formula reflects the well-known intuition that the marginal value of insurance depends on the difference in marginal utilities between the good and bad states.

Figure 5 plots $MWG(Z)$ for an agent with and without commitments. Without commitments ($k = 0$), $MWG(Z)$ increases with Z because $v'(W_0 - Z)$ increases monotonically with Z. Insuring large shocks is more valuable than insuring small shocks when all goods are freely adjustable. Commitments alter this reasoning. As shown in Proposition 1, commitments make $v'(W_g - Z)$ a non-monotonic function of Z, because the marginal utility of consumption falls discontinuously at $W = s$. As a result, the marginal value of insuring some large shocks may be smaller than the marginal value of insuring some moderate shocks. In the example shown in the figure, the marginal dollar of insurance for a \$10,000 loss raises welfare more than twice as much as the marginal dollar of insurance for a loss of \$40,000.

The mechanism underlying the difference between the marginal welfare cost of small vs. large shocks is that commitments are more likely to be abandoned when households are hit by large shocks. To assess the empirical relevance of this mechanism, we compare consumption responses to small vs. large shocks in the PSID. We classify an unemployment shock as a "large" shock if the head's wage income loss due to unemployment exceeds 33%, i.e. total wage income in year 0 is at least 33% less than total wage income in the year prior to the shock. We classify a shock as "small" if the wage income loss is between 0 and 33%. By this classification, there are 455 "large" shocks and 517 "small" shocks in the data.

Table 3 reports the results of this analysis. When hit by a small shock, households that choose to move reduce food and housing consumption by 10%, while non-movers reduce food consumption by 5% (and housing by 0). 31% of households move in the year that a small shock occurs. When hit by a large shock, movers reduce food and housing by 13-14%, while non-movers reduce food by 14%. 40% of households move following large shocks, showing that more households cross the boundary of the (S,s) band and abandon prior commitments following large shocks. This evidence supports the mechanism through which commitments amplify the welfare cost of moderate scale shocks relative to large shocks. For larger shocks, more individuals choose not to bear the cost of reducing food sharply, and mitigate the welfare cost of the shock by re-optimizing fully over both goods.

These results have practical implications for the optimal design of programs such as unemployment insurance (UI) and disability insurance (DI). The key tradeoff faced by a social planner in designing such policies is to balance the benefits of consumption smoothing with the moral hazard cost of distorting incentives (e.g. inducing agents to work less). Commitments can affect this tradeoff in several ways. For example, Chetty (2003) shows that amplified moderate stake risk aversion results in a significantly higher optimal benefit rate for UI in a model with commitments relative to Gruber's (1997) calibrations in the one-good model.[17] Commitments also suggest that the optimal wage replacement rate for UI may be higher than for DI – in contrast with the standard model – since disability is a large shock that could induce adjustment on many margins. More broadly, commitments create a force toward providing more insurance for short-term, moderate-stake shocks relative to long-term welfare programs. We defer formal analysis of these normative issues to future work.

[17] Consistent with this result, commitments appear to guide the choice of unemployment benefit levels in practice. In the Arizona State Benefit Adequacy Study, Burgess et. al. (1981) define "benefit adequacy" in terms of covering "the total of necessary/obligated expenses for the entire household."

III.C Labor Supply

Since the seminal studies of Woytinsky (1940) and Mincer (1962), a large literature in labor economics has studied "added worker effects" – the entry of spouses into the labor force when primary earners become unemployed (see e.g., Ashenfelter 1980; Lundberg 1985; Tano 1993; Stephens 2003). By changing the base of individuals actively seeking work, the added worker effect can affect the measurement of the unemployment rate, the shape of the aggregate labor supply curve, and the dynamics of labor force participation.

An added worker effect requires that the income elasticity of labor supply is large in the short run. However, most studies of long-run labor supply behavior find that the income elasticity of labor supply is small and perhaps even zero (see e.g., Pencavel 1986, Blundell and MaCurdy 1999). Why might income effects on labor supply be larger in the short run than the long run?

The commitments model provides a straightforward answer to this question. To show this, we extend the model in section 2 to introduce endogenous labor supply. The household maximizes

$$\mathrm{E}_0 \sum_{t=1}^{T} u(f_t, x_t, l_t)$$

where period utility $u(f, x, l)$ is defined over food, housing, and spouse's labor supply, l. To facilitate analysis of added worker effects, we assume that the primary earner's labor supply is exogenously determined and affects the household's problem only by changing its wealth. We make the following regularity assumptions about utility: (1) u is strictly increasing in f and x, strictly decreasing in l, strictly concave and twice differentiable; (2) $u_{f,x} \geq 0$, $u_{f,l} \leq 0$ and $u_{x,l} \leq 0$, conditions which guarantee that f, x, and leisure are normal goods (Chipman 1977); and (3) $\lim_{f \to \infty} u_f = 0$, the Inada condition for food.

Let w denote the spouse's wage rate, which is fixed over time. The household's dynamic budget constraint continues to be given by equation (1), with period income defined as $y_t = w l_t$. As above, we assume that there is no uncertainty after period 1. The optimal consumption profiles are therefore flat: $x_t = x_1$, $f_t = f_1$ and $l_t = l_1$ for all $1 \leq t \leq T$.

Let W denote the household's wealth (unearned income) at the beginning of period 1, excluding wage income earned in period 1 and subsequent periods. Our objective is to characterize the effect of changes in W – which may occur e.g. because the primary earner becomes unemployed – on l_1, the spouse's labor supply in period 1. At any given W, the household's problem can be solved with a two-step procedure: 1) find the optimal consumption bundle of f_1 and l_1 for any given x_1;

(2) solve for the optimal level of x_1. This motivates the following expression for the elasticity of labor supply in period 1 with respect to period 1 wealth:

$$\varepsilon_{l_1,W} = \varepsilon_{l_1,W \mid x_1} + \varepsilon_{l_1,x_1} \cdot \varepsilon_{x_1,W}$$

where $\varepsilon_{l_1,W \mid x_1}$ is the semi-elasticity of l_1 with respect to W, holding fixed x_1. Note that this formula holds both with and without adjustment costs.

Given prior commitments x_0, define $\overline{W}$ to be the wealth level at which a household that faces no adjustment costs would optimally choose $x^n(\overline{W}) = x_0$ in period 1. Let $\varepsilon^c_{l,W}$ denote the wealth elasticity of labor supply for a household with commitments and $\varepsilon^n_{l,W}$ denote the corresponding parameter for one without commitments. Since $\varepsilon_{l_1,W \mid x_1}$ does not depend on the adjustment cost k at $W = \overline{W}$, the difference in the wealth elasticities in the commitment and no-commitment cases at $W = \overline{W}$ is

$$\begin{aligned} \varepsilon^c_{l,W} - \varepsilon^n_{l,W} &= -\varepsilon^n_{l_1,x_1} \cdot \varepsilon^n_{x_1,W} \\ &\Rightarrow \quad \varepsilon^c_{l,W} < \varepsilon^n_{l,W} < 0 \end{aligned} \tag{9}$$

because $\varepsilon^n_{x_1,W} > 0$ and $\varepsilon^n_{l_1,x_1} > 0$. This expression shows that the wealth elasticity of labor supply is larger in magnitude when households have commitments. Insofar as commitments are retained when households face small or temporary wealth fluctuations but are adjusted in the long run, this result implies that the wealth (unearned income) elasticity of labor supply is larger in the short run than the long run.

To understand the intuition for this result, suppose the primary earner is temporarily unemployed. If the household has commitments that it wishes to maintain, there is a strong incentive for spouses to enter the labor force to help pay the mortgage and other bills, especially in the presence of liquidity constraints. In contrast, a household that experiences a large, permanent change in wealth will reoptimize on all margins of consumption in the long run, reducing the pressure to make large adjustments on l. In the extreme case where utility is quasilinear in x, the no-commitment (long run) income effect is zero, while the short-run income effect is positive.

The fact that commitments amplify income effects in the short run also has implications for the wage elasticity of labor supply, which includes both an income and substitution effect. Studies of behavioral responses to taxation generally find small wage elasticities for households with incomes

below $100,000 (e.g., Saez 2004). The inability to fully reoptimize consumption in the short run may dampen responses to tax reforms because of temporarily amplified income effects. In the short run, households may be reluctant to cut labor supply in response to a tax increase if they have made prior commitments. However, taxes may still have significant effects on labor supply in the long run, when short-run income needs due to commitments are diminished. Hence, empirical studies that focus on short-run changes in behavior could understate the distortionary effects of taxation.

