Online Appendix

for

Incidental Bequests and the Choice to Self-Insure Late-Life Risks

Lee M. Lockwood

A   Appendix

A.1   Health Uncertainty: Mortality, Long-Term Care, and Acute Medical Care Risks

This section reports the results of descriptive regressions about long-term care and long-term care insurance and describes the key elements of the model of health risks. I test the robustness of the results to a variety of changes in these risks in Section V.D.

A.1.1   Long-Term Care and Long-Term Care Insurance: Descriptive Regressions

Table 1 reports results from descriptive regressions of long-term care usage and long-term care insurance ownership on key demographic and economic variables. These regressions are based on a sample of people 65 and older in the 1998 wave of the Health and Retirement Study. The regressions of long-term care usage further restrict the sample to people who report difficulties with at least two activities of daily living (ADLs). The table shows the estimated marginal effects from probit regressions.

Use of formal care is much greater among people with more ADL limitations, is slightly greater among single people and people without children, and is perhaps slightly greater among people with greater income, though the income results are statistically insignificant and the point estimates are non-monotonic in income quartile. Long-term care insurance ownership is strongly increasing in wealth but is otherwise not well predicted by the other demographic variables, including whether someone is single and whether he or she has children.

A.1.2   Health Transitions

An individual's future health depends probabilistically on the individual's current age and health status as well as on the individual's sex (s) and (permanent) income (y), $Pr(h_{t+1} = h'|h_t, t; s, y)$. I base the model of health transitions on a model developed by Friedberg et al. (2014). This model makes a number of important improvements on the widely-used Robinson model of long-term care requirements (Robinson, 2002), including using updated data and more robust procedures.

I make three adjustments to the Friedberg et al. (2014) model in order to cater it to my

application. First, I convert the monthly health transitions calculated by Friedberg et al. (2014) into annual transitions. This choice is driven both by computation time considerations and by data limitations, since the Health and Retirement Study and many other datasets measure medical spending and other variables at lower frequencies (e.g., every year or every second year). Second, I use the estimated transition matrices for women as the baseline transition matrices for both the single men and the single women in my sample. The care usage patterns of women likely provide a closer approximation to the usage patterns of single people, whether male or female, because women receive a much smaller share of their long-term care from their spouses than men do. Women receive less informal care from their spouses because their spouses (predominantly men) tend to get sick and die at earlier ages than they do. As a result, a smaller share of women's care episodes occur when their husbands are alive and well enough to provide them with informal care. Of course, the care use patterns of women in the general population—including married women—suffers from this same problem and so tends to understate the care needs of singles, but the bias is less severe than it is for men.

The third set of adjustments I make is to adjust the Friedberg et al. (2014) transition probabilities to match De Nardi, French and Jones's (2010) estimates of life expectancy conditional on reaching age 70 for different sex and income groups. A t-year-old in sex-income quintile group (s, q) faces the Friedberg et al. (2014) transition probabilities of a $(t + \Delta(s, q))$-year-old female, where $\Delta(s, q)$ is chosen to minimize the difference between predicted life expectancy at age 70 and De Nardi, French and Jones's (2010) estimates of life expectancy at age 70.

Table 2 shows the age adjustments, $\Delta(s, q)$, and the resulting life expectancies of each group. The differences in life expectancies at age 70 across sex and income groups are substantial: Women live more than five years longer than men in the same income quintile, and men and women in the top income quintile live almost four years longer than their counterparts in the bottom quintile. Each group's adjusted life expectancy is within 0.3 years of De Nardi, French and Jones's (2010) estimate. It is important that the model of health risk be consistent with this substantial heterogeneity in life expectancy, since life expectancy can have an important impact on saving and insurance decisions (De Nardi, French and Jones, 2009).

Table 3 presents statistics related to the unconditional and conditional probabilities of being in different health states in the original and adjusted Friedberg et al. (2014) models. The adjusted model preserves the essential character of the original model in terms of the key determinants of saving and insurance decisions: the expected share of remaining life spent in different health states. The key difference is that males spend a greater share of

their remaining lives in nursing homes under the adjusted model than under the original model. This is due to the much greater supply of informal care to married than to single men and is why I base the model of health transitions for single males on the Friedberg et al. (2014) model for females. The other main differences have to do with time aggregation. Using yearly rather than monthly transitions reduces the probability of ever experiencing a nursing home stay and of leaving a nursing home alive, since yearly transitions rule out the possibility of stays of less than one year in duration. Although it would be desirable to base the model of health transitions on a model specifically estimated to match the heterogeneous experiences of single men and women with different levels of income, such a model is not available, and, as discussed in Section V.D, the conclusions are robust to many alternative assumptions and are unlikely to be affected by plausible changes in the model of health risk.

A.1.3 Long-term Care Prices

The cost of the individual's long-term care is a deterministic function of the individual's health, age, sex, and income quintile, $ltc(h_t, t, s, q)$. Part of this heterogeneity could reflect differences in the prices that people face for the same care, due, for example, to differences in prices across different locations. Other sources of heterogeneity could include unmeasured and un-modeled differences in the quantity or quality of the long-term care services consumed by different groups, conditional on their health status. For example, higher-income people might purchase higher-quality (and so costlier) long-term care.

To estimate $ltc(h_t, t, s, q)$, I combine two sources of data. The first is data from a MetLife survey about long-term care prices (MetLife Mature Market Institute, 2002a,b). This reports average prices for different long-term care services, including stays in nursing homes and assisted living facilities and skilled and unskilled home care. According to this survey, average prices in 2002 were \$52,195 per year in a nursing home, \$26,280 per year in an assisted living facility, \$18 per hour for unskilled home care, and \$37 per hour for skilled home care.

The second source of data is the National Long-Term Care Survey (NLTCS). This is a longitudinal survey of Americans age 65 and older with detailed information about health and health-related expenditures, including information about the prices of any long-term care services that surveyed individuals consume. I use the NLTCS data to estimate the following regression:

$$\frac{p_i}{\bar{p}} = \alpha + \beta \text{female}_i + \gamma \text{age}_i + \sum_{q=2}^{5} \delta_q \text{income quintile q}_i + \varepsilon_i,$$

4

where p_i is the price per month of care in i's nursing facility, $\bar{p}$ is the average price per month of care in facilities, and female$_i$ and income quintile q$_i$ are indicators for whether i is a female and in income quintile q, respectively. I use the predicted values from this regression to scale the average prices of each long-term care service (nursing homes, assisted living facilities, skilled home care, and unskilled home care).[1]

A summary of the results is presented in Table 4. Females pay slightly higher prices than males (about 6 percent) and higher-income people pay slightly higher prices than lower-income people (the top income quintile pays about 12 percent more than the bottom). Conditional on the type of care being used, age has little effect on long-term care prices (the coefficient estimate is a precise zero). The biggest source of heterogeneity is between people in the bottom income quintile and everyone else; the prices that people in the bottom income quintile pay are between 1.6 and 11.5 percent lower than the prices paid by people in higher income quintiles. But a striking feature of the results is how little heterogeneity there appears to be on average across different sex, age, and income groups. None of the individual coefficients is significant at conventional confidence levels, and the covariates taken as a whole are not significant either.

