Strategic Surveys and the Bequest Motive

John Ameriks[*], Andrew Caplin[†], Steven Laufer[†], and Stijn Van Nieuwerburgh[‡]

December 15, 2005

Abstract

Strong bequest motives can explain slow asset run-down in retirement. Yet strong precautionary motives provide an equally compelling explanation, and it is hard to tease these two motives apart empirically. We develop a series of "strategic" survey questions that allow for bequest and precautionary motives to be separately identified. Our questions are developed in the context of a model that stresses long term care costs and aversion to Medicaid as the chief drivers of precautionary savings.

1 Introduction

The life cycle model of Modigliani and Brumberg [1954] predicts asset decumulation in the retirement period, and as a result low intergenerational transfers. Yet the elderly dissave little, and intergenerational transfers account for a substantial proportion of total wealth.[1] These facts led Kotlikoff and Summers [1981] to argue that the bequest motive is the primary driver of wealth accumulation. Yet even now, more than two decades later, the issue of how important are bequest motives remains largely unresolved.

Current empirical strategies for measuring the importance of bequest motives exploit individual differences. That such differences are profound is confirmed by the fact that while many wealthy households fail to pursue such obvious tax avoidance strategies as intervivos giving (McGarry [1999]; Poterba [2001]), others take great pains to maximize bequests.[2] Further confirmation is provided by responses to survey questions on the subjective importance of bequests (e.g. Laitner

[*]The Vanguard Group

[†]Corresponding author, New York University, Department of Economics, 110 Fifth Ave #432, New York, NY, 10011, acaplin@nyu.edu, Tel: (212) 998-8950

[‡]New York University, Stern School of Business, Department of Finance, svnieuwe@stern.nyu.edu. Preliminary and incomplete. Comments welcome. We thank Mark Gertler, Miles Kimball, and the participants of the NYU macroeconomics seminar for helpful comments. The views expressed herein are those of the authors rather than of their institutions.

[1]Findings on wealth dynamics in retirement and the proportion of wealth that is accounted for by bequests are summarized briefly in the next section. See Kopczuk and Lupton [2005] for a more thorough review.

[2]Strategies employed apparently include not only using inter vivos transfers to decrease tax liabilities involved in transferring funds to heirs (Bernheim et al. [2001]; Page [2003]; Bernheim et al. 2004; Joulfaian [2004]), but also offsetting increased public transfers by purchasing life insurance and selling annuities (Bernheim [1991]), and even delaying (or accelerating) death to take advantage of changes in estate-tax law (Slemrod and Kopczuk [2001]).

and Juster [1996]).[3] In an effort to exploit these differences for estimation purposes, Hurd [1987] studies differences in consumption profiles for otherwise similar retired households with and without children. He finds that they decumulate wealth at roughly the same pace, casting doubt on the power of the bequest motive. This apparently negative finding has been challenged by Kopczuk and Lupton [2005], who argue against the strategy of identifying bequest motives with the presence or absence of children. His alternative strategy identifies the extent of these motives indirectly via differences in the extent of asset run down in retirement for demographically similar households. His results suggest that the majority of the elderly population have strong bequest motives, and that they spend some 25% less than do those with low such motives.

One limitation of Kopczuk and Lupton's approach is that bequest motives are not the only possible explanation for low spending in retirement. Precautionary motives provide an equally compelling explanation for low spending among the elderly. As individuals age their earning options diminish, and they may want to hold on to their wealth both to guard against the possibility of a longer than expected life span (Yaari [1965]) and to guard against potentially high medical and long term care expenses (Palumbo [1999]). This same identification problem impacts a second line of research on bequest motives, in which they are inferred indirectly by noting the low use made of annuities (e.g. Bernheim, Shleifer, and Summers [1985]).[4] The lack of success in past efforts to distinguish between bequest and precautionary motives has led such prominent researchers as Dynan, Skinner, and Zeldes [2002] to argue that they may be empirically indistinguishable.

> These motives for saving are overlapping and cannot generally be distinguished. A dollar saved today simultaneously serves both a precautionary life-cycle function (guarding against future contingencies such as health shocks or other emergencies) and a bequest function because, in the likely event that the dollar is not absorbed by these contingencies, it will be available to bequeath to children or other worthy causes. (274)

In this paper we expand upon the insight of Dynan, Skinner and Zeldes concerning the limitations of behavioral data in separating bequest and precautionary motives, and develop a "strategic" survey methodology designed to shed new light on their relative importance (Barsky, Juster, Kimball, and Shapiro [1997], Kimball and Shapiro [2003] and Kimball, Sahm, and Shapiro [2005] have done much to blaze this trail). This methodology exploits the fact that the decision maker's spending strategy is potentially far more revealing than is behavior alone. This allows a well-designed strategic survey question to explore anticipated behavior in particularly revealing contingencies. We introduce methods to aid in the design of such questions, and illustrate their applicability in

[3]If anything, these subjective answers suggest that bequest motives rank below life-cycle and precautionary considerations as motivations for saving. Dynan, Skinner, and Zeldes [2004] report results from the 1998 Survey of Consumer Finance households in which households were asked to choose up to five from a larger list of motives for saving. Retirement was reported by 45 percent of all households as a reason for accumulation. Saving for emergencies or illness also figured prominently, particularly among the elderly, where 40 percent listed one or both of these reasons as a motivation. In stark contrast, saving for one's estate or children was mentioned by only 12 percent of retired households.

[4]Computations in sections 5 and 7 below show that precautionary motives may equally well explain the apparently low interest in annuities.

the context of the debate concerning bequest motives. The most subtle aspect of survey design is that the questions can only be written after the complete model has been specified and the information content of the behavioral data has been characterized.

The current state of our understanding of retiree spending strategies and their connection to models of life cycle consumption is outlined in section 2 below. Section 3 presents our model of spending in the retirement period. We follow Brown and Finkelstein [2004] in assigning great importance to potentially high long term care costs as motivators of precautionary saving. Section 4 establishes the limitations of behavioral data in separating bequest and precautionary motives. Sections 5 and 6 outline our strategic survey methodology and illustrate methodological issues in the context of the bequest debate. Section 7 indicates that our model has powerful implications for consumer interest in annuities and long term care insurance. While low consumer interest in annuities may accurately reflect underlying consumer preferences, well-designed long term care insurance policies offering consumers appropriate safeguards would easily appear to pass the market test.

2 Background

2.1 Asset Rundown and Bequests

The fact that wealth runs down slowly if at all in retirement is solidly established. In fact cohort analysis suggests that wealth actually rises in early retirement (Menchik and David [1983]). While some estimates using panel data suggest that household wealth increases with age (e.g. Menchik and David [1983]), other studies have shown a decline (e.g. Hurd [1990]). The most recent of these panel studies by Anderson, French, and Lam [2004] uses panel data from the AHEAD Survey (Assets and Health Dynamics of the Oldest Old). The AHEAD is a sample of non-institutionalized individuals who were of age 70 or higher in 1993. A total of 8,222 individuals in 6,047 households were interviewed for the AHEAD survey in 1993. These individuals were interviewed again in 1995, 1998, and 2000. The survey follows individuals until they die. End of life information on deceased respondents is obtained by interviewing proxy respondents in an exit interview. Using the panel structure of the AHEAD data, they are able to document changes in wealth over time for members of different cohorts. When tracking assets of households within a cohort, they find that there is a systematic tendency for wealth to *increase* over the length of the panel. For example, assets increase about 37 percent between 1993 and 2000 for the cohort aged 70-74 in 1993. They use simple techniques to estimate how much of this increase in wealth may have been driven by the historically high return on assets in the late 1990's, and conclude that this reduces the scale of the increase in wealth without changing the qualitative conclusion.

While the facts concerning slow asset rundown are clear, the importance of bequests in total wealth is a far more contentious area. Three different estimation methods have been used. One early technique involved measuring bequests and inferring the impact on wealth. Kotlikoff and Summers (1981) applied this method to argue that as much as 46% of household wealth is accounted for by bequests, while in Modigliani's (1988) hands, the number was 17%. A second technique that produces even more variable estimates is based on comparing life-cycle saving with total wealth. A third technique using data on inheritances and intended bequests produces estimates in the range of 15% to 30% (Menchick and David, 1983; Modigliani, 1988; Gale and Scholz, 1994; Juster and Laitner, 1996).

2.2 Long Term Care Costs and Precautionary Motives

Where possible, estimates of the importance of precautionary savings rely on objective data on health dynamics and medical expense risks. For example the treatment of longevity risk derives largely from application of appropriate life tables. Data on health expense risks have focussed largely on actual out of pocket health expenses and health transitions among the elderly (see French and Jones, 2004; Palumbo, 1999; and Feenberg and Skinner, 1994). French and Jones [2004] estimate the stochastic process that determines both the distribution and dynamics of health care costs based on data from the Health and Retirement Survey and the AHEAD. The French and Jones findings indicate that 1 percent of all households incur a medical expense shock that costs $44,000 over their lifetimes and .1 percent of all households incur a medical expense shock that costs $125,000 over their lifetimes each year.

We believe that the health expense computations conducted to date understate the magnitude and behavioral salience of precautionary motives. In particular, they do not adequately reflect long term care costs and differential attitudes to private as opposed to public long term care options. Using the model of Robinson [2002], Brown and Finkelstein [2004] produce estimates of long term care risks for 65-year old men and women. Their model predicts that a 65 year-old man has a 27 percent chance of entering a nursing home at some future point. The risk is even higher for women; a 65 year-old woman has a 44 percent chance of ever entering a nursing home. When stays occur, they can be long-lasting, in particular for women. Men who enter a nursing home spend on average 1.3 years there, while women spend on average 2 years. There is a considerable right-tail to the distribution of nursing home utilization. Of individuals who enter a nursing home, 12 percent of men and 22 percent of women will spend more than 3 years there; one-in-eight women who enter a nursing home will spend more than 5 years there. Hence it is not unreasonable for a married household to anticipate the need to finance stays of up to five years between them.

The price for obtaining care, if it is required, is potentially massive. Brown and Finkelstein report data from the Metlife Market Survey national data (MetLife 2002a, MetLife 2002b) indicating that the national average daily cost of nursing home care in 2002 is $143 per day for a semi-private room, and that private rooms are far more expensive. In current dollars, the cost of a possible joint stay of five years' length in long term care in a semi-private room would be some $250,000. Not only is this higher than the tail medical events as detailed by French and Jones, but it is also far more likely. A couple may face not only the (annual) .1 percent risk identified by French and Jones of a medical expense shock that costs $125,000 over their lifetimes, but also a (lifetime) 20% or so chance of a long term care shock with the potential to cost at least $250,000.

2.3 Medicaid and Medicaid Aversion

The data above suggest that long term care risk dominates expense risk in later life. The potential for expenditures of this size would appear to be a strong motive for saving as well as for the use of private insurance arrangements that could share and reduce individual risk. However, an unusual aspect of long term care risk is the existence of a public insurance program, Medicaid, as a payer of last resort. Pauly [1990] argues that the presence of Medicaid explains the very limited nature of the long term insurance market. Brown and Finkelstein likewise see Medicaid as reducing private insurance incentives. They estimate that for the median male (female), 60 percent (75 percent) of the benefits from a private policy are redundant of benefits that Medicaid would otherwise have paid.

While Medicaid availability may reduce saving incentives for some, there are countervailing forces. First, an individual who qualifies for Medicaid coverage is allowed to keep very little in the way of income and assets to finance non-care consumption or to bequeath. Married households are allowed to retain only their housing wealth, while single households must essentially deplete all assets before qualifying for Medicaid.[5] Perhaps equally significant is anecdotal evidence concerning the relatively limited nature and low perceived quality of Medicaid as opposed to privately provided long term care. Nightmare stories about patients dying undignified deaths in Medicaid facilities abound.[6] We believe that a realistic model must follow Brown and Finkelstein in allowing for the possibility that an individual will choose to maintain high assets late in life in order to avoid ending up a ward of the state.[7] In fact, we conducted a preliminary survey strongly indicative of the importance of Medicaid aversion.

The survey was conducted in February of 2005 in cooperation with Greenfield Online, a major provider of web-based surveys. Greenfield Online posted the questionnaire on its website and notified a random sample of its "online panel" of members (more than 1 million individuals) that the survey was ready to be filled out. Respondents were screened to include only those who: were 55 years of age or older; expected $20,000 or less in earnings from work in the current year and in all future years; had no dependents other than (possibly) a partner living with them; and reported being the primary (or co-primary) financial decision maker in their household. An undisclosed financial incentive was offered to those who completed the survey. In the space of two weeks, we obtained 1,004 completed responses to the on-line survey.[8]

[5] State Medicaid programs impose a 3 to 5 year look back period on assets to make it more difficult for individuals to hide assets by transferring them to a spouse or children.

[6] The most recent of these include evidence of inadequate protection from fire and other hazards at U.S. nursing home facilities; see U.S. Government Accountability Office [2003].

[7] While the current paper focuses more on the strategic survey methodology and bequest motives, we believe that the lumpy nature of long term care costs and Medicaid aversion make them worthy of independent analysis. In particular, they may combine to give rise to such apparently unusual behavior as a non-monotonic relationship between wealth and spending. When wealth is low, it may be best to give up on any chance of privately affording long term care, getting rid of this motive for saving. With slightly higher wealth, saving for this reason may offer more hope of avoiding Medicaid, and therefore be worthwhile. It is conceivable that this pattern may help explain why the distribution of wealth is much more skewed than the distribution of income (Castaneda et al. [2003]), and also "why the rich save more" (Dynan, Skinner, and Zeldes, 2004).

[8] There are some significant differences in demographic characteristics of our sample relative to the sub-population

We outlined a scenario in which 20 years from now the respondent has $400,000 in assets left and knows that long term care will be required for the remaining years of life. The question posits that the respondent has $400,000 available and that the following three options are available:[9]

- Receive care at a Medicaid facility, and leave a $400,000 estate.
- Receive care at a private facility and leave a $200,000 estate
- Receive care at home and leave a $100,000 estate.

The majority of respondents (62%) indicated they would elect to receive care at home, with 19% each for the other two options. A quantitative follow-up question explored how much respondents would be willing to reduce their estates in order to receive care at the private facility rather than a Medicaid facility. The average willingness to pay among all respondents was $143,000, and was $80,000 even among those who selected Medicaid in the given scenario.

If aversion to Medicaid is such an important factor, why is there so little use of private long term care insurance (over one third of long-term care expenditures are paid for out of pocket, which is nearly double the proportion of expenditures in the health sector as a whole that are paid for out of pocket (CBO [2004], National Center for Health Statistics [2002])? There are several potential direct explanations. Brown and Finkelstein [2004] discuss in detail the many imperfections in the private market for LTC insurance. Prices are marked up substantially above expected claims, with loads on typical policies about 18 cents on the dollar. Maximum daily benefits are typically constant in nominal terms, and thus declining in real terms over time. Given that standard estimates project 1.5 percentage point annual real growth in care costs, one would need to purchase a policy for substantially more than current costs to be protected against future increases in the costs of long-term care. In addition, underwriting requirements for LTC insurance are stringent; pre-existing conditions can mean ineligibility for coverage.

While arguments based on cost may explain reluctance on the part of some consumers to take out LTC insurance, there are a number of more subtle barriers that may be at least as important. Currently available LTC policies do not produce the type of insurance envisioned in most theoretical models. The potential for technological innovation in health care and the long-term nature of the insurance contract make it particularly difficult to pre-specify an adequate range of contingencies under which insurance benefits would have value. Which one of us with experience arguing with insurers about the appropriate treatment of our medical bills would feel sanguine about our chances of winning these arguments as our mental faculties dwindle?

in the Survey of Consumer Finances that meets our age and income screens. Our sample is skewed toward singles, females, and renters: The survey sample is 66% females, while the population in the SCF is roughly 55% female; we have 41% renters, while in the comparable SCF population, 21% rent; and only 31% of our sample are married or living with a partner, while in the SCF 48% are married or living with a partner. While mean wealth is far lower in our sample than in the SCF, the general patterns across demographic group are similar (i.e., owners are wealthier than renters, single renters have very little wealth overall). Moreover there is a close correspondence between the means in our sample and medians in the SCF. Respondents in our survey report earnings levels that are quite similar in both pattern and level to that reported by the more representative SCF sample.

[9] Note that this is an early draft of an "end of life" strategic survey question detailed in section 5 below.