Consistent with the basic arguments made above, Fortin (1995) and Del Boca and Lusardi (2003) present evidence that mortgage commitments increase secondary earners' labor supply. Examining whether commitments affect behavioral responses to changes in unearned income and wages could shed further light on the added worker effect and the efficiency costs of income taxation.

III.D Other Applications

Consumption commitments could shed light on a number of other issues, including the following:

[Portfolio Choice] By amplifying risk aversion, commitments induce investors to hold safer portfolios. Since risk aversion over moderate risks varies with the location of W in the (S,s) band (Proposition 1), households in the middle of the band should particularly safe portfolios. Consistent with this prediction, Chetty and Szeidl (2005a) present empirical evidence that households who moved recently for exogenous reasons and are therefore close to the middle of the band hold safer portfolios, suggesting that the model can be useful in understanding heterogeneity in portfolio choice behavior. The model also offers normative lessons for asset allocation – e.g., institutions with significant financial commitments should hold safer portfolios.

[Asset pricing] Flavin and Nakagawa (2003) analyze asset pricing in a two-good adjustment cost model with complementarity between food and housing. They show that this model outperforms existing models in matching consumption dynamics and asset prices because it generates sluggish adjustment of consumption in response to shocks, as in habit formation models.

[Habit formation] Chetty and Szeidl (2005b) aggregate the commitments model with separable utility in an economy of agents with heterogeneous commitments. They show formally that aggregate dynamics in this economy coincide with those of a representative-consumer model with habit formation preferences. In this sense, commitments can provide non-psychological, neoclassical microfoundations for habit formation preferences widely used in macroeconomics and finance.

[Gambling] The model may help explain why racetrack bettors prefer skewed bets that have long

odds and large payoffs (Golec and Tamarkin 1998; Jullien and Salanie 2000), and why individuals are more likely to participate in some lotteries that have very large payoffs (Clotfelter and Cook 1987). Because commitments are endogenous to wealth, the model can generate gambling motives at all wealth levels, in contrast with an exogenous non-concave utility specification as in Friedman and Savage (1948).

[Wage rigidities] Committed individuals may prefer a gamble in which they get fired with a small probability rather than take a reduction in wages with certainty, explaining workers' disdain for small wage cuts as discussed in Bewley (1999). Postlewaite, Samuelson, and Silverman (2004) show formally that firms may choose to lay off workers instead of reducing wages in equilibrium in a model with commitments.

IV Conclusion

When agents have consumption commitments, risk aversion is context-specific, and in particular varies with the scale of the risk. Commitments increase risk aversion over moderate shocks relative to large shocks by forcing households to concentrate moderate shocks on a subset of consumption goods. Agents may also be risk-seeking in certain ranges because commitments induce non-concavities in utility over wealth. Empirical evidence on consumption responses to unemployment shocks shows that housing commitments force households to concentrate smaller shocks on adjustable goods such as food, as the model predicts.

In this paper, we explored applications of commitments to models where forward-looking agents maximize the present discounted value of expected utility. A growing literature in psychology and economics has argued that individuals may have present-biased preferences, e.g. because of myopia, self-control problems, or biased expectations. Such biases could lead to over-commitment ex-ante, further amplifying the effect of commitments on the welfare cost of shocks. Present-biased individuals may face a tradeoff between the benefits of commitments as a saving device (as in Laibson 1997) and the welfare costs of inflexibility when shocks occur ex-post. Present-biases could also augment the gambling motive for some agents with commitments. It would be interesting to explore the interface between commitments and models with imperfect optimization in future work.

References

[1] Arrow, K. J. (1965) "Aspects of the Theory of Risk-Bearing." Helsinki: Yrjö Hahnsson Foundation.

[2] Ashenfelter, Orley (1980). "Unemployment as Disequilibrium in a Model of Aggregate Labor Supply", Econometrica, vol. 48, pp. 547-64.

[3] Attanasio, Orazio (2000), "Consumer Durables and Inertial Behavior: Estimation and Aggregation of Ss Rules for Automobiles," Review of Economic Studies 67, 667-696.

[4] Bewley, Truman (1999), *Why Wages Don't Fall During a Recession*, Cambridge: Harvard University Press.

[5] Blundell, Richard and Thomas MaCurdy (1999), "Labor Supply: A Review of Alternative Approaches," in Ashenfelter, Orley and David Card, eds., *Handbook of Labor Economics.* Vol. 3. Amsterdam: North-Holland.

[6] Burgess, Paul, Jerry Kingston, Robert St. Louis, and Joseph Sloan, (1981), "Changes in Spending Patterns Following Unemployment," U.S. Department of Labor Employment and Training Administration Occasional Paper 81-3, Washington D.C.

[7] Clotfelter, Charles and Phillip Cook, (1987) "Implicit Taxation in Lottery Finance," National Tax Journal, 533-547.

[8] Chetty, Raj, (2003) "Consumption Commitments, Unemployment Durations, and Local Risk Aversion," NBER Working Paper 10211.

[9] Chetty, Raj and Adam Szeidl (2005a) "Marriage, Housing, and Portfolio Choice: A Test of Grossman-Laroque," UC-Berkeley mimeo.

[10] Chetty, Raj and Adam Szeidl (2005b) "Consumption Commitments: Neoclassical Foundations for Habit Formation," UC-Berkeley mimeo.

[11] Chetty, Raj and Adam Szeidl (2006), "Consumption Commitments and Risk Preferences," NBER Working Paper.

[12] Chipman, John S. (1977) "An empirical implication of Auspitz-Lieben-Edgeworth-Pareto complementarity," Journal of Economic Theory 14, 228-231

[13] Cicchetti, C. J. and J. A. Dubin (1994). "A Microeconometric Analysis of Risk Aversion and The Decision to Self-Insure." Journal of Political Economy 102:1, pp. 169-86.

[14] Del Boca, Daniela and Annamaria Lusardi (2003) "Credit Market Constraints and Labor Market Decisions," *Labour Economics* 10 pp. 681-703.

[15] Eberly, J. (1994) "Adustment of Consumers' Durables Stocks: Evidence from Automobile Purchases," *Journal of Political Economy* 102 , 403–436.

[16] Fortin, Nicole (1995), "Allocation Inflexibilities, Female Labor Supply and Housing Assets Accumulation: Are Women Working to Pay the Mortgage?," Journal of Labor Economics, 13: 524-557.

[17] Flavin, Marjorie and Shinobu Nakagawa, (2003) "A Model of Housing in the Presence of Adjustment Costs: A Structural Interpretation of Habit Persistence," University of California, San Diego mimeo.

[18] Fratantoni, M. C. (2001). "Homeownership, Committed Expenditure Risk, and the Stockholding Puzzle," *Oxford Economic Papers* 53: 241–259.

[19] Friedman, Milton and L.J. Savage, (1948) "The Utility Analysis of Choices Involving Risk," Journal of Political Economy 56, 279-304.

[20] Golec, Joseph and Maury Tamarkin, (1998) "Betters Love Skewness, Not Risk, at the Horse Track" Journal of Political Economy 106, 205-225.

[21] Gollier, Christian, (2005) "Does flexibility enhance risk tolerance?" Univ. of Tolouse mimeo.

[22] Grossman, Sanford J. and Guy Laroque (1990) "Asset Pricing and Optimal Portfolio Choice in the Presence of Illiquid Durable Consumption Goods," Econometrica 58, 25-52.

[23] Gruber, Jonathan (1997), "The Consumption Smoothing Benefits of Unemployment Insurance," American Economic Review 87, 192-205.

[24] Hansson, B. (1988), "Risk Aversion as a Problem of Conjoint Measurement," in Decision, Probability and Utility, ed. by P. Gardenfors and N.-E. Sahlin. Cambridge, Cambirdge University Press, 136-158.

[25] Hartley, Roger and Lisa Farrell (2002), "Can Expected Utility Theory Explain Gambling?" American Economic Review 92, 613-624.

[26] Himmelberg, Charles, Christopher Mayer and Todd Sinai (2005), "Assessing High House Prices: Bubbles, Fundamentals, and Misperceptions," Journal of Economic Perspectives.

[27] Jullien, Bruno and Bernard Salanie (2000), "Estimating preferences under risk: The case of racetrack bettors." The Journal of Political Economy, 108(3):503–530

[28] Kahneman, Daniel and Amos Tversky (1979), "Prospect Theory: An Analysis of Decision under Risk," Econometrica 47, 263-292.