A.1.4   Acute Medical Care Spending

The cost of an individual's acute medical care is log-normally distributed with the mean and variance depending on the individual's health, age, sex, and income quintile, $m_t \sim \log N(\mu_m(h_t, t, s, q), \sigma_m^2(h_t, t, s, q))$. That the mean and variance are allowed to depend on health, age, sex, and income quintile allows for rich heterogeneity in the spending risk facing different people.

I estimate the mean and variance of different groups' spending on acute medical care in two steps, using data from the HRS. First, I decompose total out-of-pocket spending (the variable in the RAND release of the HRS) into separate acute and long-term care components. To do this, I use disaggregated data on out-of-pocket spending by service type in 2006. For each health status (healthy, home care, and nursing home), I estimate the share of total out-of-pocket spending that is due to acute medical care (as opposed to long-term care). The sample is everyone age 65 and older whose combined previous-wave non-housing wealth and annual income is at least $100,000. I convert observed total out-of-pocket spending to acute out-of-pocket spending by multiplying observed total

---

[1]Data considerations lead me to estimate a single scaling factor to apply to all types of long-term care rather than estimating different scaling factors for different types of care. These considerations are the difficulty of distinguishing between nursing homes and assisted living facilities and the difficulty of estimating hourly prices of skilled and unskilled home care in the data.

out-of-pocket spending by the estimated shares of spending on acute care for each health status. The estimates imply that among healthy people, about 97 percent of total out-of-pocket spending is due to spending on acute medical care. Among people who require home care, this share is 72 percent. Among people who require nursing home care, this share is just 11 percent.

I restrict the sample to person-waves in which the individual's combined previous-wave non-housing wealth and annual income is at least $100,000 in order to reduce the bias from censoring by Medicaid, charities, and uncompensated care. These factors tend to limit an individual's out-of-pocket medical spending to his or her net wealth or liquid assets, which means that including low-net-worth individuals in the sample would bias downward the estimate of the risk people face.[2]

Second, I estimate the means and variances of the acute medical spending distributions by running two versions of the following regression:

$$m_{it} = \alpha + \beta \text{female}_i + \gamma \text{age}_{it} + \sum_{h \in \{hc, nh\}} \phi_h \text{health h}_{it} + \sum_{q=2}^{5} \delta_q \text{income quintile q}_i + \varepsilon_{it},$$

where health is either healthy (omitted), home care, or nursing home and the remaining variables are as defined in Section A.1.3. In one version of this regression, the dependent variable is the log of out-of-pocket acute medical spending. In the other, the dependent variable is the square of the log of out-of-pocket acute medical spending. In both cases, person-waves with zero spending, which comprise less than five percent of the sample, are dropped in order to take logs. Together, these regressions and the identity $Var(X) = E(X^2) - (E(X))^2$ identify the mean and variance of the distribution of acute medical spending facing these groups. The sample is the subset of my main sample of single retirees 65 and older whose combined previous-wave non-housing wealth and annual income is at least $100,000, in order to avoid the censoring issue discussed above.

Table 5 presents the results. The results are mostly as expected. On average people in worse health spend more than people in better health, women spend more than men, and older people spend more than younger people. Higher-income people are estimated to spend somewhat less than lower-income people, though none of the coefficients on the income quintile indicators are statistically significant.

---

[2]The proper input to the life cycle model is total medical spending net of care paid for by Medicare, not just out-of-pocket spending by the individual. The key difference between these two objects is care paid for by Medicaid, charities, and uncompensated care. The extent to which care that is not covered by Medicare is paid for by the individual rather than by Medicaid and other means-tested programs is an endogenous outcome of the model that depends in an important way on people's preferences.

The results of the main estimation are robust to large changes in the model of acute medical spending risk, including scaling mean spending up or down by 50 percent. This is because acute medical spending is the type of risk for which saving or buying long-term care insurance are not very effective. The vast majority of people spend little out-of-pocket on acute medical care, given Medicare's relatively comprehensive coverage and holdings of supplementary Medigap policies on top of that. Although people might wish to send extra resources to those rare states of the world in which out-of-pocket acute medical spending is very high, saving and buying long-term care insurance do not target these states well, so the exact model of acute medical spending risk has relatively little effect on predicted behavior.

A.2   Numerical Solution Procedure: Details and Accuracy

I solve the model numerically using value function iteration. The method proceeds by backward induction, beginning from the maximum possible age T. Because the individual dies by age $T + 1$ with probability one and leaves any remaining wealth as a bequest, the age-T value function can be found easily. To solve for the value function at younger ages, I discretize wealth into a fine grid and use piecewise cubic hermite interpolation to evaluate the value function between grid points. For each sex-income-long-term care insurance group and at each age-health-wealth node, I solve for optimal consumption.

The solution produced by such a method is necessarily an approximation, and its accuracy depends on a number of factors, including the wealth grid. The existence of means-tested programs poses a special challenge, since they cause the value function to be non-concave, which in turn means that the individual's first-order condition for optimal consumption is necessary but not sufficient for an optimum. The effects of means-tested programs on the value function are especially pronounced in the regions of the function in which wealth is relatively small. For this reason, I ensure that the wealth grid is especially fine at small values of wealth by combining (i) a grid that is equally-spaced in logs from the maximum of $1,000 and the Medicaid wealth threshold (which in the baseline specification is $0) to $6 million with (ii) a grid that is equally-spaced in levels from -$1,000 to the maximum of $1,000 and the Medicaid wealth threshold. The resulting grid has 196 distinct values.

I turn now to tests of the accuracy of the numerical solution. The tests are based on the Euler equation, the fundamental condition for intertemporal optimization. The first-order

condition for optimal consumption spending is

$$u'(c_t) = \beta \left\{ Pr(h_{t+1} = d | h_t, t; s, y) E_t[(1 + r_t)v'(b_{t+1})] \right.$$

$$\left. + Pr(h_{t+1} \neq d | h_t, t; s, y) E_t \left[(1 + r_t) \frac{\partial V_{t+1}(\hat{x}_{t+1}, h_{t+1}; s, y, ltci)}{\partial \hat{x}_{t+1}} \right] \right\}.$$

This equation is a necessary but not sufficient condition for optimal consumption spending away from corners, $c \in (c_m(h, Pub), c_m(h, Pub) + x)$. The condition about consumption not being at a corner involves one more element than the usual case because of the consumption value of long-term care, $c_m(h_t, Pub_t)$. The usual corner solution is when borrowing constraints bind, i.e., when the marginal utility of consumption today when borrowing as much as possible exceeds the expected discounted marginal utility of resources tomorrow. The consumption value of facility-based care creates another type of corner solution. In certain circumstances, people in facilities might wish they could save some of the consumption that is bundled together with their long-term care, which is not possible. In states of the world in which this is true, the Euler equation does not hold, since the marginal utility of consumption today (when consuming only the goods and services bundled together with long-term care) is strictly less than the expected discounted marginal utility of resources next period. States of the world in which consumption is at a corner are excluded from Euler equation-based tests of numerical solution accuracy.

