Asset Bubbles and the Cost of Economic Fluctuations

Citation

Chauvin, Kyle, David I. Laibson, and Johanna Britta Mollerstrom. 2011. Asset bubbles and the cost of economic fluctuations. *Journal of Money, Credit and Banking* 43(S1): 233–260.

Published Version

doi:10.1111/j.1538-4616.2011.00416.x

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:9938146>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Open Access Policy Articles, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#OAP>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Asset bubbles and the cost of economic fluctuations

Kyle Chauvin, David Laibson, Johanna Mollerstrom*

May 26, 2010

Abstract: Lucas (1987, 2003) estimates that the cost of fluctuations is less than 0.1% of consumption. In other words, a social planner would pay no more than 0.1% of (permanent) consumption to eliminate all future business cycle fluctuations. The current paper extends Lucas' calculations by studying the costs of fluctuations arising from asset bubbles. We estimate two classes of costs: consumption volatility due only to asset price volatility, and consumption volatility due to asset trading interacted with price volatility. We show that the magnitude of these welfare costs is driven by heterogeneous household portfolios. If assets are held proportionately across the population, the welfare costs fall by an order of magnitude. 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 3 percent reduction in the level of national consumption. Our calculations are sensitive to the details of the calibration, including the degree of balance sheet and 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 percent of (permanent) national consumption.

*Harvard University. The authors gratefully acknowledge helpful comments from Alberto Alesina, Robert Barro, John Campbell, Richard Cooper, Vince Crawford, Ben Friedman, John Leahy, Greg Mankiw, Klaus Schmidt, Andrei Shleifer, and seminar participants at Harvard, Gerzensee, UC Berkeley, and UCSD. Prepared for a conference at Gerzensee Switzerland on the integration of macroeconomic and microeconomic perspectives. Laibson acknowledges financial support from the NSF (0527516) and the NIA (R01-AG-021650, R01-AG-1665, P30-AG-012810). Corresponding author: Johanna Mollerstrom.

1 Introduction

Consumption is approximately the annuity-value of wealth, so volatile wealth dynamics can generate volatile consumption dynamics.¹ The current paper models and calibrates the linkages from asset price movements to consumption and estimates the consequences of asset bubbles for social welfare.

Because of active market timing, lifecycle portfolio adjustment, or other trading motives, some households are net buyers of assets and some households are net sellers. Such transactions produce inter-household transfers when asset prices deviate from fundamentals. We refer to such inter-household transfers as an *asset trading effect*.

Households that own bubble-priced assets perceive that they are wealthier than they actually are. Consequently, they raise consumption during a 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. In other words, during the bubble the agent overconsumes relative to the true annuity value of wealth. After the bubble, the agent correspondingly underconsumes relative to the pre-bubble annuity value of wealth. We refer to this dynamic pattern as a consumption *boom/bust effect*.

Asset trading effects and consumption boom/bust effects jointly generate excess consumption volatility. In the current paper we model these effects and calculate the resulting welfare costs. We provide exact calculations and we provide Taylor approximations that express these effects with simple closed form equations. Four key expressions emerge from the Taylor approximations.

First, there is an asset trading welfare cost. The welfare cost arising from this channel – expressed as a fraction of permanent consumption – is given

¹An important exception is the case in which wealth fluctuations are driven by interest rate fluctuations. For evidence on the propensity to consume out of variation in housing wealth, see Greenspan and Kennedy (2007, 2008) and Bosworth and Smart (2009). Carroll et al. (2006) and Campbell and Cocco (2007) estimate a housing-wealth MPC of nine percent.

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). The asset-trading effect is increasing in risk aversion, in the square of the magnitude of the asset bubble, in the cross-sectional variance of normed asset trading, and in the growth of the bubble. In our benchmark calibration, the asset trading effect represents a welfare cost that is equivalent to 1.6% of consumption.