[29] Kandel, S. and Stambaugh, R.F. (1991), "Asset Returns, Investment Horizons and Intertemporal Preferences," Journal of Monetary Economics 27, 39-71.

[30] Koszegi, Botond and Matthew Rabin (2005), "Reference-Dependent Risk Attitudes," UC-Berkeley mimeo.

[31] Laibson, David (1997), "Golden Eggs and Hyperbolic Discounting," The Quarterly Journal of Economics 112, 443-477.

[32] Li, R. (2003) "Household Portfolio Choice with Precommitment." Working paper, University of Toronto.

[33] Lundberg, Shelly (1985). "The Added Worker Effect", Journal of Labor Economics, vol. 3, pp. 11-37.

[34] Martin, Robert (2003), "Consumption, Durable Goods, and Transaction Costs," FRB International Finance Discussion Papers No. 756.

[35] Mankiw, N. Gregory and Stephen P. Zeldes (1991) "The Consumption of Stockholders and Nonstockholders," Journal of Financial Economics 29, 97-112.

[36] Markowitz, H. 1952. "The Utility of Wealth." Journal of Political Economy. 60, pp. 151-58.

[37] Mas-Colell, Andreu, Michael Whinston and Jerry R. Green, "Microeconomic Theory" Oxford University Press, 1995.

[38] Mincer, Jacob (1962) "Labor Force Participation of Married Women," in *Aspects of Labor Economics.* National Bureau of Economic Research.

[39] Pencavel, John. "Labor Supply of Men: A Survey." in Ashenfelter, Orley and Richard Layard, eds., *Handbook of Labor Economics.* Vol. 1. Amsterdam: North Holland, 1986.

[40] Postlewaite, A., Samuelson, L. and Silverman, D. (2004), "Consumption Commitments and Preferences for Risk." Working Paper, Penn Institute for Economic Research.

[41] Pratt, John W. (1964) "Risk Aversion in the Small and in the Large," Econometrica 32, 122-136.

[42] Rabin, Matthew and Richard Thaler (2001). "Anomalies: Risk Aversion." Journal of Economic Perspectives, 15, 219-32

[43] Rabin, Matthew (2000), "Risk Aversion and Expected-Utility Theory: A Calibration Theorem." Econometrica 68, 1281-1292.

[44] Saez, Emmanuel (2004), "Reported Incomes and Marginal Tax Rates, 1960-2000: Evidence and Policy Implications," in Ed. J. Poterba, Tax Policy and the Economy 18, 2004, Cambridge: MIT Press.

[45] Shore, Stephen and Todd Sinai (2005), "Commitment, Risk, and Consumption: Do Birds of a Feather Have Bigger Nests?," U-Penn Wharton mimeo.

[46] Stephens, Melvin (2003), "Worker Displacement and the Added Worker Effect," Journal of Labor Economics.

[47] Tano, Doki (1993). "The Added Worker Effect", Economic Letters, vol. 43, pp. 111-17.

[48] Woytinsky, Wladimir S. (1940). "Additional Workers and the Volume of Unemployment", Journal of Political Economy, vol. 48.

[49] Zeldes, Stephen P. (1989) "Consumption and Liquidity Constraints: An Empirical Investigation," Journal of Political Economy 97, pp 305-46.

Appendix: Proofs

Proof of Lemma 1

Consider equation (2) for the value function. We first establish that there is a range of wealth levels for which it is optimal not to change x (i.e., not to move). The maximization problem of a consumer who moves in period 1 is equivalent to maximizing the utility of a consumer with initial wealth $W - kx_0$ and no adjustment costs over goods. By A2, both f and x are normal goods for this consumer (Chipman 1977). As leftover wealth $W - kx_0 \to 0$, optimal consumption $f, x \to 0$ by the budget constraint. When $W \to \infty$, by A1, optimal consumption $f, x \to \infty$. These results show that the optimal choice of x following a move is a strictly increasing function of W that maps onto the set of positive reals and is continuous by the theorem of the maximum. It follows that there exist wealth levels $W_A < W_B$ such that the optimal choice of x following a move is $x_0(1 - k/T)$ when $W = W_A$ and x_0 when $W = W_B$.

For all $W_A \leq W \leq W_B$, the consumer optimally chooses not to move. This is because the optimal choice of x following a move must satisfy $x_0(1 - k/T) \leq x \leq x_0$ for W in this range. Hence, for wealth levels between W_A and W_B lifetime commitment spending is at least $Tx_0 - kx_0$. Given the leftover budget $W - kx_0$ following a move, the consumer can allocate at most $W - Tx_0$ on food consumption, which is the amount he would allocate if he decides not to move. To summarize, when $W_A \leq W \leq W_B$, it must be that following a move both food and commitment consumption fall As a result, staying must be optimal in this range.

The optimal policy of the consumer is to move if and only if $v^m(W, x_0) > v^0(W, x_0)$. We now establish that v^m and v^0 intersect exactly twice, at $s < W_A$ and at $S > W_B$. First note that the envelope theorem implies $v_1^m(W, x_0) = T \cdot u_1(f^m, x^m)$ where f^m and x^m are the optimal choices when moving, and that $v_1^0(W, x_0) = T \cdot u_1(W/T - x_0, x_0)$. For $W < W_A$, we have seen that optimal housing after a move satisfies $x^m < x_0(1 - k/T)$. Since $f^m + x^m = (W - kx_0)/T$, this implies $f^m > W/T - x_0$. Therefore

$$v_1^m(W, x_0) = Tu_1(f^m, x^m) < Tu_1(W/T - x_0, x_0) = v_1^0(W, x_0)$$

where we used $u_{12} \geq 0$ which holds by assumption. Similarly, when $W > W_B$, the optimal housing choice satisfies $x^m > x_0$ and hence $f^m < W/T - x_0$ by the budget constraint, which implies

$$v_1^m(W, x_0) = Tu_1(f^m, x^m) > Tu_1(W/T - x_0, x_0) = v_1^0(W, x_0).$$

These inequalities imply that for $W < W_A$, v^0 is steeper than v^m, and that for $W > W_B$, v^0 is not as steep as v^m. We also know from the analysis above that when $W_A \leq W \leq W_B$, the function v^0 lies above v^m and staying is optimal. Taken together, these results imply that v^m and v^0 intersect at most once in region$W < W_A$, and they intersect at most once in region $W > W_B$ as well. To show that the optimal policy is described by an (S, s) band, we need to prove that intersections do occur in both of these regions. We begin with $W < W_A$. For W small, A1 implies that $u(f, x_0)$ will be arbitrarily close to the smallest possible utility level $\inf_{f,x} u(f, x)$. By definition this must be smaller than period utility when moving, hence $v^m(W, x_0) - v^0(W, x_0)$ is positive for W sufficiently small, implying that an intersection exists in this range. For $W > W_B$, the Inada conditions in A1 can be used to establish that $v^m(W, x_0) - v^0(W, x_0)$ becomes positive as W grows without bound.

To establish (iii), note that an increase in k reduces $v^m(W, x_0)$ for each value of W and x_0, while it does not affect $v^0(W, x_0)$. As a result, when k increases to k', the consumer is no longer indifferent between staying and moving at wealth level $s(k)$. The utility of moving has decreased, and the consumer strictly prefers to stay. This implies that $s(k') < s(k)$, and an analogous argument can be used for S.

Proof of Proposition 1

(i) Under A3, the optimal choice of both f and x is proportional to W for a household without adjustment costs. Therefore $v^n(W)$ is proportional to $W^{1-\gamma}$ and hence $\gamma^n(W) = \gamma$. By homogeneity, for any f and x

$$u_1(f,x)f + u_2(f,x)x = (1-\gamma)u(f,x)$$

and differentiating with respect to f yields

$$\frac{-u_{11}(f,x)}{u_1(f,x)}f - \frac{u_{12}(f,x)}{u_1(f,x)}x = \gamma$$

showing that $\gamma_f(f,x) \geq \gamma$ for all f and x by A2. But $\gamma^c = \gamma_f \varepsilon^c_{f,W} > \gamma$ because $\varepsilon^c_{f,W} = W/(Tf) > 1$.