The marginal increase in future utility from a marginal increase in $\hat{x}_{t+1}$ is

$$\frac{\partial V_{t+1}(\hat{x}_{t+1}, h_{t+1}; s, y, ltci)}{\partial \hat{x}_{t+1}} = \begin{cases} 0, & \text{if } \hat{x}_{t+1} < \bar{x}(h_{t+1}, ltci_i); \\ u'(c_{t+1}), & \text{if } \hat{x}_{t+1} \geq \bar{x}(h_{t+1}, ltci_i) \text{ and } \hat{c}_{t+1} > 0; \\ \Delta \gtrless u'(c_{t+1}), & \text{otherwise.} \end{cases}$$

The marginal increase in future utility is zero if $\hat{x}_{t+1} < \bar{x}(h_{t+1}, ltci_i)$, since any savings simply reduce transfers from means-tested programs one-for-one. The marginal increase in future utility is $u'(c_{t+1})$ if $\hat{x}_{t+1} \geq \bar{x}(h_{t+1}, ltci_i)$ and $\hat{c}_{t+1} > 0$, by the Envelope theorem. The marginal increase in future utility is

$$\Delta \equiv \beta \left\{ Pr(h_{t+2} = d | h_{t+1}, t; s, y) E_{t+1}[(1 + r_{t+1})v'(b_{t+2})] \right.$$

$$\left. + Pr(h_{t+2} \neq d | h_{t+1}, t; s, y) E_{t+1} \left[(1 + r_{t+1}) \frac{\partial V_{t+2}(\hat{x}_{t+2}, h_{t+2}; s, y, ltci)}{\partial \hat{x}_{t+2}} \right] \right\} \gtrless u'(c_{t+1})$$

if $\hat{x}_{t+1} \geq \bar{x}(h_{t+1}, ltci_i)$ and $\hat{c}_{t+1} = 0$. Δ could be less than $u'(c_{t+1})$ due to borrowing constraints; the individual might wish she could borrow in period $t + 1$. Or Δ could exceed $u'(c_{t+1})$ due to the consumption value of facility-based care; the individual might wish she

could sell some of the consumption that comes bundled with her care. If next-period consumption spending is strictly positive, $\hat{c}_{t+1} > 0$, the right-hand side of the first-order condition can be calculated using the optimal consumption function to calculate $u'(c_{t+1})$. This is the idea behind the Euler equation test.

The Euler equation test measures the closeness of the approximate (numerical) solution to the exact solution that satisfies the Euler equation. I follow Judd (1992) and Fella (2014) in calculating Euler equation errors in units of current consumption:

$$EE(s) = \left| 1 - \frac{c^*(s)}{\bar{c}(s)} \right|,$$

where s is the state vector, $c^*(s)$ is the analytical solution of the Euler equation (the exact consumption level at which the marginal utility of consumption equals the expected discounted marginal utility of resources in the next period), and $\bar{c}(s)$ is the (approximate) optimal consumption rule delivered by the numerical solution algorithm. I calculate Euler equation errors for each member of the simulation sample in each year of the sample period in which he or she is alive and not at a corner.

The results suggest that the numerical solution method is performing well. The average and maximum error among everyone in the sample are 0.001 (-6.7 in natural log units) and 0.039 (-3.3 log units), respectively. The average and maximum error among people within five years of the maximum age, at which point errors have accumulated, are 0.002 (-6.5 log units) and 0.004 (-5.5 log units), respectively. These compare favorably with the results reported by Fella (2014) in tests of his endogenous grid method against value function iteration.

A.3   Asymptotic Distribution of the MSM Estimator and Over-identification Tests of the Model's Fit

Pakes and Pollard (1989) and Duffie and Singleton (1993) show that the MSM estimator, $\hat{\theta}$, is consistent and asymptotically normally distributed under regularity conditions satisfied here. The variance-covariance matrix of $\hat{\theta}$ is

$$\Omega_\theta = (G'_\theta W G_\theta)^{-1} G'_\theta W \left[\Omega_g + \frac{N_d}{N_s} \Omega_g + G_\chi \Omega_\chi G'_\chi \right] W G_\theta (G'_\theta W G_\theta)^{-1},$$

where G_θ and G_χ are the gradient matrices of the moment conditions with respect to θ and χ, Ω_g is the variance-covariance matrix of the second-stage moment conditions, Ω_χ is the variance-covariance matrix of the first-stage parameter estimates, and N_d and N_s are the

empirical sample size and the simulation sample size, respectively. I approximate the derivatives in the gradient matrices numerically. The square roots of the diagonal entries of Ω_θ are the standard errors of the second-stage parameter estimates, $\hat{\theta}$.

The number of second-stage moment conditions exceeds the number of second-stage parameters, so over-identification tests of the model are possible. If the model is correct, the (scalar) statistic

$$\hat{\varphi}(\hat{\theta}, \chi_0)'\ R^{-1}\ \hat{\varphi}(\hat{\theta}, \chi_0)$$

converges in distribution to a chi-squared random variable with degrees of freedom equal to the number of second-stage moments less the number of second-stage parameters. In this formula, $\hat{\varphi}(\hat{\theta}; \chi_0)$ is the vector of moment conditions and

$$R = P\left(\frac{\Omega_g}{N_d} + \frac{\Omega_g}{N_s} + G_\chi \Omega_\chi G'_\chi\right) P,$$

where $P = I - G_\theta(G'_\theta W G_\theta)^{-1} G'_\theta W$, except if $W = \Omega_g^{-1}$, in which case $R = \left(\frac{\Omega_g}{N_d} + \frac{\Omega_g}{N_s} + G_\chi \Omega_\chi G'_\chi\right)$. I use this matrix for all of the results in the paper.

I estimate Ω_g and W from the data. Because I adopt many of the first-stage parameter values from other sources rather than estimating them, I treat χ as if it were known with certainty, $G_\chi = 0$. Excluding the correction for the uncertainty in the first-stage parameters tends to make the second-stage parameter estimates appear more precise than they actually are and the fit of the model (as measured by the chi-squared test statistic) appear worse than it actually is. To estimate G_θ, I follow the procedure for analyzing moment conditions of non-smooth functions (Pakes and Pollard, 1989; Newey and McFadden, 1994; Powell, 1994), since the functions inside the moment conditions $\varphi(\theta; \chi)$ are non-differentiable at certain points. This involves estimating the derivatives of the simulated moments with respect to the parameters θ. The procedure approximates the change in the share of people with wealth no larger than a threshold level by assuming that the density of the wealth distribution is constant within a small neighborhood of that threshold.

