

The Cost of Unanticipated Household Financial Shocks: Two Examples

Kartik Athreya and Urvi Neelakantan

Households sometimes experience unexpected negative changes to their financial circumstances. In this article, we quantify the consequences of two representative types of unanticipated financial shocks. By “unanticipated,” we mean that households in our experiments are modeled as ignoring even the *possibility* that the shock could occur. We are thus interested in the cost of an event that comes as such a surprise to the household that its previous consumptions-savings decisions in no way prepared it for such an eventuality. Our analysis is therefore exactly analogous to a standard form of experiment in business cycle contexts, e.g., the impulse response of an unanticipated fiscal or monetary policy shock that agents know is permanent as soon as it occurs (see, for example, Baxter and King [1993]).

For each shock, our calculations tell us how costly it is for households to live in a world where the shock occurs compared to a world in which it does not. Why might such costs be useful to study? If the household (or a policymaker) could pay—say through investment in financial education—for information that would enable it to avoid the shock or mitigate its effect, our calculations may provide an upper bound on how much it might be willing to pay. The reason is that the cost of the shock depends not just on its magnitude but also its likelihood. For a shock of a given size, households will be less willing to pay to avoid it as its likelihood falls. We assume that the shock is

■ We thank Juan Carlos Hatchondo, Marianna Kudlyak, Ned Prescott, Pierre Sarte, Max Schmeiser, Jonathan Tompkins, John Weinberg, Kim Zeuli, seminar participants at the Federal Reserve Bank of Richmond and the University of Illinois at Urbana-Champaign, and participants at the Georgetown Center for Economic Research and the Southern Economics Association conferences for helpful comments and suggestions. We are solely responsible for any errors. The views expressed in this paper are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Richmond or the Federal Reserve System. E-mail: kartik.athreya@rich.frb.org; urvi.neelakantan@rich.frb.org.

completely unanticipated or seen as one with zero likelihood. If its likelihood is truly close to zero, this makes it not worth doing much about, all else equal. Moreover, if the household is incorrect in assigning zero or near-zero likelihood to the shock, that is a belief that maximizes the amount by which households “underestimate” the risks. If the household instead knew that the shock could occur with positive probability, it would take actions (to the extent warranted by the magnitude and likelihood of the shocks, and the household’s aversion to risk) to reduce its severity. By contrast, our model features households that will, by their unawareness, have made no provisions at all at the onset of either of the shocks we consider. Our cost calculations will also allow us to compare shocks, that is, to point out which shocks are costlier (assuming equal likelihood) and therefore worthy of greater attention.

The two types of shocks we consider here are (1) an unanticipated drop in net worth and (2) an unpredicted increase in borrowing costs for all forms of unsecured debt. Each is meant to represent the occurrence of an empirically plausible scenario. The first provides insight into the cost borne by those who are surprised by declines in the value of assets in their portfolio. Consider, for example, a household that has a net worth that is largely composed of equity in its home, and for which the recent decline in U.S. house prices came as a shock. It is evident that many commentators and experts placed little probability on a widespread decline in home prices.¹ The second case is that of a sudden, widespread increase in the cost of rolling over debt and captures the effects of general credit market tightening as might occur in the midst of a severe recession that was a priori assigned zero probability. Note that both shocks are fully persistent.

The size of a shock is an inadequate measure of its importance to a household, in particular because the cost is likely to vary across households. Thus, quantifying the cost requires a model of household financial decision making. Households make consumption-savings decisions with the goal of smoothing consumption over their lifetime. A consequence of this hypothesis is that households’ financial positions (and, hence, the cost of the shock to them) will differ by age. Moreover, to the extent that households face other, more predictable forms of risk throughout their lives, they will also differ from each other at any given age. In turn, the cost of a shock will vary across households of any given age as well. The economic model we use is a fairly standard version of a life-cycle model of consumption and savings, and follows Athreya (2008). We use the model and, in particular, the optimal value function of the household, to quantify the effects of the shocks. Specifically, we use the

¹ Freddie Mac’s “Rent or Buy” calculator provides anecdotal evidence of the lack of concern with house price declines. The calculator did not allow users to analyze the effects of negative realizations for home prices, even though the same device allowed for increases in house prices of up to 100 percent (Joffe-Walt and Davidson 2010).

model to determine the amount of annual consumption that a household would be willing to give up to avoid facing the shock.

The reader will no doubt see that our article is highly stylized. Importantly, it abstracts from portfolio choice and focuses instead on a simple scalar measure of net worth. In its current form it therefore cannot speak directly to particular kinds of financial decisions, such as house purchases or any other leveraged purchase of risky assets. In particular, our focus on net worth effectively precludes us from being able to assess the impact of decisions whose effects derive primarily from their impact on the gross financial positions of households—as well as on any attendant changes in the periodic payment obligations—while leaving net worth essentially unchanged. Our model also abstracts from the labor supply decision, which could mitigate the cost of the shock by allowing households to simply “work their way” out of a reduction in net worth. However, this is not wholly unreasonable because the shocks we consider are most relevant to recessionary settings, in which labor markets could plausibly preclude such adjustments.

Finally, our model embodies a strong assumption with respect to the information that households possess: We assume that the shocks that the household faces are completely unanticipated. It is possible, instead, that households are aware of the existence of the kinds of shocks we analyze in this article, but wrong about the exact probabilities with which they could occur. Nonetheless, while strong, this assumption allows us to determine what are likely to be upper bounds on the consequences of such shocks. After all, any information received in advance about the likelihood of such events can only make the eventual shock, if it occurs, easier to deal with, as households will have consumed and saved in anticipation of such possible outcomes. In addition, the current work is simply a small first step, and we have indeed begun to incorporate each of these features in ongoing work (Athreya, Ionescu, and Neelakantan 2011) that we hope will shed greater light on the questions addressed here.

With the preceding in mind, we describe the model in Section 1. Section 2 describes how each shock is introduced within this framework. Section 3 reports the results in terms of the costs of each shock. Section 4 concludes.