Under A4, we have $\varepsilon^n_{f,w} < 1$. To see why, note that the within period first order condition for a consumer without adjustment costs implies

$$\varepsilon^n_{f,W} = \frac{\gamma_x}{\gamma_f}\varepsilon^n_{x,W} < \varepsilon^n_{x,W}$$

and since a weighted average of $\varepsilon^n_{f,W}$ and $\varepsilon^n_{x,W}$ equals 1, it must be that $\varepsilon^n_{f,W} < 1$. To compute risk aversion, recall $\gamma^n(W) = \gamma_f \varepsilon^n_{f,W} < \gamma_f$ a constant, while $\gamma^c = \gamma_f \varepsilon^c_{f,W} > \gamma_f$ because $\varepsilon^c_{f,W} = W/(Tf) > 1$.

(ii) The proof follows from (i). Under A3, $\gamma^c(W') = \gamma W'/(W' - kx_0)$ while $\gamma^c(W) \geq \gamma W/(W - Tx_0)$ as shown above. Given that $W \leq S$, $W'/(W' - kx_0) \leq W/(W - Tx_0)$ holds as long as $W' > kS/T$. Under A4, $\gamma^c(W') < \gamma_f W'/(W' - kx_0)$ while $\gamma^c(W) = \gamma_f W/(W - Tx_0)$ and the result follows when $W' > kS/T$.

(iii) We have $v^c_1(s-, x_0) = v^m_1(s, x_0)$ and $v^c_1(s+, x_0) = v^0_1(s, x_0)$. By the proof of Lemma 1, $v^m_1(s, x_0) < v^0_1(s, x_0)$ since $s < W_A$. As a result, $v^c_1(s-, x_0) < v^c_1(s+, x_0)$. Now consider $W = s$, and let $\widetilde{Z}$ be a gamble that pays ϵ or $-\epsilon$ with equal probabilities. Then $W + \widetilde{Z}$ will be preferred to a sure payment of W if and only if

$$v^c(s+\epsilon, x_0) + v^c(s-\epsilon, x_0) > 2v^c(s, x_0)$$

or equivalently if

$$\frac{v^c(s+\epsilon, x_0) - v^c(s, x_0)}{\epsilon} - \frac{v^c(s, x_0) - v^c(s-\epsilon, x_0)}{\epsilon} > 0.$$

In this expression, as $\epsilon \to 0$ the left hand side converges to $v^c_1(s+, x_0) - v^c_1(s-, x_0) > 0$, hence for ϵ small enough the inequality will be satisfied. A similar argument establishes the claim for $W = S$.

There exist other wealth levels besides s and S for which the claim is true. To see why, consider some $\epsilon > 0$ which satisfies the above inequality. By continuity, for ζ small enough in absolute value we must have

$$\frac{1-\zeta}{2} \cdot v^c(s+\epsilon, x_0) + \frac{1+\zeta}{2} \cdot v^c(s-\epsilon, x_0) > v^c(s - \zeta\epsilon, x_0)$$

which shows the claim for $W = s - \zeta\epsilon$. Since ζ can be either positive or negative, there exists an interval of initial wealth levels around s for which one can find fair gambles that are attractive for the household. A similar statement holds for S as well.

Proof for Section 1.3: Borrowing constraints

Let superscript BC denote an environment in which the consumer faces borrowing constraints, and assume that A1 and A2 are satisfied.

Claim *Let $u(f,x)$ satisfy either A3 and the assumption that γ_f non-increasing in f, or A4 and the assumption that $\varepsilon^n_{f,E}$ is non-increasing in E. Then, as long as $s^{BC} < W < W^{BC}$, i.e., for wealth levels where the borrowing constraint binds but moving is not optimal:*

$$\gamma^{c,BC}(W,x_0) - \gamma^{n,BC}(W) \geq T \cdot \left[\gamma^c(W,x_0) - \gamma^n(W)\right].$$

Proof. Let $E = f_1 + x_1$ denote total expenditure in period 1, and suppose the household does not find it optimal to move. Then we can write the CRRA without and with adjustment costs as

$$\begin{aligned} \gamma^n(W) &= \left[\gamma_f \cdot \varepsilon^n_{f,E} - \varepsilon_{u_f,x} \cdot \varepsilon^n_{x,E}\right] \cdot \varepsilon_{E,W} \\ \gamma^c(W,x_0) &= \gamma_f \cdot \varepsilon^c_{f,E} \cdot \varepsilon_{E,W}. \end{aligned} \tag{10}$$

These equations follow directly from (3) and (4), using the fact that $\varepsilon_{f,E} \cdot \varepsilon_{E,W} = \varepsilon_{f,W}$. The equations are valid expressions for the CRRA irrespective of whether the consumer faces borrowing constraints or not. Borrowing constraints affect risk aversion by altering the $\varepsilon_{E,W}$ term on the right hand side of both expressions in (10). When the borrowing constraint does not bind, $\varepsilon_{E,W} = 1$; when it binds, $\varepsilon_{E,W} = W/E$. Note that when the borrowing constraint binds, $W/E > T$, because the consumer is forced to spend less than a share $1/T$ of lifetime wealth W on period 1 expenditures.

Under A3, (10) implies that $\gamma^{n,BC} = \gamma \cdot \varepsilon^{BC}_{E,W}$ and $\gamma^n = \gamma$. Using the fact that $\varepsilon^{BC}_{E,W} \geq T$, the claim follows if we can show

$$\gamma^{c,BC}_f \varepsilon^{c,BC}_{f,E} - \gamma \geq \gamma^c_f \varepsilon^c_{f,E} - \gamma.$$

Because the borrowing constraint binds, we have $E^{c,BC} \leq E^c$ and hence $\varepsilon^{c,BC}_{f,E} \geq \varepsilon^c_{f,E}$. Moreover, $f^{c,BC} \leq f^c$ and thus $\gamma^{c,BC}_f \geq \gamma^c_f$

Under A4, we need to show that

$$\gamma_f \left[\varepsilon^{c,BC}_{f,E} - \varepsilon^{n,BC}_{f,E}\right] \geq \gamma_f \left[\varepsilon^c_{f,E} - \varepsilon^n_{f,E}\right].$$

Clearly, $\varepsilon^{c,BC}_{f,E} \geq \varepsilon^c_{f,E}$ because $E^{c,BC} \leq E^c$. Moreover, $\varepsilon^{n,BC}_{f,E} \leq \varepsilon^n_{f,E}$ holds by assumption since $E^{c,BC} \leq E^c$. ■

Proof of Proposition 2

(i) Let $v^{k,c}(W,x_0)$ denote the value function and let $W^{k,CE}$ be the certainty equivalent of the wealth gamble $\widetilde{W}$ when adjustment costs equal k. We have $v^{k_1,c}(W,x_0) \geq v^{k_1,c}(W,x_0)$ for all W, and therefore

$$v^{k_1,c}(W^{k_1,CE},x_0) = \mathrm{E}v^{k_1,c}(\widetilde{W},x_0) \geq \mathrm{E}v^{k_2,c}(\widetilde{W},x_0) = v^{k_2,c}(W^{k_2,CE},x_0).$$

Since $W^{k_1,CE} \in (s^{k_1}, S^{k_1})$, we have $v^{k_1,c}(W^{k_1,CE},x_0) = v^{k_2,c}(W^{k_1,CE},x_0)$ which combined with the above chain of inequalities gives $W^{k_1,CE} \geq W^{k_2,CE}$ and hence $\pi(\widetilde{W},x_0,k_1) \leq \pi(\widetilde{W},x_0,k_2)$.