In addition to contractual incompleteness, the cost of insurance premiums are typically not guaranteed for the life of contract. Finally, as with other insurance contracts, there remains a possibility of default by the insurer. With LTC insurance, this introduces a risk that may have many of the same characteristics as the LTC expenditure risk itself (low probability of a disastrous, high-cost outcome).

A recent article in Consumer Reports summarizes the problems (Consumer Reports [2003]),

> "A CR investigation, for which we reviewed 47 policies, reveals that for most people, long-term-care insurance is too risky and too expensive. As with health insurance, you must keep paying to keep it in force. If premiums rise, you may have to drop the coverage, possibly losing everything that you've paid. The policy's benefits may cover only a portion of the total expense. Many policies are packed with catches that can keep you from collecting. Finally, there's no guarantee that long-term-care insurers, some of which have weak balance sheets, will be around 20, 30, or 40 years from now when you need them to pay."

3 The Model

The model describes a simple consumption-savings problem in retirement. It introduces two motives for saving: a precautionary savings motive to guard against longevity risk, medical risk, and long-term care risk, and a bequest motive. We use this model as a laboratory to analyze the trade-off between both savings motives.

3.1 Utility

For simplicity, the unit of analysis is the household consisting of a single individual who has just retired. The first period of observation occurs when the individual is m years old and entering retirement. The model consists of a series of one-year periods, starting at the age of retirement and ending at the year of death, which is restricted to occur by maximum age M. The maximum length of the retirement period is $T = M - m$. Periods are indexed by t, the number of years into the retirement period, starting at zero at age m, so that overall $0 \leq t \leq T$. There is a stochastic death rate δ_t in year t of retirement that evolves in a matter defined below.

The agent maximizes a standard time-separable utility function with exponential discounting. In each period of life, agents receive utility from consumption in excess of a subsistence level, c^{SUB}:

$$u(c_t) = \frac{(c_t - c^{SUB})^{1-\gamma}}{1-\gamma}. \tag{1}$$

Agents also receive end-of-life utility from bequests defined by the function $v(b)$. Hence the agent maximizes,

$$E_0 \sum_{t=0}^{T} \beta^t \left(\prod_{j=0}^{t-1} (1-\delta_j) \right) \{(1-\delta_t)u(c_t) + \delta_t v(b_t)\}. \tag{2}$$

This method of modeling the utility from the bequest matches the 'warm glow' specification of Andreoni [1989] with a CES parameter matching that for consumption rather than the dynastic altruistic formulation implied by concern with childrens' utility per se.[10] With respect to functional form, we follow De Nardi [2004] in parameterizing the bequest utility with two parameters, one to control the strength of the bequest motive (ϖ) and one to measure the degree to which bequests are a luxury good (ϕ). However, we redefine the place of these parameters in the bequest utility function so as to allow for a clear interpretation of their values. An agent leaving a bequest b receives direct utility:

$$v(b) = \frac{\varpi}{1-\gamma} \left((\phi - c^{SUB}) + \frac{b}{\varpi} \right)^{1-\gamma}. \tag{3}$$

If wealth is negative upon death, the agent is credited with having left a bequest of zero.

To understand the motivation for this choice of $v(b)$, consider a simple model in which an agent starts with wealth X dollars, lives for exactly n years and then dies. In each year of life, the agent consumes c dollars, deriving annual utility $u(c) = (c - c^{SUB})^{1-\gamma}/(1-\gamma)$. Upon death, the agent bequeathes the remaining $b = X - nc$, receiving the utility specified by equation (3). The agent's problem is to choose the bequest that maximizes total utility. The solution is to choose an annual consumption c^* such that bequest satisfies $b^* \equiv X - nc^* = \varpi(c^* - \phi)$. In other words, the agent leaves an inheritance to cover ϖ years of spending at an annual expenditure level $(c^* - \phi)$, the amount by which life time consumption exceeded the threshold ϕ. If X is insufficient to allow the agent to consume more than ϕ dollars each year, no bequest is left.

The parameter ϕ plays a role similar to one introduced by Henin and Weitzenbaum [2003] who use the form,

$$v(b) = \varpi_1 \left(\bar{\phi} + \frac{b - t(b)}{\varpi_2} \right)^{1-\rho}, \tag{4}$$

where $t(b)$ is the estate tax, which is absent in our model, and $\bar{\phi}$ is the expected annual consumption of the heir. Our ϕ parameter mirrors their $\bar{\phi}$, but we do not restrict ourselves to this interpretation of the parameter's value. Our choice to use the same parameter ϖ where they have two parameters, ϖ_1 and ϖ_2, is a simplification suggested by De Nardi [2004] and also motivated by the explanation in the preceding paragraph.

3.2 Wealth and Income

Households enter retirement with wealth $X_0 \geq 0$, and wealth at the beginning of time t is denoted X_t. We assume a deterministic stream of annual income y_t for as long as the retiree lives, and

[10] Kopczuk and Lupton [2005] provides reasons for researchers' preference for direct utility of bequest models over altruistic models, such as the finding by Altonji, Hayashi and Kotlikoff [1997] that parents do not offset inter-vivos transfers given an increase in their children's permanent income.

taxes are ignored. There is no income in the year of death. We assume that there is one composite riskless asset in which the household can invest and which yields a rate of interest r. Households are not allowed to take a negative position in assets (no-borrowing constraint).

3.3 Health Dynamics and Health Costs

Our treatment of health dynamics and death is crucial to the precautionary motive, given the high expenses associated with bad health. There are four health states modeled. State 1 is a state of good health. State 2 is a state in which there are medical problems but no need for long term care. State 3 is a state in which long term care of some form is required, and state 4 is death. In period 0, the individual is in health state $s_0 \in \{1, 2, 3\}$. The health state follows a Markov chain with age-varying one-period state transition matrix $\mathcal{P}(t)$. In each year, this is a 4×4 matrix. Retirees reaching age $M - 1$ die with probability one the following year.

Together the initial health state and the Markov transition matrices $\mathcal{P}(t)$ enable us to compute future probabilities attached to all health states, including death. Given the initial health state s_0, the transition matrix is applied repeatedly to derive the probability $\pi_t(s_t)$ that a retiree is in health state $s \in \{1, 2, 3, 4\}$ at time $t \geq 1$. This means that the death probability δ_t can be computed as $\delta_t = \pi_t(4)$.

We have not included the health state directly in the utility function. Rather, we focus on the costs associated with the various health states. Each live state $s \in \{1, 2, 3\}$ has associated with it a necessary and deterministic health cost, $h(s)$. Paying these costs entirely removes any utility penalty that would otherwise be associated with the health state. Death expenses in state 4 are also deterministic, at level $h(4)$, and are subtracted from the bequest.

3.4 Bankruptcy and Medicaid

Given the risk of substantial medical expenses which may exceed available wealth, there is need to include a bankruptcy mechanism. We model bankruptcy as affecting only a single period, in which the agent's consumption and end-of-period wealth are determined as described below. In the period following bankruptcy, the agent's income continues on its deterministic path and there are no further implications of having been previously bankrupt. Agents are forced to declare bankruptcy when they cannot afford to pay for medical costs and a subsistence level of consumption c^{SUB}, but they may choose to declare bankruptcy in any year. In practice, however, there is only a very narrow window in which the agent has sufficient assets to avoid bankruptcy but declaring bankruptcy is the optimal decision.

What happens in bankruptcy depends on the medical state. In states 1 and 2, an individual who declares bankruptcy is left with sufficient assets to consume at a minimum level $c^{BR} > c^{SUB}$, with end of period wealth remaining at zero. In state 3, the long-term-care state, treatment of bankruptcy is related to the institutional reality of Medicaid. An individual declaring bankruptcy in the long term care state forfeits all wealth to the government (end of period wealth is zero) and enters a Medicaid facility, receiving in that period the Medicaid level of consumption $c^{MED} > c^{SUB}$.

The Medicaid level of consumption is an important parameter in what follows, since its level reflects Medicaid aversion. As Pauly suggests, if the Medicaid consumption level is very close to subsistence, this will produce a strong incentive for households to retain sufficient wealth to retain the private care option. If it is closer to annual consumption in the pre-Medicaid period, then the incentive will be to run down wealth and use the Medicaid subsidy in place of savings. The value of c^{MED} therefore has powerful impact on the strength of the precautionary motive.

3.5 Specification of Optimization Problem

The timing of events is as follows:

- The household enters the period t with health state s_t and wealth state X_t. If $s_t = 4$ so that the individual is deceased, no income is received, health costs $h(4)$ are paid and the bequest b_t equals the remaining net resources, down to a minimum of zero,

$$b_t = \max[X_t - h(4), 0]. \tag{5}$$

  Otherwise, if $s_t < 4$, period t income of y_t is then accrued, and the health costs $h(s_t)$ are incurred.

- If $s_t < 4$, the consumption decision is made. The agent may choose any level of consumption c_t that exceeds the subsistence level c^{SUB} and satisfies the budget constraint,

$$X_t + y_t - h(s_t) - c_t > 0 \tag{6}$$

  Alternatively, the agent may declare bankruptcy. If no consumption level $c_t > c^{SUB}$ satisfies Equation 6, bankruptcy is the only option. If $s_t = 1$ or 2, bankruptcy means consuming $c_t = c^{BR}$. If $s_t = 3$, the agent must receive care under Medicaid and $c_t = c^{MED}$.

- At the end of the period, the agent is left with the unspent portion of assets, which earn a risk-free return r. If bankruptcy was declared in the period, wealth in the next period is zero. Letting I_t^{BR} be the indicator variable for bankruptcy in period t, the following period's wealth level obeys:

$$X_{t+1} = \begin{cases} (X_t + y_t - h(s_t) - c_t)r & \text{if } I_t^{BR} = 0; \\ 0 & \text{if } I_t^{BR} = 1. \end{cases} \tag{7}$$

- Finally, the new health state s_{t+1} is drawn according to the state transition probabilities $\mathcal{P}_{t+1}(s_{t+1}|s_t)$. If $t + 1 = T$, the final period, $s_{t+1} = S$.

The household maximizes expected utility of the remaining life time consumption (2) subject to the budget constraint (6) along with the bankruptcy option.

The Bellman equation is

$$V_t(s_t, X_t) = \begin{cases} max_{\mathcal{C}_t} \left\{ u(c_t) + \beta E_t V_{t+1}(s_{t+1}, X_{t+1}) \right\} & \text{if } s_t \neq S \\ v(b_t) & \text{if } s_t = S \end{cases} \tag{8}$$

subject to equations (3)-(7). The choice set is

$$\mathcal{C}_t = \left(c_t, I_t^{BR}\right) = \begin{array}{c} \left(c_t \in (c^{SUB}, X_t + y_t - h(s_t)], I_t^{BR} = 0\right) \\ \cup \left(c_t = c^{Bankrupt}(s_t), I_t^{BR} = 1\right) \end{array} \tag{9}$$

where $c^{Bankrupt}(s_t = 1, 2) = c^{BR}$ and $c^{Bankrupt}(s_t = 3) = c^{MED}$.

3.6 Parameter Calibration

The model has two very hard to gauge preference parameters, the Medicaid consumption (c^{MED}) parameter which impacts precautionary savings, and the bequest motive, (ϖ). The central issue of this paper is how to separate out the precautionary from the bequest aspects. In the next section, we argue that this is not possible by looking at moments of behavior alone. In these exercises we fix all preference parameters apart from the bequest and Medicaid aversion parameters at conventional values.

We start our retirees at age 62 in good health ($s_0 = 1$). Each period of the model represents one year and individuals die with probability one at age 100 ($T = 38$). All wealth, consumption, and income figures are in thousands of dollars. For wealth and income values, we approximate median values from the 2001 Survey of Consumer Finances (SCF) for households in the early years of retirement. We set $X_0 = 180$, since the median net-worth of households with heads ages 55-64 was \$181.5K. The median pre-tax income of retirees was \$22K per year, which we approximate in our model as a constant after-tax income of $y = 19$. We use a gross real risk-free asset return of $r = 1.03$ and a corresponding discount rate $\beta = 1/1.03$.

Payments under the government's Supplemental Security Income (Office of Beneficiary Determinations and Services [2005]) were \$579 per month in 2005, or approximately \$7K per year. For a subsistence level of consumption, we choose a value slightly below this, $c^{SUB} = 5$ or \$5K per year. For the level of consumption available under bankruptcy, we use a value slightly higher than the SSI figure: $c^{BR} = 8$. So as not to introduce an additional parameter value, we choose this same value for ϕ, which measures the consumption level above which one considers bequests.[11]

Standard values for the coefficient of relative risk aversion parameter in life-cycle models are between 2-6. Based on the life-cycle of risky asset positions, some research has argued that older investors are more risk averse (Morin and Suarez [1983]), but there is debate about their findings (Wang and Hanna [1997] and Bajtelsmit and Bernasek [2001]). We follow Brown and Finkelstein

[11] We are sensitive to the possibility that this value may be too low, especially under the interpretation of Henin and Weitzenblum [2003] that this parameter corresponds to the expected annual consumption of ones' heirs. We intend to explore the effects of larger values for ϕ on our findings.

Table 1: Calibration of Health Transition Probability Matrix

The first column shows the moment, the second column the target from the data, and the last column shows our calibrated value at the chosen parameters. The first 8 moments capture aspects related to long-term care (LTC); the data are from Brown and Finkelstein [2004] Table 1 for males. The next 4 moments relate to longevity; the data are from the National Center for Health Statistics, Vital Statistics (1999), Table 2 for males. The last 4 moments show features of the distribution of medical costs. These are not used in the calibration. Details of the calibration exercise are in the appendix. The small discrepancies between the simulation and the data in row 3 arises from the fact that our model is cast in years. The data on the other hand were compiled on a monthly basis. We interpret more than one year as at least two years, and that leads to an upward bias in the average.

	Moment	Data	Calibration
	Long-Term Care		
1	Probability ever use LTC	.40	.39
2	Average age of first use (among users)	80	78.14
3	Cond. Avg. years spent in care	2.90	3.45
4	Cond. Prob. use more than 1 year	.77	.73
5	Cond. Prob. use more than 3 year	.37	.37
6	Cond. Prob. use more than 5 year	.17	.18
7	Cond. Prob. ever exit to non-death state	.33	.36
8	Cond. Avg. number of spells	1.20	1.20
	Longevity		
9	Life expectancy at age 62	17.81	18.23
10	Life expectancy at age 75	9.82	9.81
11	Life expectancy at age 85	5.44	6.00
12	Life expectancy at age 95	2.99	3.21
	Total Medical Expenses during Retirement		
13	Avg. lifetime medical expenses ($k)		72.92
14	Median lifetime medical expenses ($k)		18.00
15	Prob lifetime medical expenses > $100k		.25
16	Prob lifetime medical expenses > $250k		.08

[2004] who cite a number of papers that rely on "a long line of simulation literature" that use $\gamma = 3$ as a baseline value. We also consider the effects of using other values.

The role of medical costs is central our analysis, especially the possibility of high costs associated with long-term care. The distribution of these costs in our model is controlled by the medical costs associated to each health state and by the one-period 4×4 state transition matrix $\mathcal{P}(t)$. This matrix is parameterized by twelve parameters, nine that determine the value of $\mathcal{P}(0)$ (of the sixteen elements, four are fixed by the death state being absorbing and there are three further restrictions so that each row sums to one) and three that control the flow of probability from greater health to poorer health as t increases. We select values for these parameters to match age-dependent mortality rates and many of the statistics on long-term-care utilization from Brown and Finkelstein [2004]. This calibration is described in detail in the appendix and the longevity and long-term care moments that are matched are listed in rows 1-12 in Table 1.

To model the medical costs associated to each health state, we identify the mean annual out-of-pocket medical costs for non-institutionalized seniors and calculate the annual costs of long term care for a senior insured under Medicare. The National Center for Health Statistics reports that in 2004, the average out-of-pocket medical expenses for non-institutionalized individuals over age 65 was \$600. Using our calibrated health-transition matrix, we find that among the periods our simulated retirees spend out of long term care (health states 1 and 2), they spend 10.5% in state 2 so that h(1)=0 and h(2)=6 reproduces this average. For the long-term-care state, we use Brown and Finkelstein's estimation that a semi-private room in a private LTC facility costs \$143 per day. In 2004, Medicare covered the full cost of LTC for 20 days each year and the daily costs in excess of \$109.50 for an additional 80 days. This leaves an annual out-of-pocket expense for a full year of LTC at \$46,700. We take h(3)=46. We ignore costs associated with death by setting $h(4) = 0$.