1. A LIFE-CYCLE MODEL OF CONSUMPTION AND NET WORTH

The economy is that of Athreya (2008), and consists of a continuum of J overlapping generations of working households. Households value consumption, do not value leisure, and therefore supply labor inelastically.

Table 1 Model Parametrization

Parameter	Value	Parameter	Value	Parameter	Value	Parameter	Value
α	2	R^f	1.01	γ	0.99	$\underline{\tau}$	\$7,600
β	0.96	ψ	3.4%	σ_u^2	0.063	$\underline{\tau}^R$	\$8,600
x_1	0	λ	0.9	σ_η^2	0.0275		
				$\sigma_{\eta 1}^2$	0.22		

Preferences

The household chooses consumption, $\{c_j\}_{j=1}^J$, and retirement wealth, x_R , to solve

$$\sup_{(\{c_j\}, x_R) \in \Pi(\Psi_0)} E_0 \sum_{j=1}^J \beta^j \frac{c_j^{1-\alpha}}{1-\alpha} + \frac{x_R^{1-\alpha}}{1-\alpha}. \quad (1)$$

Here, $\Pi(\Psi_0)$ denotes the space of all feasible combinations $(\{c_j\}, x_R)$ given initial state Ψ_0 , α denotes risk aversion, and β is the discount factor. In the calibration, risk aversion and the discount factor are set at the standard values of $\alpha = 2$ and $\beta = 0.96$. (See Table 1 for all model parameter values, which follows Athreya [2008].)

Income

Households have three potential sources of income: labor income, means-tested transfer income, and retirement income, with labor income being subject to shocks drawn from a probability structure that is known perfectly by the agent.

Labor Income

The model period is one calendar year. Households begin working life at age 21 and retire at age 65. Households face uncertainty in their labor income because of stochastic productivity shocks to their labor supply. Following the literature (e.g., Hubbard, Skinner, and Zeldes 1995; Huggett and Ventura 2000; Storesletten, Telmer, and Yaron 2004), the evolution of log income, $\ln y_j$, is modeled as

$$\ln y_j = \mu_j + z_j + u_j, \quad (2)$$

where μ_j is an age-specific mean of log income, z_j is the persistent shock, and u_j is the transitory shock.

The profile $\{\mu_j\}_{j=1}^J$ is parameterized using data on the median earnings of U.S. males from the 2000 Census.²

The persistent shock, z_j , is given by

$$z_j = \gamma z_{j-1} + \eta_j, \quad \gamma \leq 1, \quad j \geq 2, \quad \eta_j \sim i.i.d. N(0, \sigma_\eta^2). \quad (3)$$

We set $\gamma = 0.99$ and $\sigma_\eta^2 = 0.0275$ to capture the facts that, in the data, the cross-sectional variance in log income increases substantially, and roughly linearly, over the life cycle; it is roughly 0.28 among 21-year-olds and roughly 0.90 among new retirees. The transitory shock, u_j , is distributed as $u_j \sim i.i.d. N(0, \sigma_u^2)$ and is independent of η_j .

To capture initial heterogeneity across households, it is assumed that they draw their first realization of the persistent shock from a distribution with a different variance than at all other ages. That is,

$$z_0 = 0, \text{ and } \eta_1 \sim N(0, \sigma_{\eta_1}^2). \quad (4)$$

In the above, $\sigma_{\eta_1}^2 = 0.22$.

Note that the assumption that households supply labor inelastically restricts them from using a smoothing mechanism that could be particularly useful in the face of unanticipated shocks. However, not only does this assumption keep the model parsimonious, it is in keeping with the usefulness of providing an upper bound on the costs of the shocks we study.

Means-Tested Transfer Income

Following Hubbard, Skinner, and Zeldes (1995), means-tested transfers $\tau(\cdot)$ are specified as a function of age, j , net worth, x_j , and income, y_j , as follows:

$$\tau(j, x_j, y_j) = \max\{0, \underline{\tau} - (\max(0, x_j) + y_j)\}. \quad (5)$$

Social insurance in the United States aims to provide a floor on consumption and the specification in equation (5) captures this feature. The transfer scheme provides households with a minimum of $\underline{\tau}$ units of the consumption good at the beginning of the period. In the calibration, $\underline{\tau} \cong \$7,600$ to match data on the asset accumulation of households in the lower percentiles of the wealth distribution.

Retirement Income

Household utility at retirement is evaluated as $\frac{x_R^{1-\alpha}}{1-\alpha}$. Retirement wealth, x_R , is the sum of household personal savings, x_{J+1} , and a baseline retirement benefit, $x_{\underline{\tau}R}$:

$$x_R = x_{J+1} + x_{\underline{\tau}R}. \quad (6)$$

² Since income is lognormally distributed, the mean of log income equals the log of median income. Therefore, the log of median earnings is used to generate the profile.

The baseline retirement benefit, $x_{\underline{\tau}^R}$, yields an annual income of $\underline{\tau}^R$ when annuitized using discount rate R^f . That is, the baseline retirement benefit solves

$$\sum_{k=1}^K \frac{\underline{\tau}^R}{(R^f)^k} = x_{\underline{\tau}^R}. \quad (7)$$

Here, $\underline{\tau}^R$ represents the societal minimum amount of consumption at retirement. This amount is not means tested and is intended to represent the sum of welfare programs, Social Security, and Medicare.³ The interest rate, $R^f > 0$, is the risk-free rate of return on savings and is exogenously given.

In the calibration, the minimum amount of consumption at retirement, $\underline{\tau}^R$, is set equal to \$8,600 and $R^f = 1.01$.