(ii) We first show that for σ small, the Arrow-Pratt approximation holds for the commitments utility function. For simplicity, use the notation that $W^{CE}_\sigma = W^{CE}(\widetilde{W}_\sigma,x_0,k)$ and $v(W) = v^c(W,x_0)$. Let $M^3_\sigma = \mathrm{E}\left[\left|\widetilde{W}_\sigma - \mathrm{E}\widetilde{W}\right|^3\right]$ denote the third absolute central moment of $\widetilde{W}_\sigma$; by assumption $M^3_\sigma = o\left(\sigma^2\right)$. Note that $v(W)$ is three times continuously differentiable, strictly increasing

and strictly concave in an open neighborhood of EW. We have

$$\begin{aligned} Ev\left(\widetilde{W}_\sigma\right) &= \Pr\left[W_\sigma \in (s,S)\right] \mathrm{E}\left[v\left(W_\sigma\right) | W_\sigma \in (s,S)\right] + \Pr\left[W_\sigma \notin (s,S)\right] \mathrm{E}\left[v\left(W_\sigma\right) | W_\sigma \notin (s,S)\right] \\ &= \Pr\left[W_\sigma \in (s,S)\right] \mathrm{E}\left[v\left(W_\sigma\right) | W_\sigma \in (s,S)\right] + o\left(\sigma^2\right) \end{aligned} \tag{11}$$

because $v(W)$ has an upper bound which is linear in W, $W_\sigma > kS/T$ and $\Pr\left[W_\sigma \notin (s,S)\right] \cdot \mathrm{E}[W_\sigma | W_\sigma \notin (s,S)] = O\left(M_\sigma^3\right) = o\left(\sigma^2\right)$. A second-order Taylor approximation of v around EW yields

$$v(W) = v(EW) + v'(EW) \cdot [W - EW] + \frac{1}{2} v''(EW) \cdot [W - EW]^2 + o\left([W - EW]^2\right)$$

for all $W \in (s,S)$ because the third derivative is bounded in the (S,s) band. Substituting in (11) yields

$$Ev\left(\widetilde{W}_\sigma\right) = v(EW) + v'(EW) \cdot \left[\mathrm{E}\left[W_\sigma\right] - EW\right] + \frac{1}{2} v''(EW) \cdot \mathrm{E}\left[W - EW\right]^2 + o\left(\sigma^2\right)$$

where we also used that $\Pr\left[W_\sigma \notin (s,S)\right] = O\left(M_\sigma^3\right)$, $\Pr\left[W_\sigma \notin (s,S)\right] \mathrm{E}[W_\sigma | W_\sigma \notin (s,S)] = O\left(M_\sigma^3\right)$ and $\Pr\left[W_\sigma \notin (s,S)\right] \mathrm{E}\left[W_\sigma^2 | W_\sigma \notin (s,S)\right] = O\left(M_\sigma^3\right)$. Simplifying,

$$Ev\left(\widetilde{W}_\sigma\right) = v(EW) + \frac{1}{2} v''(EW) \cdot \sigma^2 + o\left(\sigma^2\right). \tag{12}$$

We can relate this to the certainty equivalent using a first-order Taylor approximation

$$v\left(W_\sigma^{CE}\right) = v(EW) + v'(EW) \cdot \left[W_\sigma^{CE} - EW\right] + o\left(\sigma^2\right) \tag{13}$$

where the bound on the error term follows from $\left|W_\sigma^{CE} - EW\right| \leq \left[v(EW) - Ev\left(W_\sigma\right)\right] / v'(EW) = O\left(\sigma^2\right)$. Combining (12) and (13) with $Ev\left(\widetilde{W}_\sigma\right) = v\left(W_\sigma^{CE}\right)$ yields the Arrow-Pratt approximation

$$2 \frac{\pi\left(\widetilde{W}_\sigma, x_0, k\right)}{\mathrm{var}\left[\widetilde{W}_\sigma / \mathrm{E}\widetilde{W}_\sigma\right]} = -\frac{v''(EW) \cdot EW}{v'(EW)} + o(1)$$

which states that for small risks, the proportional risk premium per unit risk is approximately equal to the coefficient of relative risk aversion.

Now we proceed with the proof. For each σ, compute the risk premium per unit risk $\zeta_\sigma = 2\pi(\widetilde{W}_\sigma, x_0, k) / \mathrm{var}\left[\widetilde{W}_\sigma / \mathrm{E}\widetilde{W}_\sigma\right]$, then $\lim_{\sigma \to 0} \zeta_\sigma = \gamma^c(\mathrm{E}\widetilde{W}, x_0)$. Moreover, ζ_σ only depends on the certainty equivalent and the variance of the risk, and hence for each σ it equals the risk premium per unit risk computed for a power utility investor who has CRRA equal to $\gamma\left(\widetilde{W}_\sigma, x_0, k\right)$. By the Arrow-Pratt approximation for power utility with CRRA equal to $\gamma\left(\widetilde{W}_\sigma, x_0, k\right)$, we have $\zeta_\sigma = \gamma\left(\widetilde{W}_\sigma, x_0, k\right) + o(1)$ (note, the error in this expression is uniformly bounded for a range CRRA functions). Thus $\lim_{\sigma \to 0} \gamma\left(\widetilde{W}_\sigma, x_0, k\right) = \lim_{\sigma \to 0} \zeta_\sigma = \gamma^c\left(\mathrm{E}\widetilde{W}, x_0\right)$ which gives the desired result for small σ.

When σ grows without bound, $\widetilde{W}_\sigma$ will assign probability close to 1 to the region outside the interval (s,S) and the certainty equivalent will also be outside this range. By (ii) of Proposition 1, outside the (S,s)-band relative risk aversion is always less than $\gamma^c(\mathrm{E}\widetilde{W}, x_0)$. This implies that

in the range $W \notin (s, S)$ and $W > kS/T$, the utility function $v^m(W, x_0)$ is everywhere less risk averse than power utility with curvature $\gamma^c(\mathrm{E}\widetilde{W}, x_0)$. Because a utility function that is everywhere less risk averse has higher certainty equivalents for all gambles, for σ large enough, the certainty equivalent for $v^c\left(\widetilde{W}_\sigma, x_0\right)$ will be larger than it is for power utility with curvature $\gamma^c(\mathrm{E}\widetilde{W}, x_0)$. As a result, equivalent relative risk aversion will be less than $\gamma^c(\mathrm{E}\widetilde{W}, x_0)$ for σ large.

Miscellaneous proofs

Proof that equivalent relative risk aversion is well defined

Consider a gamble $\widetilde{W}$ that has certainty equivalent $W^{CE} \leq \mathrm{E}\widetilde{W}$ for some consumer. Let $W^{CE}(\gamma)$ be the certainty equivalent of $\widetilde{W}$ under CRRA utility with coefficient γ. Note that $W^{CE}(\gamma)$ is well-defined because the support of $\widetilde{W}$ is bounded away from zero and $\mathrm{E}\widetilde{W}$ is assumed to be finite. Moreover, $W^{CE}(0) \geq W^{CE}$ by assumption, and that $\lim_{\gamma\to\infty} W^{CE}(\gamma) > W^{CE}$ because the certainty equivalent for an infinitely risk averse consumer is the lower bound of the support of $\widetilde{W}$, which is smaller than the certainty equivalent for any investor with concave utility. Because $W^{CE}(\gamma)$ is continuous, there must be some γ^* where $W^{CE}(\gamma^*) = W^{CE}$. Also, $W^{CE}(\gamma)$ is strictly decreasing, because a CRRA utility function with higher curvature is everywhere more risk averse, which means that it has a strictly higher risk premium for all non-infinitesimal risks. Hence equivalent relative risk aversion γ^* is uniquely determined.

Proof that $\overline{W}$ is well-defined in Section III.C

Since x is a normal good, $x^n(W)$ is increasing. For W sufficiently low (perhaps negative) the optimal choice of x is close to zero since l is bounded from above; as W increases, the optimal level of x grows without bound because the marginal utility of food goes to zero. For some intermediate value $\overline{W}$ we have $x^n(\overline{W}) = x_0$.

Data Appendix

Expenditure shares (Table 1). We compute the expenditure shares for all categories using the summary expenditure variables in the CEX FMLY 1990Q1-1999Q4 files. We convert all variables to real 2000 dollars using the CPI. We define "total expenditure" as the sum of all expenditures minus mandatory contributions to pensions (TOTEXP-PERINS). To avoid bias from selective attrition in the construction of expenditure shares, we include data from only the first interview (CEX interview 2) for each household in the share calculations, summing the PQ and CQ variables to obtain quarterly expenditure measures. We define aggregate expenditure as the sum of total expenditure across all households in their first interview. Similarly, we define aggregate expenditure on category j as the sum of expenditures on category j by all households in the dataset in their first interview. Finally, we define the aggregate expenditure share of good j as aggregate expenditure on j divided by aggregate expenditure.

We define "shelter" as total expenditures on housing (HOUS) minus expenditures on utilities, furniture/appliances, and household operations. "Household operations" (HOUSOP) includes expenses such as cleaning, daycare, etc. We define "other transportation" as total expenditures on transportation (TRANS) minus expenditures on gas, maintenance, and public transportation.

