

Asset bubbles and the cost of economic fluctuations

Kyle Chauvin, David Laibson, Johanna Mollerstrom*

October 15, 2009

Preliminary draft

*Harvard University. Corresponding author: Johanna Mollerstrom. Prepared for a conference at Gerzensee Switzerland on the integration of macroeconomic and microeconomic perspectives. Laibson acknowledges financial support from the NSF and the NIA.

Abstract: Lucas (1987) calculates the cost of economic fluctuations to be less than 0.1% of consumption. In other words, we should be willing to pay no more than 0.1% of our (permanent) consumption to eliminate all future business cycle fluctuations. The current paper shows that Lucas' calculations miss the costs of fluctuations arising from asset bubbles. We identify two novel types of costs: consumption booms/busts due to asset price volatility, and consumption volatility due to asset trading during bubble periods. We show that these effects rely on heterogeneity. If assets are spread proportionately over the population, the resulting welfare costs scale down to the magnitudes derived by Lucas. We calibrate our model using the Health and Retirement Survey. Our benchmark calibration, which assumes a coefficient of relative risk aversion of 3, implies that the asset bubbles of the last decade generated a social welfare cost equal to a permanent 4% reduction in the level of national consumption. As expected, our calculations are sensitive to the details of the calibration, including the degree of balance sheet/trading heterogeneity, the coefficient of relative risk aversion, and the magnitude of the asset bubble. Our specifications with reasonable parameter values generate welfare costs ranging from 1% to 10% of (permanent) national consumption.

1 Introduction

Unlucky never trusted the stock market and held all of his retirement savings in fixed income accounts. But repeated nudges from his friends (telling him about the historical equity premium) led him to eventually shift into stocks in 2007. Because of his uncanny bad timing, Unlucky's financial assets had a *total* real return of -30% from 1995 to 2008.¹ Unlucky was also a late-comer to the housing market. He bought his first house near the peak of the housing bubble. He also paid transaction costs, including search costs, mortgage points, legal fees, and moving costs. Because of the subsequent decline in prices, Unlucky effectively realized a capital loss equal to his downpayment.²

In contrast, Lucky benefited from the asset bubbles. Because of her approaching retirement, Lucky shifted her retirement savings from stocks to bonds in the late 1990's. Because she sold at the peak of the Tech bubble, she pocketed the capital gains of the 1990's. The Treasury bonds and Treasury Inflation Protected Securities that she subsequently bought had a great run from 2000 to 2008. For her financial assets, the total real return from 1995 to 2008 exceeded 400%. Moreover, in 2006 Lucky sold her house and moved so she could live near her grandchildren. Since she wasn't familiar with the new city, she decided to rent. Hence, Lucky also benefited from the upside of the housing bubble and (accidentally) escaped the reversal.³

During the bubble years, Typical's asset dramatically appreciated and he consequently ramped up his consumption by 10%. To pay for this extra consumption he refinanced his home and extracted equity. When the stock and housing markets crumbled in 2008, he adjusted his consumption down accordingly. He had built up a substantial amount of collateralized debt during the bubble years and he now needed to reduce his consumption to rebuild his net worth. His new steady state consumption was about 2% below his initial starting point before the bubble began.

These three agents illustrate how volatile asset prices affect individual households in different ways. Naturally, these stories are not representative but the mechanisms described above did affect many U.S. households. The

¹Laibson takes responsibility for helping one friend reallocate most of his retirement savings to equities in 2007.

²Laibson's sibling owned a small condominium in the 1990's. She sold it in 2001 and acquired three homes (which she still owns, one for living and two for investment) between 2001 and 2006.

³Another of Laibson's relatives.

current paper models and calibrates these mechanisms and estimates their impact on U.S. consumers.

In essence, this paper studies two basic economic mechanisms. Unlucky and Lucky illustrate an *asset trading effect*. Because of lifecycle dynamics, active market timing, or other trading motives, some households are net buyers of assets and some households are net sellers of assets. Such transactions produce inter-household transfers when asset prices deviate from fundamentals.

Typical illustrates a different mechanism, which we refer to as a *boom/bust effect*. Households that own bubble-priced assets perceive that they are wealthier than they actually are. Consequently, they raise consumption during the bubble period: i.e., borrow more and lower active savings. When asset prices eventually return to their fundamental values, these households need to reduce their consumption to reflect their new net worth. This consumption reduction necessarily overshoots the *initial* consumption increase, since the households need to implicitly “pay back” the fraction of consumption during the bubble years that is discovered (ex-post) to be above the annuity value of their assets. This is just a dynamic implication of their budget constraint. In other words, during the bubble the agent overconsumes relative to the true annuity value of wealth. After the bubble, the underconsumes relative to the pre-bubble annuity value of wealth.

Asset trading effects and consumption boom/bust effects jointly generate excess consumption volatility. In the current paper we estimate these effects and calculate the resulting welfare costs. Three key expressions emerge in our analysis.

First, there is an asset trading effect, which we calculate using a Taylor expansion. The welfare cost arising from this channel – expressed as a fraction of permanent consumption – is given by

$$\frac{\gamma}{2} \zeta^2 \text{Var} \left(\frac{k_i}{W_i} \right) e^{\rho N}. \quad (1)$$

In this expression, γ is the coefficient or relative risk aversion, ζ is the magnitude of the asset bubble as a fraction of the fundamental value of the capital stock, ρ is the annual discount rate, and N is the duration of the bubble (years). Finally, k_i is the net value of capital (using pre-bubble prices) purchased by household i during the bubble, and W_i is the total net worth of household i (including human capital). In our benchmark calibration, the asset trading effect represents a cost that is equivalent to about 2.8%

of consumption. Note that the asset trading effect has no first-order consequences, since every gain from selling the overpriced asset is offset by a loss from buying the overpriced asset. However, second order effects survive since resulting consumption volatility lowers welfare. Concavity in the utility function causes the gains in marginal utility to be more than offset by the losses in marginal utility. Also, note that the asset trading effect vanishes as heterogeneity is reduced – i.e. as $Var\left(\frac{k_i}{W_i}\right)$ goes to zero.