Technology and Market Arrangement

At each age j , households choose whether to save ($x_{j+1} > 0$) or borrow ($x_{j+1} < 0$). Savings earn the exogenous risk-free rate of return $R^f > 0$. The interest rate on borrowing is $R(\cdot)$, which incorporates credit risk (because households can default on the debt next period) and transaction costs, ψ , arising from resources used in intermediation. Default is costly and reduces household utility by λ in the period in which debts are repudiated. This cost includes, but is not limited to, the cost of legal representation and court fees. It is meant to capture all costs deemed relevant by households, and will be calibrated to help the model match default-related behavior.

Recursive Formulation

The household's problem is recursive in a state vector that includes age, j , beginning-of-period net worth, x_j , current-period realization of the persistent shock, z_j , and current-period realization of transitory income, u_j .

Value Functions

Households that enter a period with debt must decide whether or not to default. The value function when repaying debts is $W^R(\cdot)$, which solves

$$W^R(j, x_j, z_j, u_j) = \sup_{x_{j+1}} \left\{ \frac{c_j^{1-\alpha}}{1-\alpha} + \beta E_{z_{j+1}|z_j} V(j+1, x_{j+1}, z_{j+1}, u_{j+1}) \right\}, \quad (8)$$

³ This approach follows Huggett (1996) and Gourinchas and Parker (2002).

subject to

$$c_j + \frac{x_{j+1}}{R(j, x_{j+1}, z_j)} \leq y_j + \tau(j, x_j, y_j) + x_j, \quad (9)$$

where $R(j, x_{j+1}, z_j)$ is the interest rate associated with the level of savings or borrowing, x_{j+1} , chosen by the household of age j and current realization of the persistent shock z_j .

The value of defaulting is given by $W^D(\cdot)$, which solves

$$W^D(j, x_j, y_j, u_j) = \sup_{x_{j+1}} \left\{ \frac{c_j^{1-\alpha}}{1-\alpha} - \lambda + \beta E_{z_{j+1}|z_j} V(j+1, x_{j+1}, z_{j+1}, u_{j+1}) \right\}, \quad (10)$$

subject to

$$c_j + \frac{x_{j+1}}{R^f} \leq y_j + \tau(j, x_j, y_j), \quad (11)$$

$$x_{j+1} \geq 0. \quad (12)$$

The debt obligation in the right-hand side of (9) does not appear in (11) because the household defaults. The household pays the utility cost, λ , associated with defaulting. In the parametrization, λ is set to 0.9. Along with the other parameters in the model, this targets the Chapter 7 filing rate of 0.5 percent and the mean net worth of Chapter 7 bankruptcy filers of \$16,815.⁴

The household is not allowed to borrow in the period in which it defaults. Thus, net worth chosen in the current period must be non-negative and earns the risk-free rate of interest, R^f . The household can borrow in all subsequent periods.

The beginning-of-period value function must therefore satisfy

$$V(j, x_j, z_j, u_j) = \max [W^R(j, x_j, z_j, u_j), W^D(j, x_j, z_j, u_j)]. \quad (13)$$

Once borrowing or savings is chosen, the period ends.

Loan Pricing

In the market for loans, creditors are assumed to be competitive and to hold a sufficiently large number of loans of any given size for the law of large numbers to guarantee them a deterministic rate of return on loans of that size. They pay transactions costs, ψ , in exchange for which they can observe all factors needed to forecast the risk of default one period ahead. In the model, these factors are age, j , the persistent shock, z_j , and household debt, x_j . Creditors expect to break even on each loan by pricing contingent on these factors. Let $\pi^D(j, x_{j+1}, z_j)$ denote the probability of default on a loan of

⁴This data is as of 1991 in order to be consistent with the timing of the income and consumption data.

size x_{j+1} , made to a household of age j , with persistent income shock z_j . Let $I(j+1, x_{j+1}, z_{j+1}, u_{j+1})$ be the indicator function over whether or not a household with debt x_{j+1} and shocks z_{j+1} and u_{j+1} will choose to default. That is,

$$I(j+1, x_{j+1}, z_{j+1}, u_{j+1}) = 1,$$

if and only if

$$W^D(j+1, x_{j+1}, z_{j+1}, u_{j+1}) > W^R(j+1, x_{j+1}, z_{j+1}, u_{j+1}).$$

Therefore, $\pi^D(\cdot)$ is calculated at each age j as follows:

$$\pi^D(j, x_{j+1}, z_j) = \iint I(j+1, x_{j+1}, z_{j+1}, u_{j+1}) f(z_{j+1}, u_{j+1} | z_j) dz_{j+1} du_{j+1}. \quad (14)$$

Given $\pi^D(\cdot)$, the interest rate function, $R(j, x_{j+1}, z_j)$, is determined as follows:

$$R(j, x_{j+1}, z_j) = \frac{R^f + \psi}{(1 - \pi^D(j, x_{j+1}, z_j))}. \quad (15)$$

2. UNANTICIPATED SHOCKS AND THEIR SIZES

We now introduce unanticipated shocks to households in the above framework. We capture the effect of the shock on the “representative” household of any given age, as described by the age-specific median value of wealth.

As mentioned earlier, we quantify the effect of such shocks in terms of annual consumption. There are several ways in which we could do this. For ease of interpretation, we express all quantities in terms of *constant* consumption levels. We now describe the two scenarios under study and detail the particular calculations needed to derive the costs in terms of equivalent constant consumption levels under each.

Case 1: An Unanticipated Drop in Net Worth

The first case analyzes the consequences of an unanticipated drop in net worth. The empirical parallels we have in mind are unexpected decreases in house prices or stock prices. Since wealthier households are likely to have more expensive homes and larger stock portfolios, we assume that the shock is proportional to net worth.