The asset trading effect has no first-order consequences for welfare, since every gain from selling the overpriced asset is offset by a loss (in another household) from buying the overpriced asset. However, second order effects don't cancel. Concavity in the utility function causes the gains in marginal utility to be more than offset by the losses in marginal utility. Finally, the asset trading effect vanishes as heterogeneity is reduced – i.e. as $\text{Var} \left(\frac{k_i}{W_i} \right)$, the cross-sectional variance of asset purchases, 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

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

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). The boom/bust effect is increasing in risk aversion, in the square of the magnitude of the asset bubble, in the concentration of holding of the capital stock, in the square of the normed size of the capital stock, in the cross-sectional variance of normed capital holding, and in the growth of the bubble. In our benchmark calibration, boom/bust effects account for welfare costs equal to 1.3% of permanent consumption.

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 again survive due to concavity. Finally, 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 0.5% 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.

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

The covariance effect is increasing in risk aversion, in the square of the magnitude of the asset bubble, in the cross-sectional covariance of normed capital holding and normed capital trading, and in the growth of the bubble. In our calibration the covariance effect represents a cost that is equivalent to -0.6% of permanent consumption. In other words, this third effect turns out to partially *offset* the other two. Mean reversion in household portfolio allocations induce a negative covariance between $\frac{K_i}{W_i}$ (household i 's original holdings of the capital stock) and $\frac{k_i}{W_i}$ (household i 's net purchases of the capital stock). This negative covariance 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 covariance effect vanishes in the homogeneous case.

Fourth, we identify an aggregation effect, which arises since the social welfare cost is not a linear weighted average of the individual welfare costs. Specifically, because the utility function is concave, welfare costs aggregate in a way that overweights the households with the largest welfare losses. Specifically, the social welfare cost is equal to the average value of the household-level welfare costs (the sum of the three terms above) plus an adjustment term for aggregation:

$$\frac{\gamma}{2} \frac{V(\Lambda_i)}{1 - E\Lambda_i},$$

where Λ_i is the welfare cost of household i . Since $E\Lambda_i$ is close to zero, this aggregation effect is essentially $\gamma/2$ times the variance of the individual households' welfare costs. In our calibration this aggregation effect is equivalent to 1.4% of permanent consumption. This aggregation effect also vanishes in the homogeneous case.

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 and the Survey of Consumer Finances. In the empirical analysis, we provide a new distributional result: the ratio of equity+housing wealth to total net worth is well-approximated in the cross-section by a log-normal distribution plus mass at zero. Section 6 presents our welfare cost results – both exact results and approximations. Section 7 discusses potential extensions of our model and section 8 concludes.

2 Literature Review

Real Business Cycle models imply that policy-makers should not adopt counter-cyclical policies, since fluctuations are optimal responses to changing fundamentals (e.g., Kydland and Prescott 1982; Prescott, 1986).² However, some economic models do imply that policy-makers should adopt counter-cyclical policies. Lucas (1987) calculates an upper bound for the benefits of such counter-cyclical policies. Lucas considers a representative agent with constant relative risk aversion. Lucas models consumption as a log linear trend with noise:

$$c_t = A \exp(gt + \varepsilon_t - \sigma^2/2).$$

Here A is a scaling parameter, g is the long-run growth rate, ε is an iid Gaussian random variable with standard deviation σ . This implies that the welfare cost of fluctuations is approximately $(1/2)\gamma\sigma^2$, where γ is the coefficient of relative risk aversion.

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

²Tirole (1985) shows that asset price bubbles can be rational and may even increase welfare. See also Caballero et al., (2009).

to give up as fraction of consumption in order to set $\sigma = 0$. Lucas shows 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% of consumption.³

The most frequent critique of this result assumes heterogeneous agents and incomplete markets. Such assumptions can raise the welfare costs relative to Lucas' estimates, but the results are mixed and the average welfare cost is still generally estimated to be no more than 1% (e.g., Atkeson and Phelan 1994, Imrohoroglu 1989, Krusell et al., 2009, Krusell and Smith, 1999, Krebs, 2007, Mukoyama and Sahin, 2006 and Storesletten et al., 2001).