Second, we identify a boom/bust effect. The welfare cost arising from this channel – expressed as a fraction of permanent consumption – is given by

$$\frac{\gamma}{2}\zeta^2(1-p)^{-1}\left(\frac{K}{W}\right)^2\exp\left[Var\left(\ln\frac{K_i}{W_i}\middle|K_i>0\right)\right](e^{\rho N}-1). \quad (2)$$

The new variables include: p , the fraction of households that hold zero claims to the aggregate capital stock; K , the aggregate stock of capital; K_i , the net claim to the capital stock of household i ; and W , the aggregate net worth of the agents in the economy (including human capital). In our benchmark calibration, boom/bust effects account for welfare costs equal to 2.3% of consumption. Note that the boom/bust effect also has no first-order consequences, since every first-order gain from overconsuming (relative to the true annuity value of wealth) during the bubble is offset by a first-order loss from underconsuming (relative to the pre-bubble annuity value of wealth) after the bubble bursts. However, second order effects survive since intertemporal consumption volatility that results from the boom/bust cycle lowers welfare. Finally, note that the boom/bust effect drops by an order of magnitude as heterogeneity is reduced – i.e. as p goes to zero and as $Var\left(\ln\frac{K_i}{W_i}\middle|K_i>0\right)$ goes to zero. At this homogeneous limit, the boom bust effect is only 0.4% of consumption in our benchmark calibration. Hence, heterogeneity is also a key contributor to this effect.

Third, we identify a covariance effect, which arises because of interactions between the previous two effects.

$$\gamma\zeta^2 Cov\left(\frac{k_i}{W_i}, \frac{K_i}{W_i}\right)(e^{\rho N}-1)$$

In our calibration this covariance effect represents a cost that is equivalent to about -1.1% of consumption. In other words, this third effect turns out to partially *offset* the other two. Mean reversion in portfolio allocation induces

a negative covariance between $\frac{k_i}{W_i}$ and $\frac{K_i}{W_i}$, and this reduces the welfare costs. Intuitively, households that tend to allocate the greatest portfolio share to domestic capital, K_i , when the bubble begins, and hence are likely to raise their consumption the most during the bubble, are also the households that probabilistically sell the most capital during the bubble period, thereby partially cushioning the fall in consumption when the bubble bursts.

The paper proceeds as follows. In Section 2 we review the related literature. In Section 3 we describe our model. In Section 4, we work out the welfare costs of an asset bubble, using both exact methods and a second-order Taylor expansion. In Section 5, we calibrate the model, including a discussion of micro-level data from the Heath and Retirement Study. Section 6 presents our welfare cost results – both exact results and second-order approximations. Section 7 concludes.

2 Literature Review

Real Business Cycle models imply that fluctuations do not reduce welfare, since fluctuations are optimal responses to changing fundamentals (e.g., Prescott, 1986). Along similar lines, Tirole (1985) shows that fluctuations in the form of asset price bubbles can be rational and sometimes even increase welfare (see also Caballero et al, 2009).

However, most economic models imply that macroeconomic fluctuations reduce welfare. The starting point for the measurement of these costs is a seminal contribution by Lucas (1987).⁴ Lucas considers a representative agent with constant relative risk aversion. The agent’s consumption stream is represented as $C_t = \lambda^t(1 + \varepsilon_t)C_0$ where $\lambda > 1$, where the innovations, ε_t , are independently and identically distributed. Lucas uses US data from the period after World War II to calibrate the model and to estimate how much the representative agent would be willing to give up as fraction of consumption in order to set $\varepsilon_t = 0$, in other words, to eliminate all fluctuations. Lucas concludes that, for reasonable values of the coefficient of risk aversion, the welfare cost of economic fluctuations is very low, in fact not more than 0.1%

⁴A related conclusion was reached by Cochrane (1989) who studied the welfare costs for consumers when they make small mistakes and therefore deviate from the optimal consumption path. Cochrane estimated that these “near-rational mistakes” do not incur costs that are bigger than \$1 per quarter.

of consumption.⁵

Numerous papers have studied variations of Lucas' models, to evaluate the robustness of his result. There are several papers that study different classes of preferences, including Krep-Porteus or Epstein-Zin preferences that admit non-time separability. However, this has generally not generated higher welfare costs (e.g., Dolmas 1998, and Otrok 2001).

A second approach has assumed that agents are heterogeneous, for example with respect to risk aversion, wealth or position in the labor market. Such heterogeneity raises the welfare costs relative to Lucas' estimates, but the average welfare cost is generally estimated to be about 1% (e.g., Krusell et al, 2009, Krusell and Smith, 1999, Krebs, 2007, Mukoyama and Sahin, 2006 and Storesletten et al, 2001).

Lucas' estimates are based on an assumption of complete markets, which admits the diversification of all idiosyncratic risk. A third approach is to try to generate sizable welfare effects by allowing for incomplete markets and liquidity constraints. These models have been built on, among other things, an inability to insure against labor market risk, see e.g. Imrohoroglu, 1989 and Atkeson and Phelan, 1994. However, the introduction of incomplete markets does not generate large welfare losses. The estimates of Imrohoroglu are similar to Lucas' and the estimates of Atkeson and Phelan are even smaller than those of Lucas.

Fourth, there have been attempts to use asset-pricing models to estimate the welfare costs of fluctuations. Tallarini (2000) estimates welfare costs that are several orders of magnitude larger than those of Lucas. However, these methods implicitly assume unreasonably high risk aversion.

Finally there are a few studies that look at the potential link between growth and economic fluctuations. In Lucas' paper, the constant λ , which captures the trend in growth, is assumed to be unaffected by fluctuations. It is possible, however, that the trend in growth would be higher with less fluctuations. Barlevy (2004) builds a model in which investment is assumed to be less efficient if the economy is more volatile. In this model consumers are willing to sacrifice as much as 10 percent of consumption in order to eliminate fluctuations.