Consumption growth rates (Figure 1 and Table 1). We use the CEX FMLY files to compute annual nominal consumption growth for all goods except food, housing, and cars. For food and housing, we use data from the PSID 1967-1997 family files. Details on the definition of food and housing in this dataset are given in the description of the event studies below. We compute changes in car consumption using data from the CEX OVB files 1990Q1-1999Q4. In all three datasets, we make no exclusions and include every observation with non-missing data in the sample.

We define the consumption growth rate of food as the log difference in nominal expenditures between year t and year $t-1$ for household i in year t in the PSID. For other nondurables (utilities, housing operations, other transportation, health insurance, other health, entertainment, personal care, reading, education, tobacco, miscellaneous), we define the consumption growth rate of category j for household i as the log difference in nominal expenditures between interview 5 and interview 2 for household i in the CEX. Note that individuals who report zero expenditures are omitted from our analysis since their consumption growth rates are undefined.

For housing (shelter), we follow the procedure described in the text using nominal PSID data on rent and housing values, with an adjustment for missing data. In 8.73% of the observations, the household head reports moving within the past year but the housing growth rate is undefined because (a) the individual switches from renting to owning or vice versa, (b) the individual splits off from a previous household and starts a new household, or (c) the home value or rent data is missing. Since these observations likely involve growth rates of housing consumption different from zero, the original histogram drawn with the non-missing observations overstates the fraction of zeros. We correct for this bias by scaling down the height of the bar centered at zero in the original histogram by 8.73% to arrive at the histogram shown in Figure 1b.

For cars, we define the change in car consumption as 0 for those who do not report buying or selling a car in the OVB files while they are in the CEX sample. For those who do report buying or selling, we define the change in car consumption as follows. We first construct a measure of the initial value of the car stock when the household enters the sample. We do so by summing the reported purchase prices of all cars owned by the household. Since purchase prices of older cars are frequently missing, we impute the values of these cars by replacing all missing observations with the average purchase price reported in that calendar year for a car of the reported model year. We compute the value of the car stock at the end of the sample by adding (or subtracting) the

purchase price of all purchased (or sold) cars during the sample period to the imputed value of the initial car stock. Finally, we compute growth in car consumption as the log difference between the nominal values of the end-of-sample car stock and initial car stock.

For apparel and furniture, we compute the level change in consumption as total reported expenditures over the four quarters in the CEX minus SALEINCX, which gives data on total sales of all small durables reported by the household. We deflate this net expenditure measure using the CPI to convert all the level changes to real 2000 dollars. Since SALEINCX includes sales of all small durables, it overstates sales of apparel or furniture alone. Thus our estimates give an upper bound on the frequency of consumption reductions in these categories.

The histograms in Figure 1 show the distribution of growth rates and level changes for six categories. The histograms in Figures 1a-1e re drawn with a bin size of 0.1, with a range of -1.05 to +1.05. The histogram in Figure 1f is drawn with a bin size of $200, with a range of -$2,000 to $2,000. We compute the fraction of households that reduce consumption of each category as the share of observations with negative net expenditures for apparel and furniture, and as the share of observations with negative growth rates for all other categories. For housing, we assume that half of the movers with missing housing growth data reduce housing consumption, and add this fraction (8.73%/2) to the fraction of observed reductions to obtain the statistic reported in Table 1.

Event Studies of Unemployment (Figure 2 and Table 3). We obtain data on the head of household's employment status at the time of each interview from the PSID family files. We define an individual as unemployed if they report working at the previous interview and being unemployed or temporarily laid off in the current interview. Prior to 1976, one category (code 1) includes both unemployed and temporarily laid off individuals. After 1976, this category is split into one for unemployed and another for temporarily laid off. We combine these two categories (codes 1 and 2) after 1976 to obtain a consistent definition of unemployment across years.

Food consumption is defined in all years as the sum of food consumed in the home, food consumed outside the home, and the subsidy value of food stamps, following Zeldes (1989). From 1968 to 1974, the question on food stamps asks about the amount received last year. After 1975 the variable is defined as the amount last month, which we scale up to an annual value. For some years in the 1970s, some observations include food stamps in the "food at home" variable; in these cases, we subtracted food stamps from food at home to avoid double counting. After 1994, the raw family data are available in a different form, which gives data on amounts and time periods. We multiply these and scale up to annual values to obtain a food consumption series that is consistent across years. Food growth is defined as the log change in food expenditure from year t-1 to year t. As above, housing consumption growth is defined as zero for non-movers, log change in annual rent for renters who move, and log change in house value for homeowners who move.

The food, rent, and home value data are deflated using food and housing price deflators from the CPI to obtain real growth rates. Though the consumption data are reported at an annual frequency, the framing of the consumption questions refers to the point of the interview. As pointed out by Zeldes and Gruber, this justifies the use of these variables as measures of consumption during the time of the interview rather than measures of total consumption over the past year.

We make three exclusions on the PSID dataset to arrive at our final sample for the analysis of unemployment shocks: (1) we include only heads of household between the ages of 20 and 65, (2) we include only those who report exactly one unemployment spell during the panel, (3) we exclude observations with changes in the number of people in the household to avoid introducing noise in the consumption measures because of changes in household composition. Figure 2 and Table 3 are constructed using this sample.

TABLE 1

Expenditure Shares and Adjustment Frequencies by Consumption Category

Consumption category	Share of total expenditures	Fraction of households actively reducing consumption
Shelter	22.2%	8.7%
Cars (excluding gas+maint)	14.7%	10.6%
Apparel	5.1%	0.3%
Furniture/Appliances	4.4%	0.5%
Health Insurance	3.0%	32.9%
Food and Alcohol	18.1%	42.8%
Utilities	8.2%	45.8%
Other Transportation	7.3%	49.2%
Entertainment	6.1%	48.7%
Out-of-Pocket Health	3.0%	47.8%
Education	2.0%	45.2%
Housing Operations	1.9%	44.3%
Personal Care	1.0%	41.0%
Tobacco	0.9%	36.6%
Reading Materials	0.6%	45.7%
Miscellaneous	1.5%	39.3%

NOTE--First column in table shows aggregate expenditure shares for consumption categories in the CEX, following methodology described in the data appendix. Second column reports fraction of households who actively reduce consumption (beyond depreciation) of each category from first quarter to last quarter in CEX. For apparel and furniture, households that reduce consumption are defined as those with negative net expenditures. For all other categories, households that reduce expenditure are those with negative nominal growth ($g_{cit} < 0$). See text and data appendix for the definition of g_{cit}. Categories above dotted line are classified as "commitments" by frequency-of-adjustment definition.

TABLE 2

Risk Preferences over Small and Large Stakes: Calibration Results

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	g = 1,025		**g = 1,050**		**g = 1,100**		**g = 1,250**	
Loss (L)	**CRRA**	**Commit**	**CRRA**	**Commit**	**CRRA**	**Commit**	**CRRA**	**Commit**
100	100.2	100.2	100.5	100.5	101	101	102	102
200	201	201	204	202	204	204	208	208
500	506	506	524	512	524	524	555	555
1,000	**1,025**	**1,025**	**1,050**	**1,050**	**1,100**	**1,100**	**1,250**	**1,250**
2,000	2,103	2,103	2,211	2,211	2,446	2,446	3,373	3,370
5,000	5,695	5,695	6,573	6,570	9,388	9,359	∞	14,879
10,000	13,243	13,235	19,615	19,387	∞	23,388	∞	23,841
15,000	23,805	23,732	98,027	29,741	∞	30,149	∞	30,290
20,000	39,975	34,003	∞	36,660	∞	37,005	∞	37,239
25,000	69,671	43,003	∞	43,860	∞	44,388	∞	44,746
30,000	175,622	50,354	∞	51,591	∞	52,360	∞	52,884
40,000	∞	66,569	∞	68,897	∞	70,372	∞	71,393
50,000	∞	85,166	∞	89,184	∞	91,789	∞	93,623
Curvature over food (γ_f)		3.7		7.1		13.6		29.9
ERRA for L= 50,000 gamble		2.47		2.63		2.72		2.78

NOTE--Table shows risk preferences over gambles of various stakes for agents who are indifferent between accepting and rejecting a 50-50 lose \$1,000/gain \$g gamble at wealth level \$300,000, where g varies across columns. Entries in the table show values \$G such that agent indifferent about the lose \$1000/gain \$g gamble is also indifferent between accepting and rejecting a 50-50 lose \$L/gain \$G gamble, where L varies across rows. Odd columns show G values for agent with CRRA utility; even columns for consumer with commitments. Period utility for agent with commitments is given by equation (6) in text, with γ_x=1 (log utility) T=5, lifetime wealth W=\$300,000, adjustment cost k=0.2, and prior commitment x_0 =\$36,000. Parameters γ_f and μ are chosen jointly so that (1) optimal unconstrained budget share of commitments is 50% at wealth \$300,000 and (2) agent is indifferent with respect to lose \$1000/gain \$g gamble. Last row shows equivalent relative risk aversion (see definition in text) for the L=\$50,000 gamble.