Other analyses of Lucas' result have studied Krep-Porteus or Epstein-Zin preferences (e.g., Dolmas, 1998; and Otrok, 2001), asset-pricing kernels to generate a shadow-cost for fluctuations (Tallarini 2000), fluctuations that are persistent instead of being transitory (Obstfeld 1994; and Barro 2006, 2009), and growth that is negatively linked to fluctuations (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

³A 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" incur costs that are less than \$1 per quarter.

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

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

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.

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 date that the anticipated productivity gains fail to be realized. In this way our model is similar to that of Beaudry and Portier (2007) where news about the future changes expectations in a way that impacts consumption, investment and growth. Even more relevant is Christiano et al (2007), which investigates what happens when such expectations turn out to be wrong.

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 a 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 at date $\tau \in [0, N)$ borrows $B_\tau - B = [1 + \exp(r\tau)\zeta] k$ on the domestic market to do so. We assume that the agent uses her income from k to finance this new domestic debt. However, this income is not sufficient to finance all of the new domestic debt, since the period of higher productivity has not yet arrived. We assume that the residual domestic debt is rolled over, so the domestic debt is equal to

$$B_\tau - B = [1 + \exp(r\tau)\zeta] k.$$

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

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.

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 a sub-period $[0, \tau]$ of the bubble period (where $\tau \leq N$).

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

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

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

During the bubble, assets at the household level can be decomposed into domestic assets valued at

$$\zeta (K + k) \exp(r\tau) + (K + k),$$

debt to foreign agents valued at

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

and debt to domestic agents valued at

$$[1 + \exp(r\tau)\zeta] k + B$$

Note that the net wealth is

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

When the bubble bursts (at date N), the household is left with net assets:

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

So consumption falls from

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

to

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

In other words, consumption falls by

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

We can decompose this into three effects:

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

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 accumulation effect is the consequence of growing the bubble at rate r over an interval of length N . As the duration of the bubble goes to zero, the accumulation effect ceases to matter: $\lim_{N \rightarrow 0} \exp(rN) = 1$.

3.1 Distributional assumptions

We now study inter-household differences. Consequently, we will start using household subscripts.

We assume 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 (see the calibration section for empirical evidence that validates this assumption). Specifically,

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

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 the economy in section 5. In that 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, we first characterize 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 ergodic distribution. 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 use this ergodic assumption to (numerically) back out the distribution of the transaction variable

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

4 Welfare calculations

We first characterize the welfare of a single agent in this economy. We then show how to aggregate these agents into a social welfare cost.

An individual 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 or output). 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

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

To gain intuition for this magnitude, it is helpful to rewrite it as

$$\lambda^{1-\gamma} = \int \lambda_i^{1-\gamma} di. \tag{3}$$

It follows that λ is the certainty equivalent ‘consumption’ of an ‘agent’ with constant relative risk aversion γ and stochastic ‘consumption’ λ_i .

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,

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

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

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

This can be broken down into three terms, using the fact that $C_i = rW_i$. First, there is a consumption boom/bust effect:

$$\frac{\gamma}{2} \zeta^2 (1 - p)^{-1} \left(\frac{K}{W} \right)^2 \exp(\sigma_\varepsilon^2) (e^{rN} - 1). \quad (4)$$

Second, there is a covariance effect:

$$\gamma \zeta^2 Cov\left(\frac{k_i}{W_i}, \frac{K_i}{W_i}\right) (e^{rN} - 1). \quad (5)$$

Third, there is an asset trading effect:

$$\frac{\gamma}{2} \zeta^2 Var\left(\frac{k_i}{W_i}\right) e^{rN}. \quad (6)$$

Finally, we need to aggregate these effects. As already noted, aggregation is not linear. Concavity implies that the aggregate λ is the certainty equivalent of the individual λ_i (cf equation 3). Using a second-order Taylor expansion,

$$\lambda \approx E\lambda_i + \frac{\gamma}{2} \frac{V(\lambda_i)}{E\lambda_i}.$$

In other words, the aggregate social welfare cost is the sum of equations 4, 5, and 6, plus an aggregation effect,

$$\frac{\gamma}{2} \frac{V(\Lambda_i)}{1 - E\Lambda_i}.$$

Since $E\Lambda_i$ is close to zero, this aggregation effect is essentially $\gamma/2$ times the variance of the individual households' welfare costs.