⁵See also Lucas (2003) and Barlevy (2004).

3 Model

We analyze a continuous-time, small open economy that faces fixed world factor prices (cf. Laibson and Mollerstrom 2009). Heterogeneous households, with an index i that is temporarily suppressed, maximize an exponentially weighted integral of utility flows:

$$\int_0^{\infty} \exp(-\rho t) u(C_t) dt,$$

subject to non-stochastic dynamics

$$dK_t = Y_t^L - C_t + r(K_t - D_t - B_t) + dD_t + dB_t.$$

Here ρ is the exponential discount rate, K_t is domestic capital, Y_t^L is fixed labor income, C_t is consumption, r is the real interest rate, D_t is net foreign debt (so $-D_t$ is net foreign assets), and B_t is net domestic debt. Across households, B_t adds up to zero. We assume that $r = \rho$, which is a standard steady state restriction. Finally, we assume that households have constant relative risk aversion, γ .

$$u(C_t) = \frac{C_t^{1-\gamma} - 1}{1-\gamma}.$$

We assume that an asset bubble begins an instant after date $t = 0$ and ends at date N . In other words, the bubble exists when $t \in (0, N]$. The pre-bubble state of the economy will be the benchmark to which we keep referring. Hence, we adopt special notation for all variables at date zero. Specifically, whenever we drop the time subscript, we are implicitly referring to the date 0 (pre-bubble) value of the variable. For example, we let $C = C_0$.

We assume that the asset bubble immediately raises the notional value of fixed capital by an increment ζK . We conceptualize the bubble as the discounted value of productivity gains that are anticipated to occur N periods away.⁹ Agents expect a unit of domestic capital to yield returns of r from date 0 to date N , and returns of

$$r [1 + \exp(rN)\zeta]$$

⁹These productivity gains need to be anticipated to occur in the future, to enable the bubble to persist in the meantime. The bubble bursts on the data that the anticipated productivity gains fail to be realized.

thereafter. So a physical unit of domestic capital has price $(1 + \zeta)$ when the bubble begins an instant after date $t = 0$.

The marginal propensity to consume out of wealth is r , so that the ζK increase in notional wealth leads consumption to rise by $r\zeta K$.

Since, $r = \rho$, pre-bubble consumption is equal to the annuity value of wealth:

$$C = Y^L + r(K - D - B).$$

Bubble consumption is the annuity value of bubble-inclusive wealth:

$$C' = Y^L + r(K - D - B) + r\zeta K.$$

Recall that assets appreciate at rate r , so throughout the bubble period households hold wealth with notional value $\zeta K + K - D - B$. Capital gains and dividends are exactly offset by consumption. Without loss of generality we set B at date zero equal to 0.

We also allow our agents to trade assets. This will only matter during the bubble period: agents who buy domestic capital will be harmed (since the asset is over-valued) and agents that sell domestic capital will benefit.

To track this trading, we introduce a new variable k , which represents the net change in the physical units of bubble asset accumulated by an agent. Negative values of k represent net reduction in the physical units of the bubble asset. If we norm the pre-bubble real price of the bubble asset to 1, then an agent who pays $1 + \zeta$ dollars to buy domestic capital during the bubble period will take possession of 1 extra physical unit of K . For this illustrative example, $k = 1$. Since domestic capital is only held by domestic agents, it follows that the average value of k , across all households, is 0.

Agents who buy domestic capital (during the bubble period) are left with a rude shock when the bubble bursts. They experience an additional capital loss $\exp(rN)\zeta k$ at the time the bubble bursts. This is the additional capital loss that arises from acquiring k physical units of domestic capital. Without loss of generality we assume that domestic agents finance purchases of incremental domestic capital with domestic borrowing. Likewise, a domestic agent who sells an incremental unit of domestic capital uses the proceeds of this sale to make domestic loans. For example, an agent who buys k units of physical capital borrows $\Delta B = (1 + \zeta)k$ on the domestic market to do so. Assume that the agent rolls this debt over during the bubble, so the difference in B is equal to $\exp(rN)(1 + \zeta)k$.

During the bubble period, households have a gap between their desired level of consumption $C' = Y^L + r(K - D - B) + r\zeta K$ and the level of physical income from domestic assets $Y^L + r(K - D - B)$. The gap, $r\zeta K$, is borrowed as a flow from abroad.¹⁰ The resulting change to the trade deficit (which is the same as the initial change to the current account deficit) is

$$r\zeta K.$$

since the k 's average out to zero.

The trade deficit continues at this level throughout the bubble period. By contrast, the current account deficit grows, since the foreign debt is growing. Households must also pay interest on the accumulating shortfalls. Integrating these flows yields the net accumulation of foreign debt during any sub-period, $(t_0, t_0 + \tau]$, of the bubble period (where $\tau < N$).

$$\begin{aligned} D_{t_0+\tau} - D_{t_0} &= \int_0^\tau r\zeta K \exp(r[\tau - s]) ds \\ &= [\exp(r\tau) - 1] \zeta K \end{aligned}$$

So the (change in the) current account deficit from date 0, an instant before the bubble starts, to date $\tau < N$, is

$$|CA_\tau - CA| = r\zeta K + r[\exp(r\tau) - 1] \zeta K.$$

Just before the bubble bursts, at time N , assets at the household level can be decomposed into domestic assets valued at

$$\zeta (K + k) \exp(rN) + (K + k),$$

debt to foreign agents valued at

$$D + [\exp(rN) - 1] \zeta K.$$

and debt to domestic agents valued at

$$\exp(rN)(1 + \zeta)k.$$

¹⁰We assume that households borrow from foreign agents rather than selling the foreign agents over-valued domestic assets. The domestic agents have no reason to sell the domestic assets since they don't recognize that they are overvalued. Moreover, if the over-valued assets have value that is best-realized by local owners (e.g., residential real estate), then there are good reasons to expect that the foreign agents will primarily acquire fixed income claims.