With these values, the median value for life-time medical expenses is \$18K, while the mean is \$73K (rows 13 and 14 of Table 1). Long-term care costs dominate the model, making up 85% of all medical expenses. For the 61% of individuals who do not enter long term care, the mean life-time cost is only \$7.9K. Agents face a 25% chance of facing life-time costs greater than \$100K and an 8% chance of costs greater than \$250K. Results of simulations not shown here suggests that behavior may be quite sensitive to the probabilities of realizing these unlikely but expensive scenarios. We note that the health transition probabilities are determined only by the agent's age and current health state, so that ones expected future costs do not depend on ones past medical history.

For the two remaining parameters, Medicaid consumption (c^{MED}) and the bequest motive, (ϖ), we consider a range of possible values. Medicaid consumption must lie above subsistence ($c^{SUB} = 5$) and we consider vales from 5.25 (highly Medicaid averse) to 9.5 (hardly Medicaid averse), above which there is little effect for increasingly larger values. For the bequest motive (ϖ), we consider values from zero up to 17.

To compute optimal strategies, we first discretize the state space and the control space. The model is then solved by backwards induction. At time T (age 100), the household dies with probability one. Its value function is the instantaneous utility over bequests. $V_T(\mathcal{S}_T) = v(b_T)$. In every prior period t, the Bellman equation (8) is used. We use linear interpolation to compute continuation values for points that are not on the grid. The choice variables ruled out by the budget constraint (6) are given large negative values.

4 The Inference Problem

It is clear that both bequest and precautionary motives will cause a household to decrease its consumption. Further, our model allows both of these motives to be present to varying degrees through the (continuous) parameters ϖ and c^{MED}. Simulations confirm that with no bequest motive ($\varpi = 0$), consumption decreases as c^{MED} does (as the precautionary motive increases). They also confirm that with small precautionary motives ($c^{MED} > 8$), consumption is smaller for larger values of ϖ (larger bequest motive). Yet, because consumption changes continuously with both parameters, we observe similar behavior for large bequest motives with small precautionary motives as we do for small bequest motives with strong precautionary motives, and for a range of values in between. It is therefore difficult to infer from observed behavior the relative importance of the two motives.

4.1 Likelihood Functions

To formalize and quantify the inference problem we encounter, we imagine observing the sequence of consumption decisions $\{c_t\}$ made by a single individual, for whom we know all preference parameters except ϖ and c^{MED}. For each value of some statistic (or set of statistics) of these consumption choices $f(\{c_t\})$, we construct a likelihood function for the parameters ϖ and c^{MED}: $\mathcal{L}\left(\varpi, c^{MED}; f(\{c_t\})\right)$.

Because the model cannot be solved analytically, this likelihood function is found through simulation. Fixing all other parameters of the model, we consider a grid of possible values for the pair (ϖ, c^{MED}): c^{MED} from 5.25 to 9.5 by increments of 0.25, and values of ϖ from 0 to 17 by increments of 1. This set of values represents the prior distribution, uniform over the range $(\varpi, c^{MED}) \in [0, 17] \times [5.25, 9.5]$.

We perform 2000 simulations with random health shocks for each pair of parameter values (ϖ, c^{MED}). We separate the resulting values of the statistic of interest $f(\{c_t\})$ into bins, which we refer to by their centers. To simplify our discussion of method, we consider a case in which $f(\{c_t\})$ is a singleton. This creates a three-dimensional histogram, with each bin identified by a unique value of the triplet $(c^{MED}, \varpi, f(\{c_t\}))$. The two-dimensional histogram defined by fixing the value of $f(\{c_t\})$ approximates the posterior likelihood function $\mathcal{L}\left(\varpi, c^{MED}; f(\{c_t\})\right)$ based on the uniform prior.[12]

4.2 Mean Consumption

As a first example we choose $f(\{c_t\})$ to be the mean of annual consumption over the first ten years, conditional on survival (i.e. dropping cases in which the agent does not survive ten years). We denote this mean by $\bar{c}$.

To fix ideas, Figure 1 plots the optimal consumption choice for all different parameter pairs

[12] To see this, consider that the above procedure is equivalent to the following exercise. We start with a population of 2000 x 18 x 18 people consisting of 2000 identical individuals of each type defined by a pair of values for bequest motive and Medicaid consumption (there are 18 x 18 such pairs: 18 different bequest motives and 18 different Medicaid consumptions). For each member of this population, we measure $f(\{c_t\})$. Consider the segment of the population for whom this value falls in a certain bin (e.g. the mean of consumption is between \$20.5K and \$21.5K; we refer to this bin by its center of \$21K). At random, select a single individual from this segment. Now ask: Given the value of $f(\{c_t\})$, what is the likelihood that this individual has a certain bequest motive and Medicaid consumption? The answer is that it is the distribution of these parameters within the segment of the population with this value of $f(\{c_t\})$. Equivalently, draw a person at random from the entire population and measure $f(\{c_t\})$. Given the measured value, you know that this individual is part of this segment of the population with that value of $f(\{c_t\})$. Since you don't know anything else, this individual has essentially been drawn randomly from this segment and the likelihood of a particular pair of parameters is the distribution of the parameters within that segment. This distribution is the two-dimensional histogram defined above.

$(\varpi, c^{MED}) \in [0,17]\times[5.25,9.5]$ of an agent at age 62, in good health and with wealth equal to \$180K. Annual consumption is higher for a household with a low bequest motive and low precautionary savings motive (low aversion to medicaid or high c^{MED}).

Figure 1: Optimal Consumption Policies

This figure plots the optimal consumption choice of a 62-year old agent in good health ($s_0 = 1$) with wealth X=180 and annual income y=19 for a range of bequest and Medicaid aversion parameters (ϖ, c^{MED}): c^{MED} from 5.25 to 9.5 by increments of 0.25, and ϖ from 0 to 17 by increments of 1. All other parameters are at their benchmark values.

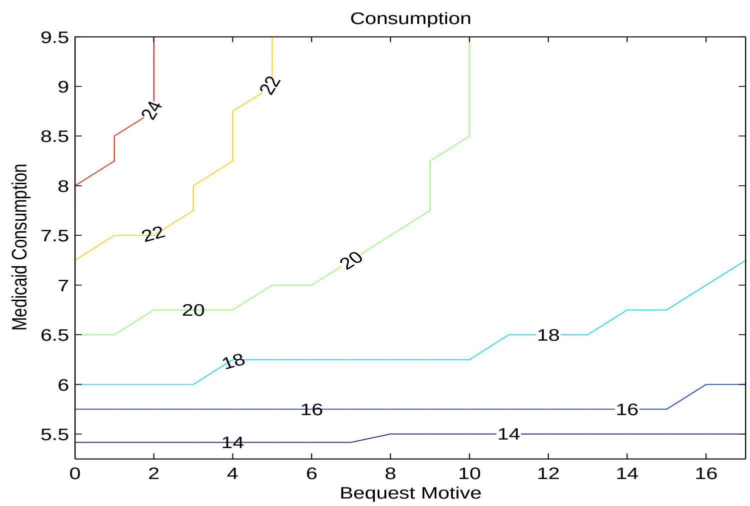

Since we are interested in the average annual consumption between ages 62-72, we then simulate the model for each parameter pair 2000 times, i.e. for 2000 different 10-year histories of health shock realizations. The means from the simulations are divided into bins of width 1, centered at integer values. To show the likelihood function $\mathcal{L}\left(\varpi, c^{MED}; \bar{c}\right)$ we draw a separate graph for each value of $\bar{c}$, with each graph showing contours of the likelihood function. In order of increasingly warmer colors, these contours enclose regions of probability 0.3, 0.4 … 0.9 and 0.95. Values in between the points on the grid are computed by linear interpolation. Especially where these regions intersect the edge of the plotting area, these numerical values are sensitive to our choice of a prior distribution for the parameters, in particular the range of values for ϖ and c^{MED} under consideration.[13]

These plots are shown in Figure 2. As an example of the discussion above, we can observe from the middle-right plot that a mean annual consumption of \$20K is equally likely to be observed in an agent with a strong bequest motive ($\varpi = 10$) with little precautionary motive ($c^{MED} > 8$) as it is in an agent with little bequest motive ($\varpi = 1$) but a strong dislike for Medicaid ($c^{MED} = 6.5$) leading to large precautionary savings. Clearly, this average annual consumption statistic does not allow us to discriminate between the two motives for wealth accumulation.

[13]Note that we are calculating likelihoods for two parameters based on a single statistic. We cannot therefore expect to discern unique maximum likelihoods for both parameters. Rather this example is useful for understanding the technique and the confounding of the two motives.

Figure 2: Likelihoods for Mean of Consumption

This figure plots likelihood contours for different parameter pairs (ϖ, c^{MED}): c^{MED} conditional on observing an average annual consumption level of \$17K (left-upper panel), \$18K (right-upper panel), ..., \$22K (right-bottom panel.

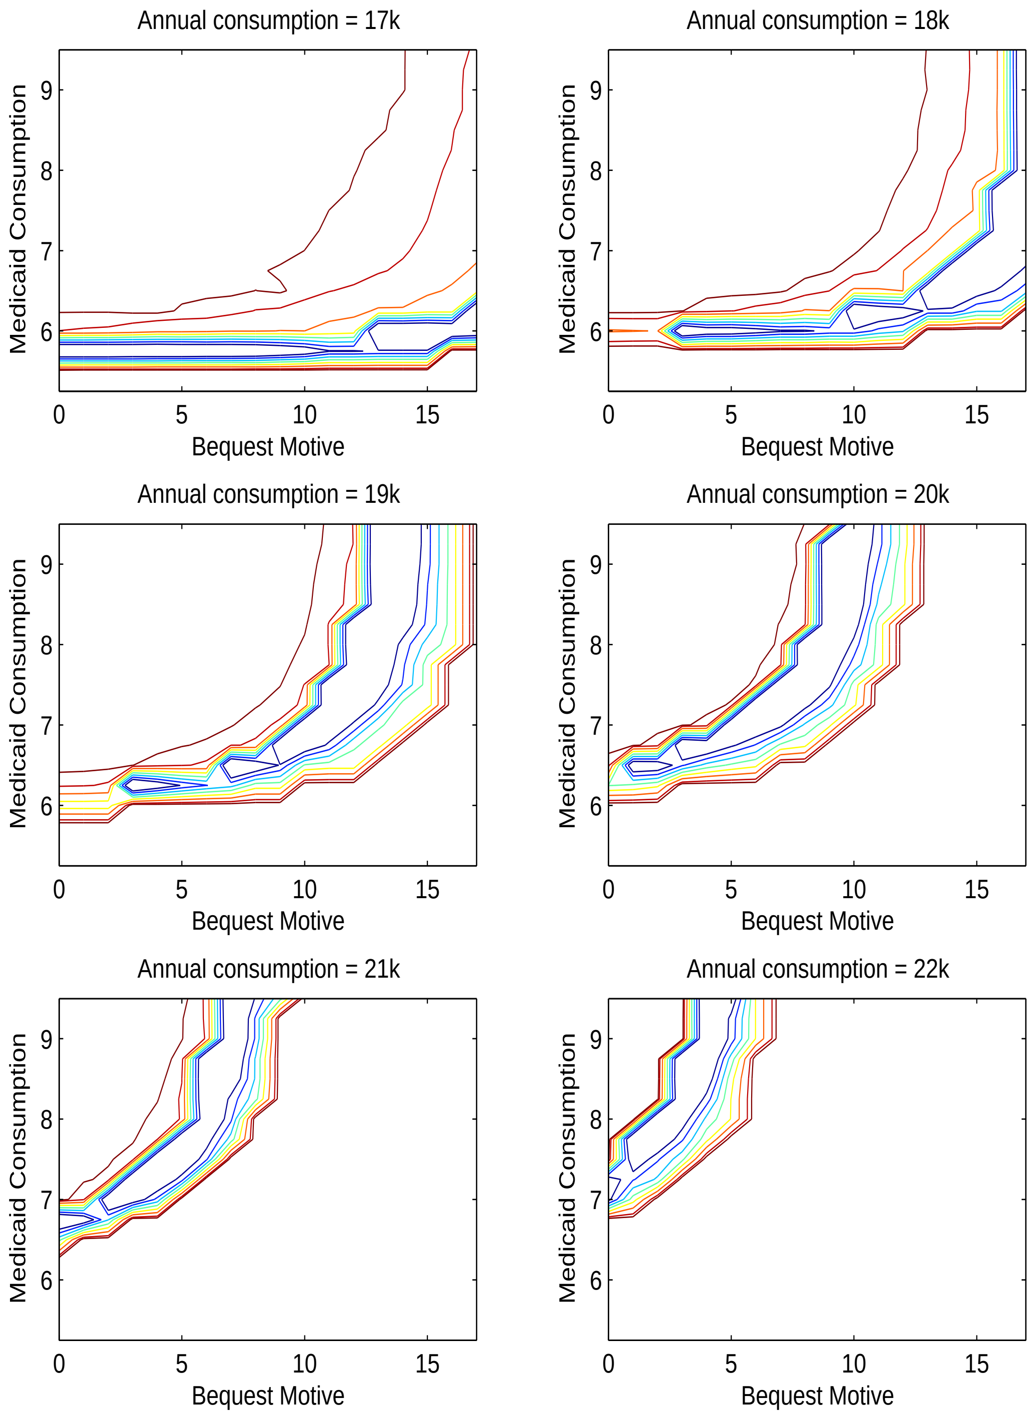

4.3 Annuities

The willingness to purchase an annuity contract is another statistic with the potential to discriminate between the two motives for wealth accumulation. An annuity allows the agent in the model to eliminate longevity risk by providing a guaranteed stream of consumption for as long as the agent lives. In theory, this might affect the trade-off between the precautionary savings motive stemming from long term care and the bequest motive. We show that in the model, it does not.

Figure 3 shows the willingness of a 62-year old agent in good health with \$19K in income and \$180K in wealth to pay for an annuity that promises \$5K annually for as long as the agent lives. It plots this willingness to pay for various types with different bequest (indexed by ϖ) and precautionay savings motives (indexed by c^{MED}). First off, the zero-load cost or 'fair value' of this annuity in the model is \$95.2K. None of the different types of agents are willing to pay this amount. Second, while different types have a different willingness to pay, a given willingness to pay cannot discriminate between the two motives for wealth accumulation. For example, an agent who is willing to pay \$74K for the policy could either have a strong precautionary savings motive (low c^{MED}) and a weak bequest motive or a strong bequest motive and a weak precautionary motive. Just as with mean consumption before, willingness to pay for an annuity is not revealing about motives. In fact, the information in Figure 3 duplicates the information in Figure 1. Further simulations indicate that only agents with less income and more wealth and with weak precautionary savings motives would be willing to pay the fair value of the annuity contract (e.g. $y = 14$ and $X = 360$). But again, we cannot distinguish between stronger bequest and very weak precautionary motives and very weak bequest and somewhat stronger precautionary motives.

4.4 Other Behavioral Statistics

In addition to the mean consumption, we consider other statistics of consumption including higher moments and change over time. In addition, we consider the actual bequest left in each simulation. Of these, some (such as the slope of consumption over time) are uninformative about the values of ϖ and c^{MED} i.e. the likelihood function is very flat. Others (such as bequests) contain information about the parameters that is duplicative when combined with mean consumption.

As an example, we consider adding information about the second moment of consumption. The standard deviation of consumption is measured over the first ten years (between ages 62 and 72), again conditional on survival, and sort the results into bins of width 0.3. In Figure 4, we plot the likelihood function conditional on the values of both mean and standard deviation: $\mathcal{L}(\varpi, C_{Med}; \bar{c}, \sigma_c)$. Mean consumption increases as we go from one row to the next; the standard deviation increases as we go across the columns of the figure. We see that the addition of this statistic does not add information about the relative strengths of the two motives. The contours of the joint likelihood function have the same shape as the contours of the likelihood functions in figure 2. Joint likelihoods using other observable data is similarly uninformative and we conclude that we are unable to distinguish the bequest and precautionary motives based on consumptive

Figure 3: Likelihoods for Willingness to Pay for Annuities.

This figure plots the willingness to pay for an annuity that promises to pay \$5,000 of a 62-year old agent in good health with wealth \$180K for a range of bequest and Medicaid aversion parameters (ϖ, c^{MED}): c^{MED} from 5.25 to 9.5 by increments of 0.25, and ϖ from 0 to 17 by increments of 1.

