

Why Do Life Insurance Policyholders Lapse?

The Roles of Income, Health and Bequest Motive Shocks*

Hanming Fang[†]

Edward Kung[‡]

March 22, 2012

Abstract

We present and empirically implement a dynamic discrete choice model of life insurance decisions to assess the importance of various factors in explaining life insurance lapsation. In order to explain some key features in the data, our model incorporates multiple dimensions of serially correlated unobservable state variables which we deal with using posterior distributions of the unobservables simulated from Sequential Monte Carlo (SMC) method. We estimate the model using the life insurance holding information from the Health and Retirement Study (HRS) data. Counterfactual simulations using the estimates of our model suggest that a large fraction of life insurance lapsations are driven by idiosyncratic shocks which are uncorrelated with health, income, and bequest motives, particularly when policyholders are relatively young. But as the remaining policyholders get older, the role of such i.i.d. shocks gets smaller, and more of their lapsation is driven either by income, health or bequest motive shocks. Income and health shocks are relatively more important than bequest motive shocks in explaining lapsation when policyholders are young, but as they age, the bequest motive shocks play a more important role. We also suggest the implications of these findings regarding the effects of the emerging life settlement market on consumer welfare.

Keywords: Life insurance lapsations, Sequential Monte Carlo Method

JEL Classification Codes: G22, L11

*We have received helpful comments and suggestions from Steven Berry, Xu Cheng, Han Hong, Aprajit Mahajan, Panle Jia, Alessandro Lizzeri, Jim Poterba, Xun Tang, Ken Wolpin, Motohiro Yogo and seminar/conference participants at New York University, Cowles Foundation Summer Conference in Structural Microeconomics (2010), AEA/CEANA (2011) and SED (2012) for many helpful comments, questions and suggestions. Fang would also like to gratefully acknowledge the generous financial support from the National Science Foundation through Grant SES-0844845. All remaining errors are our own.

[†]Department of Economics, University of Pennsylvania, 3718 Locust Walk, Philadelphia, PA 19104; and the NBER.
Email: hanming.fang@econ.upenn.edu

[‡]Department of Economics, Duke University, 213 Social Sciences Building, P.O. Box 90097, Durham, NC 27708-0097.
Email: edward.kung@duke.edu

1 Introduction

The life insurance market is large and important. According to Life Insurance Marketing and Research Association International (LIMRA International), 78 percent of American families owned some type of life insurance in 2004. By the end of 2008, the total number of individual life insurance policies in force in the United States stood at about 156 million; and the total individual policy face amount in force reached over 10 trillion dollars (see [American Council of Life Insurers \(2009, p. 63-74\)](#)).

Life Insurance Market. There are two main types of individual life insurance products, term life insurance and whole life insurance.¹ A term life insurance policy covers a person for a specific duration at a fixed or variable premium for each year. If the person dies during the coverage period, the life insurance company pays the face amount of the policy to his/her beneficiaries, provided that the premium payment has never lapsed. The most popular type of term life insurance has a fixed premium during the coverage period and is called level term life insurance. A whole life insurance policy, on the other hand, covers a person's entire life, usually at a fixed premium. In the United States at year-end 2008, 54 percent of all life insurance policies in force were Term Life insurance. Of the new individual life insurance policies purchased in 2008, 43 percent, or 4 million policies, were term insurance, totaling \$1.3 trillion, or 73 percent, of the individual life face amount issued (see [American Council of Life Insurers \(2009, p. 63-74\)](#)). Besides the difference in the period of coverage, term and whole life insurance policies also differ in the amount of cash surrender value (CSV) received if the policyholder surrenders the policy to the insurance company before the end of the coverage period. For term life insurance, the CSV is zero; for whole life insurance, the CSV is typically positive and pre-specified to depend on the length of time that the policyholder has owned the policy. One important feature of the CSV on whole life policies relevant to our discussions below is that by government regulation, CSVs do *not* depend on the health status of the policyholder when surrendering the policy.²

Lapsation. Lapsation is an important phenomenon in life insurance markets. Both LIMRA and the Society of Actuaries consider that a policy lapses if its premium is not paid by the end of a

¹The whole life insurance has several variations such as universal life (UL), variable life (VL) and variable-universal life (VUL). Universal life allows flexible premiums subject to certain minimums and maximums. For variable life, the death benefit varies with the performance of a portfolio of investments chosen by the policyholder. Variable-universal life combines the flexible premium options of UL with the varied investment option of VL (see [Gilbert and Schultz, 1994](#)).

²The life insurance industry typically thinks of the CSV from the whole life insurance as a form of tax-advantaged investment instrument (see [Gilbert and Schultz, 1994](#)).

	1998	1999	2000	2001	2002	2003	2004	2005	2006	2007	2008
By Face Amount	8.3	8.2	9.4	7.7	8.6	7.6	7.0	6.6	6.3	6.4	7.6
By Number of Policies	6.7	7.1	7.1	7.6	9.6	6.9	7.0	6.9	6.9	6.6	7.9

Table 1: Lapsation Rates of Individual Life Insurance Policies, Calculated by Face Amount and by Number of Policies: 1998-2008.

Source: [American Council of Life Insurers \(2009\)](#)

specified time (often called the grace period).³ According to LIMRA (2009, p. 11), the life insurance industry calculates the annualized lapsation rate as follows:

$$\text{Annualized Policy Lapse Rate} = 100 \times \frac{\text{Number of Policies Lapsed During the Year}}{\text{Number of Policies Exposed to Lapse During the Year}}$$

The number of policies exposed to lapse is based on the length of time the policy is exposed to the risk of lapsation during the year. Termination of policies due to death, maturity, or conversion are not included in the number of policies lapsing and contribute to the exposure for only the fraction of the policy year they were in force. Table 1 provides the lapsation rates of individual life insurance policies, calculated according to the above formula, both according to face amount and the number of policies for the period of 1998-2008. Of course, the lapsation rates also differ significantly by the age of the policies. For example, LIMRA (2009, p. 18) showed that the lapsation rates are about 2-4% per year for policies that have been in force for more than 11 years in 2004-2005.

Reasons for Lapsation Have Important Implications Regarding the Welfare Effects of Life Settlement Market. Our interest in the empirical question of why life insurance policyholders lapse their policies is primarily driven by the recent theoretical research on the effect of the life settlement market on consumer welfare. A life settlement is a financial transaction in which a policyholder sells his/her life insurance policy to a third party – the life settlement firm – for more than the cash value offered by the policy itself. The life settlement firm subsequently assumes responsibility for all future premium payments to the life insurance company, and becomes the new beneficiary of the life insurance policy if the original policyholder dies within the coverage period.⁴ The life settlement industry is quite recent, growing from just a few billion dollars in

³This implies that if a policyholder surrenders his/her policy for cash surrender value, it is also considered as a lapsation.

⁴The legal basis for the life settlement market seems to be the Supreme Court ruling in *Grigsby v. Russell* [222 U.S. 149, 1911], which upheld that for life insurance an “insurable interest” only needs to be established at the time the policy becomes effective, but does not have to exist at the time the loss occurs. The life insurance industry has typically

the late 1990s to about \$12-\$15 billion in 2007, and according to some projections (made prior to the 2008 financial crisis), is expected to grow to more than \$150 billion in the next decade (see Chandik, 2008).⁵

In recent theoretical research, Daily, Hendel and Lizzeri (2008) and Fang and Kung (2010a) showed that, if policyholders' lapsation is driven by their loss of bequest motives, then consumer welfare is unambiguously lower with a life settlement market than without; however, Fang and Kung (2010b) showed that if policyholders' lapsation is driven by income or liquidity shocks, then a life settlement market may potentially improve consumer welfare.

The reason for the difference in the welfare result is as follows. Life insurance is typically a long-term contract with one-sided commitment in which the life insurance companies commit to a specified death benefit provided that the premium payments are made, whereas the policyholder can lapse anytime. Because the premium profile is typically constant while the policyholder's mortality rate typically increases with age, the contracts are thus *front-loaded*; that is, in the early part of the policy period, the premium payments exceed the actuarially fair value of the risk insured. In the later part of the policy period, the premium payments are less than the actuarially fair value. As a result, whenever a policyholder lapses his/her policy after holding it for several periods, the life insurance company pockets the so-called *lapsation profits*, which is factored into the pricing of the life insurance policy to start with due to competition (see Gilbert and Schultz, 1994). The key effect of the settlement firms on the life insurers is that the settlement firms will effectively take away the lapsation profits, forcing the life insurers to adjust the policy premiums and possibly the whole structure of the life insurance policy under the consideration that lapsation profits do not exist. In the theoretical analysis, we show that life insurers may respond to the threat of life settlement by limiting the degree of reclassification risk insurance, which certainly reduces consumer welfare.⁶ However, the settlement firms are providing cash payments to policyholders when the policies are sold to the life settlement firms. The welfare loss from the reduction in extent of reclassification risk insurance has to be balanced against the welfare gain to the consumers when they receive payments from the settlement firms. If policyholders sell their policies due to income shocks, then the cash payments are received at a time when the marginal utility of income is particularly high, and the balance of the two effects may result in a net welfare

included a two-year contestability period during which transfer of the life insurance policy will void the insurance.

⁵The life settlement industry actively targets wealthy seniors 65 years of age and older with life expectancies from 2 to up to 12-15 years. This differs from the earlier viatical settlement market developed during the 1980s in response to the AIDS crisis, which targeted persons in the 25-44 age band diagnosed with AIDS with life expectancy of 24 months or less. The viatical market largely evaporated after medical advances dramatically prolonged the life expectancy of an AIDS diagnosis.

⁶The constant premium specified in the typical long-term life insurance contracts insures against the fluctuations in the actuarial-fair premium as the policyholder's health changes. This type of insurance is known as *reclassification risk insurance*.

gain for the policyholders. If policyholders sell their policies as a result of losing bequest motives, the balance of the two effects on net result in a welfare loss. Thus, to inform policy-makers on how the emerging life settlement market should be regulated, an empirical understanding of why policyholders lapse is of crucial importance.

What Do We Do in This Paper? For this purpose, we present and empirically implement a dynamic discrete choice model of life insurance decisions. The model is “semi-structural” and is designed to bypass data limitations where researchers only observe whether an individual has made a new life insurance decision (i.e., purchased a new policy, or added to/changed an existing policy) but do not observe what the actual policy choice is nor the choice set from which the new policy is selected. We empirically implement the model using the limited life insurance holding information from the Health and Retirement Study (HRS) data. An important feature of our model is the incorporation of serially correlated unobservable state variables. In our empirical analysis, we provide ample evidence that such serially correlated unobservable state variables are necessary to explain some key features in the data.

Methodologically, we deal with serially correlated unobserved state variables using posterior distributions of the unobservables simulated using Sequential Monte Carlo (SMC) methods.⁷ Relative to the few existing papers in the economics literature that have used similar SMC methods, our paper is, to the best of our knowledge, the first to incorporate multi-dimensional serially correlated unobserved state variables. In order to give the three unobservable state variables in our empirical model their desired interpretations as unobserved income, health and bequest motive shocks respectively, we propose two channels through which we can anchor these unobservables to their related observable counterparts. We also discuss how the additional unobservable state variables significantly improve our model fit.

Our estimates for the model with serially correlated unobservable state variables are sensible and yield implications about individuals’ life insurance decisions consistent with the both intuition and existing empirical results. In a series of counterfactual simulations reported in Table 12, we find that a large fraction of life insurance lapsations are driven by idiosyncratic shocks which are uncorrelated with health, income, and bequest motives, particularly when policyholders are relatively young. But as the remaining policyholders get older, the role of such i.i.d. shocks gets smaller, and more of their lapsations are driven either by income, health or bequest motive shocks. Income and health shocks are relatively more important than bequest motive shocks in explain-

⁷See also [Norets \(2009\)](#) which develops a Bayesian Markov Chain Monte Carlo procedure for inference in dynamic discrete choice models with serially correlated unobserved state variables. [Kasahara and Shimotsu \(2009\)](#) and [Hu and Shum \(2009\)](#) present the identification results for dynamic discrete choice models with serially correlated unobservable state variables.

ing lapsation when policyholders are young, but as they age, the bequest motive shocks play a more important role. We discuss the implications of these findings on the effects of life settlement markets on consumer welfare.