Note that the net wealth is

$$\zeta K + K - D.$$

When the bubble bursts, the household is left with net assets:

$$K - D - [\exp(rN) - 1] \zeta K - \exp(rN) \zeta k.$$

So consumption falls from

$$Y^L + r(K - D) + r\zeta K$$

to

$$Y^L + r(K - D) - r [\exp(rN) - 1] \zeta K - r \exp(rN) \zeta k$$

In other words, consumption falls by

$$r\zeta (K + k) \exp(rN).$$

We can decompose this into three effects:

$$\underbrace{r\zeta K + r\zeta k}_{\text{Direct bubble effect} + \text{Trading effect}} \times \underbrace{\exp(rN)}_{\text{Accumulated debt effect}}.$$

The direct bubble effect is the reversal in the initial consumption boom. The trading effect is the additional reduction in wealth associated with loss (or gain) on assets that have been acquired during the bubble period. The accumulated debt effect is the consequence of accumulating debt at interest rate r over an interval of length N . As the duration of the bubble goes to zero, the accumulated debt effect ceases to matter: $\lim_{N \rightarrow 0} \exp(rN) = 1$. As the duration of the bubble increases, the accumulated debt effect grows exponentially.

3.1 Distributional assumptions

We now study inter-household differences. Consequently, we will start using household subscripts. The log normal distribution that we adopt is supported by empirical analysis that follows in the calibration section. Suppose that the initial distribution of capital follows a two-part distribution.

A mass p of consumers have $K_i = 0$. The remaining mass $1 - p$ has a log normal distribution of K_i levels. Specifically,

$$K_i = \begin{cases} 0 & \text{with probability } p \\ C_i \exp(\mu + \varepsilon_i - \sigma_\varepsilon^2/2) & \text{with probability } 1 - p \end{cases}$$

where ε_i is normally distributed with mean zero and variance σ_ε^2 . Moreover, ε_i is independent of C_i and iid across households. This implies that,

$$\begin{aligned} \int K_i di &= (1 - p) \int C_i \exp(\mu + \varepsilon_i - \sigma_\varepsilon^2/2) di \\ \int K_i di &= (1 - p) \exp(\mu) \int C_i di \\ \mu &= \ln \int K_i di - \ln \int C_i di - \ln(1 - p) \end{aligned}$$

We will exploit this relationship when we calibrate our economy. In essence, μ is the natural log of the ratio of aggregate domestic capital to aggregate consumption minus the natural log of the fraction of households that hold domestic capital. In our calibration, we also truncate the right-hand-tail of the log-normal density to prevent extreme welfare costs for households with large positions in the bubble asset.

Finally, we need to characterize the distribution of trades, k . To do this, begin by thinking about the Markov process that relates K_i to K'_i . Suppose that at each iteration of this Markov process, a fraction of the population $1 - \phi$ stays at their old level of domestic capital, and the remainder ϕ adjusts their domestic capital. For simplicity, we assume that the adjusting households are randomly dropped into the same initial distribution of capital holding. Specifically,

$$K'_i = \begin{cases} K_i & \text{with probability } 1 - \phi \\ 0 & \text{with probability } \phi p_0 \\ C_i \exp(\mu + \varepsilon'_i - \sigma_\varepsilon^2/2) & \text{with probability } \phi(1 - p_0) \end{cases}$$

We can use this ergodic assumption to back out the distribution of the transaction variable

$$k_i = K'_i - K_i.$$

4 Welfare calculations

We first characterize the welfare of an agent in this economy. The agent has consumption of

$$Y^L + r(K - D) + r\zeta K = C + r\zeta K$$

during the bubble and

$$\begin{aligned} Y^L + r(K - D) - r[\exp(rN) - 1](\zeta K) - r\exp(rN)\zeta k \\ = C - r[\exp(rN) - 1](\zeta K) - r\exp(rN)\zeta k \end{aligned}$$

after the bubble. So utility is given by

$$\begin{aligned} \int_0^N \exp(-\rho t) u(C + r\zeta K) dt \\ + \int_N^\infty \exp(-\rho t) u(C - r[\exp(rN) - 1](\zeta K) - r\exp(rN)\zeta k) dt \\ = \frac{1 - \exp(-\rho N)}{\rho} u(C + r\zeta K) \\ + \frac{\exp(-\rho N)}{\rho} u(C - r[\exp(rN) - 1](\zeta K) - r\exp(rN)\zeta k) \end{aligned}$$

We can compare this to the counterfactual of a bubble-free economy. Without the bubble, lifetime utility would have been

$$\frac{1}{\rho} u(C).$$

4.1 Exact welfare calculations

We first provide an exact calculation of the welfare costs engendered by the bubble. As noted before, we are only measuring the welfare costs associated with fluctuations (and ignoring any welfare costs associated with the loss of resources). For any individual agent, we can solve for the factor λ_i that equates the welfare they received as a result of the bubble episode and the welfare they would have received had they instead simply scaled their pre-bubble consumption by λ_i . The implicit equation for λ_i is given below.

$$\begin{aligned} [1 - \exp(-\rho N)] u(C_i + r\zeta K_i) \\ + \exp(-\rho N) u(C_i - r[\exp(rN) - 1](\zeta K_i) - r\exp(rN)\zeta k_i) \\ = u(\lambda_i C_i) \end{aligned}$$

Rearranging this expression, yields a closed-form expression for λ_i :

$$\left\{ [1 - \exp(-\rho N)] \left(1 + r \zeta \frac{K_i}{C_i} \right)^{1-\gamma} + \exp(-\rho N) \left(1 - r [\exp(rN) - 1] \zeta \frac{K_i}{C_i} - r \exp(rN) \zeta \frac{k_i}{C_i} \right)^{1-\gamma} \right\}^{\frac{1}{1-\gamma}}$$