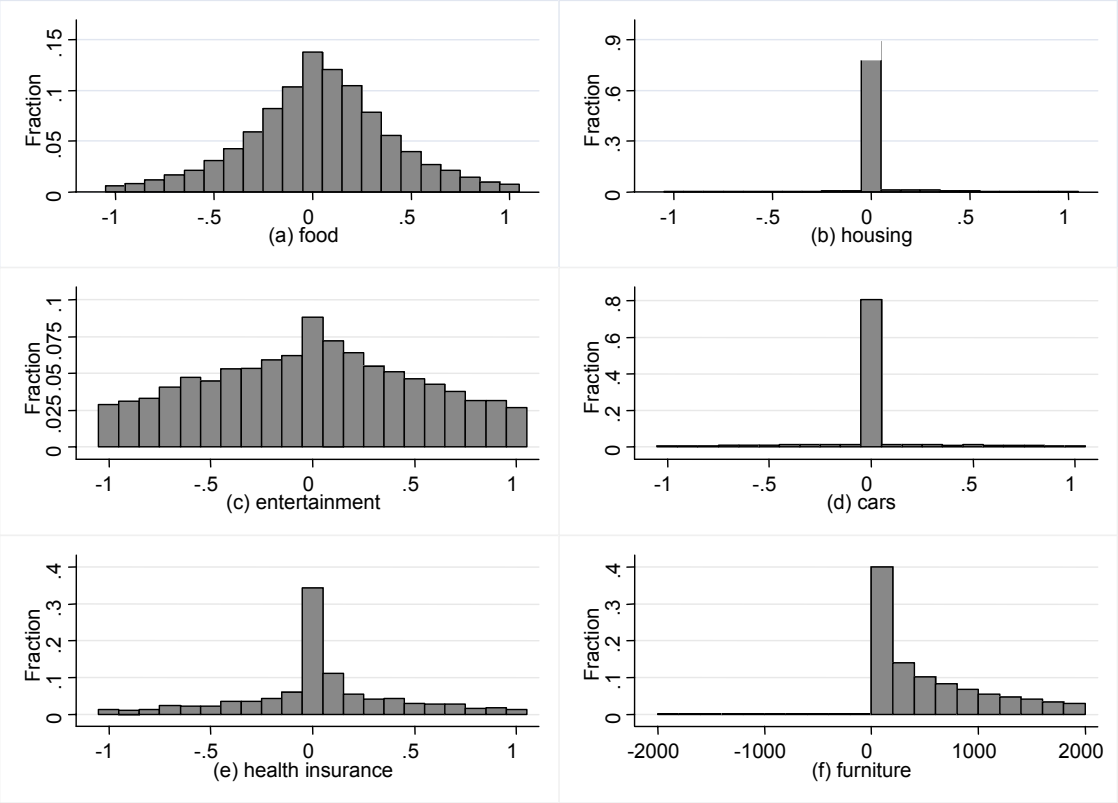
TABLE 3

Housing and Food Consumption Responses to Small and Large Shocks

	Percent Who Move	Change in Housing for Movers	Change in Food for Movers	Change in Food for Non-Movers
Small Unemployment Shock	31%	-10%	-10%	-5%
Large Unemployment Shock	40%	-14%	-13%	-14%
Year Prior to Shock	27%	13%	-2%	-1%

NOTE--Table shows fraction of movers and real annual growth rates of housing and food for movers and non-movers in response to small and large unemployment shocks. A shock is classified as large if total wage income of head in the year of job loss falls by 33% or more relative to wage income in previous year. A shock is classified as small if total wage income loss of head is between 0% and 33%. Corresponding statistics for year prior to shock are shown as a basis for comparison. The data is from the PSID; see the data appendix for definition of sample and other details. "Change in housing" statistics exclude observations when individuals switches from owning to renting or vice versa; "percent who move" statistics include all observations.

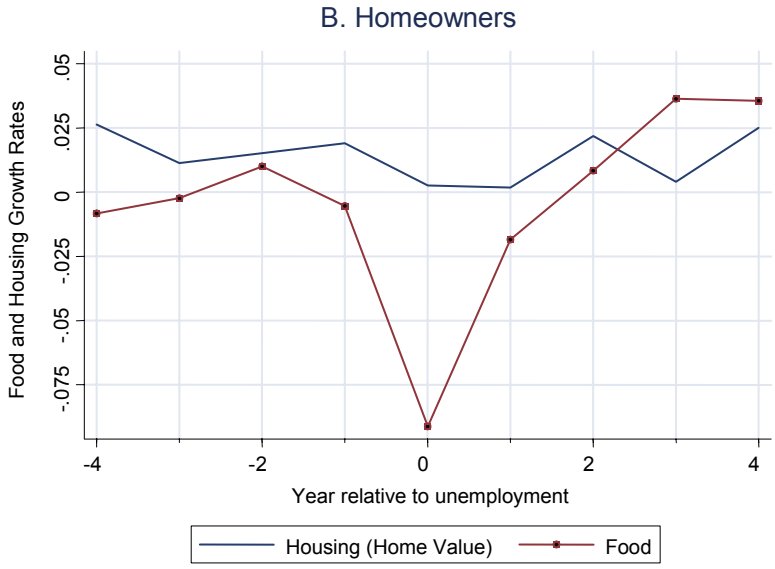
FIGURE 1
CONSUMPTION ADJUSTMENT PATTERNS



NOTE–These figures show the population distribution of annual changes in consumption of various categories. Figures 1a-1e plot histograms of the annual nominal growth rates of consumption (g_{cit}), following the methodology described in the text and data appendix. Figure 1f plots a histogram of level changes in furniture consumption (total expenditures minus sales of furniture over a one year period) in real 2000 dollars.