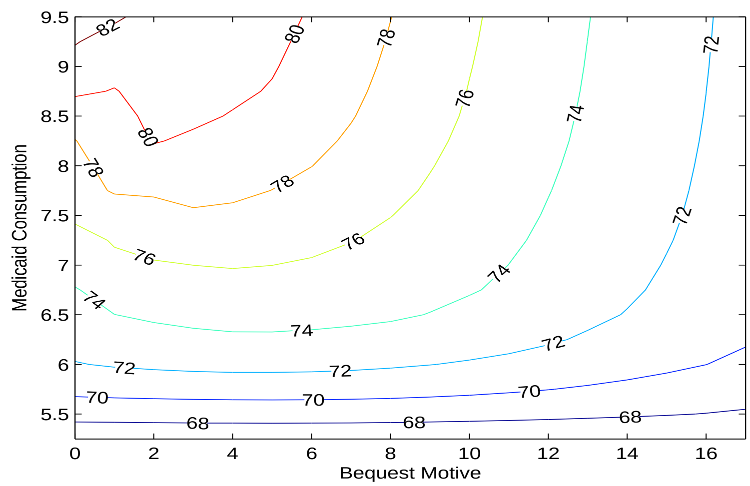

behavior alone.

5 Strategic Survey Questions

The results of the last section suggest it will be very difficult to infer bequest motives directly from consumption data, since precautionary motives produce consumption patterns similar to those we would expect from bequest motives. We propose an approach to identification that involves using survey instruments to fill in for missing evidence. A "strategic" survey question explores informative aspects of the agent's strategy that may not play out in the actual history that ensues. We derive additional statistics from these responses that combine with observed behavior to give rise to tight confidence intervals for the parameters of interest.

One feature of our approach is that we limit ourselves exclusively to questions concerning the consumer's strategy. We ask questions concerning which among a well-defined set of choices would be selected in a given contingency. In our view, seeking data on strategies is consistent with the spirit of revealed preference, according to which choice behavior is the fundamental source of behavioral insight. The strategic survey methodology is a minimum departure from standard observation of behavior, and as such represents the natural next step when purely behavioral data are insufficiently revealing.

Figure 4: Joint Likelihoods for Mean and Standard Deviation of Consumption

This figure plots likelihood contours for different parameter pairs (ϖ, c^{MED}): c^{MED} conditional on observing average annual consumption of \$17K-\$21K and a standard deviation of annual consumption of 0.6-0.9-1.2.

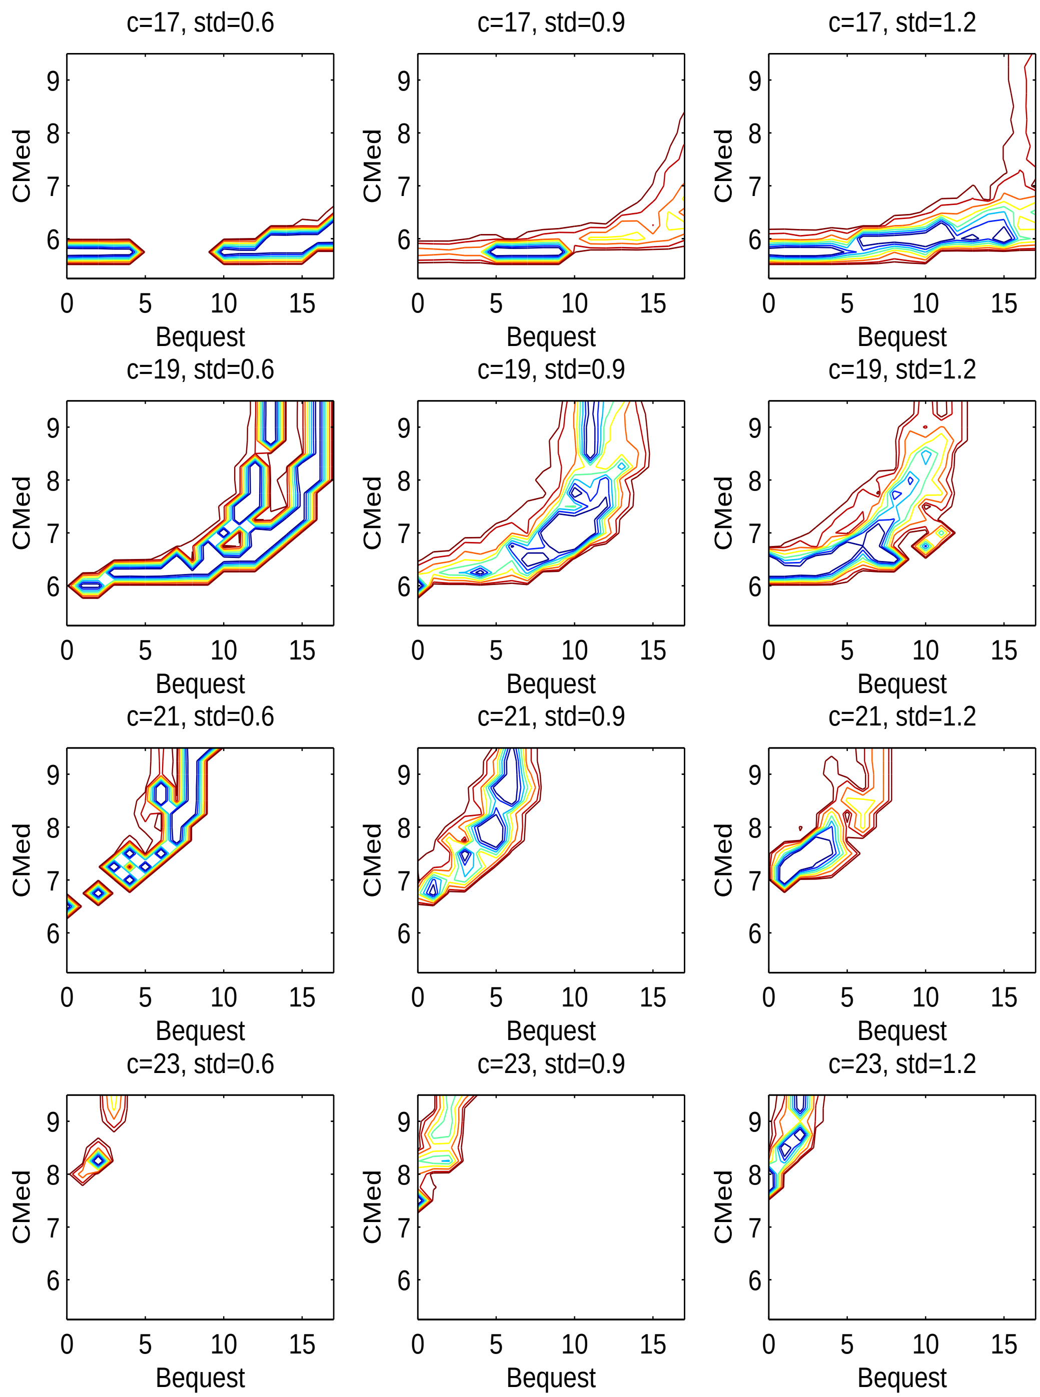

5.1 Immediate Shocks and the Continuation Strategy

There are many purely architectural necessities for a strategic survey question concerning a specific choice in a given hypothetical contingency. The first critical issue is to ensure that the contingency in question is explained to respondents with precision. The most obvious manner in which to pin down the underlying contingency in which a choice is to be made is to place all contingencies and choices in the present (immediately following survey completion). This is the approach that is taken in the papers by Barsky, Juster, Kimball, and Shapiro [1997], Kimball and Shapiro [2003], Kimball, Sahm, and Shapiro [2005] that initiated research into the information content of survey-based contingent questions.

The work of Kimball and Shapiro [2003] on the labor supply elasticity is particularly relevant. They argued that direct econometric evidence on the size of the substitution effect is muddied by several factors, such as the difficulty in finding temporary, exogenous movements in the real wage that identify movements in labor supply. To overcome this they developed the following survey instrument aimed at measuring the income effect.

> "Suppose you won a sweepstakes that will pay you [and your (husband/wife/partner)] an amount equal to your current family income every year for as long as you [or your (husband/wife/partner)] live. We'd like to know what effect the sweepstakes money would have on your life."

Follow up behavioral questions concerned whether or not the respondent would quit work entirely, and if not, whether or not they would reduce hours worked, and if so by how much: "Would you quit work entirely?" They placed the answers in the context of a model of labor supply that tightly connects income and substitution effects. They made the empirically-inspired conjecture that income and substitution elasticities are of approximately equal and offsetting magnitudes. With this, they used their survey instrument to measure the substitution effect, and to conclude that it is higher than is inferred in standard econometric procedures.

Asking questions of a given individual from the present perspective has the advantage that the contingency in which the choice is to be made is entirely familiar to the decision maker, with only the survey specified shock as a source of novelty. Yet there are two distinct problems associated with using the present situation as the base for the choice. First, individual differences are crucial. The answer depends on where the respondent is placed in terms of age, wealth, and other critical variables in the life cycle model. Second, the present choice can be largely undone by later decisions, and placing appropriate constraints on future behavior may place great strains on realism. Our model allows us to take account of the first factor, and we develop questions involving commitment to handle the second.

5.2 End of Life Choices

The limits on the information content of survey responses based on the present are shared almost in their entirety by standard choice data. As with choices based on hypothetical scenarios, the interpretation of actual choice is critically dependent on unobservable features of the future strategy about which strong assumptions must be made to draw inferences of value. We believe that it is all to the advantage of the strategic survey methodology that it opens up the possibility of placing decisions in the future and imposing natural constraints on the options for continuation.

This possibility of placing the decision maker in a hypothetical future contingency is particularly important in the arena of wealth accumulation in retirement, given that many aspects of the strategy that are of interest to us will play out only in the distant future. The questions we design play off the fact that salient contingencies take place near the end of life, at a decision node in which crucial decisions must be made that are influenced by the trade-off between bequests and precautionary motives. Focusing on decisions in these contingencies has the advantage that one can more tightly delimit future choices, and thereby provide an insight into this trade-off that is likely to be more robust. It also places individuals in a common situation, limiting the influence on the answers of individual differences.

Placing a decision in a well-defined but future contingency raises issues of its own. One opens up the issue of how remote is the proposed contingency in the decision maker's mind, and with it another possible source of survey error. At one extreme, the contingency may be entirely impossible, such as asking a respondent to consider a decision that they lies in their past, or that reflects health, wealth, and demographic contingencies that are impossible. There are many unknowns concerning exactly how well able various respondents will be to remove the details of their current lives and move into the requisite hypothetical state of mind. Clearly, a deeper understanding of response error will be invaluable in the design of strategic surveys.

Another important issue concerns how complete is the description of the contingency in which respondents are placed. For example, does it include a specification of their beliefs about the future at this point? Does it specify how they ended up at that decision node? Additional specificity adds clarity to the named contingency, yet may also reduce realism, since it opens up the issue of whether the asserted path is plausible from the respondent's viewpoint. In the current survey these issues are decided on a case-by-case basis according to an as-yet untested set of intuitions concerning the trade-off between specificity and lack of realism.

5.3 Application

In this section we lay out two sets of questions pertaining to the trade-off between bequest and precautionary motives. The entire survey is included in an Appendix to the paper. Note that there may appear to be a certain amount of redundancy in the questions we present. We do not regard this as problematic, since we show in the next section that survey error ensures that we will benefit from posing a wide variety of questions even if they are all focused on the same two parameters.

As background to the questions posed, it is important to note that the survey contains a series of questions that allow us to characterize the expected consumption profile in the ensuing period of five years. With this information gathered, we will be in a position to ask about how presently planned expenditures compare with those that would be undertaken in a scenario that neutralizes the need to set aside resources to avoid Medicaid.[14]

5.3.1 Immediate Shocks

One set of questions that are of clear interest are those that present the respondent with an immediate shock that neutralizes the incentive to retain wealth either for bequest or for precautionary reasons. The questions are preceded by the following statement: "The questions below concern how you would expect your household's average annual living expenses excluding housing, debt payments and health care (i.e. spending on food, groceries, travel and entertainment, etc.) to change if you were immediately to win a tax-free sweepstakes providing you with additional cash available for use only for restricted purposes. The various parts of the question differ only in the amount and the precise restrictions placed on the use of these winnings. The point of departure in answering these questions is the $[\cdot]$ that you estimated in question $[\cdot]$ for your household's annual living expenses over the next five years. Based on this, the questions all concern the extent to which winning this sweepstakes would result in changes to your average annual spending on living expenses over this period."

The first survey question in this block of questions is:

> "[SQ_1] Suppose that the prize that you won immediately after completing this survey was **unlimited in dollar amount** yet was tightly and accurately restricted to be used solely for purposes of **future private long term care expenses** for you and/or your partner. This money cannot be accessed except for private long term care costs, and will not become available for bequests. The one and only use to which the prize can be put is to pay for long term care. How much would you expect to change your household's annual living expenses in this scenario? "

In this scenario, there are no long term care costs for which to save so that the resulting path of saving should have nothing to do with the Medicaid aversion parameter. To solve this problem within our model, we set private long term care costs to zero, solve for the optimal consumption policies and calculate mean consumption as before. Figure 5, which shows the 'correct answer' to this question inside the model, confirms that the behavior in this scenario depends only on the value of ϖ and not at all on c^{MED}. The higher the agent's bequest motive, the lower her consumption would be. [15]

[14] Note that in practice any field survey will have some respondents with partners. However in fitting with the model we ignore that fact in the questions presented in this section.

[15] We ask questions on consumption changes to the respondents but solve for consumption levels in the model. Going from changes to levels in the model is trivial. Based on model simulations for a wide range of parameters, we found the appropriate range of answers for this question and the following two questions to be 'leave spending unchanged', 'increase spending by $1-$1,000 per year', 'increase spending by $1,001-$2,000 per year', 'increase spending by $2,001-$3,000 per year', 'increase spending by $3,001-$4,000 per year', 'increase spending by $4,001-$5,000 per year', 'increase spending by $5,001-$7,500 per year', 'increase spending by $7,501-$10,000 per year', 'increase spending by $10,001-$12,500 per year', and 'increase spending by more than $12,501 per year'. Our simulations indicate that the upper tail of the answers may only be relevant for older agents who are closer to long term care.

Figure 5: Optimal Consumption Policies in the Absence of Long Term Care Risk (SQ_1).

This figure plots the optimal consumption choice of a 62-year old agent in good health with wealth $180K for a range of bequest and Medicaid aversion parameters (ϖ, c^{MED}): c^{MED} from 5.25 to 8.25 by increments of 0.25, and ϖ from 0 to 15 by increments of 1. The cost of long term care is set to zero. All other parameters are at their benchmark values.

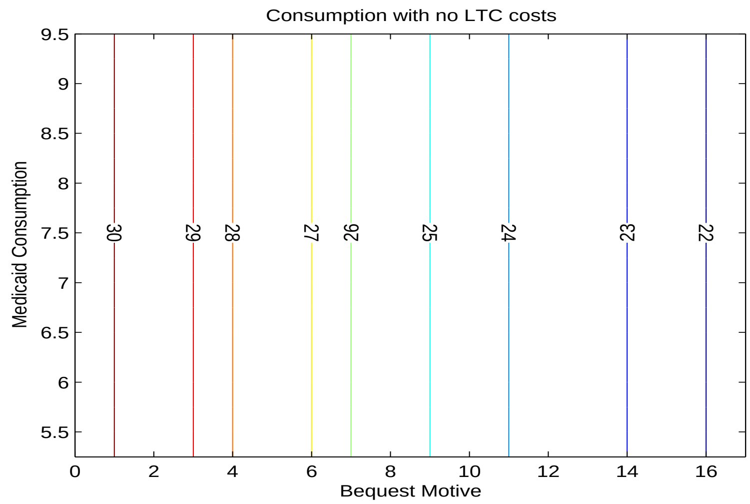

We can now compute the likelihood of a given parameter pair (ϖ, c^{MED}) conditional on a given answer to this survey question SQ_1: $\mathcal{L}\left(\varpi, c^{MED}; SQ_1\right)$. For now, we assume that these responses contain no error. In the next section we will revisit this question to show how measurement error impacts the inference. Figure 6 plots the contours of the likelihood function for each of four possible answers: $SQ_1 \in \{\$23K, \$25K, \$27K, \$29K\}$. Again, they are vertical, depending only on the bequest parameter.

One potentially attractive concept is to enrich this scenario by adding to the prize full high quality *medical* insurance. To add this form of insurance would plausibly change life expectancy due to the increased use of high quality medical resources, so that the impact on savings would depend on a more intricate set of model parameters. Such a criticism is far harder to level against a question that revolves around end of life care.