We can also relate the individual λ_i coefficients to an aggregate λ that has the property that if every agent's consumption pre-bubble were uniformly scaled down by λ , then the equilibrium path welfare (with the bubble) would equal the counterfactual welfare resulting from scaled consumption (without the bubble). More formally, λ is given by the following equation:

$$\int \frac{u(\lambda C_i)}{\rho} di = \int \frac{u(\lambda_i C_i)}{\rho} di$$

Constant relative risk aversion implies that

$$\int \lambda^{1-\gamma} \frac{u(C_i)}{\rho} di = \int \lambda_i^{1-\gamma} \frac{u(C_i)}{\rho} di.$$

Hence,

$$\lambda = \left[\frac{\int \lambda_i^{1-\gamma} u(C_i) di}{\int u(C_i) di} \right]^{\frac{1}{1-\gamma}}.$$

When all of the households are ex-ante identical, so that $C_i = C_j$ for all i, j pairs, this reduces to

$$\lambda = \left(\int \lambda_i^{1-\gamma} di \right)^{\frac{1}{1-\gamma}}.$$

4.2 Taylor Approximation of Welfare Cost

We now use a second-order Taylor expansion to calculate the welfare loss from the asset bubble. We want to find Δ such that

$$\begin{aligned} & [1 - \exp(-\rho N)] u(Y^L + r(K - D) + r\zeta K) \\ & + \exp(-\rho N) u(Y^L + r(K - D) - r[\exp(rN) - 1](\zeta K) - r \exp(rN)\zeta k) \\ & = u(\Delta + Y^L + r(K - D)) \end{aligned}$$

We are interested in solving for

$$-\frac{\Delta}{C} = \Lambda,$$

which is the permanent percent reduction in pre-bubble consumption that produces a reduction in welfare equivalent to the experience of the bubble. Hence,

$$\Lambda_i = 1 - \lambda_i$$

We expand the argument of the utility function around $C = Y^L + r(K - D)$. Since this expansion is algebraically intensive and conceptually routine, we provide the details in the appendix. For an individual household the second-order expansion yields,

$$\Lambda_i = \frac{\gamma}{2} \left(\frac{r\zeta}{C_i} \right)^2 [K_i^2 (e^{rN} - 1) + 2k_i K_i (e^{rN} - 1) + k_i^2 e^{rN}].$$

We can also integrate across households to produce an average value for Λ_i . This derivation exploits the fact that the average value of k_i is zero. The average welfare cost, $\int \Lambda_i di$, is given by

$$\frac{\gamma}{2} \left(\frac{r\zeta}{C} \right)^2 [EK_i^2 (e^{rN} - 1) + 2Cov(k, K) (e^{rN} - 1) + Var(k)e^{rN}].$$

This can be broken down into three terms. First, there is a consumption boom/bust effect:

$$\begin{aligned} \frac{\gamma}{2} (r\zeta)^2 (1-p)^{-1} \left(\frac{K}{rW} \right)^2 \exp(\sigma_\varepsilon^2) (e^{rN} - 1) \\ = \frac{\gamma}{2} \zeta^2 (1-p)^{-1} \left(\frac{K}{W} \right)^2 \exp(\sigma_\varepsilon^2) (e^{rN} - 1). \end{aligned}$$

Second, there is a covariance effect:

$$\begin{aligned} \gamma (r\zeta)^2 \frac{Cov(k_i, K_i)}{r^2 W^2} (e^{rN} - 1) &= \gamma \zeta^2 \frac{Cov(k_i, K_i)}{W^2} (e^{rN} - 1) \\ &= \gamma \zeta^2 Cov\left(\frac{k_i}{W_i}, \frac{K_i}{W_i}\right) (e^{rN} - 1). \end{aligned}$$

Third, there is an asset trading effect:

$$\begin{aligned} \frac{\gamma}{2} (r\zeta)^2 \frac{Var(k_i)}{r^2 W^2} e^{rN} &= \frac{\gamma}{2} \zeta^2 \frac{Var(k_i)}{W^2} e^{rN} \\ &= \frac{\gamma}{2} (r\zeta)^2 Var\left(\frac{k_i}{W_i}\right) e^{rN}. \end{aligned}$$

We calibrate and compare our measures of welfare loss in the next section of the paper.

5 Calibration of the model

For tractability, we study the steady state, in which $r = \rho$. We therefore set

$$1.05 = \exp(r) = R = \exp(\rho).$$

Our baseline value for the coefficient of relative risk aversion is $\gamma = 3$. We also consider CRRA values from 0 to 5.

We need to calibrate the (plausible) magnitude of the U.S. asset bubbles. Figure 1 plots the U.S. ratio of household wealth to GDP.¹¹ This series fluctuated historically in a range roughly between 3 and 3.5 units of GDP. Starting in the mid-1990's, however, the series broke from this historical range and rose sharply. At its peak in the second quarter of 2007, the series reached a value of 4.7 units of GDP. By the first quarter of 2009, the series had fallen back to just above its historical range. These comparisons imply an estimated peak bubble value of about 1 unit of GDP (\$14 trillion). However, the 2007 ratio of household wealth to GDP misses part of the value of the bubble in 2007, since it nets out the value of debt accumulated to finance consumption during the preceding bubble years. The key contributor is mortgage debt, which increased from 0.44 units of GDP in 1996:1 to 0.75 units of GDP in 2007:2. This analysis implies an additional bubble increment of 0.3 units of GDP.

We estimate the total value of the bubble as 1/3 of the value of all domestic capital (including land). The Federal Reserve Balance Sheets (B.100) report that the total value of household tangible assets and equity-based assets was nearly \$50 trillion at year-end 1998 (in 2009 dollars). Hence, we set $\zeta = 1/3$. In 1998, total U.S. Consumption was about \$10 trillion (2009 dollars), so we set


$$\frac{K}{C} = \frac{K}{\rho W} = 5.$$

We calibrate the distribution of K_i from the Health and Retirement Study. Though this database has the limitation that it only covers respondents who are middle-aged or older, the HRS has two offsetting advantages. Respondents are surveyed longitudinally every two years. Moreover, the HRS asset data is of higher quality than the asset data in the Panel Survey on Income Dynamics.⁶

¹¹The numerator is compiled by the Federal Reserve and is available back to 1952.