A.4 Anticipated and Realized Rates of Return on Wealth

Table 6 lists the historical returns data that I use to estimate the anticipated and realized rates of returns on retirees' portfolios. I follow Baker, Doctor and French (2007) and French and Benson (2011) in terms of data sources and assumptions.[3] Using data from the

---

[3]The main exception is that I use a different rate-of-return series for bonds because Baker, Doctor and French's (2007) series does not extend to 2008, the end of my sample period. I am grateful to Eric French for providing me with the historical returns data.

HRS, I classify retirees' assets into the six categories shown in the table as well as a residual "Other" category (which includes vehicles, for example) that I assume earns 0 percent real, after-tax returns. Following Baker, Doctor and French (2007), I assume that Individual Retirement Account (IRA) assets are allocated 60 percent to stocks and 40 percent to bonds and that the rate of return on business assets is a weighted average of the returns on housing and stocks, with an 85 percent weight on housing.

Anticipated returns on wealth in the model.— Individual i in income quintile q believes that she draws an (annual) rate of return on any saving she might have from the following distribution: $r_i \sim N(\mu_r(q), \sigma_r(q)^2)$, where $\mu_r(q)$ and $\sigma_r(q)$ are estimated based on (i) the average portfolio shares of individuals in income quintile q in the data and (ii) historical data on the realized rates of return on different types of assets. For each income quintile, I estimate the average shares of their portfolios held in the asset classes listed in Table 6. Then, using annual rate-of-return data for each asset class from 1960–2010, I estimate the mean and variance of the distribution of annual rates of return for each income quintile based on their portfolio shares. The resulting means and variances are similar if I instead estimate them based on returns during the time period immediately preceding the sample period (1960–1997), rather than including data through the sample period.

Realized returns on wealth in the simulation, $r_{i,t} = \sum_j \alpha_{i,j,t} r_{j,t}$.— Retiree i's realized rate of return in year t is the weighted average of the realized rates of returns on different assets j in year t ($r_{j,t}$), weighted by i's portfolio shares in that year ($\alpha_{i,j,t}$). The portfolio shares of retirees with zero or negative net wealth are set equal to the median shares among people with between \$5,000 and \$15,000 of net worth. I assume that individuals' portfolio shares are the same in years between interviews as they were in the previous year.

Allowing for differences between anticipated and realized returns and estimating person-wave-specific rates of return protect against two potential sources of bias. One potential source of bias is that the sample period, 1998–2008, was characterized by unusually high rates of return on many assets. The average real return earned by a portfolio that matches the asset allocations of retirees around the middle of the wealth distribution was about 6 percent per year over the period, compared to about 4 percent in the three-and-a-half decades leading up to the sample period. Failing to account for the unusually, and probably unexpectedly, high rates of return could bias the results; the naive estimation would attribute wealth outcomes as arising solely from purposeful saving behavior whereas unusual capital gains or losses may have been important as well (Baker, Doctor and French, 2007). The other source of bias that this procedure protects against is that retirees' portfolios vary systematically across the wealth distribution. Retirees in the middle of the wealth distribution, for example, hold more of their wealth in housing than

11

richer and poorer retirees, and the average return on housing wealth was especially high (7.9 percent per year) over the sample period. Ignoring the differences in retirees' portfolios could bias the results by leading the estimation to wrongly attribute differences in wealth as arising solely from differences in saving behavior whereas differences in realized returns may have been important as well.

A.5 Simulation Procedure

Simulated wealth moments.— The simulated wealth moments are analogous to their empirical counterparts. Given a vector of parameter values, θ, I solve the model to find optimal consumption spending, $\hat{c}_t(\hat{w}_t, h_t, t; s, y, ltci; \theta)$. I use these decision rules together with each individual's fixed characteristics, initial state, subsequent health path, and year-specific rates of return on wealth to simulate each individual's wealth as long as they live between 1999–2008. Given the simulated wealth profiles of each individual in the simulation sample, I use the same procedure to calculate the simulated wealth moments from the simulated data as I use to calculate the empirical wealth moments from the actual data.

Simulated long-term care insurance moments.— The simulated long-term care insurance moments are the long-term care insurance ownership rates by wealth quartile among the subset of the simulation sample who were 65–69 years old in 1998. Only people in good health in 1998 are allowed to buy long-term care insurance in the simulation. This is meant to capture the fact that people in bad health are prevented from buying long-term care insurance—their applications are rejected by insurers (Murtaugh et al., 1997; Hendren, 2013). Given a vector of parameter values, θ, I solve the model to find the value functions, $V_t(\hat{x}_t, h_t; s, y, ltci; \theta)$. Simulated long-term care insurance ownership by individual i is one if both (i) i is healthy in 1998 and (ii) i would be better off buying long-term care insurance given his or her state variables; it is zero otherwise:

$$ltci_i^s = \mathbb{1}(h_{i,t_i} = he) \times \mathbb{1}\left[V_{t_i}(\hat{x}_{i,t_i}, h_{i,t_i}; s, y, ltci = 1; \theta) > V_{t_i}(\hat{x}_{i,t_i}, h_{i,t_i}; s, y, ltci = 0; \theta)\right].$$

The simulated aggregate long-term care insurance ownership rates are the averages of the individual ownership indicators among individuals in each wealth quartile. Simulated long-term care insurance ownership depends on θ through the value functions' dependence on θ.

Because long-term care insurance premiums depend on the age at which long-term care insurance is purchased, and because the model must be solved separately for each

long-term care insurance premium schedule, I simulate the demand for long-term care insurance only among healthy 65–69-year-olds and treat them for this purpose as if they were all 67 years old, the average age at which people buy long-term care insurance (Brown and Finkelstein, 2007). Everyone who can buy long-term care insurance therefore faces the same load (proportional markup over actuarial cost); there is no adverse selection in the model once insurance rejections are accounted for.[4] The assumption that people face a one-time decision about whether to buy long-term care insurance—which I make to economize on computation time—is a rough approximation to the fact that people most often purchase long-term care insurance in their 60s (America's Health Insurance Plans, 2007), with an average purchasing age of 67 (Brown and Finkelstein, 2007).

A.6 Roles of Different Features of the Data in Determining the Parameter Estimates

This section discusses the extent to which different features of the data are informative about the key parameters of the model and the sensitivity of the parameter estimates to changes in the first-stage parameter values and second-stage moments.

A.6.1 Bequest Motives, (ϕ, c_b)

Retirees' saving and long-term care insurance choices, when interpreted in standard life cycle models, are highly informative about bequest motives. As reported in Tables 3 and 5, across a wide range of first-stage parameter values and second-stage estimating moments, the estimates imply that bequests are a luxury good, that bequest motives significantly increase saving and decrease holdings of long-term care insurance and annuities, and that versions of the model without bequest motives are highly inconsistent with retirees' choices. The estimated bequest motive is pinned down relatively sharply and is not very sensitive to changes in the first-stage parameter values and second-stage moments. Across the wide range of specifications in Tables 3 and 5, $\hat{c}_b$ is always between \$12,500 and \$30,000 and usually between \$15,000 and \$20,000, $\hat{\phi}$ is always between 0.93 and 0.99 and usually between 0.95 and 0.96, and the restriction implicit in nested versions of the model without bequest motives is always strongly rejected (in all cases $p \ll 0.01$).