Another attractive scenario involves removing bequest motives by offering a prize that provides fully for these motives. However, the nature of these needs is hard to specify, since there are essentially limitless desires. There is also the intricacy that by providing for one's children's needs, one may in fact liberate them to pass more consumption back in an implicit deal. The following question attempts to take account of these issues in two ways: by providing a cash limit of some form on the prize and by attempting to rule out feedback effects of the higher bequest on own consumption.

> "[SQ_2] Suppose instead that the prize comprised**$135,000 tax free dollars** and was tightly and accurately restricted to be used for **bequests upon your death**. The prize

Figure 6: Likelihoods Given Survey Response to no-LTC Risk Question (SQ_1)

This figure plots likelihood contours for different parameter pairs (ϖ, c^{MED}): c^{MED} conditional on observing an answer to the no-LTC risk survey question stated above. This answer is expressed as an average annual consumption level, and we plot the contours for four possible answers: \$23K (left panel) - \$29K (right panel. The survey response is assumed to contain no error.

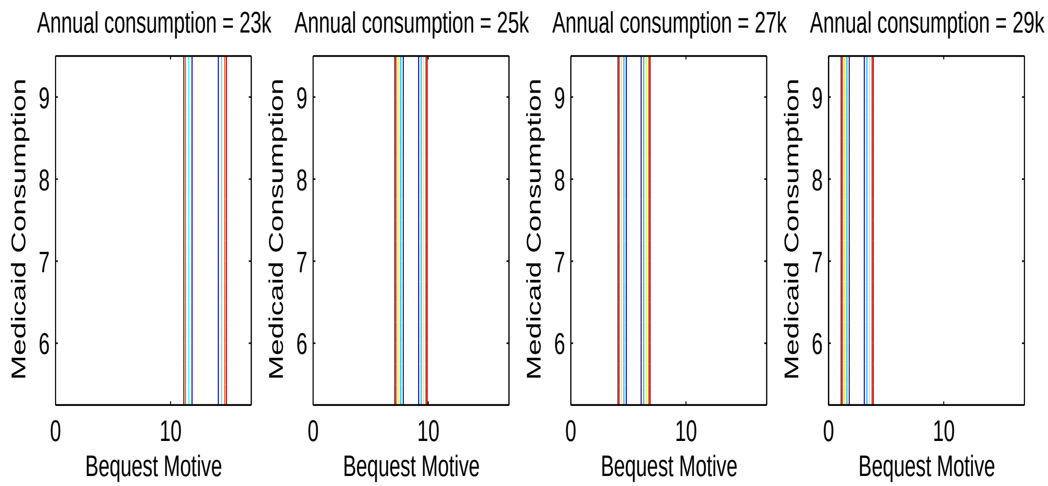

cannot be accessed either by you or by your beneficiaries until you have passed, and until that time it yields a rate of return equal to the inflation rate, so that the money will keep its current value when it becomes available. Note that while you can inform beneficiaries immediately of having won the prize, there is a prohibition against their passing this additional money back to you in the form of increased transfer payments. In other words, you are to assume that the gift has absolutely no impact on the resources that are available to you while alive, passing on in whole to your beneficiaries when you have passed. How much would you change your annual living expenses in this scenario?"

Question SQ_2 is intended to form a partial converse to the long term care question, neutralizing in large part bequest motives. Note that the question involves a prize level of \$135k, which was chosen on the basis of model simulations for a large variety of income and wealth levels. Figure 7 illustrates the 'correct' answer to this question in the model for various types characterized by (ϖ, c^{MED}). Comparing to figure 1, the agent of type (7,7) consumes now \$22K instead of \$20K, because his/her bequest motive is largely satisfed by the locked box. The lines in Figure 7 are more horizontal than in Figure 1, reflecting the fact the prize has partially neutralized the bequest motive.

A follow-up to the above question can be used to gain insight into the relative impact on consumption of the bequest and Medicaid aversion motives.

"[SQ_3] Now suppose that this same **\$135,000 tax free dollars** won in the sweepstakes was dedicated tightly and accurately to **private long term care costs** for you and/or

Figure 7: Optimal Consumption Policies with a \$135K Lock Box for Bequests (SQ_2)

This figure plots the optimal consumption choice of a 62-year old agent in good health with wealth \$180K for a range of bequest and Medicaid aversion parameters (ϖ, c^{MED}): c^{MED} from 5.25 to 9.5 by increments of 0.25, and ϖ from 0 to 17 by increments of 1. This agent has access to a \$135K bequest upon death. All other parameters are at their benchmark values.

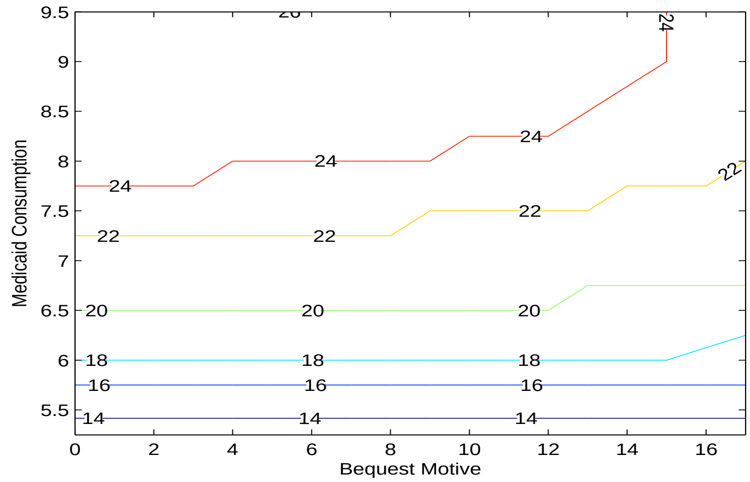

your partner rather than for the bequest. You should assume that this money will be sufficient to pay for **three years** of long-term care costs in the facility of your choice. As before, the account is a locked box that cannot be accessed except for private long term care costs, and will not become available for bequests. Again as before, in the periods before it is needed, the capital grows at the rate of inflation, so that the money will keep its current value when it becomes available. How much would you change your annual living expenses in this scenario?"

We decided to tie the prize to a fixed number of years of care to avoid differences in the perceived amount of long term care that \$135K would cover. In the model these are tightly linked (1 year is assumed to cost \$45K), and we want to be able to interpret the answers under the same restriction. Figure 8 plots the 'correct' answer to question SQ_3 in the model. The diagonal lines of figure 1 have become more vertical, closer to those in figure 5 where all LTC risk was eliminated. For each level of Medicaid aversion c^{MED}, the agent consumes more.

In addition to looking to understand the bequest and Medicaid aversion motives separately, we are interested in questions that allow us to compare the relative power of these two motives. We follow up further on questions concerning the use of the reward as follows.

"[SQ_4] Finally, suppose that all \$135,000 tax free dollars won in the sweepstakes could be used for **either** long term care or bequests upon your passing. One additional requirement is that you must **immediately decide how much to dedicate to each**

Figure 8: Optimal Consumption Policies with a \$135K Lock Box for Private Long Term Care (SQ_3)

This figure plots the optimal consumption choice of a 62-year old agent in good health with wealth \$180K for a range of bequest and Medicaid aversion parameters (ϖ, c^{MED}): c^{MED} from 5.25 to 9.5 by increments of 0.25, and ϖ from 0 to 17 by increments of 1. This agent has access to a \$135K provision to cover long-term care expenses. All other parameters are at their benchmark values.

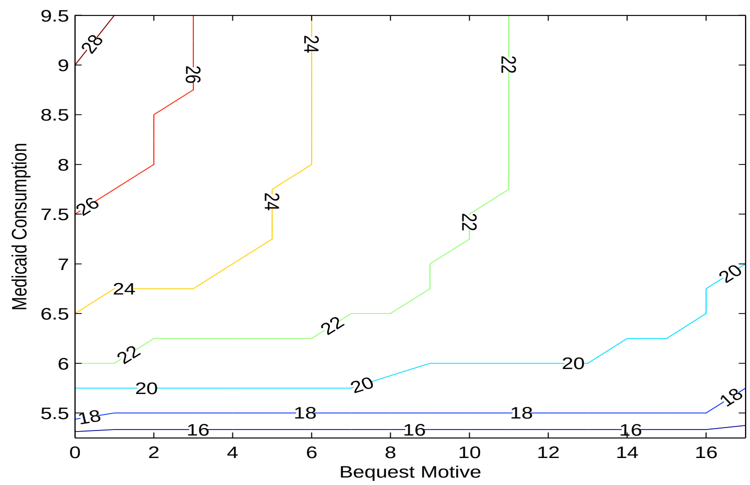

of long term care and the bequest. How would you divide the \$135,000 prize between these two locked boxes, the long term care box and the bequest box? ”

Figure 9 shows how a typical agent in the model would answer question SQ_4 with a prize of \$100K. An agent with a strong bequest motive and a weak precautionary savings motive (high c^{MED}) is in the upper right-hand side corner, and would put only \$20K into long term care. The person with opposite motives is the the lower left-hand side corner, and would use all of the \$100K towards long-term care.

5.3.2 End of Life Choice

Our imposition of commitment devices (‘the locked box’) is intended to increase the information content of current decisions. Yet the consumers in our model and in the real world are still facing many future uncertainties, and these have a significant impact on the choices they make. For this reason it is worthwhile also to pose a question that places the decision maker close to the end of life, at the moment at which the decision on long term care must be taken. The question below refers to such a narrow section of the model that the ‘correct’ answer can be computed analytically. It is not sensitive, for example, to uncertainty in medical costs or longevity. In essence, this is the end of dynamic programming problem with few problematic contingencies left to deal with.

“[SQ_{5a}] We move now to an entirely different scenario. Consider a situation in which you have arrived at age 85, you are the sole surviving partner (if applicable), and are

Figure 9: Choice Between Lock Boxes (SQ_4)

This figure plots the optimal choice of a 62-year old agent in good health with wealth \$180K for a range of bequest and Medicaid aversion parameters (ϖ, c^{MED}): c^{MED} from 5.25 to 8.25 by increments of 0.25, and ϖ from 0 to 15 by increments of 1. This agent has access to a \$100K lock box which she can split as stated in the question above. All other parameters are at their benchmark values.

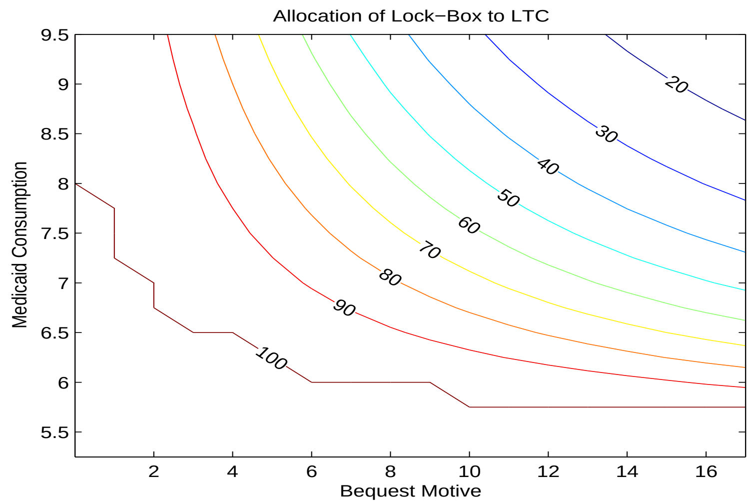

in need of long term care. You know that you have exactly one year left to live and will need to sell the home (should you own one at this stage) and enter a long term care facility for that year. Assume there has been no inflation, and that you have total net worth of \$100,000 (this includes the value of the home where relevant). Finally, suppose that you have the following two options:

(A) Place all \$100,000 in a tax-free trust so that it will all be in your bequest, and go on Medicaid, receiving the services of a Medicaid facility with no income to pay for any other consumption.

(B) Use your money to pay for private long term care and stay off Medicaid. For that year, you have pay \$45,000 for your private room and for the medical care in a private facility. In this option, you would use your income to pay for your other living expenses (e.g. food, entertainment etc.), leaving \$55,000 for the bequest.

Which of these two options would you choose?"

Option A: go to Medicaid facility and leave \$100,000 bequest to your beneficiaries

Option B: go to private facility and leave \$55,000 bequest to your beneficiaries

The amount of \$100K is chosen based on model-simulations so that it is maximally informative for all agents within a broad set of income and wealth levels. This question is then followed up by a question that determines the cost of a private facility that makes the respondent indifferent between the two alternatives.

> "[SQ_{5b}] In the same scenario, what part of a possible bequest of \$100,000 would you choose to forgo to avoid long term care in a Medicaid facility?"

The answers take the form: 'Nothing, I would choose to go to a Medicaid facility and leave a bequest of \$100,000', 'I would forgo 10% of the bequest and leave a bequest of \$90,000,' ..., ' I would forgo the entire bequest to avoid the Medicaid facility.'[16]

Figure 10 plots the 'correct' answer to the question SQ_{5b} about the willingness to pay for a private facility for agents in the model with different parameter pairs (ϖ, c^{MED}). This willingness to pay decreases as the bequest motive gets stronger as consumption in the Medicaid facility increases (i.e. as Medicaid is seen as a better alternative).

Figure 10: Willingness to Pay for Private Long Term Care Facility (SQ_{5b})

This figure plots the answer of a 62-year old agent in good health with wealth \$100K for a range of bequest and Medicaid aversion parameters (ϖ, c^{MED}) to the above question about the willingness to pay for 1 year of long-term care in a private facility. All other parameters are at their benchmark values.

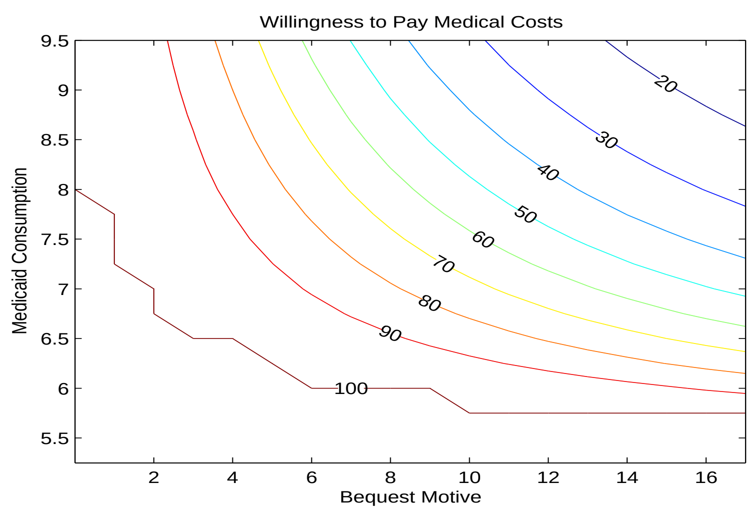

5.4 Combining Choice Data and Strategic Survey Data

Section 4 argued that behavioral data do not identify the two parameters of interest because they are determined by the combination of the two motives and do not speak to their relative

[16]Note that within our model, this is not the final decision made, even if it is believed that the long term care is terminal. Specifically, while entry into the Medicaid facility pins down the level of consumption spending in our model (c^{MED}), the same is not true for private long term care, which is treated as leaving the subject free to choose the level of consumption. This survey question SQ_{5a} pins down consumption in the private facility to remove its possible impact on the answer.

strength. We do not propose to abandon behavioral data, but rather, we hope to supplement them with survey questions that are specifically designed to identify the relative strength of the two motives. Thus, the most informative survey questions are those that provide information that is complementary to the behavioral data. Even though strategic survey questions may be measured with larger error than behavioral data, they are useful to separate out motives. In our approach, combining behavioral and survey data entails constructing likelihood functions conditional on the values of *both* observations. Hence in what follows we are asking joint likelihoods to be maximally informative. As an example, we combine survey question 5b with mean consumption data and show that the likelihood contours are concentric circles. This proves that the method is able to identify a narrow range for each parameter.

To set the stage for this exercise, we first plot the contours of the likelihood function over (ϖ, c^{MED}), conditional on the answer to the above question: $\mathcal{L}\left(\varpi, c^{MED}; SQ_{5b}\right)$. Each panel of Figure 11 represents a different answer to question SQ_{5b}. We assume in this plot that there is some response error around that answer, which we discuss in more detail in the next section. What is important for now is that the contour plots are downward sloping. Based on this question alone, a respondent who is willing to pay \$60K to avoid the Medicaid facility is equally likely to have a high bequest motive and a strong precautionary savings motive $(\varpi, c^{MED}) = (15, 6.5)$ as a weak bequest motive and a weak precautionary savings motive $(\varpi, c^{MED}) = (7, 9)$.