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 $r = \rho = 0.05$. Our baseline value for the coefficient of relative risk aversion is $\gamma = 3$. We also consider CRRA values from 1 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 minus government debt⁴ all divided by GDP.¹¹ This series fluctuated historically in a range roughly between 2.5 and 3 units of GDP. Starting in the late-1990's, however, the series broke from this historical range and rose sharply. At its peak in the 4th quarter of 2006, the series reached a value of 4.1 units of GDP. By the 4th quarter of 2008, the series had fallen back to its historical range. These comparisons imply an estimated peak bubble value of about 1.1 units of GDP. However, the ratio of household wealth to GDP misses part of the value of the bubble, since it nets out the value of debt accumulated to finance consumption during the bubble years. Household debt increased from 0.3 units of GDP from the late 1990's to 2006:4. This analysis implies that had U.S. households not consumed some of their bubble wealth, the ratio of economy-wide net worth would have risen 1.65 units of GDP. Since,

$$\begin{aligned} \zeta &= \left(\frac{\Delta K}{K} \right)_{1998} \\ &= e^{(g-r)(2006-1998)} \left(\frac{\Delta K}{Y} \right)_{2006} \left(\frac{Y}{C} \right)_{1998} \left(\frac{C}{K} \right)_{1998} \\ &= 0.33. \end{aligned}$$

⁴Government debt includes federal, state, and local sources.

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

Figure 1: Household net worth minus total government debt all divided by GDP (1951:4 – 2009:3).

Sources: Federal Reserve Board flow of funds balance sheets, Bureau of Economic Analysis NIPA accounts, and authors' calculations.

For this calculation we take $g = 0.02$, $r = 0.05$, $(\frac{\Delta K}{Y})_{2006} = 1.1$, $\frac{Y}{C} = \frac{3}{2}$, and $\frac{C}{K} = \frac{1}{5}$. Here g is the real growth rate in GDP. This calculation implies that the total value of the bubble is 0.33 of the value of all domestic capital (including land). To err on the side of conservatism, we set $\zeta = 0.3$ in our benchmark calibration.

We calibrate the distribution of K_i from the Health and Retirement Study (HRS) and the Survey of Consumer Finances (SCF). Though the HRS has the limitation that it only covers respondents who are middle-aged or older, the HRS has an offsetting advantage. Respondents are surveyed longitudinally every two years, whereas the SCF is purely cross-sectional.⁵

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.⁶ 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).⁷ 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 $\frac{K_i}{C_i} \leq 30$. Furthermore, we distinguish between the mass-at-zero and greater-than-zero sub-distributions by partitioning the data at the 18th (for the HRS waves) and 30th (for the SCF waves) percentiles. This is done to rule out extreme welfare losses (for households with extreme values of K).

Our natural log plots from the HRS have associated standard deviations

⁵The HRS asset data is of higher quality than the asset data in the Panel Survey on Income Dynamics.

⁶Both the HRS and the SCF bias down this ratio by collapsing private businesses into a net equity summary statistic. For example, a family business with \$1 million of assets and \$800,000 in debt would be recorded in these data sets as a private equity position of \$200,000. This understates the household's leverage.

⁷The data in Figures 2 and 4 are displayed between ratio values of 0 and 1.5. In Figures 3 and 5 we display all data points with log ratio values within 2.5 of the sub-distribution mean.