⁶In our next draft of this paper, we will also use the Survey of Consumer Finances, which is a representative sample of the entire U.S. population.

Figure 1: Household wealth divided by GDP 1952:1-2009:2, US




Source: Federal Reserve Board, BEA, and authors calculations.

First, we use the HRS to construct a household-level estimate of $\frac{K_i}{W_i}$, where W_i is total household resources, including an estimate of the value of human capital. In our model it should be the case that $\frac{K_i}{W_i} = \frac{\rho K_i}{C_i}$. This procedure is explained in our second appendix. Figure 2 plots the distribution of $\frac{K_i}{W_i}$, for seven different waves of HRS surveys (starting in 1992/93 and ending in 2006). The distributions all have two distinct components. First, there is substantial mass at zero (about 15% of the HRS households have no K assets). Second, a log-normal distribution fits the data that is greater than zero. To confirm this second parametric property, we analyze the households with $K > 0$, and plot the natural log of their $\frac{K_i}{W_i}$ ratios. Figure 3 plots these distributions for the same seven waves of HRS surveys. Figure 3 also superimposes a Gaussian density to confirm our parametric assumptions. In our actual simulations (for exact calculations of welfare losses), we truncate the log normal distribution so that $\frac{K_i}{C_i} \leq 30$. This is done to rule out extreme welfare losses (for households with extreme values of K).


Our natural log plots have associated standard deviations that range from 0.66 to 0.74 (depending on the Wave of the HRS). We therefore adopt a benchmark standard deviation, σ_ε , of 0.70 for our calculations. The U.S. Census Bureau reports that 65% of U.S. households own their own home. We therefore assume that $p = 0.30$: i.e. 30% of households own neither a home nor equity.⁷

Finally, we turn to asset ownership dynamics. To provide a simple but workable calibration, we assume that $1 - \phi = 0.5$ of the households did not change their physical claims K_i during the duration of the bubble period. The remaining households (mass ϕ) trade their claims during the N -year bubble period, replacing their initial ratio K_i/W_i with a new iid draw from the stable distribution of K_i/W_i . This assumption produces a simulated correlation between K_i/W_i and K'_i/W_i , of 0.50 (a numerical coincidence). Note that K'_i/W_i is unit claims to domestic capital of household i at the end of the N -year bubble divided by initial net worth of household i . The actual empirical correlation (using the HRS) is much lower: 0.26. If we raised ϕ to match this empirical correlation, our imputed welfare costs would be even higher (since more trading increases the magnitude of the asset trading effect). However, we believe that the low empirical correlation is partly due to measurement error in the HRS. For this reason we would be biasing our imputed welfare


⁷About half of U.S. households own equity either directly or indirectly through a mutual fund (SCF). The overwhelming majority of these households own their own home.




(a) 1992/3: Mean = 0.1792, SD = 0.1725




(b) 1994/5: Mean = 0.1839, SD = 0.1717




(c) 1998: Mean = 0.2206, SD = 0.2110




(d) 2000: Mean = 0.2256, SD = 0.2120



(e) 2002: Mean = 0.2146, SD = 0.2031




(f) 2004: Mean = 0.2002, SD = 0.1895

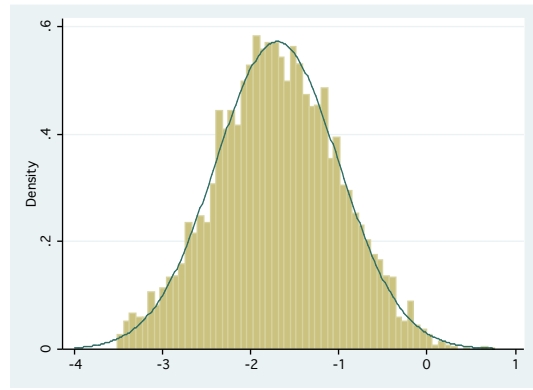


(g) 2006: Mean = 0.1945, SD = 0.1764


Figure 2: Distribution of K/W .




(a) 1992/3: Mean = -1.7125, SD = 0.7030




(b) 1994/5: Mean = -1.6976, SD = 0.6966




(c) 1998: Mean = -1.4924, SD = 0.7311




(d) 2000: Mean = -1.4804, SD = 0.7228



(e) 2002: Mean = -1.5128, SD = 0.7061



(f) 2004: Mean = -1.5600, SD = 0.6979



(g) 2006: Mean = -1.5732, SD = 0.6597

Figure 3: Distribution of $\ln(K/W)$ for non-zero households.

costs up if we picked an ϕ value that was high enough to match the empirical correlation of 0.26.

6 Results

Table 1 reports our benchmark calibration values for the key parameters that we will vary. Table 1 also reports a low/high range for each variable. Table 2 reports three additional variables that we will hold fixed in all of our simulations.

Table 1: Critical Parameters		low	benchmark	high
γ	CRRRA	1	3	5
ζ	bubble magnitude	1/6	1/3	1/2
N	bubble duration	8	10	12
σ_ε	standard deviation $\ln(K_i/W_i)$, $K_i > 0$	0.5	0.7	0.9
ϕ	Fraction of households trading	0.3	0.5	0.7

Table 2: **Fixed Parameters**

$\rho = r$	discount rate and real interest rate	0.05
$K/C = K/(\rho W)$	aggregate ratios	5
p	fraction of households with no K	0.3

Table 3 reports welfare costs using the benchmark values and varying each variable independently. Several properties stand out. First, our derived welfare costs are typically one to two orders of magnitude larger than Lucas' welfare cost. However, this comparison is somewhat arbitrary, since our welfare costs are not discounted. We derive the welfare evaluation from the instant before the bubble begins.

Second, our welfare costs are highly sensitive to the calibration values of the key parameters. The size of the bubble turns out to be particularly important. For example, a bubble equal in size to half of the value of the capital stock is cataclysmic, since a bubble of this magnitude causes a small mass of households to lose most of their net worth as a result of such an event.