Second, Figure 12 combines the behavioral data, more precisely the average annual consumption of section 4.2, with the strategic survey question of section 5.3.2 (SQ_{5b}). It plots the contour plots of the likelihood function conditional on both observations: $\mathcal{L}\left(\varpi, c^{MED}; \bar{c}, SQ_{5b}\right)$. Graphically, it overlays figures 2 and 11. Because the former has upward sloping contours and the latter has downward sloping contours, the likelihood conditional on both observations has nice concentric contours. For example, an agent who consumes \$20K per year and answers that he/she would be willing to pay \$80K to avoid a Medicaid facility has a bequest motive around $\varpi = 7$ and a precautionary savings motive around $c^{MED} = 7$. The posterior likelihood reaches a peak at this parameter pair. The concentric contours arise because (1) uncertainty over health shock realizations generates a distribution over annual average consumption, and (2) survey error generates a distribution over the 'correct' survey answer. This strategic survey question was specifically designed to complement the consumption data, and so the degree to which these plots identify the parameters is evidence of the usefulness of our approach.

6 Response Error and Inference

Expecting survey answers to be error-free is unrealistic. In practice, there will be inconsistencies not only between the answers to the various questions, but also as between these answers and the behavior itself. A crucial issue is how to model the various sources of error. Section 6.1 outlines the basic model of survey response error that will be used in extracting information from the data to narrow the set of likely parameter values. One simplification is that we assume that error is

Figure 11: Likelihoods Given Survey Response to Willingness to Pay for Private Long Term Care Facility (SQ_{5b})

This figure plots likelihood contours for different parameter pairs (ϖ, c^{MED}): c^{MED} conditional on observing an answer to the willingness to pay to avoid Medicaid survey question stated above. This answer is expressed as an annual dollar amount, and we plot the contours for six possible answers: \$0K (left-top panel) - \$100K (right-bottom panel). The survey response is assumed to contain a \$5K response error.

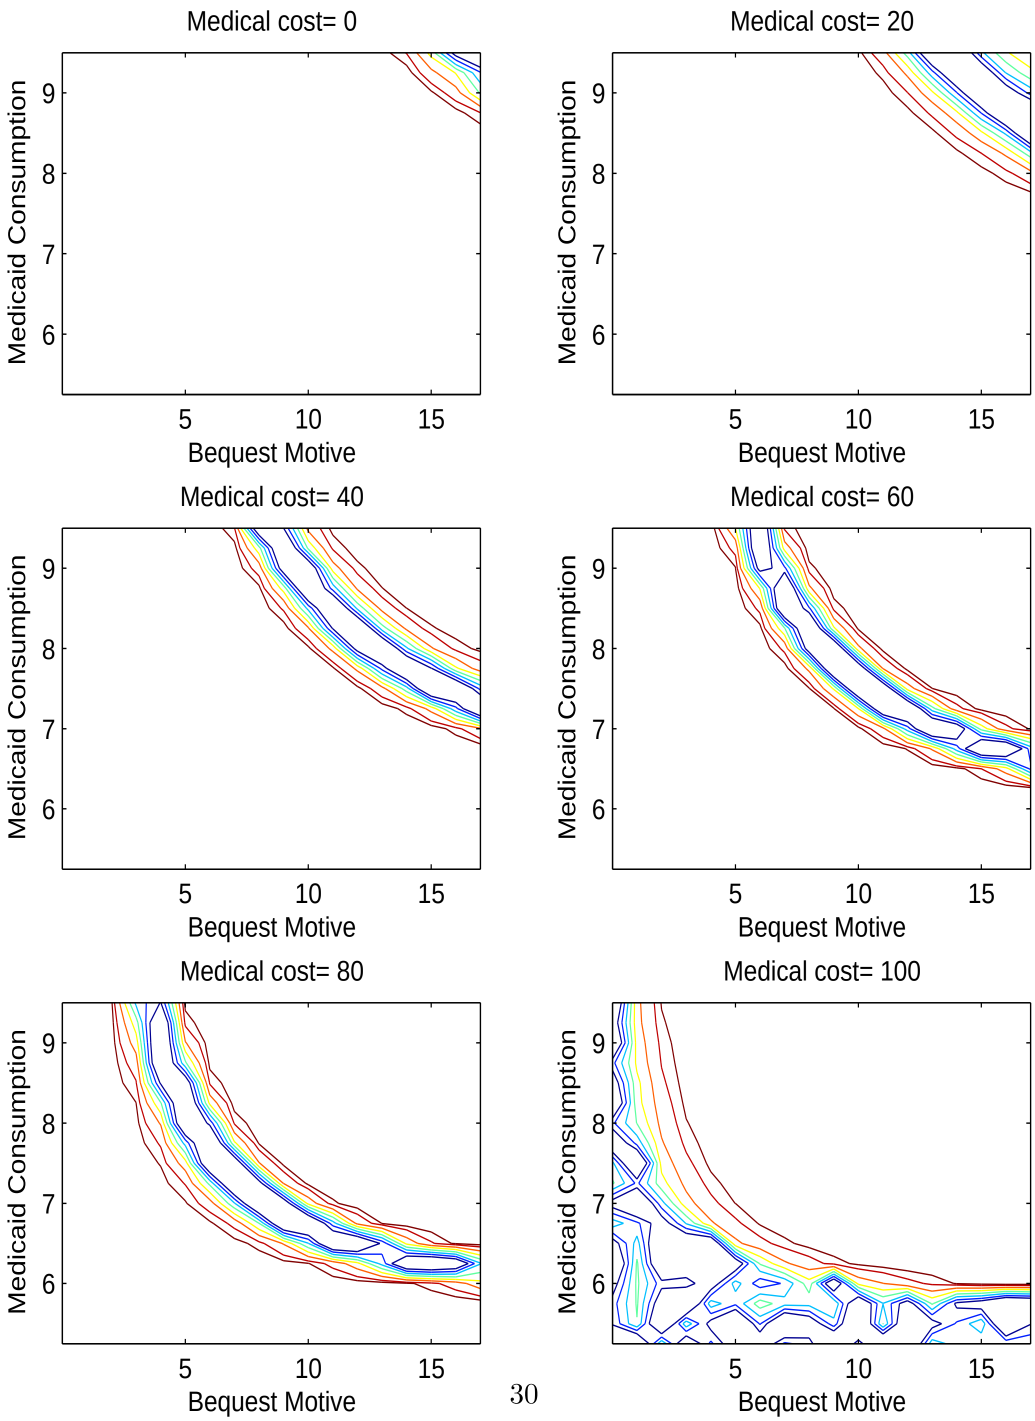

Figure 12: Joint Likelihoods Given Actual Annual Consumption and Survey Response to Willingness to Pay for Private Long Term Care Facility

This figure plots likelihood contours for different parameter pairs (ϖ, c^{MED}): c^{MED} conditional on observing both the average annual consumption level and an answer to the willingness to pay to avoid Medicaid survey question stated above. This answer is expressed as an annual dollar amount, and we plot the contours for six possible answers: \$0K (left-top panel) - \$100K (right-bottom panel). The survey response is assumed to contain a \$5K response error.

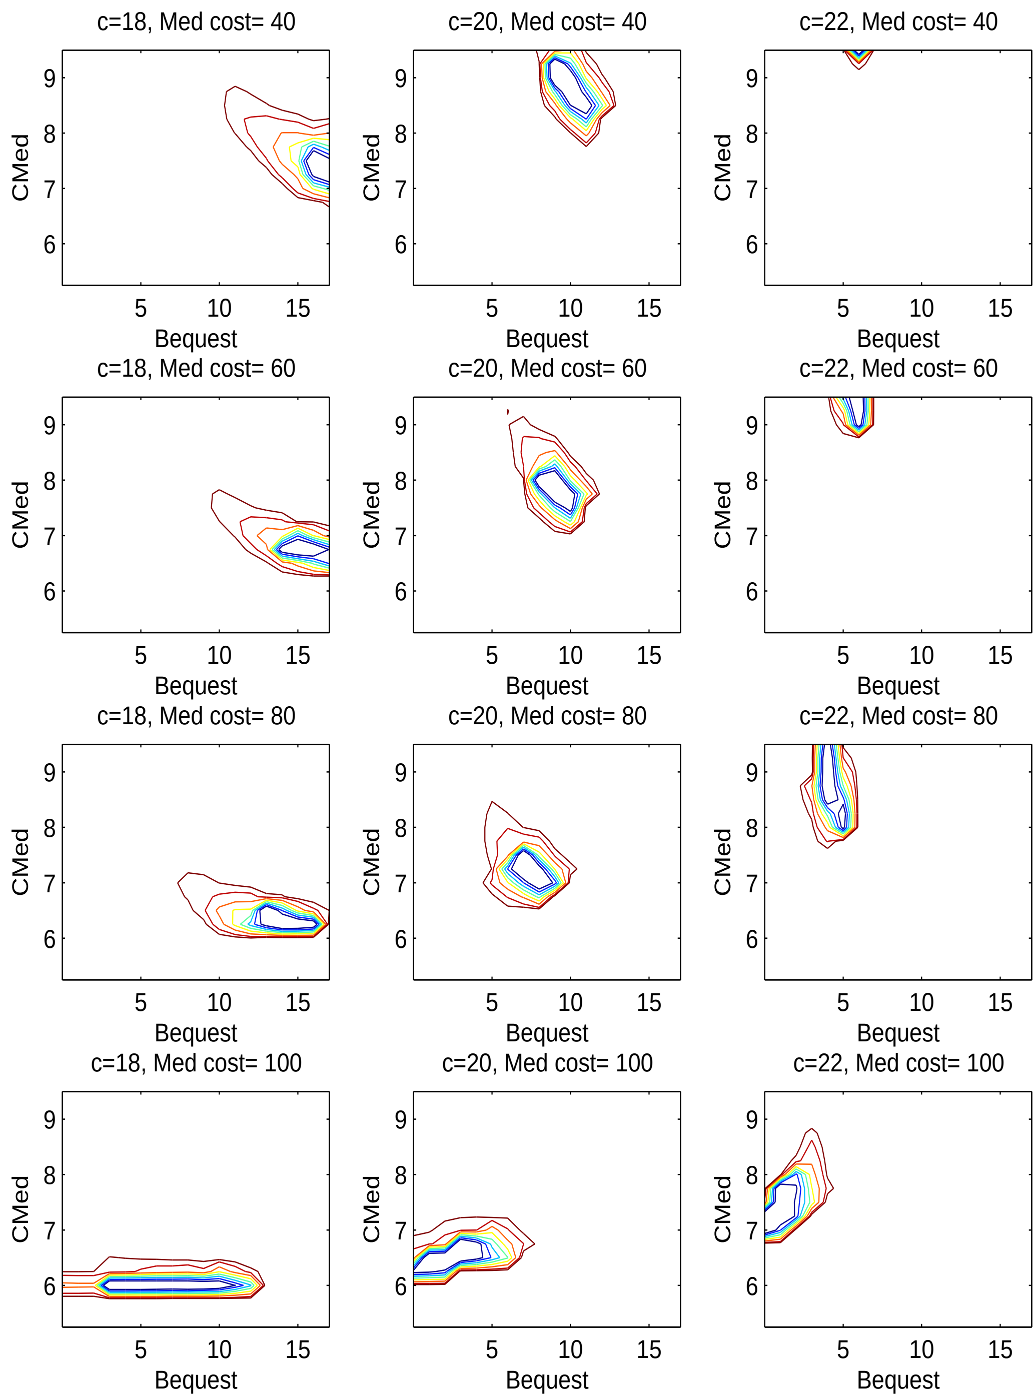

present only in the responses to the strategic survey questions, while data on actual and expected annual spending is error free.[17] We assume in addition that answers to all numerical questions are subject to mean zero normally distributed errors, and that the structure of the error term is common across individuals. Section 6.2 shows how the structure of the variance-covariance matrix impacts the interpretation of given survey responses. Section 6.3 presents an estimation procedure in conceptual form.

It must be understood that the error analysis we outline below is only the first word on a subject that we see as critical to the future advances in survey methodology. In particular much future effort will have to be devoted to developing and estimating models of survey error that are respectful of the psychology of survey response. We close the section with a few comments on specific aspects of this broad agenda that are already in our sights (section 6.4).

6.1 The Multivariate Normal Model

Given that expenditure and intended spending data are assumed to be error free, we will be able to use data on these behaviors to pin down one particular "iso-expenditure" contour in the parameter space, as in Figure 1 above. This will leave only the position on this contour line to be inferred from the strategic survey data. The exact nature of this inference depends on the precise assumptions that we make concerning the nature of the survey response error not only at the level of the individual response to a given question, but also across individuals and across questions.

Formally, we define the actual answer given by individual i on an open numerical question k to be impacted by survey error:

$$\widehat{SQ^i}_k = SQ^i_k + \varepsilon^i_k,$$

where the left-hand side is the answer that is given, and the right-hand side is the truth as defined in our model (SQ^i_k) plus an additive survey error ε^i_k.[18] The fundamental question is how to model the error terms ε^i_k. In this paper, we assume that there is no systematic bias in response to any question, so that the mean of ε^i_k is zero for all individuals i and questions k. We assume also that the distribution of the vector of error terms for respondent i is multivariate normal, and does not depend either in any way on the respondent's characteristics:

$$\varepsilon^i \sim N(0, \Sigma).$$

With respect to the correlation in errors across the various questions, we assume only that these are characterized by a fixed six-by-six variance-covariance matrix Σ.[19]

[17] The ideal measure of behavior is the actual path of spending. However it will not always be possible to follow actual behavior over time. In our survey we ask instead about income and expenditure in the past year, and also the expected level five years from now.

[18] Formulating the model in this way exposes already some aspects of the estimation strategy. First, we are assuming that numerical response error is the appropriate metric for measuring error. Second, we are assuming that these numerical errors are additive. Alternative formulations are the subject of future research.

[19] Note that it is in principle easy to allow for cross sectional differences in the error term. It is also easy in principle to add cross sectional connections between individuals beyond those in the error term. An extreme version of this

A Precedent The challenge in estimation is to identify both the parameters of the error distribution and the posterior over each individual's preference parameters. The inference process used by Kimball, Sahm, and Shapiro [2005] serves as an informative precedent in this regard. They analyze survey error in a case with one survey question with mean zero normal error that is assumed identical across individuals. They do this in the context of questions that were posed on the HRS to measure risk aversion. They assume that the response error of individual i, ε^i, is a purely random measurement error that is normally distributed with mean zero and unknown standard deviation σ, again assumed common across individuals. They exploit the fact that many respondents answered the same question on risk aversion on two different waves of the HRS and observe that a certain proportion of the answers changed between waves. Each individual's preference parameter is then calculated by averaging that individual's responses in both waves. Their estimation procedure for σ rests on the assumption that the underlying risk aversion parameter did not change between waves, so that response switches provide information on the variance of the response error. For any assumed level of σ, it is therefore possible to compute the likelihood of the observed pattern of responses, and in particular of response changes. Maximizing this likelihood produces an estimated value for $\hat{\sigma}$. With this estimate that derives from the universe of respondents, confidence intervals are provided for the risk aversion parameter.

6.2 The Variance-Covariance Matrix and Inference

The analysis of Kimball, Sahm, and Shapiro [2005] has the simplifying feature that each individual's responses are sufficient to identify the mean estimate of the parameter of interest. In our case, this separability disappears and even our mean estimate of the bequest motive depends on the estimated structure of the error term. This difference arises because of the multivariate response error that we are analyzing. While they analyze two responses to the same question that are spaced in time, we analyze simultaneous responses to six different questions. Our questions differ in difficulty and meaning, so that error variances may differ. They are asked simultaneously and have sufficiently strong common elements so that one might expect errors to be correlated. Rather than prejudge the direction of the cross-questional correlation, we outline an estimation procedure for the variance-covariance matrix Σ below.