The remainder of the paper is structured as follows. In Section 2 we describe the data set used in our empirical analysis and describe how we constructed key variables, and we also provide the descriptive statistics. In Section 3 we present the empirical model of life insurance decisions. In Section 4 we estimate the dynamic model without serially correlated unobservable state variables, and show via simulations that the dynamic model without serially correlated unobservable state variables fails to replicate some important features of the data. In Section 5 we extend the dynamic model to incorporate serially correlated unobserved state variables, describe the SMC method to account for them in estimation, and provide the estimation results. In Section 6 we report our counterfactual experiments using the model with unobservables. In section 7 we conclude.

2 Data

We use data from the Health and Retirement Study (HRS). The HRS is a nationally representative longitudinal survey of older Americans which began in 1992 and has been conducted every two years thereafter.⁸ The HRS is particularly well suited for our study for two reasons. First, the HRS contains rich information about income, health, family structure, and life insurance ownership. If family structure can be interpreted as a measure of bequest motive, then we have all the key factors motivating our analysis. Second, the HRS respondents are generally quite old: between 50 to 70 years of age in their first interview. As we described in the introduction, the life settlements industry typically targets policyholders in this age range or older, so it is precisely the lapsation behavior of this group that we are most interested in.

Our original sample consists of 4,512 male respondents who were successfully interviewed in both the 1994 and 1996 HRS waves, and who were between the ages of 50 and 70 in 1996. We chose 1996 as the period to begin decision modeling because the 1996 wave is *the first time* the HRS began to ask questions about whether or not the respondent lapsed any life insurance policies and whether or not the respondent obtained any new life insurance policies since the last interview. As we will explain below, these questions are used prominently in the construction of the key decision variable used in our structural model. We use only respondents who were also interviewed in 1994 so we can know whether or not they owned life insurance in 1994.

We follow these respondents until 2006. Any respondent who missed an interview for any

⁸See <http://hrsonline.isr.umich.edu/concord> for the survey instruments used in all the waves of HRS.

Table 2: Sample Selection Criterion and Sample Size

Selection Criterion	Sample Size
All individuals who responded to both 1994 and 1996 HRS interviews	17,354
... males who were aged between 50 and 70 in 1996 (wave 3)	4,512
... did not having any missing interviews from 1994 to 2006	3,696
... did not have any missing values for reported life insurance ownership status from 1994 to 2006	3,567
... reported owning life insurance at least once from 1996 to 2006	3,324

Note: The selection criteria are cumulative.

reason other than death between 1996 and 2006 was dropped from the sample. Any respondent with a missing value on life insurance ownership any time during this period was also dropped. This leaves us a sample of 3,567 males. We also dropped 243 individuals who never reported owning life insurance in any wave of HRS data. Our final analysis sample thus consists of 3,324 in wave 1996 and the survivors among them in subsequent waves, 3,195 in wave 1998, 3022 in wave 2000, 2,854 in wave 2002, 2,717 in wave 2004 and 2,558 in wave 2006. Table 2 describes how we come to our final estimation sample.

Construction of Variables Related to Life Insurance Decisions. Here we describe the questions in HRS we use to construct the life-insurance related variables.

- For whether or not an individual owned life insurance in the current wave, we use the individual’s response to the following HRS survey question, which is asked in all waves:

[Q1] “Do you currently have any life insurance?”

- For whether or not an individual obtained a policy since the previous wave, we use the individual’s response to the following HRS question, which is asked all waves starting in 1996:

[Q2] “Since (previous wave interview month-year) have you obtained any new life insurance policies?”

If the respondent answers “yes,” we consider him to have obtained a new policy.

- For whether or not an individual lapsed a policy since last wave, we use the individual’s response to the following HRS question, which is asked all waves starting in 1996:

[Q3] “Since (previous wave interview month-year) have you allowed any life insurance policies to lapse or have any been cancelled?”

We also use the response to another survey question, which is also asked all waves starting in 1996:

[Q4] “Was this lapse or cancellation something you chose to do, or was it done by the provider, your employer, or someone else?”

If the respondent answers “yes” to the first question *and* answers “my decision” to the second question, we consider him to have lapsed a policy.

In the notation of the model we will present in the Section 3, we construct an individual’s period- t (or wave- t) decisions as follows:

- For the individual who reported not having life insurance in the previous wave ($d_{t-1} = 0$), we let $d_t = 0$ if the individual reports not having life insurance this wave; and $d_t = 1$ if the individual reports having life insurance this wave (“yes” to Q1). Because the individual does not own life insurance in period $t - 1$ but does in period t , we interpret that he chose the optimal policy in period t given his state variables at t .
- For the individual who reported having life insurance in the previous wave ($d_{t-1} \geq 1$), we let $d_t = 0$ if the individual reports not having life insurance this wave (“no” to Q1); and we let $d_t = 1$ (i.e. the individual re-optimizes his life insurance) if the individual reports having life insurance this wave (“yes” to Q1) *and* he obtained new life insurance policy (“yes” to Q2), **OR** if the individual answered “yes” to Q1, reported lapsing (i.e. answered “yes” to Q3) and reported that lapse was his own decision (answered “my own decision” to Q4). Note that under this construction, we have interpreted the “lapsing or obtaining” of any policies as an indication that the respondent re-optimized his life insurance coverage. Finally, we let $d_t = 2$ (i.e. he kept his previous life insurance policy unchanged) if the individual reports having life insurance this wave (“yes” to Q1) **AND** the individual reported no to obtaining new policy (“no” to Q2) **AND** the individual did not lapse any existing policy (either reported “no” to Q3 or reported “yes” to Q3 but did not report “my decision” to Q4).

Information About the Details of Life Insurance Holdings in the HRS Data. HRS also has questions regarding the face amount and premium payments for life insurance policies. However, there are several problems with incorporating these variables into our empirical analysis. First, the questions differ across waves. In the 1994 wave, questions were asked regarding the total face amount and premiums for term life policies; but for whole life policies only total face amount was collected.⁹ In the 1996 and 1998 waves, information about lapsation and face amounts, but not

⁹The questions in 1994 wave related to premium and face amount are: [W6768]. About how much do you pay for (this term insurance/these term insurance policies) each month or year? [W6769]. Was that per month, year, or what?

premiums. From 2000 on, the HRS asked about the combined face value for all policies, combined face value for whole life policies, and the combined premium payments for whole life policies. Note that information about the premiums for term life policies were not collected from 2000 on. Second, there is a very high incidence of missing data regarding life insurance premiums and face amounts. In our selected sample, 40.3% of the respondents have at least one instance of missing face amount in waves when he reported owning life insurance. The incidence of missing values in premium payments is even higher. Third, even for those who reported face amount and premium payments for their life insurance policies, we do not know the choice set they faced when purchasing their policies.

For these reasons, we decide to only model the individuals' life insurance decisions regarding whether to re-optimize, lapse or maintaining an existing policy, and only use the observed information about the above decisions in estimating the model.

2.1 Descriptive Statistics

Patterns of Life Insurance Coverage and its Transitions. Table 3 provides the life insurance coverage and patterns of transition between coverage and no coverage in the HRS data. Panel A shows that among the 3,324 live respondents in 1996, 88.1% are covered by life insurance; among the 3,195 who survived to the 1998 wave, 85.7% owned life insurance, etc. Over the waves, the life insurance coverage rates among the live respondents seem to exhibit a declining trend, with the coverage rate among the 2,558 who survived to the 2006 wave being about 78.6%.

Panel B and C show, however, that there is substantial transition between coverage and no coverage. Panel B shows that among the 512 individuals who did not have life insurance coverage in 1994, almost a half (47.5%) obtained coverage in 1996; in later waves between 25.6% to 33.7% of individuals without life insurance in the previous wave ended up with coverage in the next wave. Panel C shows that there is also substantial lapsation among life insurance policyholders. In our data, between-wave lapsation rates range from 4.6% to 10.2%. Considering that our sample is relatively old and the tenures of holding life insurance policies in the HRS sample are also typically longer, these lapsation rates are in line with the industry lapsation rates reported in the introduction (see Table 1).

Panel D shows that even among those individuals who own life insurance in both the previous wave and the current wave, a substantial fraction has changed the terms of their coverage, or in the words of our model, re-optimized. Between 6.0% to 9.5% of the sample who have insurance coverage in adjacent waves reported changing the terms of their coverages.

[W6770]. What is the current face value of all the term insurance policies that you have? [W6773]. What is the current face value of (this [whole life] policy/these [whole life] policies?)

Table 3: Life Insurance Coverage and Transition Patterns in HRS: 1996-2006

	Wave					
	1996	1998	2000	2002	2004	2006
Panel A: Life Insurance Coverage Status						
Currently covered by life insurance	2,927	2,739	2,524	2,313	2,187	2,011
	88.1%	85.7%	83.5%	81.0%	80.5%	78.6%
No life insurance coverage	397	456	498	541	530	547
	11.9%	14.3%	16.5%	19.0%	19.5%	21.4%
Total live respondents	3,324	3,195	3,022	2,854	2,717	2,558
Panel B: Life Insurance Coverage Status Conditional on No Coverage in Previous Wave						
Life insurance coverage this wave	243	125	130	150	163	123
	47.5%	33.4%	31.9%	33.7%	32.7%	25.6%
No life insurance coverage this wave	269	249	277	295	336	357
	52.5%	66.6%	68.1%	66.3%	67.3%	74.4%
Total live respondents with no coverage last wave	512	374	407	445	499	480
Panel C: Life Insurance Coverage Status Conditional on Coverage in Previous Wave						
Life insurance coverage this wave	2,684	2,614	2,394	2,163	2,024	1,888
	95.4%	92.7%	91.5%	89.8%	91.3%	90.9%
No life insurance coverage this wave	128	207	221	246	194	190
	4.6%	7.3%	8.5%	10.2%	8.7%	9.1%
Total respondents with coverage last wave	2,812	2,821	2,615	2,409	2,218	2,078
Panel D: Whether Changed Terms of Coverage Conditional on Coverage in Both Current and Previous Waves						
Did not change terms of coverage	2,430	2,395	2,233	2,034	1,881	1,769
	90.5%	91.6%	93.3%	94.0%	92.9%	93.7%
Changed terms of coverage	254	219	161	129	143	119
	9.5%	8.4%	6.7%	6.0%	7.1%	6.3%
Total live respondents with coverage in both waves	2,684	2,614	2,394	2,163	2,024	1,888

Table 4: Summary Statistics of State Variables.