Third, the Taylor approximation is usually about $2/3$ as large as the exact calculation. So the Taylor expansion should only be used as a pedagogical tool and not as a good numerical approximation.

Fourth, the asset trade effect is usually a little bit bigger than the boom/bust effect, but in general the two effects are generally of similar magnitude.

Fifth, the covariance effect is about half as large as the boom/bust effect.

7 Conclusion

To be written.

Table 3: Welfare loss for different calibrations

	Welfare loss as % of permanent consumption		Decomposition as percent of total		
	Exact	Taylor	Asset trade effect	Boom/bust effect	Covariance effect
Benchmark case	3.9%	2.7%	70.7%	57.1%	-27.8%
Low CRRA ($\gamma=1$)	1.0%	0.9%	70.7%	57.1%	-27.8%
High CRRA ($\gamma=5$)	12.0%	4.5%	70.7%	57.1%	-27.8%
Small bubble ($\zeta=1/6$)	0.7%	0.7%	70.7%	57.1%	-27.8%
Large bubble ($\zeta=1/2$)	95.6%	6.2%	70.7%	57.1%	-27.8%
Short bubble (N=8)	3.0%	2.3%	74.2%	50.3%	-24.5%
Long bubble (N=12)	5.7%	3.1%	67.8%	62.8%	-30.6%
Low K heterogeneity ($\sigma=0.5$)	2.6%	2.1%	67.5%	59.0%	-26.6%
High K heterogeneity ($\sigma=0.9$)	5.3%	3.6%	66.6%	59.6%	-26.2%
Low trading ($\phi=0.3$)	2.9%	2.2%	51.1%	69.0%	-20.1%
High trading ($\phi=0.7$)	4.9%	3.2%	84.5%	48.7%	-33.2%

8 References

- Atkeson, A. and Phelan, C. (1994) Reconsidering the Costs of Business Cycles with Incomplete Markets. NBER Macroeconomics Annual, p. 187-207.
- Barlevy, G. (2004) The Cost of Business Cycles Under Endogenous Growth. American Economic Review, Vol 94, Issue 4, p. 964-990.
- Caballero, R., Farhi, E. and Hammour, M. (2009) Speculative Growth: Hints from the U.S. Economy. American Economic Review, Vol. 96, Issue 4, p. 1159-1192.
- Cochrane, J.H. (1989) The Sensitivity of Tests of the Intertemporal Allocation of Consumption to Near-Rational Alternatives. American Economic Review, Vol 79, Issue 3, p. 319-337.
- Dolmas, J. Risk Preferences and the Welfare Cost of Business Cycles. Review of Economic Dynamics, Vol 1, Issue 3, p. 646-676.
- Imrohoroglu, A. (1989) The Cost of Business Cycles with Indivisibilities and Liquidity Constraints. Journal of Political Economy, Vol 97, Issue 6, p. 1364-1383.
- Krebs, T.G. (2007) Job Displacement Risk and the Cost of Business Cycles. American Economic Review, Vol 97, Issue 3, p. 664-686.
- Krusell, P. et al (2009) Revisiting the Welfare Effects of Eliminating Business Cycles. Review of Economic Dynamics, Vol 12, Issue 3, p. 393-402.
- Krusell, P. and Smith, A. (1999) On the Welfare Effects of Eliminating Business Cycles. Review of Economic Dynamics, Vol 2, Issue 2, p. 245-272.
- Laibson, David and Johanna Mollerstrom (2009) "Capital Flows, Consumption Booms and Asset Bubbles: A Behavioral Alternative to the Savings Glut Hypothesis," Harvard mimeo.
- Lucas, R. (1987) Models of business cycles. Oxford: Basil Blackwell.
- Lucas, R. (2003) Macroeconomic Priorities. American Economic Review, Vol 93, Issue 1, p. 1-14.

- Mukoyama, T. and Sahin, A. (2006) Cost of Business Cycles for Unskilled Workers. *Journal of Monetary Economics*, Vol 53, Issue 8, p. 2179-2193.
- Otrok, C. (2001) On Measuring the Welfare Cost of Business Cycles. *Journal of Monetary Economics*, Vol 47, Issue 1, p. 61-92.
- Prescott, E.C. (1986) Theory Ahead of Business Cycle Measurement. *Federal Reserve Bank of Minneapolis Quarterly Review*, Vol 10, p. 9-22.
- Storesletten, K. et al (2001) The Welfare Costs of Business Cycles Revisited: Finite Lives and Cyclical Variation in Idiosyncratic Risk. *European Economic Review*, Vol 45, p. 1311-1339.
- Tallarini, T. (2000) Risk-Sensitive Real Business Cycles. *Journal of Monetary Economics*, Vol 45, Issue 3, p. 507-532.
- Tirole, J. (1985) Asset Bubbles and Overlapping Generations. *Econometrica*, Vol 53, Issue 6, p. 1499-1528.

Appendix A

We now use a second-order Taylor expansion to calculate the welfare loss from the asset bubble. We want to find Δ such that

$$\begin{aligned} & [1 - \exp(-\rho N)] u(C + r\zeta K) \\ & + \exp(-\rho N) u(C - r[\exp(rN) - 1](\zeta K) - r \exp(rN)\zeta k) \\ & = u(\Delta + C) \end{aligned}$$

where

$$C = Y^L + r(K - D).$$

We expand the argument of the utility function around C . Hence,

$$u(C + r\zeta K) = u(C) + u'(C)(r\zeta K) + \frac{1}{2}u''(C)(r\zeta K)^2,$$

and

$$\begin{aligned} & u(C - r[\exp(rN) - 1](\zeta K) - r \exp(rN)\zeta k) = \\ & u(C) + u'(C)(-r[\exp(rN) - 1](\zeta K) - r \exp(rN)\zeta k) \\ & + \frac{1}{2}u''(C)(r[\exp(rN) - 1](\zeta K) + r \exp(rN)\zeta k)^2. \end{aligned}$$