---

[4]In practice, insurance companies limit adverse selection by denying coverage to people with certain health conditions (Murtaugh et al., 1997; Hendren, 2013) and by front-loading premiums to minimize policy lapsation by people who remain healthy (Hendel and Lizzeri, 2003). In long-term care, Finkelstein and McGarry (2006) find that average long-term care usage is roughly equal for the insured and uninsured population, though Finkelstein, McGarry and Sufi (2005) find that people who become healthier than average are more likely than others to drop their coverage.

Figure 1 plots four versions of the classical minimum distance objective as a function of the bequest motive parameters, ϕ and c_b, holding fixed the other parameters at their baseline estimates. Each of the four objective functions is based on a different set of moment conditions: the baseline set of wealth and long-term care insurance moments, only the wealth moments, only the long-term care insurance and median wealth moments, and only the median wealth moments. With the exception of the objective function based on the median wealth moments alone, the objective functions are well behaved and imply that the underlying data are highly informative about bequest motives. The objective functions all feature a single, small "valley" in the same location, centered on the estimate, the "hills" around which increase steeply as ϕ and c_b move away from their estimated values in any direction. These imply that the wealth moments alone, the long-term care insurance and median wealth moments together, and, especially, the full set of long-term care insurance and wealth moments are all much more consistent with models in which bequests are valuable luxury goods than with other configurations, including those with no bequest motive.

The median wealth moments alone, by contrast, are relatively uninformative about the bequest motive parameters: Many combinations of ϕ and c_b are similarly consistent with these moments. This illustrates the lack of power in the saving choices of retirees at a particular point in the wealth distribution to discriminate between different underlying preferences. But as the other figures show, taking into account a broader set of patterns—the saving of retirees with different levels of wealth or retirees' saving and long-term care insurance choices together—is a powerful way to discriminate between different underlying preferences. Both broader sets of patterns are highly inconsistent with versions of the model without bequest motives but are matched well by the model with bequest motives.

A.6.2 Non-Bequest Motive Parameters, $(c_{pub}, \bar{x}_{comm}, \beta, \sigma)$

Retirees' saving and long-term care insurance choices are less informative about the other, non-bequest motive parameters. As reported in Tables 3 and 5, many changes in the first-stage parameter values and second-stage moments have non-negligible effects on the estimates of the discount factor, the coefficient of relative risk aversion, and the consumption values of means-tested programs. Across the specifications reported in Tables 3 and 5, for example, $\hat{\beta}$ varies from 0.80 to 0.95, $\hat{\sigma}$ from 2.0 to 6.1, $\hat{c}_{pub}$ from \$4,000 to \$19,100, and $\hat{\bar{x}}_{comm}$ from \$1,100 to \$5,200.

Figure 2 plots the baseline objective function as a function of different combinations of

parameters, holding fixed the other parameters at their baseline estimates. The objective function is well behaved, but in many cases it is not very informative about the non-bequest motive parameters. Panel (a) shows that a wide range of β and $\bar{x}_{comm}$ values are similarly consistent with retirees' saving and long-term care insurance choices. Panel (b) compares the extent to which β and c_b are pinned down by the estimation. β is extremely poorly pinned down, as values of β from 0.75 to 0.95 are similarly consistent with the data. c_b is much more tightly pinned down, between about \$13,000 and \$23,000, despite its range being increased by interaction effects with β given β's large range. Panel (c) shows that β is also much less well pinned down than σ. Panel (d) shows that c_{pub} is poorly pinned down as well, as values between about \$10,000 and \$20,000 are similarly consistent with retirees' saving and long-term care insurance choices as a whole.

A.6.3   Why Bequest Motives are Pinned Down More Sharply than the Other Parameters

Retirees' saving and long-term care insurance choices are highly informative about bequest motives and less informative about the other parameters because bequest motives affect saving and long-term care insurance choices in a way unlike those of any of the other second-stage parameters (or, indeed, any of a large set of plausible changes one might make to the model, including, for example, if people over- or underestimate health spending risk or life expectancy), whereas the other parameters tend to affect saving and long-term care insurance choices in ways that are similar to one another.

All of the non-bequest motive parameters affect saving and long-term care insurance in the same direction. Saving and long-term care insurance are monotonically increasing in β and σ and monotonically decreasing in c_{pub} and $\bar{x}_{comm}$. As a result, a change in one or more of these parameters can be roughly offset by changes in others, so many different combinations of values of these parameters have similar implications for retirees' saving and long-term care insurance choices. The result is that the estimates of these parameters are not pinned down very sharply; they are sensitive to changes in the first-stage parameter values and the relative weights of different second-stage moments in the estimation. In other words, retirees' saving and long-term care insurance choices are relatively uninformative about the non-bequest motive parameters because these parameters all have similar effects on saving and long-term care insurance.[5] This is an example of the common

---

[5]Most of the difficulty lies in pinning down the values of these parameters jointly, not individually given values of the others. The standard errors on the estimates of these parameters, with the partial exception of c_{pub}, tend to be small. Across specifications, σ tends to be negatively related to β and positively related to c_{pub} and $\bar{x}_{comm}$, presumably because the effects of a given increase in σ can be roughly offset by a decrease in β or an increase in c_{pub} or $\bar{x}_{comm}$.

finding that risk aversion and time preferences are often not sharply pinned down in estimated life cycle models, since changes in risk aversion and time preferences have similar effects on many behaviors. As a result, one should not draw strong conclusions about the value of the non-bequest motive parameters from this evidence.

Bequest motives in which bequests are a luxury good, by contrast, tend to increase saving but reduce long-term care insurance by reducing the opportunity cost of precautionary saving. That is why retirees' saving and long-term care insurance choices are highly informative about bequest motives; bequest motives play a key role in allowing the model to match observed behavior in which many retirees hold much of their wealth well into retirement yet do not buy annuities or long-term care insurance. This is also why the bequest motives are pinned down well even though the other parameters are not.

A.6.4   Why $\hat{\beta}$ Tends to be Low

While most of the parameter estimates take standard or plausible-seeming values, the estimates of the discount factor, β, tend to be unusually low. Across a wide range of specifications, $\hat{\beta}$ ranges from values around 0.95, a typical value in the literature, down to values as low as 0.80, which implies strong impatience. The large range indicates that β is not well pinned down by this evidence, so it would be wrong to conclude that retirees' saving and long-term care insurance choices are strongly indicative of a low discount factor. But the estimates of β tend to be lower than is often the case in this literature, with a central tendency around 0.9, so it is useful to discuss why this might be.[6]

With many estimating moments and parameters, determining the relative roles of different features of the data in driving a particular parameter estimate is not straightforward; all of the parameter estimates are determined jointly by all of the moments. But a variety of tests suggest that the low values of β arise from the difficulty of matching the wealth holdings of the poor, especially in combination with the low rates of long-term care insurance ownership throughout the wealth distribution.