To illustrate the importance of correlation in the errors between questions, we provide two examples. First, we consider two questions relating to our two parameters of interest (ϖ, c^{MED}). The first question is a continuous version of the mixed locked box question at the end of section 5.3.1 (SQ_4). The answer is the amount of dollars the respondent would lock into the long term care box when given a locked box of \$100K for mixed use. Of course, \$100 minus this number would go into the bequest box. The second question is the end-of-life question from section 5.3.2 (SQ_{5b}). It

is to assume that in fact all individuals have the same preference parameters, and thereby to assign all individual differences in the estimates to response error. A less extreme version is to assume that various sub-populations have their preference parameters drawn from a distribution with common mean and specific variance, or that there be some explicit function connecting the distribution of these parameters to observables. However the findings concerning heterogeneity outlined in the introduction suggest that it may be best to be cautious in this regard.

too is a question about how much an individual would pay for private long-term care when wealth is \$100K, but now the respondent has only one more year to live. Both questions are essentially about the trade-off between bequest and precautionary savings motive. The two full lines in each of the panels of Figure 13 give the curve of types of agents characterized by different (ϖ, c^{MED}) pairs but the same pair of 'correct' answers $(SQ_4, SQ_{5b}) = (\$80, \$70)$. However, we don't observe the true answers since both are assumed to contain survey error:

$$\begin{aligned} \widehat{SQ}_4 &= SQ_4 + \varepsilon_4, \\ \widehat{SQ}_{5b} &= SQ_{5b} + \varepsilon_{5b}, \end{aligned}$$

with known variance-covariance error Σ^{ε}. Each panel of Figure 13 makes a different assumption on the correlation between te errors. In the left panel, the correlation between ε_4 and ε_{5b}, $\frac{\Sigma^{\varepsilon}_{12}}{\sqrt{\Sigma^{\varepsilon}_{11}}\sqrt{\Sigma^{\varepsilon}_{22}}}$, is assumed to be -1, in the middle panel it is 0, and in the right panel it is +1. Their variances are held constant across the panels. We compute the likelihood function $\mathcal{L}\left(\varpi, c^{MED}; \widehat{SQ}_4, \widehat{SQ}_{5b}\right)$ under these three assumptions on the correlation. Comparing the middle to the right plot, we notice how much smaller the confidence interval over the parameters is if the answers are negatively correlated (as opposed to independent or positively correlated). Questions with negatively correlated errors are valuable, because they lead to tighter likelihood contours, an insight reminiscent of portfolio theory. Another consequence of negative correlation is that such questions place tighter restrictions on the parameter estimates and the error variances. When two answers appear inconsistent with each other (the two estimates are far apart), negative correlation necessitates a larger error variance on one or both questions to reconcile both answers. Thus the correlations will affect our estimates of errors on individual questions. We return to this in the next section.

Because the likelihood plots in figure 13 are parallel rather than perpendicular, these questions alone cannot identify the two parameters of interest. Rather, the figure illustrates how the information contained in two survey questions leads to differences in inference depending on the correlation structure of the errors. In the second example, we revisit these same two survey questions (SQ_4 and SQ_5) but complement them with a behavioral question that asks for next year's consumption. This allows us to study the effect of correlation on the inference about the two parameters of interest. The three solid lines in 14 plot the parameter pairs corresponding to the following three answers: the respondent would put \$80K in LTC box (SQ_4), pay \$60K for private long-term care ($SQ = 5$), and consume \$19K next year. The two survey questions have a downward slope (the upper one is the end-of-life question SQ_5), while the consumption question is upward sloping. We then assume that the answers to all three questions contain errors that are normally distributed with standard deviations 15K, 15K, and .35K respectively. The consumption question is assumed to have a small error, in line with our previous discussion. The figure shows the likelihood contours for the preference parameters conditional on these three answers. The left-most plot is for uncorrelated errors. In the middle plot, the error in the consumption question is perfectly positively correlated with the end-of-life question and uncorrelated with the lock box question. In the right-most plot, the

Figure 13: Understanding Correlated Survey Error - 2 Strategic Survey Questions

This figure plots likelihood contours for different parameter pairs (ϖ, c^{MED}) conditional on observing an answer to survey question 4 and survey question 5b. We assume that both answers $\widehat{SQ}_4$ and $\widehat{SQ}_{5b}$ contain measurement error around the truth SQ_4 and SQ_{5b}. We plot the contours for three assumptions on the correlation of the error process. The left panel assumes that the errors are perfectly positively correlated, the middle panel assumes that they are uncorrelated, and the right panel assumes that they are perfectly negatively correlated.

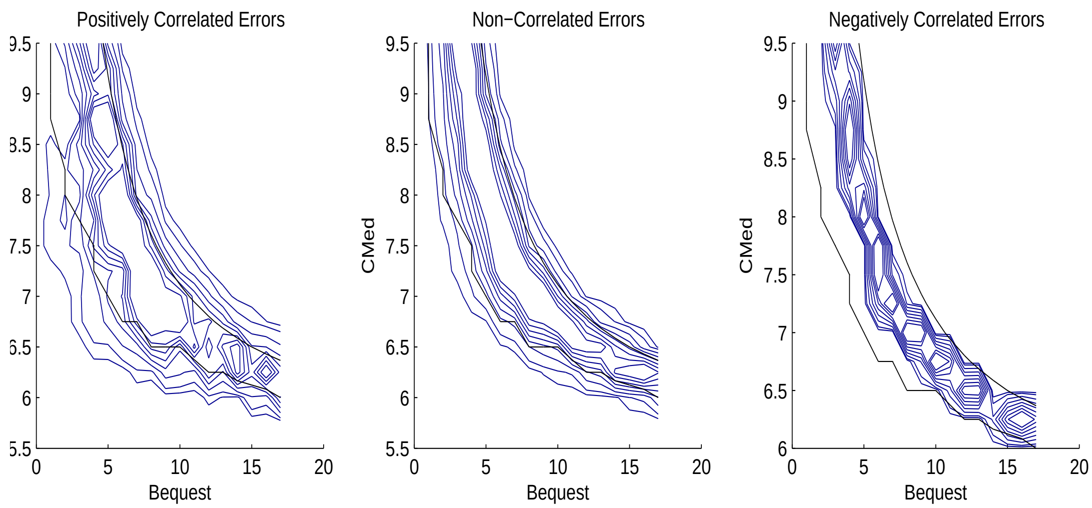

consumption error is perfectly positively correlated with the lock-box question and uncorrelated with the end-of-life question. We see that depending on the correlation structure, the estimates of the bequest parameter change from 7 to 5 to 9.

Finally, we can analyze the inference when the variance of the error is greater for some questions than for others, for a given correlation structure between questions. Naturally, the likelihood contours will be closer to the answers that are measured with more precision.

These examples illustrate the potentially profound impact that the variance-covariance matrix of errors may have on the interpretation of a given set of survey answers. We believe that these insights provide important guidance in survey design, in that there is really a portfolio effect at work. Ultimately, the goal of the survey designer will be to identify a portfolio of questions such that the likely structure of the error term offers most hope of resolving uncertainties about critical parameters.

6.3 The Estimation Problem

The key insight in Kimball, Sahm, and Shapiro [2005] that generalizes to our case is that the analysis of preference parameters revealed by means of survey answers is holistic. The answers of one individual can be interpreted only in light of the universe of responses.

We are interested in the joint likelihood of preference parameters $\{\varpi_j, c_j^{MED}\}_{j=1}^J$ for all J

Figure 14: Understanding Correlated Survey Error - With Behavioral Question

This figure plots likelihood contours for different parameter pairs (ϖ, c^{MED}) conditional on observing an answer to survey question 4, survey question 5b, and conditional on observing the consumption for next year. We assume that both answers $\widehat{SQ}_4$ and $\widehat{SQ}_{5b}$ contain normally distributed measurement error around the truth SQ_4 and SQ_{5b} with mean 0 and standard deviation of \$15K each. Expected consumption next year has measurement error with a standard deviation of \$0.35K. We plot the likelihood contours for three assumptions on the correlation of the error process. The left panel assumes that the errors are uncorrelated. In the middle plot, the error in the consumption question is perfectly positively correlated with the end-of-life question (SQ_{5b}) and uncorrelated with the lock box question (SQ_4). In the right-most plot, the consumption error is perfectly positively correlated with the lock-box question and uncorrelated with the end-of-life question.

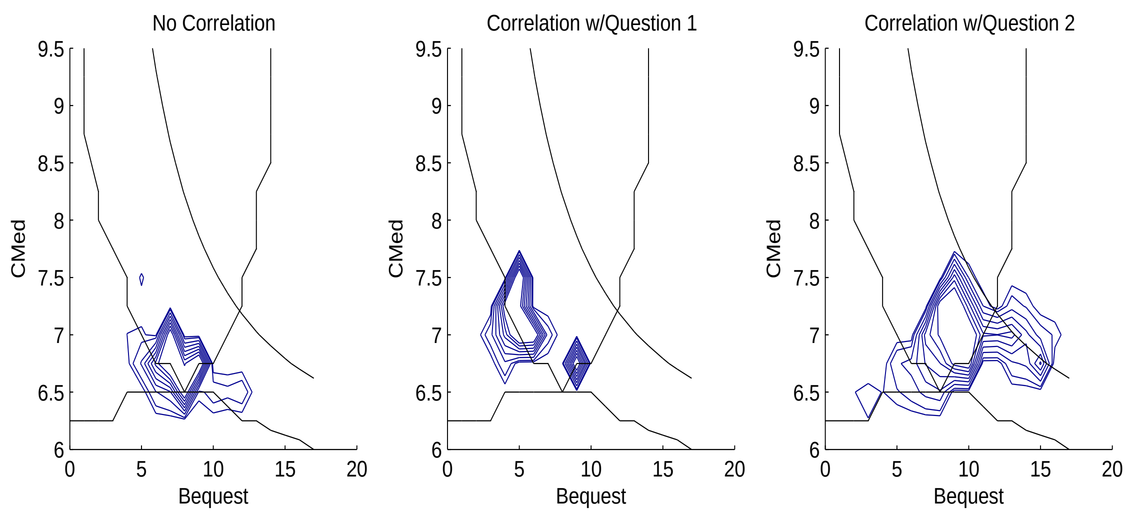

respondents and the variance covariance matrix of the survey error $\widehat{\Sigma}$. Our identification assumption is that each respondent's parameters only depend on the other respondents' answers through the variance-covariance matrix. Denote the set of answers or respondent j to all survey questions by $\{\widehat{SQ}_j\}$. Formally, the joint likelihood of all parameters and the covariance matrix, conditional on all answers can be factored as follows:

$$\mathcal{L}\left(\{\varpi_j, c_j^{MED}\}_{j=1}^J, \widehat{\Sigma}; \{\widehat{SQ}\}_{j=1}^J\right) = \mathcal{L}\left(\widehat{\Sigma}; \{\widehat{SQ}\}_{j=1}^J\right) \times \prod_{j=1}^{I} \mathcal{L}\left(\varpi_j, c_j^{MED}; \widehat{\Sigma}, \widehat{SQ}_j\right) \tag{10}$$

A Bayesian estimation procedure would use the prior joint distribution $\mathcal{L}^p$ and the survey responses to maximize the posterior likelihood function. Using Bayes' law, this posterior can be written as:

$$\mathcal{L}\left(\{\varpi_j, c_j^{MED}\}_{j=1}^J, \widehat{\Sigma}; \{\widehat{SQ}\}_{j=1}^J\right) = \frac{\Pi_{j=1}^J \mathcal{L}\left(\widehat{SQ}_j; \varpi_j, c_j^{MED}, \widehat{\Sigma}\right) \mathcal{L}^p\left(\{\varpi_j, c_j^{MED}\}_{j=1}^J, \widehat{\Sigma}\right)}{\int \Pi_{j=1}^J \mathcal{L}\left(\widehat{SQ}_j; \varpi_j, c_j^{MED}, \widehat{\Sigma}\right) d\mathcal{L}^p\left(\{\varpi_j, c_j^{MED}\}_{j=1}^J, \widehat{\Sigma}\right)} \tag{11}$$

One simplifying assumption would be to make the prior over the preference parameters independent across respondents, and independent from the prior over the covariance matrix. In that case, equation (11) further simplifies to:

$$\frac{\mathcal{L}^p\left(\widehat{\Sigma}\right) \Pi_{j=1}^J \mathcal{L}\left(\widehat{SQ}_j; \varpi_j, c_j^{MED}, \widehat{\Sigma}\right) \mathcal{L}^p\left(\varpi_j, c_j^{MED}\right)}{\int \mathcal{L}^p\left(\widehat{\Sigma}\right) \Pi_{j=1}^J \mathcal{L}\left(\widehat{SQ}_j; \varpi_j, c_j^{MED}, \widehat{\Sigma}\right) \mathcal{L}^p\left(\varpi_j, c_j^{MED}\right) d\left(\{\varpi_j, c_j^{MED}\}_{j=1}^J\right) d\widehat{\Sigma}}$$

The details of the estimation procedure will be worked on in parallel with the fielding of the survey.

6.4 Psychological Sources of Measurement Error

There are several psychological factors that could potentially cause systematic biases in the survey responses. In this section we discuss a few of these factors in more detail. Rather than analyze them all, we will add more intricate forces based on the failure of a suggestive specification test for the basic normal model. We see a more in-depth analysis of the psychological sources of survey error as a crucial feature of future research. The work of Ljunqvist [1993], who uses formal techniques to study the impact of a specific utility-based response bias, is an important precedent. Our approach to psychological biases will be involve formally incorporating them into the model, and estimating the underlying parameters in just this spirit.

6.4.1 Order Effects

We wish to allow for the possibility that the order of presentation of questions impacts responses. Order effects are a clean example where a psychological bias in survey response may be corrected based on a formal model of the error process. The issue of ordering arises in the context of the relative assignment to long term care and bequests in the locked box question (SQ_4). The most natural presentation involves specifying 100% in one box at the top of the list, and 100% in the other box at the end. One may suspect that respondents will be influenced by the order and be more likely to pick an answer high in the list than one that is lower down.

To understand how order impacts response probabilities, we rotate the order of answers for different respondents at random. We then test whether or not there is evidence of order effects that give rise to systematically different substantive responses depending on the order of the list. If so, and if the difference is in the direction of greater selection of higher options, we will fit a model with a simple geometric decay parameter to correct for order bias. In the spirit of the previous section, this parameter is just another parameter to be fit with the survey response data.

6.4.2 Anchoring

A well-known bias that may impact responses to our survey is the 'anchoring and adjustment' heuristic of Kahnemann and Tversky [1973]. The chief impact of anchoring is likely to be on the answers to survey questions SQ_1 through SQ_3 that call for respondents to report how various prizes would change their level of spending. We suspect that spending estimates will be anchored in the initial answer, so that answers to questions concerning changes will both be biased to zero and will be bunched close to one another. Again, we design our survey so that different respondents are presented with a different ordering of these questions at random.

If the anchoring and adjustment heuristic is in play, then we will expect to find that the errors in the answer to how much spending will increase based on various wealth shocks in the locked box question will be biased toward zero. In addition, we may expect them to display an intricate mutual dependence, in which subsequent answers are impacted by the numerical values of earlier answers. If these patterns are in evidence, then modeling responses in a manner that captures this behavior will allow us to improve the information content of the responses. As before, our approach will involve developing tightly parametrized models to gauge the extent of any biases we uncover, and therefore to render actual responses as informative as possible.

6.4.3 Planning and the Error

How planners differ from non-planners with regard to wealth accumulation strategies is of considerable interest (Lusardi [2003] and Ameriks, Caplin, and Leahy [2003]). Our survey will measure various general individual planning attributes, as well as analyze information gathering and planning activities in the specific areas of bequests and long term care. The goal is not only to understand how these planning propensities and planning behaviors play out in later life, but also to gain insights into the nature of survey response error. A natural conjecture is that those who plan heavily will respond to the survey in a manner that involves smaller errors. We will be able to explore this issue in the context of our statistical model of the error term.

7 Implications for Financial Innovation

The financial choices of the elderly are poorly understood. In particular, many questions have been raised concerning the lack of interest among the elderly in annuities and in long term care insurance. These are two products aimed at reducing the impact of major financial risks faced by the elderly: longevity risk and long term care expense risk. While we do not yet have complete survey evidence, it is already clear that the model has entirely different implications for the potential future growth of these two markets. Specifically, the model can help explain why there may be a limited market for actuarially fair annuities, yet a vibrant market for fair long term care insurance. If policy makers are interested in promoting markets that would benefit the elderly, they would do well to consider how best to resolve the many problems that currently bedevil the long term care insurance market.

7.1 Interest in Fair Annuities

At median wealth and income levels, the agents in the model do not want to purchase annuities because they already have insurance against longevity risk in the form of income. Purchasing annuities would take away resources needed for both other savings motives (bequests and long term care). However, for higher values of wealth and lower values of income, those with small bequest and precautionary motives would buy annuities at fair value.