Variable Description	Wave					
	1996	1998	2000	2002	2004	2006
Age of respondent	61.102 (4.3535)	63.035 (4.3441)	64.925 (4.3070)	66.852 (4.2861)	68.789 (4.2858)	70.676 (4.2512)
Household income (\$1,000s)	62.632 (73.846)	61.978 (73.512)	65.187 (87.389)	63.246 (76.716)	66.278 (87.937)	66.549 (80.344)
Log household income	10.582 (1.3013)	10.587 (1.2086)	10.623 (1.2051)	10.611 (1.2045)	10.688 (1.0377)	10.727 (0.9150)
Whether ever diagnosed with high blood pressure	0.4025 (0.4905)	0.4372 (0.4961)	0.4765 (0.4995)	0.5172 (0.4998)	0.5683 (0.4954)	0.6197 (0.4856)
Whether ever diagnosed with diabetes	0.1439 (0.3510)	0.1594 (0.3661)	0.1755 (0.3805)	0.2062 (0.4046)	0.2273 (0.4191)	0.2523 (0.4344)
Whether ever diagnosed with cancer	0.0599 (0.2373)	0.0811 (0.2731)	0.1007 (0.3009)	0.1287 (0.3349)	0.1591 (0.3659)	0.1874 (0.3903)
Whether ever diagnosed with lung disease	0.0710 (0.2569)	0.0777 (0.2677)	0.0791 (0.2700)	0.0880 (0.2834)	0.1083 (0.3108)	0.1170 (0.3215)
Whether ever diagnosed with heart disease	0.1902 (0.3926)	0.2145 (0.4106)	0.2351 (0.4241)	0.2672 (0.4426)	0.3050 (0.4605)	0.3435 (0.4750)
Whether ever diagnosed with stroke	0.0494 (0.2167)	0.0579 (0.2337)	0.0652 (0.2470)	0.0712 (0.2572)	0.0829 (0.2757)	0.0986 (0.2982)
Whether ever diagnosed with psychological problem	0.0599 (0.2373)	0.0714 (0.2575)	0.0788 (0.2695)	0.0873 (0.2823)	0.0932 (0.2907)	0.1002 (0.3003)
Whether ever diagnosed with arthritis	0.3904 (0.4879)	0.4460 (0.4972)	0.4831 (0.4998)	0.5312 (0.4991)	0.5790 (0.4938)	0.6205 (0.4854)
Sum of above conditions	1.3672 (1.2409)	1.5453 (1.2912)	1.6940 (1.3103)	1.8969 (1.3372)	2.1230 (1.3941)	2.3392 (1.4204)
Whether married	0.8504 (0.3567)	0.8431 (0.3638)	0.8444 (0.3626)	0.8447 (0.3623)	0.8438 (0.3631)	0.8326 (0.3734)
Whether has children	0.9419 (0.2340)	0.9402 (0.2372)	0.9394 (0.2386)	0.9386 (0.2400)	0.9366 (0.2436)	0.9335 (0.2492)
Age of youngest child	24.832 (10.059)	26.813 (10.020)	28.764 (9.959)	30.892 (9.839)	32.933 (9.779)	34.894 (9.718)
# of live respondents	3,324	3,195	3,022	2,854	2,717	2,558

Note: Standard deviations are in parenthesis.

Table 5: Reduced-Form Logit Regression on the Probability of Buying Life Insurance, Conditional on Having No Life Insurance in the Previous Wave

Variable	Coefficient	Std. Error
Constant	0.9224	0.6039
Age	-0.0407***	0.0084
Logincome	0.0929***	0.0315
Number of health conditions	-0.0810***	0.0297
Married	0.0831	0.1007
Has children	0.2704	0.1713
Age of youngest child	0.0006	0.0043
Observations	2,707	
Log likelihood	-1,713.4	

Notes: *, **, and *** respectively represent significance at 10%, 5% and 1% levels.

Summary Statistics of State Variables. Table 4 summarizes the key state variables for the sample used in our empirical analysis. It shows that the average age of the live respondents in our sample is 61.1. The means of household income in our sample are quite stable around \$62,000 to \$66,000; and similarly the means of log household incomes are stable around 10.58 to 10.73. The next eight rows report the mean of the incidence of health conditions, including high blood pressure, diabetes, cancer, lung disease, heart disease, stroke, psychological problems and arthritis. It shows clear signs of health deterioration for the surviving samples over the years. The sum of the above eight health conditions steadily increase from 1.37 in 1996 to 2.34 in 2006. Finally, the marital status of the surviving sample seems to be quite stable, with the fraction married being in the range of 83% to 85%.

Tables A.1 and A.2 in the appendix summarize the mean and standard deviation of the state variables by the life insurance coverage status. There does not seem to be much of a difference in ages between those with and without life insurance coverage, but the mean log household income is significantly higher for those with life insurance than those without and life insurance policyholders are much more likely to be married than those without. We also find that those with life insurance tend to be healthier than those without life insurance.

Table 6: Reduced-Form Multinomial Logit Regression on the Probability of Lapsing, Changing Coverage, or Maintaining Coverage, Conditional on Owning Life Insurance in the Previous Wave

Variable	Change existing coverage		Maintain existing coverage	
	Coefficient	Std. Err.	Coefficient	Std. Err.
Constant	-0.1119	0.9182	1.6361***	0.6333
Age	-0.0789***	0.0099	-0.0480***	0.0069
Logincome	0.4542***	0.0578	0.2522***	0.0396
Number of health conditions	-0.0493	0.0391	-0.0411	0.0267
Married	0.3358**	0.1372	0.3092***	0.0908
Has children	-0.0912	0.2027	0.0464	0.1366
Age of youngest child	0.0097*	0.0051	0.0075**	0.0034
Δ Age	-0.0956	0.0623	0.2067***	0.0439
$(\Delta$ Age) ²	0.0090*	0.0048	-0.0082**	0.0034
Δ Logincome	-0.1406***	0.0457	-0.0386	0.0305
$(\Delta$ Logincome) ²	0.0174***	0.0053	0.0102***	0.0037
Δ Conditions	0.1268	0.1362	0.0159	0.0877
$(\Delta$ Conditions) ²	-0.0843	0.0514	-0.0178	0.0285
Δ Married	0.3218	0.1998	-0.0954	0.1391
Observations			14,951	
Log likelihood			-7,565.6	

Notes: (1). Conditional on owning life insurance, the three choices are: (a). to lapse all coverage; (b). to change the existing coverage; and (c). to maintain the existing coverage. The base outcome is set to choice (a). For any variable x , Δx is the difference between the current value of x and the value of x which occurred during the last period in which the respondent changed his coverage.

(2). *, **, and *** respectively represent significance at 10%, 5% and 1% levels.

Reduced-Form Determinants of the Life Insurance Decisions. Table 5 presents the coefficient estimates of a Logit regression on the probability of purchasing life insurance among those who did not have coverage in the previous wave. It shows that the richer, younger, healthier and married individuals are more likely to purchase life insurance coverage than the poorer, older, unhealthier and widowed or unmarried individuals. Table 6 presents the estimates of a multinomial Logit regression for the probability of lapsing, changing coverage, or maintaining the previous coverage. The omitted category is lapsing all coverage. The estimates show that richer individuals are more likely to either maintain the current coverage or changing existing coverage than to lapse all coverage; individuals who experienced negative income shocks are more likely to lapse all coverage; individuals who are either divorced or widowed are more likely to lapse all coverage; finally, individuals who have experienced an increase in the number of health conditions are somewhat more likely to lapse all coverage, though the effect is not statistically significant.

3 An Empirical Model of Life Insurance Decisions

In this section we present a dynamic discrete choice model of how individuals make life insurance decisions, which we will later empirically implement. Our model is simple, yet rich enough to capture the dynamic intuition behind the life insurance models of [Hendel and Lizzeri \(2003\)](#) and [Fang and Kung \(2010a\)](#).

Time is discrete and indexed by $t = 1, 2, \dots$. In the beginning of each period t , an individual i either has or does not have an existing life insurance policy. If the individual enters period t without an existing policy, then he chooses between not owning life insurance ($d_{it} = 0$) or optimally purchasing a new policy ($d_{it} = 1$). If the individual enters period t with an existing policy, then, besides the above two choices, he can additionally choose to keep his existing policies ($d_{it} = 2$). If an individual who has life insurance in period $t - 1$ decides not to own life insurance in period t , we interpret it as lapsation of coverage. As we describe in Section 2, the choice $d_{it} = 1$ for an individual who previously owns at least one policy is interpreted more broadly: an individual is considered to have re-optimized his existing policy if he reported purchasing a new policy or choosing to lapse one of the existing policies (while he continues to hold at least one other policy). The key interpretation for choice $d_{it} = 1$ is that the individual re-optimized his life insurance holdings.

Flow Payoffs from Choices. Now we describe an individual's payoffs from each of these choices. First, let $x_{it} \in \mathcal{X}$ denote the vector of *observable* state variables of individual i in period t , and let

$z_{it} \in \mathcal{Z}$ denote the vector of *unobservable* state variables.¹⁰ These characteristics include variables that affect the individual’s preference for or cost of owning life insurance, such as income, health and bequest motives. We normalize the utility from not owning life insurance (i.e., $d_{it} = 0$) to 0; that is,

$$u_0(x_{it}, z_{it}) = 0 \text{ for all } (x_{it}, z_{it}) \in \mathcal{X} \times \mathcal{Z}. \quad (1)$$

The utility from optimally purchasing a new policy in state (x_{it}, z_{it}) , i.e., $d_{it} = 1$, regardless of whether he previously owned a life insurance policy, is assumed to be:

$$u_1(x_{it}, z_{it}) + \varepsilon_{1it}, \quad (2)$$

where ε_{1it} is an idiosyncratic choice specific shock, drawn from to a Type-I extreme value distribution. In our empirical analysis, we will specify $u_1(x_{it}, z_{it})$ as a flexible polynomial of x_{it} and z_{it} .

Now we consider the flow utility for an individual i entering period t with an existing policy which was last re-optimized at period \hat{t} . That is, let $\hat{t} = \sup \{s | s < t, d_{is} = 1\}$. Let $(\hat{x}_{it}, \hat{z}_{it}) = (x_{i\hat{t}}, z_{i\hat{t}})$ denote the state vector that i was in when he last re-optimized his life insurance. We assume that the flow utility individual i obtains from continuing the existing policies purchased when his state vector was $(\hat{x}_{it}, \hat{z}_{it})$ is given by

$$u_2(x_{it}, z_{it}, \hat{x}_{it}, \hat{z}_{it}, \varepsilon_{2it}) = u_1(x_{it}, z_{it}) - c((x_{it}, z_{it}), (\hat{x}_{it}, \hat{z}_{it})) + \varepsilon_{2it}, \quad (3)$$

where $c(\cdot, \cdot)$ can be considered as a *sub-optimality penalty function*, which may also include adjustment costs (see discussion below in Section 3.1), that depends on the “distance” between the current state (x_{it}, z_{it}) and the state in which the existing policy was purchased $(\hat{x}_{it}, \hat{z}_{it})$. The adjustment cost can be positive or negative, depending on the factors that have changed. For example, if the individual was married when he purchased the existing policy but is not married now, then, all other things equal, the adjustment cost is likely to be negative; he would have less incentive to keep the existing policy. On the other hand, if the individual’s health has deteriorated substantially, then obtaining a new policy could be prohibitively costly, in which case the adjustment cost is likely to be positive; he would have more incentive to keep the existing policy which was purchased during a healthier state.

To summarize, we model the life insurance choice as the decision to either: (1). hold no life insurance; (2). purchase a new insurance policy which is optimal for the current state; or (3).

¹⁰We present the model here assuming the presence of both the observed and unobserved state variables. In Section 4, we will also estimate a model with only observed state variables. In that case, we simply ignore the unobserved state vector z_{it} .

continue with an existing policy. By decomposing the ownership decision into continuation vs. re-optimization, our model is able to capture the intuition that an individual who has suffered a negative shock to a factor that positively affects life insurance ownership (such as income or bequest motive) may still be likely to keep his insurance if the policy was initially purchased a long time ago during a better health state.

Moreover, the decomposition of the ownership decision allows us to examine two separate motives for lapsation: lapsation because the individual no longer needs *any* life insurance, and lapsation because the policyholder's personal situation, i.e. (x_{it}, z_{it}) , has changed such that new coverage terms are required.

Parametric Assumptions on u_1 and c Functions. In our empirical implementation of the model, we let the observed state vector x_{it} include age, log household income, sum of the number of health conditions, marital status, an indicator for whether the individual has children, and the age of the youngest child. We let the unobserved state vector z_{it} include z_{1it} , z_{2it} and z_{3it} which respectively represent the unobserved components of income, health and bequest motives. In Section 5 below, we will describe how we anchor these unobservables to their intended interpretations and how we use sequential Monte Carlo methods to simulate their posterior distributions.