The model has trouble matching the wealth holdings of poor people. In the data, many people report holding small-but-positive amounts of wealth and fewer report holding zero wealth. Among my sample of single retirees, for example, of the roughly 27 percent of person-waves in which wealth is no greater than \$10,000, about 56 percent have strictly positive wealth. In the model, with plausible-seeming parameter values fewer people hold

---

[6]The strength of the baseline estimation's "preference" for a low value of β is moderate, much less than its preference for bequest motives but not a matter of indifference either. The p-value of the restriction that $\beta = 0.95$ is 0.015, that $\beta = 0.925$ is 0.06, and that $\beta = 0.90$ is 0.48.

small-but-positive amounts of wealth and more hold zero wealth. When $\theta = (c_{pub} = \$20,000,\ \phi = 0.95,\ c_b = \$20,000,\ \sigma = 3,\ \beta = 0.97,\ \bar{x}_{comm} = \$7,000)$, for example, the model matches well the long-term care insurance moments, the median and 75th percentile wealth moments, and even the (low and not targeted) 25th percentiles of wealth, but it over-predicts the "probability of zero wealth" moments by an average of almost 16 percentage points, 29.4 percent vs. 13.8 percent.

By increasing the strength of precautionary motives to save (i.e., increasing σ and decreasing c_{pub} and $\bar{x}_{comm}$), the model can better match the probability of zero wealth. But, absent other adjustments, this comes at the expense of dramatically over-predicting saving higher in the wealth distribution and long-term care insurance holdings. When θ is the baseline estimate except with $\beta = 0.97$, for example, the model matches pretty well the "probability wealth equals zero" moments (overstating them by 3.0 percentage points on average) but over-predicts the long-term care insurance moments by 8.5 percentage points on average, the median wealth moments by about \$16,100 on average, and the 75th percentile wealth moments by about \$46,400 on average.

Reducing β helps the model reduce the extent to which it over-predicts saving higher in the wealth distribution and long-term care insurance ownership, while only slightly worsening its over-prediction of the probabilities of zero wealth. When θ is the baseline estimate, the model matches each set of moments pretty well. It under-predicts the long-term care insurance moments by 1.4 percentage points and the median wealth moments by \$2,700 on average, and it over-predicts the 75th percentile wealth moments by \$7,300 and the zero-wealth moments by 4.3 percentage points on average.

To summarize, the estimations tend to favor a relatively low β because that tends to be the least-costly way (in terms of the objective-function penalty) to reduce the over-predictions of saving higher up in the wealth distribution and long-term care insurance ownership that result from their efforts to match the prevalence of small-but-positive wealth levels among the poor.

All of the key conclusions are highly robust to calibrating β to more standard values.

A.6.5 Why it is "Hard" for the Model to Match the Very Bottom of the Wealth Distribution

The proximate reason the model has trouble matching the prevalence of small-but-positive wealth holdings with reasonable parameter values is that holding small amounts of wealth means forgoing consumption today in exchange for small expected future benefits. The

expected future benefits are low mainly because of health spending shocks, which, given the presence of means-tested programs, effectively function as stochastic wealth taxes. The ultimate reason the model has trouble matching the prevalence of small-but-positive wealth holdings with reasonable parameter values is that there is a mismatch between the concept of wealth in the model and available measures in the data.[7]

The model does not include many of the real-world motives to hold at least small amounts of wealth, many of which arise from the desire to economize on transactions costs. For example, in reality most people who regularly consume "car services" own their own car rather than renting car services on a flow basis. There is no analogous motive for holding wealth in the model. Moreover, there is a mismatch in the timing of when wealth is measured. In the model, wealth is measured "between periods," immediately after the income and spending from one period are realized and immediately before the income and spending from the next period occur. In the data, by contrast, wealth is measured whenever the individual happens to be surveyed, which is unlikely to coincide perfectly with the analog—to the extent there even is one—of the beginning- or end-of-period timing in the model. Someone who lives "hand-to-mouth," exhausting her resources by the time the next income payment arrives, will typically have a small but positive amount of wealth when she is surveyed by the HRS, despite not saving anything from the perspective of the model.

These mismatches cause the estimations to "stretch"—i.e., adjust the values of some of the parameters away from the more standard values that they would otherwise prefer—to try to reduce the predicted share of people with zero wealth to come into closer alignment with the low share in the data. Reducing β enables it to do so without missing the other target moments too badly, especially long-term care insurance holdings.

All of the key conclusions are highly robust to a wide range of changes in the assumptions and target moments that at least partially address this mismatch, including using different measures of wealth (e.g., excluding the value of cash and vehicles or housing) and dropping the zero-wealth moments.

---

[7]This mismatch is present for other models in this literature as well, but it becomes important mainly when low wealth levels are targeted, which is rarely done.

References

America's Health Insurance Plans. 2007. "Who Buys Long-Term Care Insurance? A 15-year Study of Buyers and Non-Buyers, 1990-2005." Prepared for America's Health Insurance Plans by LifePlans, Inc.

Baker, Olesya, Phil Doctor, and Eric French. 2007. "Asset rundown after retirement: the importance of rate of return shocks." *Economic Perspectives*, 2Q: 48–65.

Brown, Jeffrey R., and Amy Finkelstein. 2007. "Why is the market for long-term care insurance so small?" *Journal of Public Economics*, 91(10): 1967–1991.

De Nardi, Mariacristina, Eric French, and John Bailey Jones. 2009. "Life Expectancy and Old Age Savings." *The American Economic Review*, 99(2): 110–115.

De Nardi, Mariacristina, Eric French, and John Bailey Jones. 2010. "Why Do the Elderly Save? The Role of Medical Expenses." *Journal of Political Economy*, 118(1): 39–75.

Duffie, Darrell, and Kenneth J. Singleton. 1993. "Simulated moments estimation of Markov models of asset prices." *Econometrica*, 929–952.

Fella, Giulio. 2014. "A generalized endogenous grid method for non-smooth and non-concave problems." *Review of Economic Dynamics*, 17(2): 329–344.

Finkelstein, Amy, and Kathleen McGarry. 2006. "Multiple Dimensions of Private Information: Evidence from the Long-Term Care Insurance Market." *American Economic Review*, 96(4): 938–958.

Finkelstein, Amy, Kathleen McGarry, and Amir Sufi. 2005. "Dynamic Inefficiencies in Insurance Markets: Evidence from long-term care insurance." *American Economic Review*, 95(2): 224–228.

French, Eric, and David Benson. 2011. "How do sudden large losses in wealth affect labor force participation?" *Chicago Fed Letter*, 282(Jan).

Friedberg, Leora, Wenliang Hao, Wei Sun, Anthony Webb, and Zhenyu Li. 2014. "New Evidence on the Risk of Requiring Long-Term Care." Center for Retirement Research Working Papers, Center for Retirement Research at Boston College wp2014-12.

Hendel, Igal, and Alessandro Lizzeri. 2003. "The Role Of Commitment In Dynamic Contracts: Evidence From Life Insurance." *The Quarterly Journal of Economics*, 118(1): 299–327.