The cost of this shock is calculated as follows. Let $V(j, \bar{x}_k, z_j, u_j)$ denote the value of arriving in a given period k with wealth \bar{x}_k . Let $V(j, \tilde{x}_k, z_j, u_j)$ denote the value of arriving in a given period k with wealth \tilde{x}_k , where, for $0 < \theta < 1$,

$$\begin{aligned} \tilde{x}_k &= (1 - \theta)\bar{x}_k \text{ if } \bar{x}_k \geq 0 \\ \tilde{x}_k &= (1 + \theta)\bar{x}_k \text{ if } \bar{x}_k < 0. \end{aligned}$$

The latter is the value associated with the discounted expected utility that a household can obtain from behaving optimally after the occurrence of the shock to net worth. Given this, we define \bar{c} and \tilde{c} as the constant values for consumption that, when received over an entire lifetime, generate discounted utility equal to $V(j, \bar{x}_k, z_j, u_j)$ and $V(j, \tilde{x}_k, z_j, u_j)$, respectively. Thus, \bar{c} and \tilde{c} solve, respectively,

$$V(j, \bar{x}_k, z_j, u_j) = \sum_{k=j}^{J+25} \beta^{j-k} \frac{\bar{c}^{1-\alpha}}{1-\alpha},$$

$$V(j, \tilde{x}_k, z_j, u_j) = \sum_{k=j}^{J+25} \beta^{j-k} \frac{\tilde{c}^{1-\alpha}}{1-\alpha}.$$

The difference $\bar{c} - \tilde{c}$ represents the number of units of consumption that would make the household indifferent to facing the shock or not. The difference thus represents the cost of the shock in units of consumption, or, alternately, the amount the household would pay in units of consumption to avoid facing the shock. (See the Appendix for calculation details.)

What are empirically plausible sizes of the net worth shocks? We may arrive at upper bounds by assuming that the household's entire net worth is composed of a single asset and attribute the shock to a drop in the value of that asset. This asset might be the household's equity in its home or its stock portfolio. To obtain an approximation of the upper bound on the size of shocks to house and stock prices, we use available data. The recent house price bust serves as a case in point. The largest annual decline of nearly 19 percent in the S&P/Case-Shiller Home Price Index of U.S. National Values since 1987 came between the first quarter of 2008 and the corresponding quarter of 2009.⁵ If we look at stocks, the shock could correspond to the worst one-year performance of diversified mutual funds such as Vanguard's S&P 500 Index or Total Stock Market Index, in which their value fell by roughly 37 percent. Different values may be appropriate for households with different profiles, but we carry out the exercise for a 40 percent drop in asset values, i.e., $\theta = 0.4$, to serve as an upper bound.

Case 2: An Unexpected Tightening of Credit Markets

Our second scenario aims to measure the cost imposed on a household by a sudden change in borrowing premia, and so intends to be reflective of dysfunction in credit markets very generally. We capture this in the model by comparing the maximal value that is attainable to an agent under the initial

⁵ See

www.standardandpoors.com/indices/sp-case-shiller-home-price-indices/en/us/?indexId=spusa-cashpidff-p-us.

interest rate function with that attainable from living in an environment where loans are costlier than before. Specifically, we model the shock as raising the interest rate on credit to $R(\cdot) + \iota$, $\iota > 0$. Importantly, we assume that agents will be faced with the tighter credit conditions *for the rest of their lives*. As a result, our calculations will likely represent an upper bound on the cost of such credit market tightening.

The net worth shock did not change household value functions. This is because the shock did not alter the subsequent uncertainty or costs in the household's environment. In this case, by changing the interest rate faced by households for the rest of their lives, the shock does change the maximal value of attainable utility coming from any given wealth position. Let $V(j, x_k, z_j, u_j)$ be the value function before the shock, when the interest on credit is $R(\cdot)$. Let $\tilde{V}(j, x_k, z_j, u_j)$ be the value function after the shock, derived from solving the household's optimization problem over its remaining life under the new credit-pricing function.

In this case, \bar{c} and \tilde{c} solve

$$V(j, x_k, z_j, u_j) = \sum_{k=j}^{J+25} \beta^{j-k} \frac{\bar{c}^{1-\alpha}}{1-\alpha},$$

$$\tilde{V}(j, x_k, z_j, u_j) = \sum_{k=j}^{J+25} \beta^{j-k} \frac{\tilde{c}^{1-\alpha}}{1-\alpha}.$$

As before, cost of shock is $\bar{c} - \tilde{c}$.

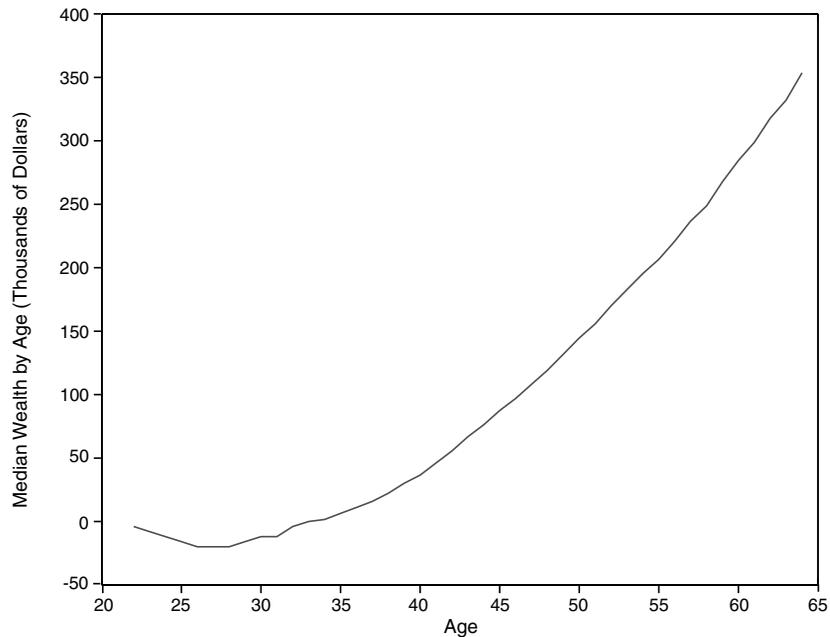
We increase borrowing costs at all debt levels by 300 basis points, which corresponds to among the largest spreads observed between mortgage market interest rates and 10-year Treasury securities.