In section 4.3 we showed that the model was unable to generate a demand for annuities. We computed the actuarially fair value of an annuity that pays \$5K while alive and compared it to the willingness to pay for such contract of a 62-year old household in good health with \$19K in annual income and \$180K in wealth. Figure 3 illustrated that none of the different types (different bequest and precautionary motives) were willing to pay the fair value. Further simulations indicate that only agents with less income and more wealth and with weak precautionary savings motives would be willing to pay the fair value of the annuity contract. For example, when annual income is $y = 14$ instead and wealth is $X = 360$, then households with $\varpi < 5$ and $c^{MED} > 7$ are willing to pay the actuarially fair value for the annuity.

7.2 Interest in Fair Long Term Care Insurance

We have emphasized the role of long-term care costs as the main driver of precautionary motives. Given that long-term care costs are the gorilla in the retirement expense room, an important question is why so few people have long-term care insurance. Extant insurance contracts and markets are riddled with imperfections. Our model suggests that long-term care insurance that is actuarially fair and not subject to these imperfections would be very valuable.

Simulations indicate that a 62-year old in good health with \$19K in annual income and \$180K in wealth would be willing to pay between \$6K and \$15K per year (depending on bequest and precautionary motives) for a long-term care insurance policy that pays full life-long costs at a private facility (see Figure 15). As before, we take these costs to be \$45K per year. The actuarially fair value insurance premium of this policy is only \$3.9K per year, so that all types are willing to

Figure 15: Likelihoods for Willingness to Pay for Long Term Care Insurance.

This figure plots the willingness to pay for a long-term care insurance contract that promises to pay all costs associated with a private long term care facility ($45K per year) of a 62-year old agent in good health with wealth $180K for a range of bequest and Medicaid aversion parameters (ϖ, c^{MED}): c^{MED} from 5.25 to 9.5 by increments of 0.25, and ϖ from 0 to 17 by increments of 1. This is willingness to pay per year (the insurance premium) for a contract whose actualially fair premium is $3.9K per year.

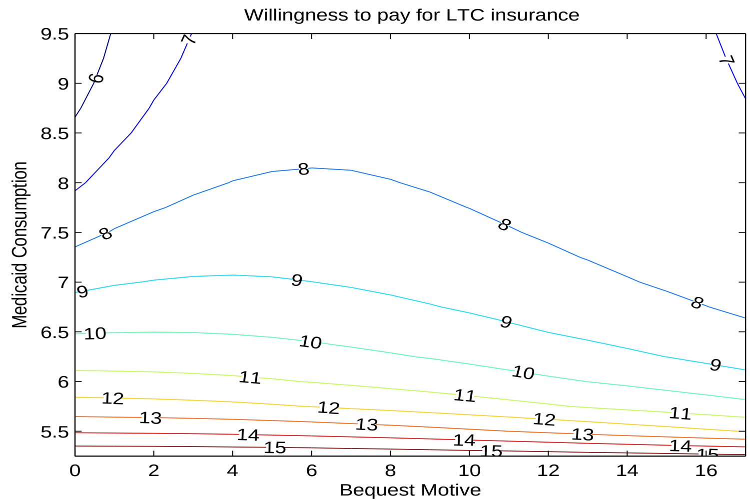

pay at least 50% more than the fair value. This suggests that the welfare gains from actuarially fair insurance contracts are potentially very large.

8 Concluding Remarks

We have outlined a strategic survey methodology designed to shed new light on the relative importance of bequest and precautionary motives for wealth accumulation in retirement. We have argued that a standard dynamic model of consumption and savings cannot separate out the two motives for wealth accumulation based on behavioral data, such as actual consumption and bequests. However, strategic survey questions that place respondents in informative contingencies, can be useful complements to behavioral data. We have formulated such contingent questions in the context of the trade-off between saving for bequests versus precautionary motives, and solved for the answer inside the model. When combined with standard behavioral data, these strategic survey questions were successful at isolating the strength of the bequest and precautionary motives.

9 Appendix: Health Transition Calibration

The distribution of medical costs in our model is controlled by the medical costs associated to each health state and by the one-period 4×4 state transition matrix $\mathcal{P}(a)$, where a denotes age in excess of 62. This matrix is parameterized by twelve parameters, nine that determine the value of $\mathcal{P}(0)$ (of the sixteen elements, four are fixed by the death state being absorbing and there are three further restrictions so that each row sums to one) and three that control the flow of probability from greater health to poorer health as age increases. We calibrate these 12 parameters to match 8 moments related to long-term care utilization (Brown and Finkelstein [2004], Table 1, males), and 4 moments related to longevity (National Center for Health Statistics, Vital Statistics [1999], Table 2 for males). Table 1 shows the moments we match, their target value, and our best fit. The last 4 rows show some features of the distribution of medical costs. For medical costs we use $h(1) = 0$, $h(2) = 4.5$, $h(3) = 30$, and $h(4) = 0$.

More precisely, the 1-period ahead transition matrix at age $62 + a$ is given by:

$$\mathcal{P}(a) = \begin{bmatrix} p_{11} & p_{12} & p_{13} & 1-p_{11}-p_{12}-p_{13} \\ p_{21} & p_{22} & p_{23} & 1-p_{21}-p_{22}-p_{23} \\ p_{31} & p_{32} & p_{33} & 1-p_{31}-p_{32}-p_{33} \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1-c_1a^e & c_1a^e\left(\frac{c_2c_3}{1+c_2+c_3c_2}\right) & c_1a^e\left(\frac{c_3}{1+c_2+c_3c_2}\right) & c_1a^e\left(\frac{1}{1+c_2+c_3c_2}\right) \\ 0 & 1-c_1a^e & c_1a^e\left(\frac{c_2}{1+c_2}\right) & c_1a^e\left(\frac{1}{1+c_2}\right) \\ 0 & 0 & 1-c_1a^e & c_1a^e \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

The second matrix is the age-adjustment. It shifts probability mass from the left (better health states) towards the right (worse health states and death), relative to the transition matrix at age 62, $\mathcal{P}(0)$. The 3 parameters c_1, c_2, and c_3 control how fast this shifting occurs. Loosely speaking, the parameter c_1 controls the transition from LTC to death as age increases; c_2 determines how much more likely death is relative to LTC when in health state 1 or 2, and c_3 determines how much likely state 2 is when in good health. The exponent e allows for faster than linear shifting as the agent becomes older. It is held fixed at $e = 1.5$. We note that there is no unique solution to the system of 12 equation and 12 parameters because the system is highly non-linear. Using a non-linear lest squares procedure, the best fit is obtained for $p_{11} = .9600$, $p_{12} = 0.0308$, $p_{13} = .0013$, $p_{21} = .3855$, $p_{22} = .6435$, $p_{23} = .0695$, $p_{31} = .0246$, $p_{32} = .1335$, $p_{33} = .7468$, $c_1 = .0016$, $c_2 = .8887$, and $c_3 = .5570$. To scale the moments to the same units, and to attach more importance to matching some moments than others, we use the following weights on the 12 moments: 100, 5, 10, 100, 100, 100, 100, 1, 4, 5, 6, and 7.

References

[1] Altonji, Joseph G., Hayashi, Fumio and Kotlikoff, Laurence J, 1997, "Parental Altruism and Inter Vivos Transfers: Theory and Evidence." Journal of Political Economy, 105(6), pp. 1121-1166.

[2] Ameriks, John, Andrew Caplin, and John Leahy, 2003, "Wealth Accumulation and the Propsensity to Plan." Quarterly Journal of Economics, 118(3), pp. 1007-1047.

[3] Anderson, Kate, Eric French, and Tina Lam, 2004, "You Can't Take It With You: Asset Run-Down At the End Of The Life Cycle." Economic Perspectives, Federal Reserve Bank of Chicago, 53, pp. 40-54.

[4] Andreoni, James, 1989, “Giving with Impure Altruism: Applications to Charity and Ricardian Equivalence.” Journal of Political Economy, Vol. 97, pp. 1447-1458.

[5] Barsky, Robert, Juster, Thomas, Kimball, Miles, and Shapiro, Matthew, 1997, “Preference Parameters and Behavioral Heterogeneity: An Experimental Approach Using the Health and Retirement Survey.” Quarterly Journal of Economics, 112, pp. 537-579.

[6] Bajtelsmit, Vickie L. and Bernasek, Alexandra, 2001, “Risk Preferences and the Investment Decisions of Older Americans.” AARP Working Paper 2001-11.

[7] Bernheim, B. Douglas, 1991, “How Strong Are Bequest Motives? Evidence Based on the Estimates of the Demand for Life Insurance and Annuities.”Journal of Political Economy, 99(5), pp. 889-927.

[8] Bernheim, Douglas B., Robert J. Lemke and John Karl Scholz, 2004, “Do Estate and Gift Taxes Affect the Timing of Private Transfers.”Journal of Public Economics, 88(12), pp. 2617-34.

[9] Bernheim, Douglas B., Shleifer, Andrei, and Summers, Lawrence H., 1995, “The Strategic Bequest Motive.” Journal of Political Economy, 93(6), pp. 1045-1076.

[10] Bernheim, Douglas B., Jonathan Skinner, and Steven Weinberg, 2001, “What Accounts for the Variation in Retirement Wealth among Households.” American Economic Review, 91(4), pp. 832-57.

[11] Brown, Jeffrey and Finkelstein, Amy, 2004, “The Interaction of Public and Private Insurance: Medicaid and the Long-Term Care Insurance Market.” NBER Working Paper.

[12] Castaneda, Ana, Javier Díaz-Giménez, and José - Víctor Ríos-Rull, 2003, “Accounting For U.S. Earnings and Wealth Inequality.” Journal of Political Economy, 111(4), pp. 818-57.

[13] Congressional Budget Office, 2004, “Financing Long Term Care for the Elderly.”

[14] Congressional Research Service, 2002, “Medicaid: Eligibility for the Aged and Disabled.”

[15] Consumer Reports, 2003, “Do You Need Long Term Care Insurance?”November.

[16] De Nardi, Christina, 2004, “Wealth Distribution, Intergenerational Links, and Estate Taxation.” Review of Economic Studies, 71(3), pp. 743-68.

[17] Dynan, Karen; Skinner, Jonathan and Zeldes, Stephen P., 2004, “Do the Rich Save More?,” Journal of Political Economy, 112(2), pp. 397-444.

[18] Dynan, Karen; Skinner, Jonathan and Zeldes, Stephen P., 2002, “The Importance of Bequests and Life-Cycle Saving in Capital Accumulation: A New Answer.”American Economic Review Papers and Proceedings, pp. 274-278.

[19] Feenberg, D., and J. Skinner, 1994, “The Risk and Duration of Catastrophic Health Care Expenditures.” The Review of Economics and Statistics, Vol. 76, No. 4, pp. 633-647.

[20] French, E., and J. Jones, 2004, "On the Distribution and Dynamics of Health Care Costs." Journal of Applied Econometrics, forthcoming.

[21] Gale, William G. and Scholz, John Karl, 1994, "Intergenerational Transfers and the Accumulation of Wealth."Journal of Economic Perspectives, 8(4), pp. 145-60.

[22] Hénin, Pierre-Yves and and Thomas Weitzenblum, 2003, "Redistribution through alternative pension reforms : A life-cycle analysis on French occupational groups." CEPREMAP Working Paper (Couverture Orange).

[23] Hurd, Michael D., 1987, "Savings of the Elderly and Desired Bequests." American Economic Review, 77 (3), pp. 298-312.

[24] Hurd, Michael D., 1990 "Research on the Elderly: Economic Status, Retirement, and Consumption."Journal of Economic Literature, 28(2), pp. 565-637.

[25] Joulfaian, David, 2004, "Gift Taxes and Lifetime Transfers: Time Series Evidence."Journal of Public Economics, 88(9-10), pp. 1917-1929.

[26] Kahneman, Daniel and Tversky, Amos, 1973, "On the Psychology of Prediction."Psychology Review, 80, pp. 237-251.

[27] Kimball, Miles S. and Shapiro, Matthew D., 2003, "Labor Supply: Are income and Substitution Effects Both Large or Both Small?"University of Michigan Working Paper.

[28] Kimball, Miles S., Sahm C. and Shapiro, Matthew D., 2005, "Using Survey-Based Risk Tolerance."University of Michigan Working Paper.

[29] Kopczuk, Wojciech and Lupton, Joseph P., 2005, "To Leave or Not To Leave: The Distribution of Bequest Motives.", NBER Working Paper 11767.

[30] Kotlikoff, Laurence J. and Summers, Lawrence H., 1981, "The Role of Intergenerational Transfers in Aggregate Capital Accumulation." Journal of Political Economy, 89(4), pp. 706-32.

[31] Laitner, John, Juster, Thomas F., 1996, "New Evidence on Altruism: A Study of TIAA-CREF Retirees."American Economic Review, 86(4), pp. 893-908.

[32] Ljungqvist, Lars, 1993, "A Unified Approach to Measures of Privacy in Randomized Response Models: A Utalitarian Perspective."Journal of the American Statistical Association, 88(421), pp. 97-103.

[33] Lusardi, Annamaria, 2003, "Planning and Saving for Retirement."Mimeo Dartmouth College.

[34] McGarry, Kathleen, 1999, "Inter Vivos Transfers and Intended Bequests." Journal of Public Economics, 73(3), 321-51.

[35] Menchik, Paul L., Martin, David, 1983, "Income Distribution, Lifetime Savings, and Bequests."American Economic Review, 73(4), pp. 672-90.

[36] MetLife Mature Market Institute, 2002a, “MetLife Market Survey of Nursing Home and Home Care Costs 2002.”

[37] MetLife Mature Market Institute, 2002b, “MetLife Survey of Assisted Living Costs 2002”

[38] Modigliani, Franco, 1988, “The Role of Intergenerational Transfers and Life-Cycle Saving in the Accumulation of Wealth.” Journal of Economic Perspectives, 2(2), pp. 15-40.

[39] Modigliani, Franco, Brumberg, Richard, 1954, “Utility Analysis and the Consumption Function: An Attempt at Integration of Cross-Section Data.”In: Post-Keynesian Economics, Kurihara, Kenneth K. (editor). New Brunswick, NJ: Rutgers University Press, pp. 388-436.

[40] Morin, R.A. and Suarez A.F, 1983, “Risk Aversion Revisited.”Journal of Finance, 38(4), pp. 1201-1216.

[41] National Center for Health Statistics, 1999. “Method for Constructing Complete Annual U.S. Life Tables.”Vital and Health Statistics, 2(129).

[42] National Center for Health Statistics, 2002. Health, United States, 2002 with Chartbook on Trends in the Health of Americans. Hyattsville, Maryland.

[43] National Center for Health Statistics, 2004. Health, United States, 2004 with Chartbook on Trends in the Health of Americans. Hyattsville, Maryland.

[44] Office of Beneficiary Determinations and Services, 2005, “Understanding Supplemental Security Income.”

[45] Page, Benjamin R., 2003, “Bequest Taxes, Inter Vivos Gifts, and the Bequest Motive.” Journal of Public Economics, 87, pp. 1219-1229.

[46] Palumbo, M., 1999, “Uncertain Medical Expenses and Precautionary Saving Near The End Of The Life Cycle.” Review of Economic Studies, Vol. 66, No. 2, pp. 395-421.

[47] Pauly, Mark, 1990, “The Rational Nonpurchase of Long-Term-Care Insurance. ”Journal of Political Economy, 98(1), pp. 153-168.

[48] Poterba, James M., 2001, “Estate and Gift Taxes and Incentives for Inter Vivos Giving.” Journal of Public Economics, 79(1), 237-64.

[49] Robinson, James, 2002, “A Long Term Care Status Transition Model.” Mimeo, University of Wisconsin, Madison.

[50] Slemrod, Joel and Kopczuk, Wojciech, 2001, “Dying to Save Taxes: Evidence from Estate Tax Returns on the Death Elasticity.”National Bureau of Economic Research (Cambridge, MA) Working Paper No. 8158.

[51] U.S. Government Accountability Office, 2003, “NURSING HOME FIRE SAFETY Recent Fires Highlight Weaknesses in Federal Standards and Oversight.”GAO-04-660.

[52] Wang, Hui and Sherman, Hanna, 1997, "Does Risk Tolerance Decrease With Age?"Association for Financial Counseling and Planning Education, Vol. 8, No. 2.

[53] Yaari, Menahem E., 1965, "Uncertain Lifetime, Life Insurance, and the Theory of the Consumer."Review of Economic Studies, 32(2), pp. 137-150.