Hendren, Nathaniel. 2013. "Private information and insurance rejections." *Econometrica*, 81(5): 1713–1762.

Judd, Kenneth L. 1992. "Projection methods for solving aggregate growth models." *Journal of Economic Theory*, 58(2): 410–452.

MetLife Mature Market Institute. 2002*a*. "MetLife Market Survey of Nursing Home and Home Care Costs."

MetLife Mature Market Institute. 2002*b*. "MetLife Survey of Assisted Living Costs."

Murtaugh, Christopher M., Peter Kemper, Brenda C. Spillman, and Barbara L. Carlson. 1997. "The Amount, Distribution, and Timing of Lifetime Nursing Home Use." *Medical Care*, 35(3): 204–218.

Newey, Whitney K., and Daniel McFadden. 1994. "Large sample estimation and hypothesis testing." In *Handbook of Econometrics*. Vol. 4 of *Handbook of Econometrics*, , ed. R. F. Engle and D. McFadden, Chapter 36, 2111–2245. Elsevier.

OECD. 2011. "An Overview of Growing Income Inequalities in OECD Countries: Main Findings." *SourceOECD General Economics & Future Studies*, 2011(14): 6–30.

Pakes, Ariel, and David Pollard. 1989. "Simulation and the Asymptotics of Optimization Estimators." *Econometrica*, 57(5): 1027–1057.

Powell, James L. 1994. "Estimation of semiparametric models." In *Handbook of Econometrics*. Vol. 4 of *Handbook of Econometrics*, , ed. R. F. Engle and D. McFadden, Chapter 41, 2443–2521. Elsevier.

Robinson, James. 2002. "A Long-term Care Status Transition Model." mimeo, University of Wisconsin, Madison.

Appendix Figures and Tables

(a) Objective function: baseline

(b) Objective function: wealth only

(c) Objective function: LTCI and median wealth

(d) Objective function: median wealth only

Figure 1: Contour plots of different versions of the objective function in (c_b, ϕ)-space with the other parameters held fixed at their baseline estimated values. Higher contours indicate greater mismatch between the simulated and empirical moments. The asterisks mark the baseline estimates. All plots use the same scale.

Figure 2: Contour plots of the baseline objective function as a function of different pairs of θ parameters with other parameters held fixed at their baseline estimated values. Higher contours indicate greater mismatch between the simulated and empirical moments. The asterisks mark the baseline estimates. All plots use the same scale.

(a) Odd-numbered cohorts

(b) Odd-numbered cohorts

(c) Even-numbered cohorts

(d) Even-numbered cohorts

Figure 3: Empirical wealth moments (solid lines) and simulated wealth moments (dashed lines) for odd- and even-numbered cohorts under the baseline model with bequest motives. Panels (a) and (c) show the 25th, 50th, and 75th percentiles of the wealth distributions among surviving members of each cohort. The 25th percentiles are not included in the estimation. Panels (b) and (d) show the share with zero wealth among surviving members of each cohort. The x-axis shows the average age of surviving members of the cohort.

Figure 4: Cumulative distribution function of wealth in the last wave in which an individual is alive among individuals who die during the sample period. Wealth in the last period in which an individual is alive is a proxy for realized bequests that is better-measured than actual bequests (see, for example, De Nardi, French and Jones, 2010). The simulated distribution is generated by the baseline model and parameter estimates.

	(1) Used formal home care since last interview	(2) Stayed in nursing home since last interview	(3) Owns long-term care insurance
Female=1	0.0494 (0.0333)	-0.0548 (0.0262)	0.0123 (0.00512)
Single=1	0.0263 (0.0402)	0.124 (0.0299)	-0.0207 (0.00564)
No kids=1	0.0310 (0.0448)	0.0905 (0.0336)	0.00535 (0.00973)
Age	0.0219 (0.0309)	0.0721 (0.0247)	-0.00438 (0.00711)
Age^2	-0.0000788 (0.000194)	-0.000360 (0.000151)	0.00000950 (0.0000468)
Number of ADLs (0-5)=3	0.0759 (0.0367)	0.0429 (0.0282)	
Number of ADLs (0-5)=4	0.187 (0.0425)	0.107 (0.0327)	
Number of ADLs (0-5)=5	0.292 (0.0431)	0.276 (0.0321)	
Income quartile=2	0.0371 (0.0379)	0.0262 (0.0281)	
Income quartile=3	0.0618 (0.0475)	0.0127 (0.0375)	
Income quartile=4	0.0391 (0.0551)	0.0719 (0.0451)	
Wealth quartile=2			0.0124 (0.00464)
Wealth quartile=3			0.0482 (0.00577)
Wealth quartile=4			0.137 (0.00776)
Observations	1083	1360	10599
ymean	0.422	0.283	0.0738
Sample restrictions	Age 65+ 2+ ADLs	Age 65+ 2+ ADLs	Age 65+

Marginal effects; Standard errors in parentheses

Table 1: Marginal effects from probit regressions, i.e., the increase in the average predicted probability of the dependent variable being one if everyone in the sample had their value of the indicator variable in question increased from zero to one or their value of the continuous variable in question (age) increased by one unit. Columns 1, 2, and 3 report results from probit regressions of indicator variables for whether the individual used any (formal) home care since the last interview, whether the individual stayed in a nursing home since the last interview, and whether the individual owns long-term care insurance, respectively. All of the columns restrict the sample to people age 65 and older. Columns 1 and 2 further restrict the sample to people who report having problems with at least two activities of daily living. The difference in the number of observations between columns 1 and 2 reflect a difference in the number of missing values of the dependent variables. Age is measured in years. Wealth quartiles are calculated based on wealth values that are adjusted for whether the individual is part of a one- or two-person household according to the widely-used square root equivalence scale (e.g., OECD, 2011) (so an individual in a couple is assigned a wealth value equal to his or her household wealth divided by $\sqrt{(2)}$ before calculating quartiles). The qualitative results are not sensitive to plausible alternatives.

	Healthy males			Healthy females		
		Life expectancy at 70			Life expectancy at 70	
Income quintile	Age adjustment to FHSWL, Δ	De Nardi et al. (2010)	Adjusted FHSWL	Age adjustment to FHSWL, Δ	De Nardi et al. (2010)	Adjusted FHSWL
1	17	7.6	7.7	6	12.8	13.0
2	15	8.4	8.4	5	13.8	13.6
3	13	9.3	9.3	3	14.7	14.8
4	11	10.5	10.3	2	15.7	15.4
5	9	11.3	11.3	0	16.7	16.8

Table 2: Adjustments to the Friedberg et al. (2014) (FHSWL) model of health transitions for females to match the life expectancy differences across sex and income groups documented by De Nardi, French and Jones (2010).