The primitives of our model are thus given by the utility function of optimally purchasing life insurance u_1 , and the sub-optimality adjustment function c . In our empirical analysis we adopt the following parametric specifications for $u_1(x_{it}, z_{it})$ and $c((x_{it}, z_{it}), (\hat{x}_{it}, \hat{z}_{it}))$:

$$\begin{aligned}
u_1(x_{it}, z_{it}) &= \theta_0 + \theta_1 \text{AGE}_{it} + \theta_2 (\text{LOGINCOME}_{it} + z_{1it}) + \theta_3 (\text{CONDITIONS}_{it} + z_{1it}) \\
&\quad + \theta_4 (\text{MARRIED}_{it} + z_{1it}) + \theta_5 \text{AGE}_{it}^2 + \theta_6 (\text{LOGINCOME}_{it} + z_{1it})^2 \\
&\quad + \theta_7 (\text{CONDITIONS}_{it} + z_{2it})^2 + \theta_8 (\text{MARRIED}_{it} + z_{3it})^2 + \\
&\quad + \theta_9 \text{HAS CHILDREN}_i + \theta_{10} \text{HAS CHILDREN}_i \times \text{AGE OF YOUNGEST CHILD}_{it}; \quad (4)
\end{aligned}$$

$$\begin{aligned}
c((x_{it}, z_{it}), (\hat{x}_{it}, \hat{z}_{it})) &= \theta_{11} + \theta_{12} (\text{AGE}_{it} - \widehat{\text{AGE}}_{it}) + \theta_{13} (\text{AGE}_{it} - \widehat{\text{AGE}}_{it})^2 \\
&\quad + \theta_{14} (\text{LOGINCOME}_{it} + z_{1it} - \widehat{\text{LOGINCOME}}_{it} - \hat{z}_{1it}) \\
&\quad + \theta_{15} (\text{LOGINCOME}_{it} + z_{1it} - \widehat{\text{LOGINCOME}}_{it} - \hat{z}_{1it})^2 \\
&\quad + \theta_{16} (\text{CONDITIONS}_{it} + z_{2it} - \widehat{\text{CONDITIONS}}_{it} - \hat{z}_{2it}) \\
&\quad + \theta_{17} (\text{CONDITIONS}_{it} + z_{2it} - \widehat{\text{CONDITIONS}}_{it} - \hat{z}_{2it})^2 \\
&\quad + \theta_{18} (\text{MARRIED}_{it} + z_{3it} - \widehat{\text{MARRIED}}_{it} - \hat{z}_{3it}) \\
&\quad + \theta_{19} (\text{MARRIED}_{it} + z_{3it} - \widehat{\text{MARRIED}}_{it} - \hat{z}_{3it})^2. \quad (5)
\end{aligned}$$

In Section 3.1 below, we will provide an interpretation of the above-specified $c(\cdot, \cdot)$ function as

a sub-optimality penalty function.

Transitions of the State Variables. The state variables which an individual must keep track of depend on whether the individual is currently a policyholder. If he currently does not own a policy, his state variable is simply his current state vector (x_{it}, z_{it}) ; if he currently owns a policy, then his state variables include both his *current state vector* (x_{it}, z_{it}) and the state vector $(\hat{x}_{it}, \hat{z}_{it})$ at which he purchased the policy he currently owns.

In our empirical analysis, we assume that the current state vectors (x_{it}, z_{it}) evolve exogenously (i.e., not affected by the individual's decision) according to a joint distribution given by

$$(x_{it+1}, z_{it+1}) \sim f(x_{it+1}, z_{it+1} | x_{it}, z_{it}).$$

In particular, for the observed state vector x_{it} , which includes age, log household income, sum of the number of health conditions, and their respective squares, marital status, whether the individual has children, and the age of the youngest child, we estimate their evolutions directly from the data. For the unobserved state vector z_{it} , we will use sequential Monte Carlo methods to simulate its evolution (see Section 5.2 below for details).

The evolution of the state vector $(\hat{x}_{it}, \hat{z}_{it})$ is endogenous, and it is given as follows. If the individual does not own life insurance at period t , which we denote by setting $(\hat{x}_{it}, \hat{z}_{it}) = \emptyset$, then

$$[(\hat{x}_{it+1}, \hat{z}_{it+1}) | (\hat{x}_{it}, \hat{z}_{it}) = \emptyset] = \begin{cases} (x_{it}, z_{it}) & \text{if } d_{it} = 1 \\ \emptyset & \text{if } d_{it} = 0 \end{cases} \quad (6)$$

where \emptyset denotes that the individual remains with no life insurance. If the individual owns life insurance at period t purchased at state $(\hat{x}_{it}, \hat{z}_{it})$, then

$$[(\hat{x}_{it+1}, \hat{z}_{it+1}) | (\hat{x}_{it}, \hat{z}_{it}) \neq \emptyset] = \begin{cases} \emptyset & \text{if } d_{it} = 0 \\ (x_{it}, z_{it}) & \text{if } d_{it} = 1 \\ (\hat{x}_{it}, \hat{z}_{it}) & \text{if } d_{it} = 2. \end{cases} \quad (7)$$

3.1 Discussion

Dynamic Discrete Choice Model without the Knowledge of the Choice and Choice Set. As we mentioned in Section 2, we do not have complete information about the exact life insurance policies owned by the individuals, and for those whose life insurance policies we do know about, we do not know their choice sets. However, we do know whether an individual has re-optimized

his life insurance policy holdings (i.e., i.e. purchased a new policy or lapsed an existing one), or has dropped all coverage.

In fact the data restrictions we face are fairly typical for many applications.¹¹ For example, in the study of housing market, it is possible that all we observe is whether a family moved to a new house, remained in the same house, or decided to rent; but we may not observe the characteristics (including the purchase price) of the new house the family moved into, or the characteristics of the house/apartment the family rented; and most likely, we are not able to observe the set of houses or apartments the family has considered purchasing or renting (see, e.g., [Kung \(2012\)](#)).

Our formulation provides an *indirect utility approach* to deal with such data limitations. Suppose that when individual i 's state vector is (x_{it}, z_{it}) , he has a choice set $\mathcal{L}(x_{it}, z_{it})$ which includes all the possible life insurance policies that he could choose from. Note that $\mathcal{L}(x_{it}, z_{it})$ depends on i 's state vector (x_{it}, z_{it}) , which captures the notion that life insurance premium and face amount typically depend on at least some of the characteristics of the insured. Let $\ell \in \mathcal{L}(x_{it}, z_{it})$ denote one such available policy. Let $u^*(\ell; x_{it}, z_{it})$ denote individual i 's *primitive* flow utility from purchasing policy ℓ . If he were to choose to own a life insurance, his choice of the life insurance contract from his available choice set will be determined by the solution to the following problem:

$$V(x_{it}, z_{it}) = \max_{\ell \in \mathcal{L}(x_{it}, z_{it})} \{u^*(\ell; x_{it}, z_{it}) + \varepsilon_{lit} + \beta E[V(x_{it+1}, z_{it+1}) | \ell, x_{it}, z_{it}]\}. \quad (8)$$

Let $\ell^*(x_{it}, z_{it})$ denote the solution. Then the flow utility $u_1(x_{it}, z_{it})$ we specified in (2) can be interpreted as the *indirect* flow utility, i.e.,

$$u_1(x_{it}, z_{it}) = u^*(\ell^*(x_{it}, z_{it}); x_{it}, z_{it}). \quad (9)$$

It should be pointed out that, under the above indirect flow utility interpretation of $u_1(x_{it}, z_{it})$, in order for the error term ε_{lit} in (2) to be distributed as i.i.d extreme value as assumed, we need to make the assumption that ε_{lit} in (8) does *not* vary across $\ell \in \mathcal{L}(x_{it}, z_{it})$. This seems to be a plausible assumption.

Interpretations of the Sub-Optimality Penalty Function $c(\cdot)$. The sub-optimality penalty function $c(\cdot)$ we introduced in (3) permits the interpretation that changing an existing life insurance policy may incur adjustment costs. To see this, consider an individual whose current state vector

¹¹[McFadden \(1978\)](#) and [Fox \(2007\)](#) studied problems where the researcher only observes the choices of decision-makers from a subset of choices. [McFadden \(1978\)](#) showed that in a class of discrete-choice models where choice specific error terms have a block additive generalized extreme value (GEV) distributions, the standard maximum likelihood estimator remains consistent. [Fox \(2007\)](#) proposed using semiparametric multinomial maximum-score estimator when estimation uses data on a subset of the choices available to agents in the data-generating process, thus relaxing the distributional assumptions on the error term required for [McFadden \(1978\)](#).

is (x_{it}, z_{it}) and owns a life insurance policy he purchased at \hat{t} when his state vector was $(x_{i\hat{t}}, z_{i\hat{t}})$. Suppose that he decides to change (lapse or modify) his current policy and re-optimize, but there is an adjustment cost of $\kappa > 0$ for changing. Thus, the flow utility from lapsing into no coverage for this individual will be

$$u_0(x_{it}, z_{it}) = 0.$$

The flow utility from re-optimizing, using the notation from (9), will be

$$u_1(x_{it}, z_{it}) = u^*(\ell^*(x_{it}, z_{it}); x_{it}, z_{it}) - \kappa.$$

And the flow utility from keeping the existing policy is

$$\begin{aligned} u_2(x_{it}, z_{it}, \hat{x}_{it}, \hat{z}_{it}) &= \overbrace{u^*(\ell^*(\hat{x}_{it}, \hat{z}_{it}); x_{it}, z_{it})}^{u_1(x_{it}, z_{it})} - \kappa \\ &\quad - \left[\overbrace{u^*(\ell^*(x_{it}, z_{it}); x_{it}, z_{it}) - u^*(\ell^*(\hat{x}_{it}, \hat{z}_{it}); x_{it}, z_{it})}^{\text{sub-optimality penalty}} - \kappa \right] \\ &= u_1(x_{it}, z_{it}) - c((x_{it}, z_{it}), (\hat{x}_{it}, \hat{z}_{it})) \end{aligned} \quad (10)$$

where

$$c((x_{it}, z_{it}), (\hat{x}_{it}, \hat{z}_{it})) \equiv [u^*(\ell^*(x_{it}, z_{it}); x_{it}, z_{it}) - u^*(\ell^*(\hat{x}_{it}, \hat{z}_{it}); x_{it}, z_{it})] - \kappa.$$

Note that in the above expression for $c((x_{it}, z_{it}), (\hat{x}_{it}, \hat{z}_{it}))$, the term in the square bracket is the difference between the flow utility the individual could have received by purchasing the optimal policy for his current state (i.e. $\ell^*(x_{it}, z_{it})$), and the flow utility he receives from his existing policy, $\ell^*(\hat{x}_{it}, \hat{z}_{it})$, which he purchased when his state was $(\hat{x}_{it}, \hat{z}_{it})$. The sub-optimality penalty therefore measures the utility loss from holding a policy $\ell^*(\hat{x}_{it}, \hat{z}_{it})$ that was optimal in state $(\hat{x}_{it}, \hat{z}_{it})$, but sub-optimal when state vector is (x_{it}, z_{it}) . But by not re-optimizing, the individual saves the adjustment cost κ . Given the presence of adjustment cost κ , we would expect that an existing policyholders will hold on to his policy until the sub-optimality penalty $[u^*(\ell^*(x_{i\hat{t}}, z_{i\hat{t}}); x_{it}, z_{it}) - u^*(\ell^*(x_{it}, z_{it}); x_{it}, z_{it})]$ exceeds the adjustment cost κ , if we ignore decisions driven by i.i.d preference shocks ε_{1it} and ε_{2it} .

It is clear from the above discussion that, in this formulation, we can also allow the adjustment cost κ to be made a function of (x_{it}, z_{it}) , though we will not be able to separate the sub-optimality penalty $[u^*(\ell^*(x_{i\hat{t}}, z_{i\hat{t}}); x_{it}, z_{it}) - u^*(\ell^*(x_{it}, z_{it}); x_{it}, z_{it})]$ from $\kappa(x_{it}, z_{it})$.¹²

¹²If the adjustment cost κ is incurred both when the individual lapses into no coverage, and when he re-optimizes, i.e., if $u_0(x_{it}, z_{it}) = -\kappa$, and $u_1(x_{it}, z_{it}) = u^*(\ell^*(x_{it}, z_{it}); x_{it}, z_{it}) - \kappa$, then we can allow κ to depend on both (x_{it}, z_{it}) and $(x_{i\hat{t}}, z_{i\hat{t}})$.