Figure 2: Distribution of K/W [HRS]

(a) Mean = 0.2173, SD = 0.4093

(b) Mean = 0.2263, SD = 0.4507

(c) Mean = 0.2561, SD = 0.2897

(d) Mean = 0.2596, SD = 0.3892

(e) Mean = 0.2537, SD = 0.2752

(f) Mean = 0.2467, SD = 0.3585

(g) Mean = 0.2354, SD = 0.4186

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

(a) Mean = -1.5844, SD = 0.7744

(b) Mean = -1.5301, SD = 0.7210

(c) Mean = -1.4601, SD = 0.8216

(d) Mean = -1.4090, SD = 0.7515

(e) Mean = -1.4312, SD = 0.7408

(f) Mean = -1.4986, SD = 0.8086

(g) Mean = -1.5050, SD = 0.8094

Figure 4: Distribution of K/W [SCF]

(a) Mean = 0.1710, SD = 0.1849

(b) Mean = 0.1717, SD = 0.1795

(c) Mean = 0.1842, SD = 0.1831

(d) Mean = 0.1933, SD = 0.1881

(e) Mean = 0.2182, SD = 0.2105

(f) Mean = 0.2297, SD = 0.2904

Figure 5: Distribution of $\ln(K/W)$ for non-zero households [SCF]

(a) Mean = -1.6512, SD = 0.7219

(b) Mean = -1.6148, SD = 0.6600

(c) Mean = -1.5257, SD = 0.6145

(d) Mean = -1.4747, SD = 0.6086

(e) Mean = -1.3325, SD = 0.5593

(f) Mean = -1.2709, SD = 0.5617

that range from 0.72 to 0.82 (depending on the wave of the HRS). In Figures 3 and 4 we plot analogous figures generated by the SCF. The SCF plots in Figure 4 have standard deviations that range from 0.56 to 0.72. We therefore adopt a benchmark standard deviation, σ_ε , of 0.70 for our calculations.

The U.S. Census Bureau reports that the household homeownership rate ranged from a low of 65.9% (1998q1) to a high of 69.2% (2004q2) during the bubble period. We assume that $p = 0.30$: i.e. 30% of households own neither a home nor equity.

Finally, we turn to asset ownership dynamics. 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 ergodic 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 costs up if we picked a 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	CRRA	1	3	5
ζ	bubble magnitude	0.2	0.3	0.4
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 social welfare costs are typically one to two orders of magnitude larger than Lucas' welfare costs. However, this comparison is misleading, since our welfare costs are not discounted. Our calculations derive the welfare evaluation from the instant before the bubble begins. In the next section, we generate a timeless perspective with a recursive argument.

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 to 40% of the value of the capital stock is cataclysmic.

Third, the Taylor approximation is usually significantly larger than the exact calculation. So the Taylor expansion should only be used as a pedagogical tool and not as a close numerical approximation.

Fourth, the asset trading, boom/bust, and aggregation effects are usually of similar magnitude.

Fifth, the covariance effect is about half as large (and of opposite sign) as the other effects.

7 The timeless perspective

To make our calculations directly comparable to those of Lucas (1987, 2003), we now analyze a timeless perspective. Let V be the value function of an economy with a constant θ probability of entering a bubble event and a welfare loss of $1 - \lambda$ fraction of permanent consumption in the event of a bubble. Then,

$$\begin{aligned}\rho V &= u(C) + \theta [\lambda^{1-\gamma} V - V] \\ V &= \frac{u(C)}{\rho - \theta [\lambda^{1-\gamma} - 1]}\end{aligned}$$

Table 3: Calibrated social welfare loss

	Welfare cost as percent of permanent consumption		Decomposition of Taylor Approximation			
	Exact	Taylor	Trade	Boom/bust	Covariance	Aggregation
Benchmark case [†]	2.9%	3.7%	1.6%	1.3%	-0.6%	1.4%
Low CRRA ($\gamma=1$)	0.8%	1.1%	0.5%	0.4%	-0.2%	0.4%
High CRRA ($\gamma=5$)	7.2%	7.3%	2.6%	2.1%	-1.0%	3.6%
Small bubble ($\zeta=0.2$)	1.1%	1.5%	0.7%	0.6%	-0.3%	0.5%
Large bubble ($\zeta=0.4$)	33.6%	31.6%	2.8%	2.3%	-1.1%	27.6%
Short bubble ($N=8$)	2.3%	3.2%	1.4%	1.0%	-0.5%	1.3%
Long bubble ($N=12$)	3.8%	4.2%	1.7%	1.6%	-0.8%	1.6%
Low K heterogeneity ($\sigma=0.5$)	2.0%	2.6%	1.1%	1.0%	-0.5%	0.9%
High K heterogeneity ($\sigma=0.9$)	3.8%	4.8%	1.9%	1.8%	-0.8%	1.9%
Low trading ($\phi=0.3$)	2.2%	2.7%	0.9%	1.3%	-0.4%	0.9%
High trading ($\phi=0.7$)	3.6%	4.6%	2.2%	1.3%	-0.9%	2.0%