	(1) Life	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
		Share of time $h =$				Ever	Years	Years	Leave $h = nh$	Leave $h = nh$
	expect	he	hhc	alf	nh	$h = nh$	$h = nh$	$h = nh \mid > 0$	alive	alive $\mid$ ever $h = nh$
Males										
Original FHSWL	16.8	0.95	0.02	0.01	0.02	0.43	0.39	0.90	0.35	0.82
Yearly transitions	17.3	0.95	0.02	0.01	0.02	0.20	0.38	1.92	0.11	0.56
By income quintile										
Bottom	9.8	0.81	0.03	0.04	0.12	0.45	1.19	2.65	0.22	0.48
Second	10.8	0.83	0.03	0.04	0.10	0.43	1.12	2.59	0.22	0.51
Third	11.9	0.84	0.04	0.03	0.09	0.41	1.05	2.56	0.22	0.53
Fourth	13.0	0.86	0.04	0.03	0.08	0.39	0.99	2.53	0.21	0.54
Top	14.2	0.87	0.04	0.03	0.07	0.37	0.94	2.52	0.21	0.55
Females										
Original FHSWL	20.3	0.91	0.03	0.02	0.04	0.57	0.83	1.46	0.46	0.80
Yearly transitions	20.8	0.91	0.03	0.02	0.04	0.33	0.82	2.47	0.19	0.56
By income quintile										
Bottom	16.1	0.89	0.04	0.02	0.05	0.35	0.88	2.50	0.20	0.56
Second	16.8	0.89	0.04	0.02	0.05	0.35	0.86	2.49	0.19	0.56
Third	18.3	0.90	0.03	0.02	0.05	0.34	0.84	2.48	0.19	0.56
Fourth	19.1	0.90	0.03	0.02	0.04	0.34	0.83	2.48	0.19	0.56
Top	20.8	0.91	0.03	0.02	0.04	0.33	0.82	2.47	0.19	0.56

Table 3: Health transitions model statistics. "Original FHSWL" corresponds to the Friedberg et al. (2014) (FHSWL) model of monthly health transitions. The other rows show simulated statistics from adjusted versions of the model. The adjustments are: switching from monthly to yearly transitions to economize on computation time and better match the frequency of the HRS data, using FHSWL's female model for males to better reflect the long-term care risk facing single retirees, and adjusting the male and female models to match the heterogeneity in remaining life expectancy from age 70 documented by De Nardi, French and Jones (2010). Column (1) shows life expectancy at age 65. Columns (2)–(5) show the expected shares of time from age 65 spent healthy, receiving home care, living in assisted living facilities, and living in nursing homes, respectively. Column (6) shows the probability of ever living in a nursing home. Column (7) shows expected years spent in nursing homes. Column (8) shows expected years in nursing homes conditional on living in a nursing home at some point. Column (9) shows the unconditional probability of ever leaving a nursing home alive. Column (10) shows the probability of leaving a nursing home alive conditional on living in a nursing home at some point. I am grateful to Wenliang Hou and Tony Webb for providing me with the transition matrices from Friedberg et al. (2014).

	Relative price of nursing home care
Female	0.0604
	(0.0669)
Age	-0.0000976
	(0.00369)
2nd income quintile	0.0783
	(0.0799)
3rd income quintile	0.0164
	(0.0782)
4th income quintile	0.0958
	(0.0849)
Top income quintile	0.115
	(0.0812)
N	536

Table 4: Regression of the relative price of nursing home care using data from the NLTCS. Standard errors in parentheses.

	(1) Log acute medical spending	(2) Square of log acute medical spending
Age	0.00707	0.0884
	(0.00499)	(0.0712)
Female	0.202	3.158
	(0.0729)	(1.004)
Home care	0.331	4.905
	(0.0711)	(1.024)
Nursing home	0.476	6.970
	(0.105)	(1.513)
2nd income quintile	-0.0112	-0.115
	(0.159)	(2.334)
3rd income quintile	-0.149	-2.237
	(0.148)	(2.146)
4th income quintile	-0.141	-1.920
	(0.142)	(2.070)
Top income quintile	-0.169	-2.186
	(0.140)	(2.052)
N	3969	3969

Table 5: Acute medical spending regressions. The omitted dummy variables are "healthy" and "bottom income quintile." The sample is the subset of my main sample (single retirees 65 and older) whose combined previous-wave non-housing wealth and annual income was at least \$100,000 and who have strictly positive spending (in order to take logs). Home care indicates whether the individual used home care since the last interview. Nursing home indicates whether the individual is living in a nursing home at the time of the interview. Standard errors in parentheses.

Asset	Data source	Taxation	Return, 1998–2008		Portfolio share
			Mean	Std. dev.	(percent)
Occupied housing	OFHEO, Baker et al. (2007)	0 percent on capital gains, 1 percent/yr property tax	7.9	3.2	54.8
Stocks	CRSP	0 percent on capital gains, 20 percent on div yield (2 percent yield)	2.6	16.9	9.3
Bonds	AAA long bonds yield to maturity	20 percent	3.2	1.0	2.2
Liquid (CDs)	Treasury	20 percent	1.2	1.4	6.9
Unoccupied housing	OFHEO	0 percent	4.3	3.2	1.5
Debt	Baker et al. (2007)	20 percent	2.4	-	-16.9

Table 6: Data sources and assumptions underlying the calculations of the expected and realized rates of return on wealth. The mean returns are the geometric averages of annual real, after-tax returns, in percent. The portfolio shares are the average shares of net wealth held in each asset in 1998 by the sample of single retirees, weighted by HRS respondent-level weights. The assumption of zero taxation of capital gains comes from the assumption that a large share of retirees' capital gains are not realized (by asset sales) during the sample period. Additional details about the data sources can be found in Baker, Doctor and French (2007).

	Average deviations, (simulated-empirical)			
	LTCI (p.p.)	Medians ($1,000s)	75th ptiles ($1,000s)	Pr(w=0) (p.p.)
Baseline	−1.4	−2.7	7.3	4.3
No BM	9.1	−11.1	−55.1	6.3
No BM, medians	33.5	2.5	−50.1	2.0
No BM, 0s and p75s	44.0	29.6	−9.9	3.3

	Average absolute deviations, |simulated-empirical|			
	LTCI (p.p.)	Medians ($1,000s)	75th ptiles ($1,000s)	Pr(w=0) (p.p.)
Baseline	2.7	10.9	27.5	5.9
No BM	9.9	15.4	62.7	6.8
No BM, medians	33.5	10.7	60.8	5.3
No BM, 0s and p75s	44.0	32.5	47.8	4.4

Table 7: Economic fit of different estimated models to each set of moment conditions. The estimations are the baseline estimation ("Baseline"), the main estimation without bequest motives ("No BM"), the estimation without bequest motives based on the median wealth moments ("No BM, medians"), and the estimation without bequest motives based on the "share with zero wealth" and 75th percentile wealth moments ("No BM, 0s and p75s"). The first set of rows shows the average excess of the simulated moments over the empirical moments of each type. The last set of rows shows the average absolute deviations of the simulated moments from the empirical moments of each type. Deviations from the long-term care insurance and probability-of-zero-wealth moments are in percentage points.