It is also worth pointing out that our parametric specifications of $u_1(x_{it}, z_{it})$ and $c((x_{it}, z_{it}), (\hat{x}_{it}, \hat{z}_{it}))$, as given in (4) and (5) respectively, are consistent with the above interpretations of the sub-optimality penalty function.¹³

Limitations of the “Indirect Flow Utility” Approach. In this paper we adopt the “indirect flow utility” approach to deal with the lack of information regarding individuals’ actual choices of life insurance policies and their relevant choice set. This is useful for our purpose of understanding why policyholders lapse their coverage (as we will demonstrate later), but it comes with a limitation. The indirect flow utilities $u_1(x_{it}, z_{it})$ and $u_2(x_{it}, z_{it}, \hat{x}_{it}, \hat{z}_{it})$, defined in (9) and (10) respectively, are derived only under the *existing life insurance market structure*. As a result, the estimated indirect flow utility functions are *not* primitives that are invariant to counterfactual policy changes that may affect the equilibrium of the life insurance market. Of course, this limitation is also present in other dynamic discrete choice models where the flow utility functions can have the interpretation as reduced-form, indirect utility function of a more detailed choice problem.¹⁴

4 Estimates from a Dynamic Discrete Choice Model Without Unobservable State Variables

In this section, we present our estimation and simulation results for the dynamic structural model of life insurance decisions presented in Section 3. However, in order to illustrate the importance of unobserved state variables in the life insurance decisions (which we turn to in the next section), we deliberately do not include any unobserved state variables z_{it} in this section. The estimation and simulation results for a dynamic discrete choice model with unobserved state variables are presented in Section 5.

As described in Section (3) the flow utilities are given by equations (1)-(3). Since we only include observed state variables in this section, we will for simplicity denote the transition distributions of the state variables x_{it} by $P(x_{it}|x_{it-1})$. Since x_{it} are observable, we estimate $P(x_{it}|x_{it-1})$ separately and then take it as given. As we mentioned in Section 3, we assume that the evolution of the state variables x_{it} does not depend on the life insurance choices analyzed in this model. However, the evolution of the “hatted” state variables does depend on the choices. Specifically, $\hat{x}_{it+1} = x_{it}$ if $d_{it} = 1$; $\hat{x}_{it+1} = \hat{x}_{it}$ if $d_{it} = 2$; and $\hat{x}_{it+1} = \emptyset$ if $d_{it} = 0$. As usual we use $\hat{x}_{it} = \emptyset$

¹³Due to the ages of the individuals in our estimation sample, we have practically no changes in the number of children. Thus the difference between AGE OF YOUNGEST CHILD_{it} and AGE OF YOUNGEST CHILD_{iit} is essentially the same as the difference between AGE_{it} and AGE_{iit}.

¹⁴For example, in many I.O. papers a reduced-form flow profit function is assumed. Presumably, the profit function is not invariant to changes in the market structure.

to denote an individual who does not own life insurance at the beginning of period t . Finally, we assume that the *yearly* discount factor is 0.9, and thus the per period (two years) discount factor in our model is $\beta = 0.81$, and we assume that the time horizon is finite. We choose age 80 as the last year in the decision horizon, because that is the oldest age in our data set. Thus, an individual of age 80 chooses myopically according to $u_1(x_{it})$ and $u_2(x_{it}, \hat{x}_{it})$.

At period t , let $V_{0t}(x_{it})$ be the present value from choosing $d_{it} = 0$ (no life insurance); and let $V_{1t}(x_{it})$ be the present value from choosing $d_{it} = 1$ (re-optimize), and let $V_{2t}(x_{it}, \hat{x}_{it})$ be the present value to choosing $d_{it} = 2$ (keep existing policy) for those who owned policies previously purchased at state \hat{x}_{it} . To derive these choice-specific value functions, it is useful to first derive the inclusive continuation values from being in a give state vector. Let $V_t(x_{it}, \hat{x}_{it})$ denote the period- t inclusive value for being in state x_{it} and having an existing policy purchased when the state vector is \hat{x}_{it} , and let $W_t(x_{it})$ denote the period- t inclusive value for being in state x_{it} and not having any existing life insurance. Under the assumption of additively separable choice specific shocks drawn from i.i.d. Type-1 extreme value distributions and using $G(\cdot)$ to denote the joint distribution of the random vector $\varepsilon_t \equiv (\varepsilon_{1t}, \varepsilon_{2t})$, $V_t(x_{it}, \hat{x}_{it})$ and $W_t(x_{it})$ can be, following Rust (1994), expressed as:

$$\begin{aligned} V_t(x_{it}, \hat{x}_{it}) &= \int \max \{V_{0t}(x_{it}), V_{1t}(x_{it}) + \varepsilon_{1t}, V_{2t}(x_{it}, \hat{x}_{it}) + \varepsilon_{2t}\} dG(\varepsilon_t) \\ &= \log \{ \exp [V_{0t}(x_{it})] + \exp [V_{1t}(x_{it})] + \exp [V_{2t}(x_{it}, \hat{x}_{it})] \} + 0.57722, \end{aligned} \quad (11)$$

$$\begin{aligned} W_t(x_{it}) &= \int \max \{V_{0t}(x_{it}), V_{1t}(x_{it}) + \varepsilon_{1t}\} dG(\varepsilon) \\ &= \log \{ \exp [V_{0t}(x_{it})] + \exp [V_{1t}(x_{it})] \} + 0.57722, \end{aligned} \quad (12)$$

where 0.57722 is the Euler constant. Then, the choice-specific present value functions can be written as follows:

$$V_{0t}(x_{it}) = \beta \int W_{t+1}(x_{it+1}) dP(x_{it+1}|x_{it}), \quad (13)$$

$$V_{1t}(x_{it}) = u_1(x_{it}) + \beta \int V_{t+1}(x_{it+1}, x_{it}) dP(x_{it+1}|x_{it}), \quad (14)$$

$$V_{2t}(x_{it}, \hat{x}_{it}) = u_2(x_{it}, \hat{x}_{it}) + \beta \int V_{t+1}(x_{it+1}, \hat{x}_{it}) dP(x_{it+1}|x_{it}). \quad (15)$$

Since we assume that age 80 is the final period, we have $V_{0,80}(x_{it}) = 0$, $V_{1,80}(x_{it}) = u_1(x_{it})$ and $V_{2,80}(x_{it}, \hat{x}_{it}) = u_2(x_{it}, \hat{x}_{it})$. Using this, we can solve for the choice-specific value functions at each age through backward recursion.

The choice probabilities at each period t are then given as follows. For individuals without life

insurance in the beginning of period t , their choice probabilities for $d_{it} \in \{0, 1\}$ are given by:

$$\Pr \{d_{it} = 0 | x_{it}, \hat{x}_{it} = \emptyset\} = \frac{\exp [V_{0t}(x_{it})]}{\exp [V_{0t}(x_{it})] + \exp [V_{1t}(x_{it})]},$$

$$\Pr \{d_{it} = 1 | x_{it}, \hat{x}_{it} = \emptyset\} = \frac{\exp [V_{1t}(x_{it})]}{\exp [V_{0t}(x_{it})] + \exp [V_{1t}(x_{it})]}.$$

For individuals who own life insurance in the beginning of period t , which are purchased in previous waves when state vector is \hat{x}_{it} , their choice probabilities for $d_{it} \in \{0, 1, 2\}$ are given by:

$$\Pr \{d_{it} = 0 | x_{it}, \hat{x}_{it} \neq \emptyset\} = \frac{\exp [V_{0t}(x_{it})]}{\exp [V_{0t}(x_{it})] + \exp [V_{1t}(x_{it})] + \exp [V_{2t}(x_{it}, \hat{x}_{it})]},$$

$$\Pr \{d_{it} = 1 | x_{it}, \hat{x}_{it} \neq \emptyset\} = \frac{\exp [V_{1t}(x_{it})]}{\exp [V_{0t}(x_{it})] + \exp [V_{1t}(x_{it})] + \exp [V_{2t}(x_{it}, \hat{x}_{it})]},$$

$$\Pr \{d_{it} = 2 | x_{it}, \hat{x}_{it} \neq \emptyset\} = \frac{\exp [V_{2t}(x_{it}, \hat{x}_{it})]}{\exp [V_{0t}(x_{it})] + \exp [V_{1t}(x_{it})] + \exp [V_{2t}(x_{it}, \hat{x}_{it})]}.$$

We estimate the parameters using maximum likelihood. A simulation and interpolation method is used to compute and then integrate out the inclusive value terms. The numerical solution method we employ closely follows [Keane and Wolpin \(1994\)](#). Among the state variables, two of them are allowed to be continuous, namely the current log income and the log income when the last re-optimization of life insurance occurred; the other state variables are discrete. But the size of the state space, not including log incomes, is still very large.¹⁵ We thus use Keane and Wolpin's method for approximating the expected continuation values using only a subset of the state space.¹⁶

4.1 Estimation Results

Table 7 presents the coefficient estimates for $u_1(\cdot)$ and $c(\cdot)$ for the dynamic discrete choice model without serially correlated unobservable state variables. The estimated coefficients for the function $u_1(\cdot)$ as specified in (4) are reported in Panel A. The estimates indicate that married individuals have more to gain from optimally purchasing new insurance, whereas older and less healthy individuals have less to gain. This is consistent with the interpretation that the cost of re-optimizing one's life insurance increases when one gets older and has poorer health, whereas

¹⁵The state variable "conditions" is the number of health conditions ever diagnosed, where the health conditions used are: 1. high blood pressure; 2. diabetes; 3. cancer; 4. lung disease; 5. heart disease; 6. stroke; 7. psychological problem; 8. arthritis. Each of these 8 health conditions was carried around as a binary state variable (1 or 0) and the transitions for each of these were estimated separately. So, there were 2^8 possible combinations of health conditions, but only 9 (0 through 8) possible values for "CONDITIONS" and "CONDITIONS".

¹⁶The subset we use for interpolation consists of 400 randomly drawn points in the state space. For the numerical integration over the state space, 40 random draws from the state space were used.

Table 7: Estimation Results from Dynamic Model without Serially-Correlated Unobservables

	Estimate	Std. Error
Panel A: Coefficients for $u_1(x_{it})$		
Constant (θ_0)	-2.1297***	0.0711
Age (θ_1)	-0.1026***	0.0013
Logincome (θ_2)	-0.0601***	0.0102
Conditions (θ_3)	-0.0380***	0.0052
Married (θ_4)	0.1072***	0.0112
Age ² (θ_5)	0.0006***	0.0000
Logincome ² (θ_6)	0.0126***	0.0007
Conditions ² (θ_7)	-0.0005	0.0007
Has children (θ_9)	0.0385	0.0236
Has children \times Age of youngest child (θ_{10})	0.0040***	0.0004
Panel B: Coefficients for $c(x_{it}, \hat{x}_{it})$		
Constant (θ_{11})	1.6679***	0.0214
Δ Age (θ_{12})	0.2827***	0.0084
$(\Delta$ Age) ² (θ_{13})	-0.0138***	0.0006
Δ Logincome (θ_{14})	0.0074	0.0041
$(\Delta$ logincome) ² (θ_{15})	0.0018**	0.0006
Δ Conditions (θ_{16})	0.0177	0.0134
$(\Delta$ Conditions) ² (θ_{17})	0.0077	0.0051
Δ Married (θ_{18})	-0.0346	0.0272
$(\Delta$ Married) ² (θ_{19})	-0.2396***	0.0320
Log likelihood	-9,338.59	

Notes: (1). The specifications for $u_1(x_{it})$ and $c(x_{it}, \hat{x}_{it})$ are given in 4 and 5 respectively.

(2). For any variable x , Δx is the difference between the current value of x and \hat{x} , which is the value of x at the time when the respondent changed his coverage.

(3). Married² is not included in Panel A because without unobservables, it is perfectly correlated with Married.

(4). The annual discount factor β is set at 0.9.