These expansions imply that

$$\begin{aligned} & u(C) + [1 - \exp(-\rho N)] \left[u'(C)(r\zeta K) + \frac{1}{2}u''(C)(r\zeta K)^2 \right] \\ & + \exp(-\rho N) u'(C)(-r[\exp(rN) - 1](\zeta K) - r \exp(rN)\zeta k) \\ & + \exp(-\rho N) \frac{1}{2}u''(C)(r[\exp(rN) - 1](\zeta K) + r \exp(rN)\zeta k)^2 \\ & = u(C) + u'(C)\Delta + \frac{1}{2}u''(C)\Delta^2 \end{aligned}$$

We ignore terms in Δ^2 . Adding this equation up across all agents, the first order terms vanish. The households that lose are exactly offset by the households that gain (in first order terms).

$$\begin{aligned} u'(C)\Delta & = \frac{1}{2}u''(C)(r\zeta K)^2 \left[1 - \exp(-\rho N) + \exp(\rho N) \left[1 - \exp(-rN) + \frac{k}{K} \right]^2 \right] \\ & = \frac{1}{2}u''(C)(r\zeta K)^2 \left[\exp(rN) - 1 + 2(\exp(rN) - 1) \frac{k}{K} + \exp(rN) \left(\frac{k}{K} \right)^2 \right] \end{aligned}$$

$$\begin{aligned}
\frac{\Delta}{C} &= \frac{1}{2} \frac{C \times u''(C)}{u'(C)} \left(\frac{r\zeta K}{C} \right)^2 \left[\exp(rN) - 1 + 2(\exp(rN) - 1) \frac{k}{K} + \exp(rN) \left(\frac{k}{K} \right)^2 \right] \\
&= -\frac{\gamma}{2} \left(\frac{r\zeta K}{C} \right)^2 \left[\exp(rN) - 1 + 2(\exp(rN) - 1) \frac{k}{K} + \exp(rN) \left(\frac{k}{K} \right)^2 \right] \\
&= -\frac{\gamma}{2} \left(\frac{r\zeta K}{C} \right)^2 \left[(e^{rN} - 1) \left(1 + \frac{2k}{K} \right) + e^{rN} \left(\frac{k}{K} \right)^2 \right]
\end{aligned}$$

Recall our notation:

$$-\frac{\Delta}{C} = \Lambda.$$

Hence,

$$\Lambda = \frac{\gamma}{2} \left(\frac{r\zeta}{C} \right)^2 [K^2 (e^{rN} - 1) + 2kK (e^{rN} - 1) + k^2 e^{rN}].$$

So the average welfare loss, $\int \Lambda_i di$.

$$\begin{aligned}
&\frac{\gamma}{2} \left(\frac{r\zeta}{C} \right)^2 [EK_i^2 (e^{rN} - 1) + 2Cov(k, K) (e^{rN} - 1) + Var(k)e^{rN}] \\
&= \frac{\gamma}{2} \zeta^2 \left[(1-p)^{-1} \left(\frac{K}{W} \right)^2 \exp(\sigma_\varepsilon^2) (e^{rN} - 1) + 2Cov\left(\frac{k_i}{W_i}, \frac{K_i}{W_i} \right) (e^{rN} - 1) + Var\left(\frac{k_i}{W_i} \right) e^{rN} \right]
\end{aligned}$$

Appendix B

Each household has financial assets $A = E + (A - E)$, where E is equity and $A - E$ is all other financial assets. We label financial liabilities L (including mortgages). Real assets are $Z + H$, where H is housing and Z includes all other real assets. Let K^h represent human capital, and estimate this as

$$K_t^h = \sum_{s=t}^{\infty} \frac{Y_s^L}{R^{s-t}}.$$

Here, Y^L is labor income and R is $1 + r$, where r is the interest rate. For our calibration we use $R = 1.05$. The bubble assets are H and E , where $H + E = K$. We want to study the ratio

$$\frac{E + H}{A - L + H + Z + K^h} = \frac{K}{W}.$$

We track each household as a single unit. All relevant values are normalized in terms of September 2009 US prices.

For each household and each wave of the HRS we started by calculating the discounted sum of future income. When a household dies (meaning that none of the members followed by HRS is alive) we let the income series end. For households that are still alive at the end of the survey, we estimate the future income scheme with a geometric weighting of the following form, where we set the annual probability of death, d , equal to 0.02.

Suppose a household enters the survey at t and exits at $t + \tau$. Then we impute future values of Y^L by setting $Y_{t+\tau+s}^L$ equal to average real Y^L over the course of the survey. Then

$$K_t^h = \sum_{s=t+1}^{t+\tau} \frac{Y_s^L}{R^{s-t}} + \sum_{s=t+\tau+1}^{\infty} \frac{Y_{Average}^L}{(R + d)^{s-t}}$$

We use the following variables from the Health and Retirement Survey (HRS): *total income* (h*itot), *other real estate* (h*arles), *vehicles owned* (h*atran), *businesses owned* (h*absns), *ira* (h*aira), *stocks* (h*astck), *checking accounts* (h*achck), *cds* (h*acd), *bonds* (h*abond), *other savings* (h*aother), *other debts* (h*adebt), *first house* (h*ahous), *mortgages on first house* (h*amort), *other loans on first house* (h*ahmln), *second house* (h*ahoub), and *mortgages on second house* (h*amrtb). See the RAND-HRS Data Documentation, financial section, for a complete description of the variables.

We split ira into ira_e (equity) and ira_a (other financial assets) by setting ira_e to $(100 - age)$ percent of ira . We then calculated

$$E = \text{stocks} + \text{businesses owned} + ira_e$$

$$A = \text{checking accounts} + \text{cd} + \text{bonds} + \text{other savings} + ira_a + E$$

$$L = \text{mortgage on 1st house} + \text{mortgage on 2nd house} + \\ \text{other loans on 1st house} + \text{other debts}$$

$$H = \text{first house} + \text{second house} + \text{other real estate}$$

$$Z = \text{vehicles owned}$$

$$Y^L = \text{total income}$$

These variables are used to calculate the magnitudes described above.