[†]For the benchmark case, the coefficient of relative risk aversion is $\gamma=3$, the bubble is $\zeta=0.3$ proportion of the value stock of physical capital, the duration of the bubble period is $N=10$ years, the cross-sectional household-level standard deviation of the ratio of K/C (for $K>0$) is $\sigma=0.7$, the fraction of households that trade in the N year bubble period is $\phi=0.5$, the interest rate and discount rate are both equal to 0.05, the aggregate K/C ratio is 5, and the fraction of households with no claims to the capital stock is 0.3.

Let λ_∞ represent the timeless social welfare cost of having a stationary θ likelihood of entering a bubble event. Then,

$$\begin{aligned}\frac{u(C)}{\rho}\lambda_\infty^{1-\gamma} &= \frac{u(C)}{\rho - \theta [\lambda^{1-\gamma} - 1]} \\ \lambda_\infty &= \left(\frac{\rho}{\rho - \theta [\lambda^{1-\gamma} - 1]} \right)^{\frac{1}{1-\gamma}}.\end{aligned}$$

If $\rho = 0.05$, $\theta = 0.02$, and $\lambda = 0.97$, then $\lambda_\infty = 0.987$. Hence, the timeless welfare cost of bubbles is 1.1% of permanent consumption. Moreover, this calculation overlooks the fact that welfare costs are convex in the magnitude of the bubble. Even a tiny chance of a bubble that is significantly larger than $\zeta = 0.3$ will have potentially large effects on the timeless welfare cost. Recall that a bubble of magnitude $\zeta = 0.4$ implies a welfare cost of 33% of permanent consumption. If bubbles of this magnitude are added to the distribution of bubbles, and bubbles of this magnitude occur with a probability of only 1/1000, then the total timeless welfare cost doubles to 2.3% of permanent consumption.

8 Conclusion

Lucas (1987, 2003) estimates the welfare costs of economic fluctuations and bounds them above at only 1/10 of 1% of permanent consumption. Motivated by his analysis, we estimate the welfare costs of an asset bubble. Our model identifies two types of welfare costs: consumption fluctuations due to active asset trading and consumption fluctuations due to passive asset ownership.

To calibrate the model we use the Health and Retirement Survey and the Survey of Consumer Finances. With a coefficient of relative risk aversion of 3 (our benchmark assumption) the asset bubbles of the last decade have a welfare cost equal to 3% of permanent consumption. With calibrated values for key parameters, the welfare costs generally lie between 1% and 10% percent of permanent consumption. Even from the timeless perspective, the welfare costs in our model are greater than 1% of permanent consumption.

Our model predicts that asset bubbles should give rise to increased consumption inequality in their wake.⁸ This prediction needs to be tested.

⁸ Atkinson (2009) shows that inequality in Singapore and Malaysia increased as a consequence of the Asian financial crisis in 1997 and 1998.

Our model also has several policy implications. Our results imply that the welfare costs of bubbles are highly convex. This suggests that policy makers consider leaning against the wind only in the presence of large asset price movements that may be bubbles. A small policy intervention is likely to have small (second-order) welfare costs if no bubble is present, and enormous (first-order) welfare gains in the presence of a large bubble. Our results imply that it is not important to eliminate asset bubbles, just to reduce or constrain their size.