(5). *, **, and *** respectively represent significance at 10%, 5% and 1% levels.

marriage is a bequest factor that leads to the purchase of life insurance. Log income enters $u_1(\cdot)$ nonlinearly, but in the most relevant income ranges in our data, individuals with higher incomes are more likely to purchase life insurance. This is partly because higher income makes any given amount of coverage more affordable and partly because individuals with higher incomes would like to leave more bequests to their dependents. Panel A also shows that individuals with children have higher utility from re-optimizing though the effect is not statistically significant.

Panel B of Table 7 reports the estimated coefficients for the function $c(\cdot)$ as specified in (5). The positive coefficient estimate on ΔAGE and the negative coefficient estimate on $(\Delta\text{AGE})^2$ indicate that the effect of ΔAGE – the difference in the current age and the age when the existing policy was purchased – on $c(\cdot)$ is nonlinear: when ΔAGE is small, the sub-optimality penalty is high, or in other words, it is less costly for the policyholder to re-optimize; but as ΔAGE gets larger, the cost of re-optimizing increases. The positive coefficient estimates on $\Delta\log\text{INCOME}$ and $(\Delta\log\text{INCOME})^2$ indicate that policyholders are more likely to stay with their existing policy when income has declined, but more likely to re-optimize when income has risen.

We find that in this model without serially-correlated unobservables the coefficient estimates on both $\Delta\text{CONDITIONS}$ and $(\Delta\text{CONDITIONS})^2$ are positive, which suggests that as the policyholder becomes less healthy, the sub-optimality penalty of staying with the existing policy increases, and individuals would be more likely to re-optimize, though the effect is not statistically significant. The finding that the less healthy policyholders are more likely to re-optimize is at odds with data. Indeed He (2010) found that less healthy policyholders are more likely to keep existing policies, relative to healthier policyholders, which is consistent with the predictions of adverse selection in lapsation decisions in Hendel and Lizzeri (2003). We will show in Table 9a below that once we introduce serially correlated unobservables in our model, we find that individuals who experience deterioration in health are more likely to re-optimize, thus illustrating the importance of including serially correlated unobservables.

The estimated coefficients for both $\Delta\text{MARRIED}$ and $(\Delta\text{MARRIED})^2$ are negative, which indicates that policyholders who experience changes in marital status are more likely to keep the existing policy. This is again not consistent with data; we will again show in Table 9a below that once we introduce serially unobserved state variables in our model, we will find that individuals who experience changes in bequest motives are more likely to re-optimize.

4.2 Model Fit

To assess the performance of the dynamic model without unobservable state variables, we report in Panel A of Table 8 the comparisons between the simulated model predictions regarding aggregate choice probabilities by wave and those in the data. It shows that the dynamic model

Table 8: Model Fit for Dynamic Model without Serially Correlated Unobservable State Variables

	Wave					
	1996	1998	2000	2002	2004	2006
Panel A: Aggregate Choice Probabilities by Wave						
<i>Actual Data</i>						
No life insurance coverage	0.1192	0.1422	0.1642	0.1889	0.1948	0.2136
Covered, but changed or bought new coverage	0.1493	0.1077	0.0964	0.0978	0.1123	0.0943
Covered, and kept existing coverage	0.7315	0.7500	0.7394	0.7132	0.6928	0.6921
<i>Simulation using dynamic model without serially correlated unobservables</i>						
No life insurance coverage	0.1764	0.2012	0.2073	0.2155	0.2277	0.2474
Covered, but changed or bought new coverage	0.1421	0.1179	0.1125	0.1058	0.1018	0.0992
Covered, and kept existing coverage	0.6813	0.6808	0.6801	0.6786	0.6704	0.6532
Panel B: Cumulative Outcomes for 1994 Policyholders						
<i>Actual data</i>						
Lapsed to no life insurance	0.0455	0.0899	0.1358	0.1792	0.2095	0.2340
Changed coverage amount	0.0903	0.1472	0.1846	0.2080	0.2283	0.2489
Kept 1994 coverage	0.8642	0.7315	0.6138	0.5203	0.4491	0.3812
Policyholder died	0.0000	0.0313	0.0658	0.0924	0.1131	0.1358
<i>Simulation using dynamic model without serially correlated unobservables</i>						
Lapsed to no life insurance	0.1048	0.1822	0.2341	0.2745	0.3091	0.3420
Changed coverage amount	0.0901	0.1389	0.1690	0.1902	0.2066	0.2209
Kept 1994 coverage	0.8049	0.6497	0.5392	0.4524	0.3842	0.3200
Policyholder died	0.0000	0.0290	0.0575	0.0827	0.1000	0.1169

without serially correlated unobservable state variables does a fairly good job at predicting the aggregate distribution of choices, but there appears to be a dynamic effect which our model is not capturing. Specifically, the actual data exhibits that, in the aggregate, there is a sharp increase in the likelihood over time of holding no life insurance. This pattern does not appear to be fully captured by our model. The simulation of the model predicts that the fraction of individuals without life insurance would increase from 17.64% in 1996 to 24.74% in 2006 (in contrast to an increase of 11.92% in 1996 to 21.36% in 2006 in the data).

In Panel B of Table 8, we report the comparison between the simulated model predictions and the data of the cumulative decision outcomes by wave for individuals who owned life insurance in 1994. In the actual data, the cumulative fraction of 1994 policyholders lapsing to no insurance steadily increases over time, going from 4.55% in 1996 to 23.40% by 2006. The model simulation is not able to replicate the initially low lapsation rates. The model also under-predicts by a large margin the cumulative fraction of the 1994 policyholders who kept their policies.¹⁷

5 Estimates from Dynamic Discrete Choice Model with Unobserved State Variables

As we showed in the previous section, the dynamic model without serially correlated unobservable state variables fails to match the dynamic persistence in both the aggregate choice probabilities and cumulative outcomes. We suspect that some of these idiosyncratic factors may not be so idiosyncratic at all. In particular, if there are serially correlated unobservables, such as unobserved components of bequest motive or liquidity shocks, then our model will fall short in explaining the empirical determinants of lapsation. Indeed, our simulations seem to over-forecast lapsation due to no longer needing life insurance and under-forecast lapsation due to re-optimization, indicating the presence of some persistent unobservable that positively influences the likelihood of needing life insurance. In this section, we go back to the empirical framework we presented earlier in Section 3, which explicitly takes into account the presence of unobserved state variables. In particular, we add three unobserved state variables: z_1 , z_2 and z_3 , that are meant to represent the serially-correlated unobservable components of income, health and bequest motive, respectively.

¹⁷In un-reported counterfactual simulations we find that elimination of marital shocks has a larger effect on lapsation patterns than elimination of income shock. Elimination of marital shocks reduces the share of policyholders lapsing to no insurance, as well as the share of policyholders which are re-optimized. Elimination of income shocks reduces the share of policies lapsing to none, but increases the share of policies which are re-optimized. The increase in re-optimization is due to the fact that income is generally going down for individuals of our age range.

However, our counterfactual simulations also show that elimination of either of these shocks has an economically insignificant impact on the aggregate lapsation patterns. Most of the lapsation in our data seems to be driven by unobserved, idiosyncratic factors. This is not surprising given that marital status and income are not very volatile in our data.

Below we first describe how we anchor the interpretations of these unobservable state variables.

5.1 Anchoring the Unobserved State Variables

In this specification, we would like to give the unobserved state variable z_1 the interpretation of an *unobserved liquidity (or income) shock*, and normalize its unit to the same as log income, z_2 the interpretation of an *unobserved health shock* that is normalized to the units of health conditions, and finally z_3 the interpretation as an *unobserved component of bequest motive* that is normalized to the units of marital status.

We assume that the initial distribution in 1994 (which we set to be $t = 0$) for each of these unobserved variables is degenerate and given by:¹⁸

$$z_{1i0} = \theta_{20}h_{i0}, \quad (16)$$

$$z_{2i0} = \theta_{21}h_{i0}, \quad (17)$$

$$z_{3i0} = \theta_{22}h_{i0}, \quad (18)$$

where h_{i0} is an indicator dummy for whether the individual reported owning life insurance in 1994. In order to anchor z_1, z_2 and z_3 to have the desired interpretation given above, we expect, but do not restrict, that the coefficients θ_{20}, θ_{21} and θ_{22} to be of certain signs. For example, because income is a positive factor for life insurance ownership, we expect that the sign of θ_{20} to be positive, so that individuals who owned life insurance in the initial period also have higher z_1 . This anchors the interpretation of z_1 as an unobserved component of income. Similarly, because being married is also a positive factor for life insurance ownership, we also expect the sign of θ_{22} to be positive, so that individuals who owned life insurance in the initial period also have higher values of unobserved bequest motive z_3 .

The second channel that anchors the unobserved state variables to having the desired interpretation is incorporated in our specifications for $u_1(\cdot)$ and $c(\cdot)$, as formulated in (4) and (5). Note that we restricted z_{1i} to entering both $u_1(\cdot)$ and $c(\cdot)$ in the same way as LOGINCOME, z_{2i} the same way as CONDITIONS, and z_{3i} the same way as MARRIED. These restrictions, together with the sequential Monte Carlo method (described in the next section below) we use to simulate the distributions of $\mathbf{z}_t \equiv (z_{1t}, z_{2t}, z_{3t})$, ensures that the unobserved variables \mathbf{z}_t have the desired interpretations.

If we know the distributions of the unobservable state vectors $(\mathbf{z}_t, \hat{\mathbf{z}}_t)$, solving this model is done no differently from that in section 4. Given a vector of parameters $\boldsymbol{\theta} = (\theta_0, \dots, \theta_{19})$, we can

¹⁸The assumption that the initial distribution of the unobservable state variables \mathbf{z}_0 is degenerate is for computational simplicity.

compute the value functions at each age through backward recursion. The difficulty of handling unobserved state variables comes during estimation, because we have to integrate over the unobservables when computing the likelihood. We now turn our attention to this problem.

5.2 Using Sequential Monte Carlo Method to Simulate the Distributions of the Unobserved State Variables

We use a Sequential Monte Carlo (SMC) method to simulate the distributions of the unobservable state vectors.¹⁹ SMC is a set of simulation-based methods which provides a convenient and attractive approach to computing the posterior distributions of non-Gaussian, non-linear, and high dimensional random variables.²⁰ A thorough discussion of the method, from both the theoretical and the practical perspectives, is available in [Doucet, de Freitas and Gordon \(2001\)](#). The SMC method has been widely used in fields such as speech recognition, biology, and physics, etc. Despite the obvious potential importance of serially correlated unobservable state variables, there are only few applications of SMC in the economics literature. [Fernandez-Villaverde and Rubio-Ramirez \(2007\)](#) used SMC for estimating macroeconomic dynamic stochastic general equilibrium models with serially correlated unobservable state variables using a likelihood approach. [Blevins \(2011\)](#) proposed the use of SMC to allow for serially correlated unobservable state variables in estimating dynamic single agent models and dynamic games. [Hong, Gallant and Khwaja \(2008\)](#) also discusses the method in an application to the pharmaceuticals industry. All of the papers allow for a single serially correlated unobservable state variable. In our application, we believe that there might be important serially correlated unobservable components for each of the three potential sources of lapsation, shocks to income, health and bequest motives.

Now we provide a detailed discussion about the SMC algorithm. For a given individual, we observe the sequence of choices $\{d_t\}_{t=0}^T$, observed state variables $\{x_t, \hat{x}_t\}_{t=0}^T$, and whether the individual had life insurance in 1994 h_0 . The data set is thus $\{d_t, x_t, \hat{x}_t, h_0\}_{t=0}^T$ (we have dropped the i subscript for convenience). Let $p(d_{0:T}|x_{0:T}, \hat{x}_{0:T}, h_0)$ denote the conditional likelihood of the observed data. We can write:

$$p(d_{0:T}|x_{0:T}, \hat{x}_{0:T}, h_0) = p(d_0|x_0, \hat{x}_0, h_0) \prod_{t=1}^T p(d_t|d_{t-1}, x_t, \hat{x}_t) \quad (19)$$

¹⁹SMC algorithms are also called bootstrap filters, particle filters, and sequential importance samplers with resampling.