3. RESULTS

The following section presents results for the two cases. In both scenarios, we impose the shock on households at one period and calculate the cost of the shock for households of various ages. The household does not expect, nor does it receive, any further shocks aside from the age at which we study them.

Notice that in our model, the presence of uninsurable risk will lead households to vary not only in income, but importantly, in consumption and wealth. There is, therefore, no "representative" household. This raises the issue of *whose* well-being and costs we are studying. A natural candidate, which we use, is the household with the median level of wealth and the median level of income shocks for its age group. Figure 1 shows median wealth for households at each age.

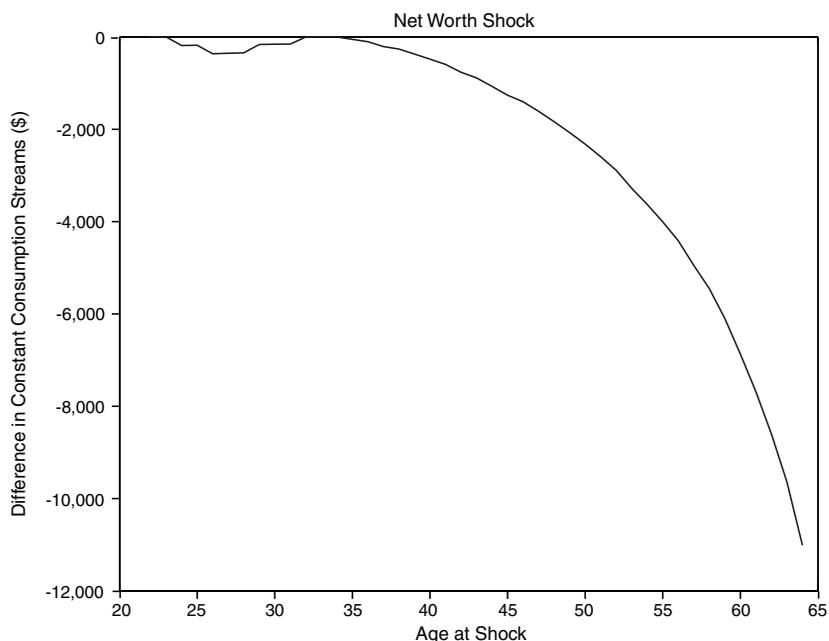
We find in general that the cost depends on the household's initial level of assets, the size of the shock, and the time period in which it occurred. Each case is discussed in detail.

Figure 1 Median Wealth by Age**Case 1: Net Worth Shock**

We calculate the cost of a 40 percent decrease in net worth to a household with the median level of assets and income shocks for its age. Figure 2 shows the cost of the shock to the household in dollars of constant annual consumption, which is calculated as $\tilde{c} - \bar{c}$. Figure 3 shows the cost as a fraction of constant annual consumption, calculated as $\frac{\tilde{c} - \bar{c}}{\bar{c}}$.

The cost as a function of age displays a U-shape and then a steep increase. The U-shape corresponds to ages at which the household has negative net worth. The proportional shock pushes the household deeper into debt, and is costliest when the household has the most debt (at age 26). Subsequently, the cost of the shock is small for a while. This is because the absolute value of wealth is low, so a proportional decline amounts to a very small number. However, the cost rapidly increases with age. As age increases, so does median wealth and the same percentage drop in net worth represents a much larger absolute loss. For example, the cost to a household that faces the shock at age 40 is \$481 in constant annual consumption for the rest of its life.⁶ For

⁶ All costs in this article are reported in 2010 dollars.

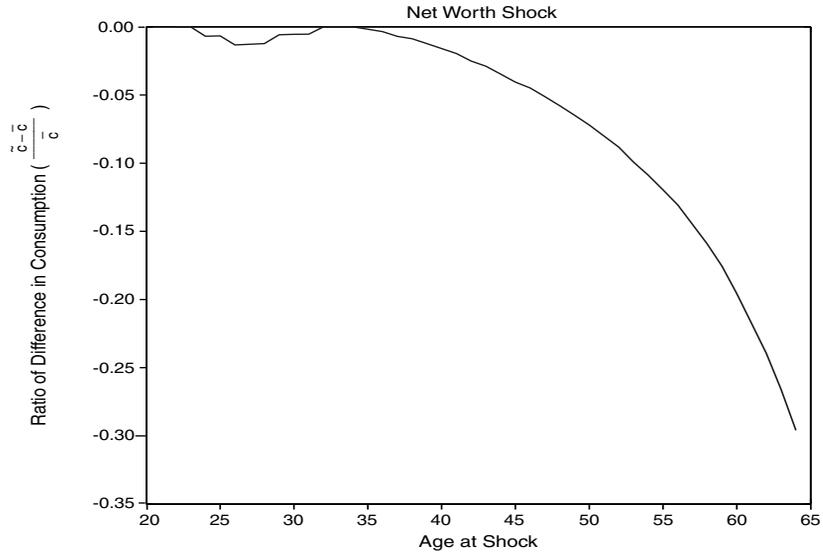
Figure 2 Difference in Consumption by Age at Shock: Case 1

a household at age 60, the cost of an unforeseen 40 percent reduction in net worth is \$6,870 per year for the rest of its life. Figure 3 represents the same cost in percentage terms. The cost to the 40-year-old household is 1 percent of annual consumption without the shock, while for the 60-year-old household, it is nearly 15 percent.

While older households face greater costs in terms of annual consumption, they also face them for fewer periods. It is therefore useful to compare the present value of the sequence $\{\bar{c} - \tilde{c}\}_{j=k}^{J+25}$ for various k , where k is the date of the shock. Calculating the present value in the first period of the model gives the perspective of one household looking ahead, while calculating the present value at date k compares the relative cost of the shock to older and younger households living in date k . Figure 4 shows the results of both calculations.⁷ Even in present value terms, the cost of the shock is highest for the oldest households.