9 References

- Atkeson, A. and Phelan, C. (1994) Reconsidering the Costs of Business Cycles with Incomplete Markets. NBER Macroeconomics Annual, p. 187-207.
- Atkinson, A.B. (2009) "Top Incomes in a Rapidly Growing Economy: Singapore". Top Incomes: A Global Perspective, edited by A B Atkinson and T Piketty, OUP, 2010.
- Barlevy, G. (2004) The Cost of Business Cycles Under Endogenous Growth. American Economic Review, Vol 94, Issue 4, p. 964-990.
- Barro, R.J. (2006) "Rare Disasters and Asset Markets in the Twentieth Century." Quarterly Journal of Economics, Vol. 121, No. 3, pp. 823-866.
- Barro, R.J. (2009) "Rare Disasters, Asset Prices, and Welfare Costs." American Economic Review, Vol. 99, No. 1, pp. 243-264.
- Beaudry, P. and Portier, F. (2007). 'When can changes in expectations cause business cycle fluctuations in neo-classical settings?' Journal of Economic Theory, vol. 135(1), pp. 458-477.
- Bosworth, B. and Smart, R. (2009) The Wealth of Older Americans and the Subprime Debacle. Mimeo, Center for Retirement Research, Boston College.
- 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.
- Campbell, J.Y. and Coco, J.F. (2007) "How do House Prices Affect Consumption? Evidence from Micro Data." Journal of Monetary Economics, Vol. 54, No. 3, pp. 591-621
- Carroll, C. et al. (2006) How large is the housing wealth effect? A new approach. NBER working paper 12746.
- Christiano, L., Ilut, C., Motto, R. and Rostagno, M. (2007). 'Monetary policy and stock market boom-bust cycles. North Western University, mimeo.

- Cochrane, J.H. (1989) The Sensitivity of Tests of the Intertemporal Allocation of Consumption to Near-Neutral 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.
- Greenspan, A. and Kennedy, J. (2007). 'Sources and uses of equity extracted from homes' Finance and Economic Discussion Series 2007-20, Washington, DC, Board of Governors of the Federal Reserve System.
- Greenspan, A. and Kennedy, J. (2008). 'Sources and uses of equity extracted from homes', *Oxford Review of Economic Policy*, vol. 24(1), pp. 120-144.
- 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.
- Kydland, F. and Prescott, E. (1982) Time to Build and Aggregate Fluctuations. *Econometrica* 50: 1345-1371.
- Laibson, David and Johanna Mollerstrom (forthcoming) "Capital Flows, Consumption Booms and Asset Bubbles: A Behavioral Alternative to the Savings Glut Hypothesis." *Economic Journal*.
- 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.
- Obstfeld, M. (1994) Evaluating Risky Consumption Paths: The Role of Intertemporal Substitutability. *European Economic Review*, Vol. 38, No. 7, pp. 1471-1486.
- 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, $\int \Lambda_i di$.

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

The first term in brackets follows from

$$\begin{aligned}
E(K_i^2) &= (1-p)E(K_i^2 | K_i > 0) \\
&= (1-p)E[C_i^2 \exp(2\mu + 2\varepsilon_i - \sigma_\varepsilon^2) | K_i > 0] \\
&= (1-p)C^2 \exp(2\mu + \sigma_\varepsilon^2) \\
&= (1-p)C^2 \exp(\sigma_\varepsilon^2) \left[\frac{K}{C(1-p)} \right]^2 \\
&= (1-p)^{-1} \exp(\sigma_\varepsilon^2) K^2.
\end{aligned}$$

Appendix B

We estimate the ratio of ownership of bubbly assets to total net worth,

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

Here, E represents equity and $A = E + (A - E)$ are total financial assets. L represents financial liabilities, including mortgages. Physical assets are denoted by H (housing) and Z (all other assets, e.g. vehicles). K^h , human capital, is the net present value of future income (excluding capital income). To calculate this ratio we use data from the Health and Retirement Survey (HRS) and the Survey of Consumer Finances (SCF).