²⁰SMC for non-linear, non-Gaussian models is the analog of Kalman filter for linear, Gaussian models. [Gordon, Salmond and Smith \(1993\)](#) is the seminal paper that proposed this algorithm, which they refer to as the *bootstrap filter*.

Because the initial distribution of \mathbf{z}_0 is degenerate and depends only on h_0 , we can write:

$$p(d_0|x_0, \hat{x}_0, h_0) = p(d_0|x_0, \hat{x}_0, \mathbf{z}_0).$$

Assuming we have solved for the value functions in a first stage, we should be able to compute $p(d_0|x_0, \hat{x}_0, \mathbf{z}_0)$.

Now, for each $t > 0$ we can write:

$$p(d_t|d_{t-1}, x_t, \hat{x}_t) = \int p(d_t|d_{t-1}, x_t, \hat{x}_t, \mathbf{z}_t, \hat{\mathbf{z}}_t) p(\mathbf{z}_t, \hat{\mathbf{z}}_t|d_{t-1}) d\mathbf{z}_t d\hat{\mathbf{z}}_t. \quad (20)$$

We know how to compute $p(d_t|d_{t-1}, x_t, \hat{x}_t, \mathbf{z}_t, \hat{\mathbf{z}}_t)$ for a given set of parameter values θ . What we need is a method to draw from $p(\mathbf{z}_t, \hat{\mathbf{z}}_t|d_{t-1})$; and we use the Sequential Monte Carlo method for this purpose.

SMC is a recursive algorithm that begins by drawing a swarm of particles approximating the initial distribution of the hidden state. The initial swarm is then used to draw a swarm for the next period, and this swarm is then filtered according to sequential importance weights.

The unobservables transition according to the following equations:

$$z_{1it} = \theta_{23} z_{1it-1} + \theta_{24} \epsilon_{z_1t}, \quad (21)$$

$$z_{2it} = \theta_{25} z_{2it-1} + \theta_{26} \epsilon_{z_2t}, \quad (22)$$

$$z_{3it} = \theta_{27} z_{3it-1} + \theta_{28} \epsilon_{z_3t}, \quad (23)$$

where ϵ_{z_1t} , ϵ_{z_2t} and ϵ_{z_3t} are i.i.d. random variables with standard normal distributions $\mathcal{N}(0, 1)$. The transition distribution of the observed state variables are given by $P(x_{it}|x_{it-1})$ as in section 4.²¹

The filtered particles are then used to draw another swarm for the next period, and so on. In the following notation, we will absorb $\hat{\mathbf{z}}$ into \mathbf{z} and use \mathbf{z} to denote any unobserved variable, including the ‘‘hatted’’ \mathbf{z} ’s.

The method proceeds as follows:

0. Set $t = 0$, draw a swarm of particles $\{\mathbf{z}_0^{(r)}\}_{r=1}^R$ from the initial distribution $p(\mathbf{z}_0)$. This distribution must be parametrically assumed, with potentially unknown parameters. In our case it is assumed to be degenerate as described by (16)-(18). Set $t = 1$.
1. For $t > 0$, use $\{z_{t-1}^{(r)}\}_{r=1}^R$ to draw a new swarm $\{\tilde{\mathbf{z}}_t^{(r)}\}_{r=1}^R$ from the distribution $p(\mathbf{z}_t|\mathbf{z}_{t-1}, d_{t-1})$. This distribution is known because we have imposed a parametric specification on it. The

²¹Thus, we do not allow the transitions of the observed state variables to depend on the realization of the unobserved state variables. This is only for simplicity and is not a limitation of the SMC algorithm.

previous period's choice, d_{t-1} , is required because that determines how the "hatted" \mathbf{z} 's evolve. The swarm of particles $\{\tilde{\mathbf{z}}_{t,r}\}_{r=0}^R$ now approximates the distribution $p(\mathbf{z}_t|d_{t-1})$.

2. For each $r = 1, \dots, R$, compute $w_t^{(r)} = p(d_t|d_{t-1}, x_t, \hat{x}_t, \tilde{\mathbf{z}}_{t,r})$. The vector $\{w_t^{(r)}\}_{r=1}^R$ is known as the vector of importance weights. We can now approximate the integral in (20) by $\frac{1}{R} \sum_{r=1}^R w_t^{(r)}$.
3. Draw a new swarm of particles $\{\tilde{\mathbf{z}}_t^{(r)}\}_{r=1}^R$ by drawing with replacement from $\{\tilde{\mathbf{z}}_t^{(r)}\}_{r=1}^R$. Use the normalized importance weights as sampling probabilities.
4. Set $t = t + 1$ and go to step 1.

Figure 1 presents a graphical representation of the SMC algorithm. The SMC starts at time $t - 1$ with an un-weighted measure $\{\tilde{\mathbf{z}}_{t-1}^{(r)}, R^{-1}\}$, which provides an approximation of $p(\mathbf{z}_{t-1}|d_{1:t-2})$. For each particle, the importance weight is computed using the information about the observed choice d_{t-1} . The importance weight is given by the model's prediction of the likelihood $p(d_{t-1}|\tilde{\mathbf{z}}_{t-1}^{(r)})$ of observing d_{t-1} when the particle is $\tilde{\mathbf{z}}_{t-1}^{(r)}$, properly re-normalized. This results in the weighted measure $\{\tilde{\mathbf{z}}_{t-1}^{(r)}, \tilde{w}_{t-1}^{(r)}\}$, which yields an approximation of $p(\mathbf{z}_{t-1}|d_{1:t-1})$. Subsequently, the re-sampling with replacement (or the selection) step selects only the fittest particles to obtain the un-weighted measure $\{\tilde{\mathbf{z}}_{t-1}^{(r)}, R^{-1}\}$, which is still an approximation of $p(\mathbf{z}_{t-1}|d_{1:t-1})$. Finally, the prediction step draws new varieties of particles from the parametric process $p(\mathbf{z}_t|\mathbf{z}_{t-1})$, resulting in the measure $\{\tilde{\mathbf{z}}_t^{(r)}, R^{-1}\}$, which is an approximation of $p(\mathbf{z}_t|d_{1:t-1})$. The measures $\{\tilde{\mathbf{z}}_{t-1}^{(r)}, R^{-1}\}$ and $\{\tilde{\mathbf{z}}_t^{(r)}, R^{-1}\}$ are the posterior distributions of the unobservables we use in the numerical integration of the choice probabilities (20).

At each iteration, we compute the per period probability of observing the data given by equation (20). Because this is done iteratively, starting from $t = 0$, we can eventually work our way up to $t = T$ and compute the entire likelihood given by (19). Repeating this process for each individual for the data gives us the entire likelihood of the data. We can then estimate the parameters via maximum likelihood.²²

We use simulated maximum likelihood to estimate the parameters of $u_1(\cdot)$, $c(\cdot)$, the initial values of z_1 , z_2 and z_3 , namely θ_{20} - θ_{22} as described in (16)-(18), as well as their AR(1) auto-correlation coefficients and variance terms θ_{23} - θ_{28} as described in (21)-(23). We compute the standard errors using a bootstrap procedure. In each iteration of the procedure, a new random seed is used to create a bootstrapped sample of individuals from the original roster. The structural parameters are

²²We employed 64 particles in each swarm to integrate out the \mathbf{z} 's when computing the conditional choice probabilities. We use 40 articles when computing the expected future value term. One evaluation of the likelihood takes about 2 minutes on an 8-core, 2.5GHz, 64 bit AMD computer while using all 8 CPUs, and the entire estimation routine took about 4 days when starting from an initial guess of all zeros.




Figure 1: A Graphical Representation of the Sequential Monte Carlo Algorithm: Adapted from Doucet, de Freitas and Gordon (2001, Chapter 1, p. 12)

then re-estimated using this bootstrapped sample. 50 bootstrapped samples were used. For each structural parameter, the standard error is calculated as the standard deviation of the estimates from the 50 bootstrapped samples.²³

5.3 Estimation Results

Tables 9a and 9b present the estimation results. Panel A shows the estimated coefficients for $u_1(x_{it}, z_{it})$ as specified in (4). There are several important changes in the estimated coefficients relative to those for the model without serially correlated unobservables as in Table 7. First, the estimated coefficient on (MARRIED + z_3) in Table 9a is negative, in contrast to that in Table 7; however, the estimated coefficient for the term (MARRIED + z_3)² is positive and statistically significant in 9a (recall that in Table 7 the term MARRIED² was perfectly correlated with MARRIED and thus had to be dropped). Thus the overall effect from the estimates in Table 9a is still that married individuals are more likely than divorced or widowed individuals to own life insurance. Second, in contrast to that in Table 7, with unobservable state variables the estimates of (LOGINCOME + z_1)

²³See Olsson and Rydén (2008) for a discussion about the asymptotic performance of approximate maximum likelihood estimators for state space models obtained via sequential Monte Carlo methods. It provides criteria for how to increase the number of particles and the resolution of the grid in order to produce estimates that are consistent and asymptotically normal.

and $(\text{LOGINCOME} + z_1)^2$ are both positive and significant, though the qualitative effect of Log income on the probability of purchasing life insurance is the same in the two models in the relevant ranges of income in our data. Third, the estimated coefficient for $(\text{CONDITIONS} + z_2)^2$ is now positive and significant in contrast to that in Table 7 where the estimate for CONDITIONS^2 was negative and insignificant. Except for these differences, the other parameter estimates are consistent across the two models.

Panel B in Table 9a presents the estimated coefficients for $c((x_{it}, z_{it}), (\hat{x}_{it}, \hat{z}_{it}))$ as specified in (5). There are several important changes in the estimates for $c(\cdot)$ when unobservable state variables are included. In contrast to the estimates in Table 7 for the model without serially correlated unobservables, the coefficient estimate for $\Delta(\text{CONDITIONS} + z_2)$ is now negative (-0.0443) and that for $(\Delta(\text{CONDITIONS} + z_2))^2$ is positive (0.0060), and both are statistically significant. This implies that individuals who experience negative health shocks tend to have lower values of the sub-optimality penalty $c(\cdot)$ and thus are more likely to keep the existing policy. As we mentioned in the previous section, this is consistent with the empirical findings in He (2010) and the theoretical predictions of Hendel and Lizzeri (2003).

Also importantly, once we incorporate serially-correlated unobservables, the coefficient estimates for both $\Delta(\text{MARRIED} + z_3)$ and $(\Delta(\text{MARRIED} + z_3))^2$ are now positive (0.0136 and 0.0284, respectively) and significant in $c(\cdot)$. This implies that changes in marital status and other unobservables related to bequest motives will increase the sub-optimality penalty and lead policyholders to adjust their life insurances by either lapsing into no coverage or re-optimizing their coverages. This is much more plausible than the implications from the estimates without unobservable state variables we discussed in the previous section. The other parameter estimates in Table 9a are similar to those in Table 7.

Panels C in Table 9b presents the estimated relationship, as specified in (16)-(18), between the initial distributions of the three unobservables z_{10} , z_{20} and z_{30} and the indicator of whether the individuals owned life insurance in 1994. The positive estimates for coefficients θ_{20} in (16) and θ_{22} in (18) indicate that those who owned life insurance policies in 1994 tend to have values of unobservable income and bequest motives; on the other hand, the negative estimate of coefficient θ_{21} in (17) indicates that the policyholders in 1994 tends to be healthier. This result is consistent with the findings in Cawley and Philipson (1999), and most likely reflects survivorship bias as explained in He (2009) and He (2010).