⁷ See the Appendix for calculation details.

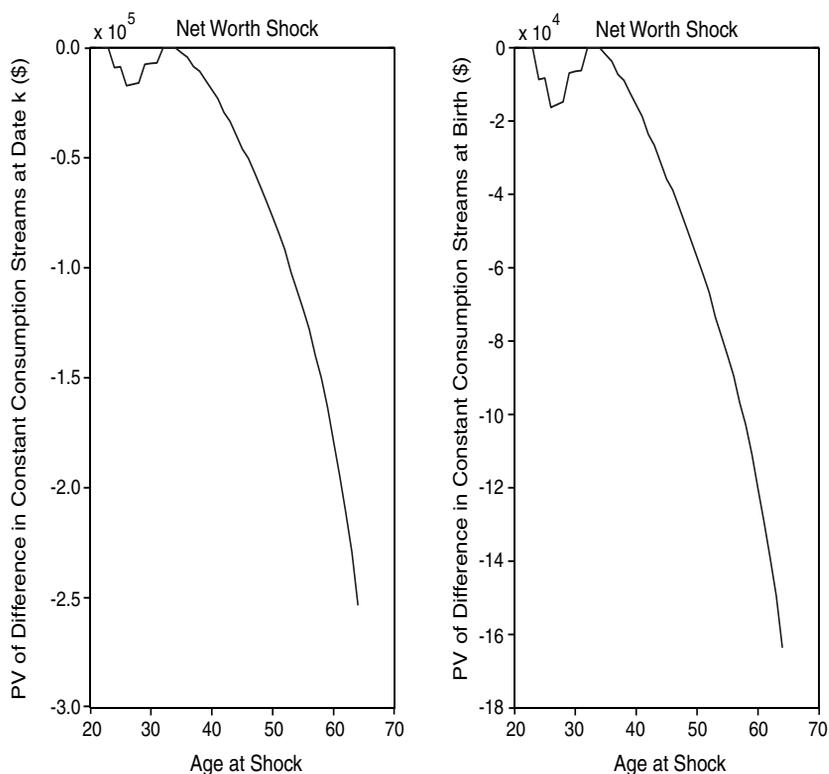
Figure 3 Percentage Difference in Consumption by Age at Shock: Case 1



Case 2: Interest Rate Shock

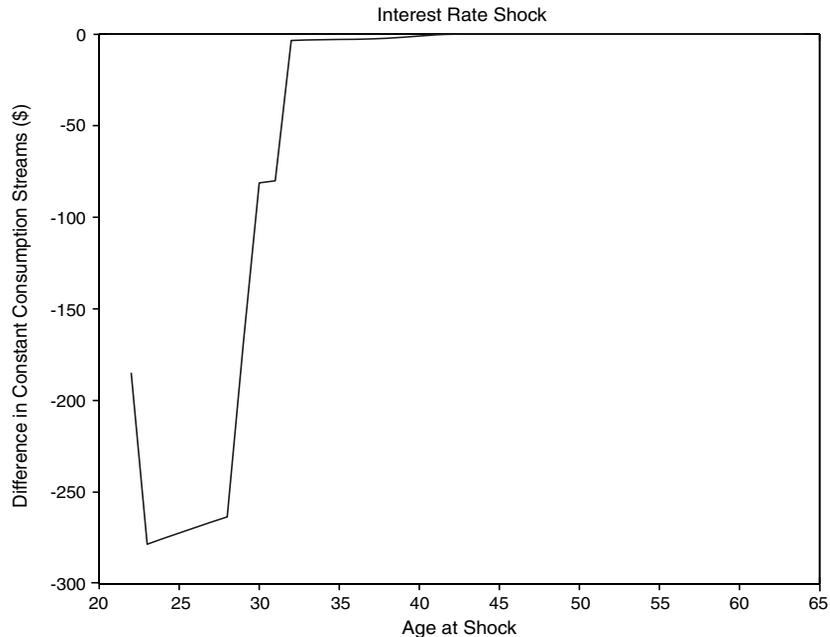
We now consider the shock to a household that comes from facing unexpectedly higher interest rates on credit for the rest of its working life. Recall that in this case, we measure the cost as arising from a 300-basis-point increase in the interest rate on all debt levels that the household might choose. The corresponding costs in terms of annual consumption, fraction of original consumption, and present value of annual consumption are shown in Figures 5–7.

The cost of this shock is very small relative to the net worth shock; it does not exceed \$300 of annual consumption for agents of any age. An important part of why this cost remains small is that households can adjust savings in the interim fairly effectively to nullify the effects of such an increase in costs. Moreover, as long as such a shock does not occur at the time when a household is holding peak debt (lowest net worth), the size of the shock itself is not large. Lastly, once households leave the first 15 years or so of their working life cycle, they are typically not in debt (have non-negative net worth), and, moreover, do not typically expect to return to such a state. Thus, contractions in credit markets will not hurt them.

Figure 4 Present Value of Difference in Consumption: Case 1

The costs we report have all been calculated under the assumption that households can declare bankruptcy and remove all unsecured debt obligations subject to a penalty. As a result, debt in the model is priced to reflect this possibility. Household net worth over the life cycle is, of course, different than what it would be in a setting where households did not have this option but instead had access to risk-free borrowing. As a check for robustness, we have shut down this option in the model by making the utility cost, λ , infinite, which effectively precludes bankruptcy.⁸ We find that this has little or no effect on the size of the costs. This is because the option to declare bankruptcy is relevant only to a subset of households—those with negative net worth (that do not have sufficient assets to pay off their debts). In our model, these are younger households. Because the debt they hold is not large on average, the proportional net worth shock translates into a small shock for them in absolute

⁸ Results are available from the authors upon request.