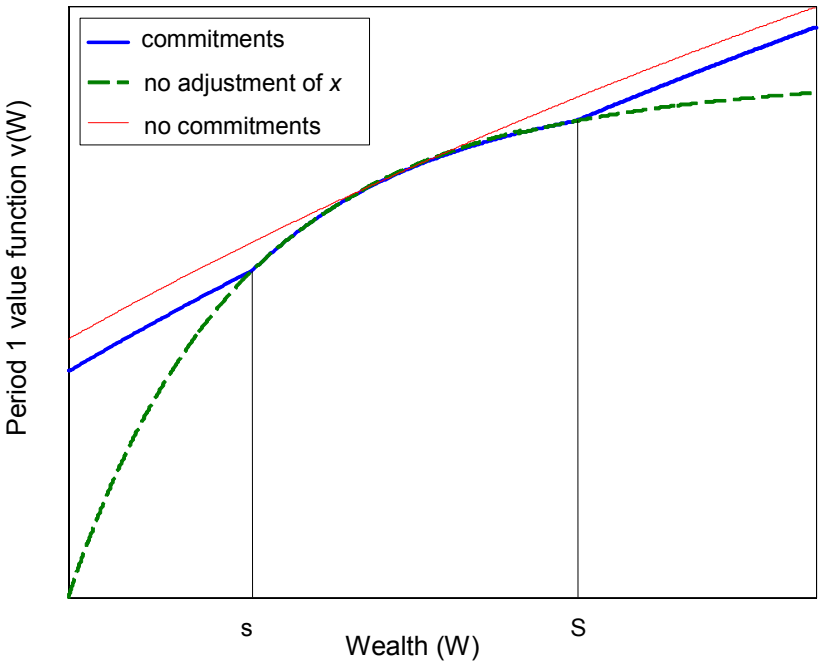
FIGURE 2
EVENT STUDY OF CONSUMPTION AROUND UNEMPLOYMENT SHOCKS





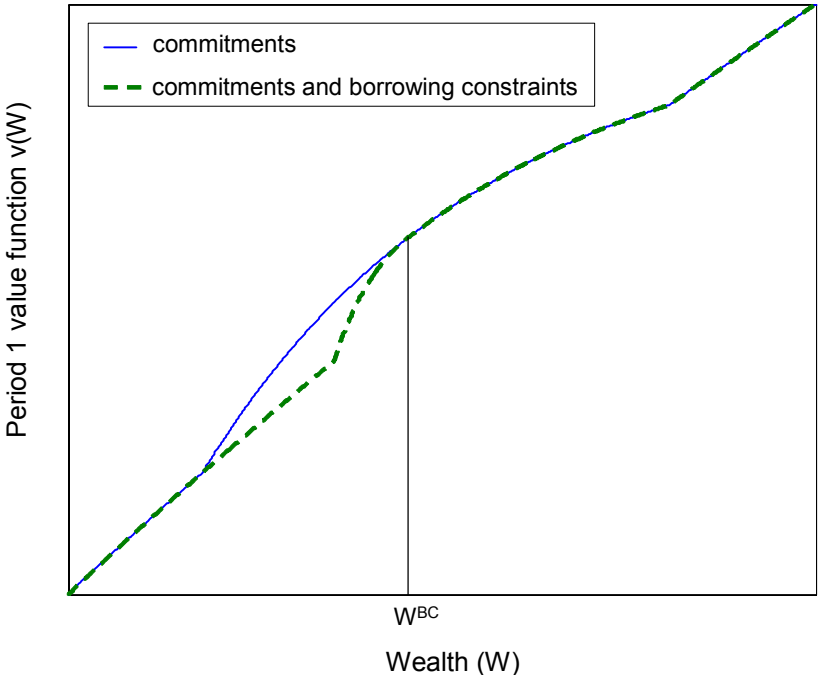
NOTE– These figures plot real annual growth rates of food and housing consumption (changes in log consumption) from year t-1 to year t. The year of job loss is normalized to t=0 for all individuals. The sample for figure A includes households who rented in year -1. The sample for Figure B includes households who owned a house in year -1. See the data appendix for sample definitions and construction of growth rates.

FIGURE 3A
EFFECT OF COMMITMENTS ON VALUE FUNCTION



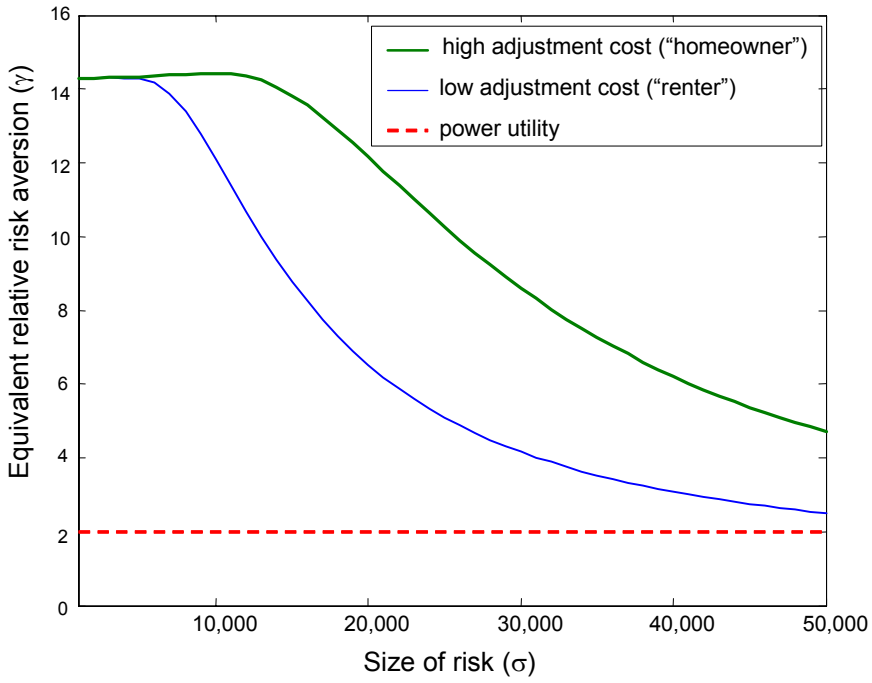
NOTE–This figure shows the value function of a consumer with consumption commitments (dark line). The adjustment cost makes it optimal not to move for $W \in (s, S)$. Outside the (S,s) band, moving is optimal, because maximized utility when the consumer does not adjust (dotted line) is below maximized utility with adjustment. The presence of commitments increases curvature inside the (S,s) band relative to a consumer with no adjustment costs (thin line). The figure is constructed using the same preferences and parameter values as column 4 of Table 2. See the notes to Table 2 for details.

FIGURE 3B
COMMITMENTS AND BORROWING CONSTRAINTS



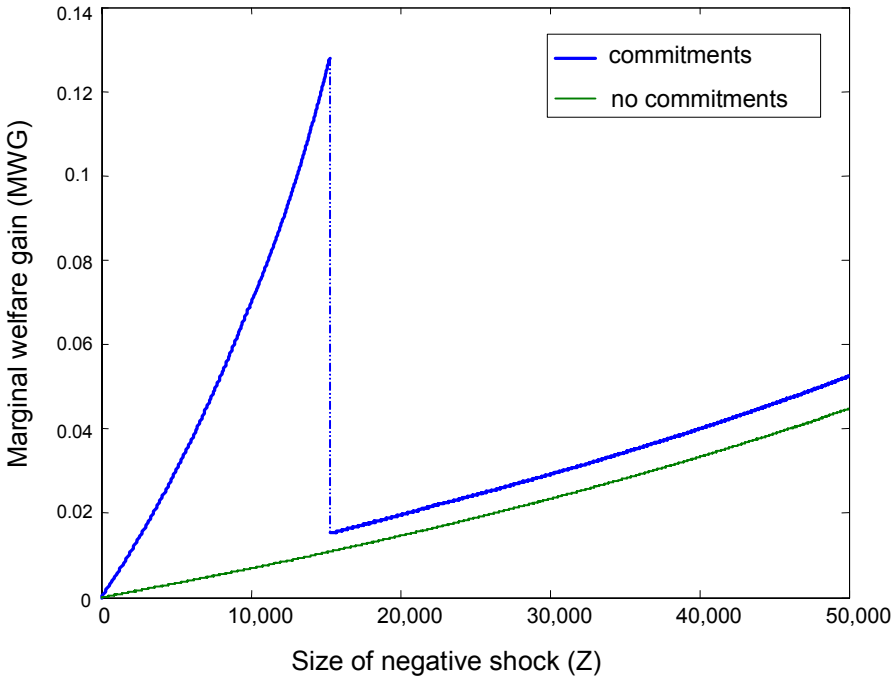
NOTE–This figure shows the value function of a consumer with consumption commitments and a borrowing constraint as specified in the text (dotted line). For high wealth realizations, utility is identical to that of a consumer without a borrowing constraint (thin line). For low wealth realizations utility falls sharply relative to the benchmark with no borrowing constraints, because the consumer is unable to smooth the impact of a shock intertemporally. The preferences and other parameters used to construct this figure are the same as those used in column 4 of Table 2.

FIGURE 4
RISK AVERSION OVER SMALL VS. LARGE RISKS



NOTE–This figure plots equivalent relative risk aversion (see text for definition) versus the size (standard deviation) of the risk. This curve is plotted for three types of individuals: one with high adjustment costs, such as a homeowner (heavy line); one with lower adjustment costs, such as a renter (thin line); and, as a basis for comparison, a consumer with no commitments and power utility preferences with risk aversion $\gamma = 2$ (dashed line). The preferences used in constructing the homeowner and renter curves are the same as those used in column 4 of Table 2; the adjustment cost is $k = 1$ for homeowners and $k = 0.2$ for renters. The shocks have lognormal distribution truncated at \$80,000, with mean \$300,000 and standard deviation shown in the figure.

FIGURE 5
MARGINAL WELFARE GAIN FROM INSURANCE



NOTE–This figure plots the marginal welfare gain from an actuarially fair insurance program that raises wealth in the bad state by $1, as a function of the size of the wealth loss in the bad state. The marginal welfare gain is a money metric, defined as the increase in expected utility from $1 of additional insurance divided by the change in expected utility from a $1 increase in wealth in the good state (see text for details). The marginal value of insurance is monotonically increasing in the case with no adjustment costs (thin line) but non-monotonic in the presence of commitments (heavy line). The preferences and parameters used in constructing figure are the same as those used in column 4 of Table 2; in particular the consumer has initial wealth of $300,000. The probability of the bad state is $p = 0.1$.