Our unit of analysis for both the HRS and SCF is a single household, which means that ownership and income values are summed over the primary respondent and, where applicable, the respondent's spouse.⁹ To calculate the above ratio, we use the following procedure. First, all variables except human capital are constructed from raw survey variables as detailed below. They are then converted into real terms using CPI data from the Department of Labor.¹⁰

Since the HRS-data are in panel format, we calculate human capital at time t as

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

where R is the discount rate, which we calibrate as 1.05. Within the survey timespan, we set Y_s^L equal to its realized value when possible. When data are missing for year s , but are available for $s + 1$ and $s - 1$ we impute

$$Y_s = Y_{s-1} + \frac{Y_{s+1} - Y_{s-1}}{2}$$

If a respondent is alive in the last survey wave, we project Y_s^L by setting it to the average of Y values recorded over the course of the survey, and

⁹In the SCF, the household, defined in this way, is already used as the unit of analysis. As for the HRS, we select the longest surviving individual member as a representative for the household.

¹⁰Monthly CPI data can be found at <ftp://ftp.bls.gov/pub/special.requests/cpi/cpia1.txt>. HRS data was obtained from <https://ssl.isr.umich.edu/hrs/files.php> and SCF data from <http://www.federalreserve.gov/pubs/oss/oss2/scfindex.html>.

increase the discount factor R to $R + d = 1.07$ to account for the probability of mortality. Therefore, if a respondent enters the survey at time t and exists at $t + \tau$, we get

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

Because the SCF is only cross-sectional, we cannot directly compute K^h through income. Instead, we first use the HRS to regress the (log) value of human capital on age, age squared, years of education¹¹, years of education squared, and the log value of current income.¹² We then take the regression coefficients and use them to predict values of K^h in the SCF by applying them to age, years of education, and current income.

¹¹In the SCF, only education of the household *head* is provided. In the HRS, because no head is designated, we compute years of education for the household as the average over individual members.

¹²This specification calculates a single regression, aggregating all waves of the HRS survey. The shape of the distribution of the bubble asset ratio is robust to predicting SCF values of K^h by regressing on the closest corresponding wave of the HRS. Also, the distributions are robust to dropping age from the regression entirely.

Table 1: Composition of ratio components in the HRS

Variable(s)	Survey Variable	Survey Variable Description
$E =$	HwASTCK	Net value of stocks and mutual funds
+	HwABSNS	Net value of businesses
+	HwAIRA	Net value of IRA and Keogh accounts (only 100 – <i>age</i> percent are counted as E)
$H =$	HwAHOUS	Value of the primary residence
+	HwAHOUB	Value of the secondary residence (not reported in 1996)
+	HwARLES	Net value of real estate, besides primary and secondary residence
$A - E =$	HwACHCK	Net value of checking, savings, and money market accounts
+	HwACD	Value of CDs, government savings bonds, and treasury bills
+	HwABOND	Net value of bonds or bond funds
+	HwAOTHR	Net value of all other savings
+	HwAIRA	(only <i>age</i> percent are counted as $A - E$)
$Z =$	HwATRAN	Net value of vehicles
$L =$	HwAMORT	Values of first and second mortgages on primary residence
+	HwAHMLN	Values of home loans other than first and second mortgages on primary residence
+	HwAMRTB	Value of a second home mortgage on secondary residence
+	HwADEBT	Value of other debt
$Y =$	HwITOT	Sum of all income in household
–	HwICAP	Business or farm income, self-employment earnings, business income, gross rent, dividend and interest income, trust funds or royalties, and other asset income

Table 2: Composition of ratio components in the SCF

Variable(s)	Survey Variable	Survey Variable Description
$E =$	NMMF	Total value of directly held pooled investment funds
+	STOCKS	Total value of directly held stocks
+	RETQLIQ	Value of IRAs, Keoghs, thrift-type accounts, and account-type pensions (only $100 - age$ percent counted as E)
$H =$	HOUSES	Total value of primary residence
+	ORESRE	Total value of other residential real estate
+	NNRESRE	Total value of net equity in nonresidential real estate
$A + H + Z - L =$	NETWORTH	Total net worth of household
$Y =$	INCOME	Total income of household
–	KGINC	Capital gain or loss income
–	INTDIVINC	Interest (taxable and nontaxable) and dividend income