Figure 5 Difference in Consumption by Age at Shock: Case 2

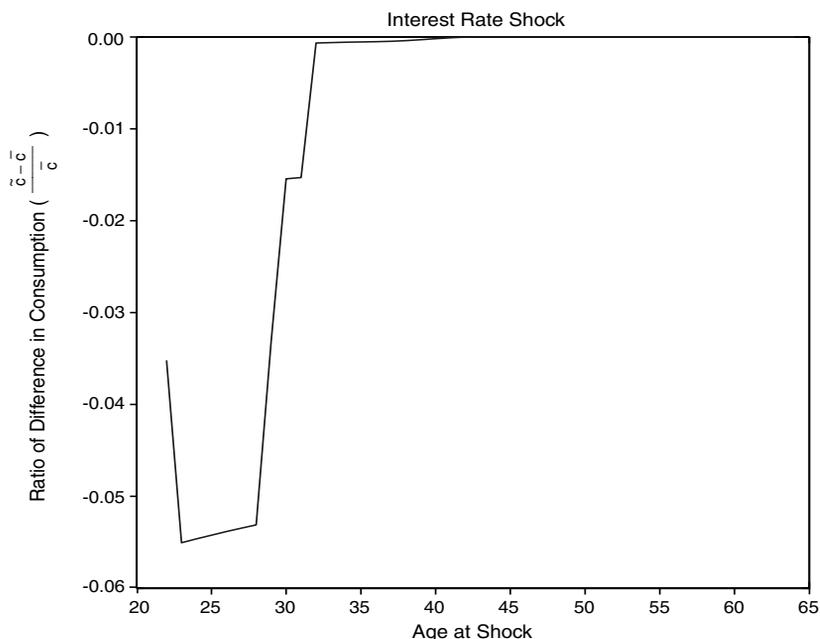
terms. Our results have shown that the cost of the interest rate shock is also small in absolute terms for these households. As a result, neither shock is large enough to make bankruptcy an important consideration for the examples we study.

We note here that the focus of the model on net worth is likely very important for the small role it assigns to the effects of an interest rate shock. In a richer setting, the fact that households are often engaged in very heavily leveraged investment (taking out mortgages to finance a home purchase) means that credit market costs could likely affect people well into late middle-age. This is simply because, while they might have positive net worth by middle-age (indeed will, in most instances), they may also owe substantial amounts on a mortgage, and the size of these obligations may be quite large and difficult to deal with. In ongoing work (Athreya, Ionescu, and Neelakantan 2011), we are considering precisely this.

4. CONCLUSION

In this article, we take a first step in measuring the cost of two particular and, we think, representative types of financial shocks. The results yield some general

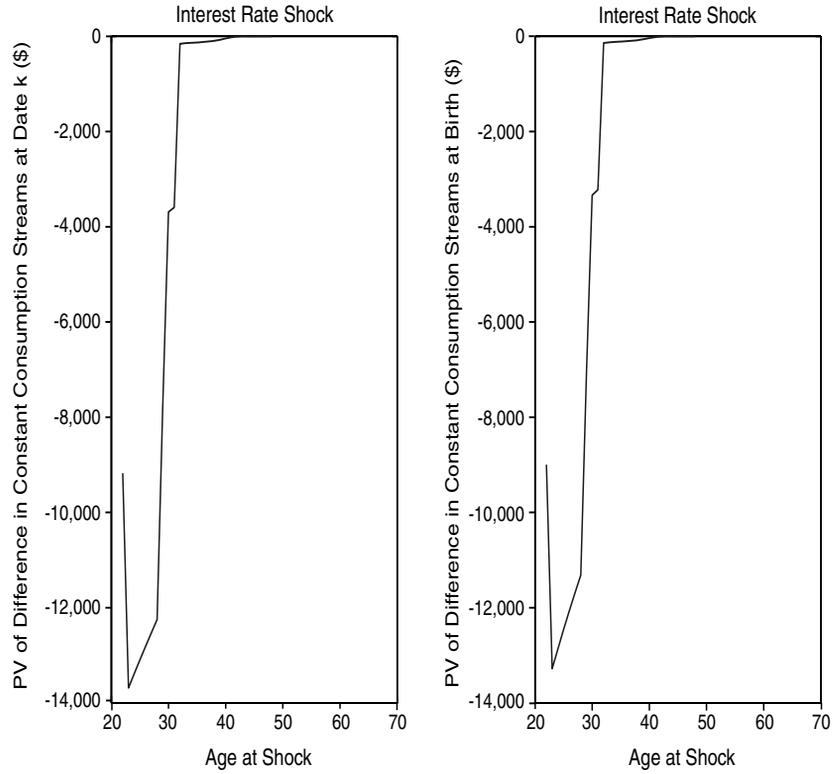
**Figure 6 Percentage Difference in Consumption by Age at Shock:
Case 2**



insights about such shocks and their costs. Comparing across households of various ages, shocks that are proportional to net worth are costliest to the oldest households for which the proportional shock translates into the largest absolute drop in net worth. Interest rate shocks, in the form of an unanticipated tightening of the credit market, are much less costly.

As mentioned at the outset, this article is stylized along several dimensions, and thus represents only a small first step in the important task of assessing the power of financial shocks to compromise household well-being. In particular, the model abstracts from portfolio choice and focuses instead on a simple scalar measure of net worth. This in turn prevents us from fully analyzing particular kinds of financial decisions, such as house purchases or any other leveraged purchase of risky assets, which can greatly change gross financial positions and periodic payment obligations while leaving net worth essentially unchanged. The model also abstracts from the labor supply decision, which could mitigate the cost of the shocks. Finally, the shocks that the household faces are completely unanticipated, something that is likely not as stark in reality. Households may well be aware of the existence of the kinds of shocks we analyze in this article, but incorrect about or unable to assess

Figure 7 Present Value of Difference in Consumption: Case 2



“true” probabilities for such events. In ongoing work (Athreya, Ionescu, and Neelakantan 2011), we consider these issues in a richer model of household portfolio choice.

APPENDIX

Solving for \bar{c} and \tilde{c}

$$V(j, \bar{x}_k, z_j, u_j) = \sum_{j=k}^{J+25} \beta^{j-k} \frac{\bar{c}^{1-\alpha}}{1-\alpha}.$$

First, we change the index of summation. Let $i = j - k$. Then

$$V(j, \bar{x}_k, z_j, u_j) = \sum_{i=0}^{J+25-k} \beta^i \frac{\bar{c}^{1-\alpha}}{1-\alpha}.$$

We use the following rule for the sum of a finite series:

$$\sum_{i=0}^n a^i = \frac{1 - a^{n+1}}{1 - a},$$

to obtain

$$V(j, \bar{x}_k, z_j, u_j) = \frac{\bar{c}^{1-\alpha}}{1-\alpha} \left[\frac{1 - \beta^{J+25-k+1}}{1 - \beta} \right].$$

Let

$$\Phi_k = \left[\frac{1 - \beta^{J+25-k+1}}{1 - \beta} \right].$$

Then

$$V(j, \bar{x}_k, z_j, u_j) = \frac{\bar{c}^{1-\alpha}}{1-\alpha} \Phi_k.$$

Solving for \bar{c} yields

$$\bar{c} = \left[\frac{V(j, \bar{x}_k, z_j, u_j)}{\Phi_k} (1 - \alpha) \right]^{\frac{1}{1-\alpha}}. \quad (16)$$

Similarly,

$$\tilde{c} = \left[\frac{V(j, \tilde{x}_k, z_j, u_j)}{\Phi_k} (1 - \alpha) \right]^{\frac{1}{1-\alpha}}. \quad (17)$$

Finding the Present Value of $\bar{c} - \tilde{c}$

To allow us to easily compare the cost of shocks at various ages, we now compute two types of present values. First, we begin by discounting *to age 0*, not just back the date at which the shock occurred (date k). The present value *at date zero* of a shock occurring at date k , given that the constant consumption equivalents are \bar{c} and \tilde{c} , is

$$PV_0(k) = \frac{\bar{c} - \tilde{c}}{(1+r)^k} + \frac{\bar{c} - \tilde{c}}{(1+r)^{k+1}} + \dots + \frac{\bar{c} - \tilde{c}}{(1+r)^{J+25}},$$

where r is the interest rate on savings. To be clear, notice that the first discounting term $\frac{1}{(1+r)^k}$ shows that we are discounting to age 0 events that begin at age k .

Let $\frac{1}{1+r} = \hat{\beta}$. Next, we'll use the known formula for the finite sum of a geometric series. We want the sum from age k to death (age $J + 25$). We

therefore first take the series from 0 to $J + 25$, $\left[\frac{1 - \hat{\beta}^{J+25+1}}{1 - \hat{\beta}} \right]$ and subtract from this the sum going from 0 to $k - 1$, $\left[\frac{1 - \hat{\beta}^k}{1 - \hat{\beta}} \right]$.

$$\begin{aligned} PV_0(k) &= (\bar{c} - \tilde{c}) \left[\frac{1 - \hat{\beta}^{J+25+1}}{1 - \hat{\beta}} - \frac{1 - \hat{\beta}^k}{1 - \hat{\beta}} \right] \\ &= (\bar{c} - \tilde{c}) \left[\frac{\hat{\beta}^k - \hat{\beta}^{J+25+1}}{1 - \hat{\beta}} \right] \\ &= (\bar{c} - \tilde{c}) \left[\frac{1 - \hat{\beta}^{J+25-k+1}}{1 - \hat{\beta}} \right] \hat{\beta}^k \\ &= (\bar{c} - \tilde{c}) \Phi_{diff_at_birth}, \end{aligned}$$

letting

$$\Phi_{diff_at_birth} = \left[\frac{1 - \hat{\beta}^{J+25-k+1}}{1 - \hat{\beta}} \right] \hat{\beta}^k,$$

we have

$$PV_0(k) = (\bar{c} - \tilde{c}) \Phi_{diff_at_birth}.$$

Similarly, the present value of a shock occurring at date k , discounted back to date k is

$$PV_k(k) = (\bar{c} - \tilde{c}) \Phi_{diff_at_k},$$

where we define

$$\Phi_{diff_at_k} = \left[\frac{1 - \hat{\beta}^{J+25-k+1}}{1 - \hat{\beta}} \right].$$

REFERENCES

- Athreya, Kartik B. 2008. "Default, Insurance, and Debt over the Life-Cycle." *Journal of Monetary Economics* 55 (May): 752–74.
- Athreya, Kartik B., Felicia A. Ionescu, and Urvi Neelakantan. 2011. "How Nasty are Financial Surprises? A View Over the Life-Cycle." Mimeo, Federal Reserve Bank of Richmond.
- Baxter, Marianne, and Robert G. King. 1993. "Fiscal Policy in General Equilibrium." *American Economic Review* 83 (June): 315–34.

- Gourinchas, Pierre-Olivier, and Jonathan A. Parker. 2002. "Consumption Over the Life Cycle." *Econometrica* 70 (January): 47–89.
- Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes. 1995. "Precautionary Saving and Social Insurance." *Journal of Political Economy* 103 (April): 360–99.
- Huggett, Mark. 1996. "Wealth Distribution in Life-Cycle Economies." *Journal of Monetary Economics* 38 (December): 469–94.
- Huggett, Mark, and Gustavo Ventura. 2000. "Understanding Why High Income Households Save More than Low Income Households." *Journal of Monetary Economics* 45 (April): 361–97.
- Joffe-Walt, Chana, and Adam Davidson. 2010. "Housing Guy Apologizes for Housing Bubble." Available at www.npr.org/blogs/money/2010/10/22/130756958/finance-guy-apologizes-for-housing-bubble.
- Storesletten, Kjetil, Christopher I. Telmer, and Amir Yaron. 2004. "Consumption and Risk Sharing Over the Life Cycle." *Journal of Monetary Economics* 51 (April): 609–33.