Panel D in Table 9b presents the estimates of the coefficients of the autoregressive processes described in (21)-(23). The estimates for coefficients θ_{23} , θ_{25} and θ_{27} are all positive and significant (both economically and statistically), suggesting that the unobservable income, health and bequest motives shocks are rather persistent, though there are sizeable variations in the unobservables,

Table 9a: Estimation Results from Dynamic Model with Serially-Correlated Unobservables

	Estimate	Std. Error
Panel A: Coefficients for $u_1(x_{it}, z_{it})$		
Constant (θ_0)	-2.4438***	0.0071
Age (θ_1)	-0.0134***	0.0001
Logincome + z_1 (θ_2)	0.0139***	0.0007
Conditions + z_2 (θ_3)	-0.0358***	0.0018
Married + z_3 (θ_4)	-1.1704***	0.0105
Age ² (θ_5)	0.0001***	0.0000
(Logincome + z_1) ² (θ_6)	0.0070***	0.0001
(Conditions + z_2) ² (θ_7)	0.0014***	0.0001
(Married + z_3) ² (θ_8)	1.1663***	0.0092
Has children (θ_9)	0.0374***	0.0022
Has children \times Age of youngest child (θ_{10})	0.0046***	0.0002
Panel B: Coefficients for $c((x_{it}, z_{it}), (\hat{x}_{it}, \hat{z}_{it}))$		
Constant (θ_{11})	1.3024***	0.0077
Δ Age (θ_{12})	0.1865***	0.0007
(Δ Age) ² (θ_{13})	-0.0074***	0.0000
Δ (Logincome + z_1) (θ_{14})	0.0103***	0.0006
(Δ (Logincome + z_1)) ² (θ_{15})	-0.0003***	0.0000
Δ (Conditions + z_2) (θ_{16})	-0.0443***	0.0014
(Δ (Conditions + z_2)) ² (θ_{17})	0.0060***	0.0002
Δ (Married + z_3) (θ_{18})	0.0136***	0.0010
(Δ (Married + z_3)) ² (θ_{19})	0.0284***	0.0025

Notes: See Table 9b.

Table 9b: Estimation Results from Dynamic Model with Serially-Correlated Unobservables:
[Table 9a Continued]

	Estimate	Std. Error
Panel C: Initial Distribution of Unoversvables		
z_1 : whether covered in 1994 (θ_{20})	2.7331**	0.1482
z_2 : whether covered in 1994 (θ_{21})	-10.3789***	0.2273
z_3 : whether covered in 1994 (θ_{22})	0.4779***	0.0358
Panel D: Transition Distribution of Unobservables		
z_1 : autocorrelation (θ_{23})	0.6957***	0.0158
z_2 : autocorrelation (θ_{25})	0.8765***	0.0143
z_3 : autocorrelation (θ_{27})	0.4997***	0.0655
z_1 : std. dev. (θ_{24})	0.2623***	0.0101
z_2 : std. dev. (θ_{26})	0.0012***	0.0001
z_3 : std. dev. (θ_{28})	0.0793***	0.0050
Log likelihood	-9,164.13	

Notes: (1). The specifications for $u_1(x_{it}, z_{it})$ and $c((x_{it}, z_{it}), (\hat{x}_{it}, \hat{z}_{it}))$ are given in 4 and 5 respectively.

(2). For any variable x , Δx is the difference between the current value of x and \hat{x} , which is the value of x at the time when the respondent changed his coverage.

(3). The annual discount factor β is set at 0.9.

(4). *, **, and *** respectively represent significance at 10%, 5% and 1% levels.

particularly the unobserved income.

5.4 Model Fit

Table 10 presents an assessment of the performance of the dynamic model with serially correlated unobservable state variables. We report in Panel A of Table 10 the comparisons between the simulated model predictions regarding aggregate choice probabilities by wave and those in the data. Relative to the model fit reported in Table 8, by incorporating serially correlated unobservable state variables, we are able to dramatically improve the model's ability to capture the dynamic patterns in the actual data. That is, the actual data exhibits that, in the aggregate, there is an increasing likelihood over time of holding no life insurance, which the model without serially correlated unobservable state variables was unable to capture (see Section 4). Indeed our simulation is able to replicate the increase in the fraction of individuals without life insurance coverage from 12.12% in 1996 to 22.86% in 2006, an increase that almost matches what is in the actual data (from 11.92% in 1996 to 21.36% in 2006).

In Panel B of Table 10, we report the comparison between the simulated model predictions and the data of the cumulative decision outcomes by wave for individuals who owned life insurance in 1994. Again, by incorporating serially correlated unobservable state variables, we are able to capture the pattern of a steadily increasing cumulative fraction of 1994 policyholders lapsing to no insurance in the actual data. In the data, this cumulative fraction went from 4.55% in 1996 to 23.39% by 2006; in our simulation, it goes from 4.59% in 1996 to 23.43% in 2006. In contrast, recall that in Table 8, the model without serially correlated unobservables is unable to replicate the initially low lapsation rates. Panel B of Table 10 also shows that by incorporating serially correlated unobservables, we are able to significantly improve the fit of the cumulative fraction of 1994 policyholders who kept their policies, whereas the model without serially-correlated unobservables under-predicts these by large margin (see Section 4).

6 Counterfactual Simulations

In this section, we report the results from a large number of counterfactual simulations to address two important questions. The first set of counterfactual simulations highlights the importance of serially-correlated unobserved state variables in explaining the patterns of life insurance decisions observed in the data. The second set of counterfactual simulations attempts to disentangle the contributions of income, health and bequest motives shocks, both observed and unobserved, in explaining the observed lapsations.

It is useful to emphasize at the outset the nature of our counterfactual simulations. When we

Table 10: Model Fit for Dynamic Model with Serially Correlated Unobservable State Variables

	Wave					
	1996	1998	2000	2002	2004	2006
Panel A: Aggregate Choice Probabilities by Wave						
<i>Actual Data</i>						
No life insurance coverage	0.1192	0.1421	0.1642	0.1889	0.1948	0.2136
Covered, but changed or bought new coverage	0.1493	0.1077	0.0963	0.0978	0.1123	0.0942
Covered, and kept existing coverage	0.7314	0.7500	0.7394	0.7131	0.6928	0.6920
<i>Simulation using dynamic model with serially correlated unobservables</i>						
No life insurance coverage	0.1212	0.1423	0.1608	0.1789	0.2005	0.2286
Covered, but changed or bought new coverage	0.1698	0.1142	0.1046	0.1029	0.1023	0.1032
Covered, and kept existing coverage	0.7089	0.7434	0.7344	0.7180	0.6971	0.6681
Panel B: Cumulative Outcomes for 1994 Policyholders						
<i>Actual data</i>						
Lapsed to no life insurance	0.0455	0.0899	0.1358	0.1792	0.2094	0.2339
Changed coverage amount	0.0903	0.1472	0.1845	0.2080	0.2283	0.2489
Kept 1994 coverage	0.8641	0.7315	0.6137	0.5202	0.4491	0.3812
Policyholder died	0.0000	0.0312	0.0657	0.0924	0.1130	0.1358
<i>Simulation using dynamic model with serially correlated unobservables</i>						
Lapsed to no life insurance	0.0459	0.0913	0.1317	0.1674	0.2003	0.2343
Changed coverage amount	0.1165	0.1798	0.2183	0.2445	0.2647	0.2816
Kept 1994 coverage	0.8375	0.6981	0.5876	0.4967	0.4229	0.3522
Policyholder died	0.0000	0.0306	0.0622	0.0913	0.1119	0.1318

remove the shocks, we are assuming that the market environment faced by consumers remain unchanged from when all shocks are present. That is, our counterfactual simulation does not allow for the market to re-equilibrate to respond to the fact that there are now fewer shocks. In particular, we must assume that the choice set of life insurance contracts that each individual faces in a given state does not change. Thus our counterfactual simulations are a partial equilibrium exercise.

6.1 The Importance of Serially-Correlated Unobserved State Variables

In this section, we report a series of counterfactual simulations to demonstrate the importance of including serially-correlated unobservable state variables. Panel A of Table 11 is identical to the bottom sub-panel of Panel B in Table 10 and it reports the model's predictions about the cumulative outcomes for 1994 policyholders.

In Panel B of Table 11, we report the predictions of the model using the coefficient estimates as reported in Panel A and B of Table 9a, but under the counterfactual assumption that the *unobservable* state variables did not change over time. It shows that without the shocks to the unobserved state variables, the model is unable to match the sharply increasing cumulative fraction of 1994 policyholders that lapse to no life insurance, and the model also over-predicts by a large margin the cumulative fraction of 1994 policyholders who kept their 1994 coverage.

In Panel C of Table 11, we report the predictions of the model using the coefficient estimates as reported in Table 9a, but under the counterfactual assumption that the *observable* state variables stayed the same as their values in 1994, except for age, while keeping the changes in the unobserved state variables. Surprisingly, assuming away the changes in the observable state variables *barely changes* the model's predictions about the cumulative outcomes for 1994 policyholders.

In Panel D, we report the predictions of the model using the coefficient estimates of the model as reported in Panel A and B of Table 9a, but under the counterfactual assumption that the *neither* the unobservable state variables *nor* the observed state variables (except for age) change over time. Only i.i.d choice specific shocks are retained in these simulations. The predictions in Panel D are very similar to Panel B where only changes in unobservable state variables are eliminated.

The counterfactual simulations in Table 11 thus provide very strong evidence for the importance of serially correlated unobservable state variables in explaining the choice patterns in the data. Qualitatively, it plays a much more important role than the variations in the observable state variables in capturing the key features in the data.

Table 11: Counterfactual Simulations Using the Estimates of the Dynamic Model with Serially-Correlated Unobservables: Cumulative Outcomes for 1994 Policyholders

Outcome	Wave					
	1996	1998	2000	2002	2004	2006
Panel A: All shocks included						
Lapsed to no life insurance	0.0459	0.0913	0.1317	0.1674	0.2003	0.2343
Changed coverage amount	0.1165	0.1798	0.2183	0.2445	0.2647	0.2816
Kept 1994 coverage	0.8375	0.6981	0.5876	0.4967	0.4229	0.3522
Policyholder died	0.0000	0.0306	0.0622	0.0913	0.1119	0.1318
Panel B: No shocks to the unobserved state variables						
Lapsed to no life insurance	0.0433	0.0651	0.0788	0.0889	0.0976	0.1071
Changed coverage amount	0.1171	0.1773	0.2150	0.2425	0.2666	0.2899
Kept 1994 coverage	0.8395	0.7267	0.6419	0.5719	0.5150	0.4571
Policyholder died	0.0000	0.0307	0.0641	0.0965	0.1207	0.1456
Panel C: No shocks to the observable state variables except for age						
Lapsed to no life insurance	0.0428	0.0832	0.1186	0.1502	0.1809	0.2137
Changed coverage amount	0.1149	0.1768	0.2150	0.2416	0.2632	0.2819
Kept 1994 coverage	0.8421	0.7091	0.6033	0.5153	0.4419	0.3697
Policyholder died	0.0000	0.0307	0.0629	0.0927	0.1138	0.1345
Panel D: Only i.i.d. choice specific shocks						
Lapsed to no life insurance	0.0353	0.0520	0.0622	0.0696	0.0761	0.0832
Changed coverage amount	0.1150	0.1728	0.2079	0.2330	0.2554	0.2778
Kept 1994 coverage	0.8496	0.7440	0.6644	0.5985	0.5445	0.4887
Policyholder died	0.0000	0.0310	0.0653	0.0987	0.1238	0.1501

Table 12: Disentangling the Contributions of Income, Health and Bequest Motive Shocks to the Lap-sations of 1994 Policyholders

	Wave					
	1996	1998	2000	2002	2004	2006
Panel A: The Role of Income Shocks						
[1] i.i.d. choice specific shocks only	0.0353	0.0520	0.0622	0.0696	0.0761	0.0832
[2] i.i.d. and income shocks only	0.0384	0.0613	0.0772	0.0897	0.1011	0.1134
[3] Incremental contribution of income shocks (%)	6.75	10.19	11.39	12.01	12.48	12.89
[4] All but income shocks	0.0415	0.0740	0.0988	0.1186	0.1365	0.1555
[5] All shocks included	0.0459	0.0913	0.1317	0.1674	0.2003	0.2343
[6] Incremental contribution of income shocks (%)	9.59	18.95	24.98	29.15	31.85	33.63
Panel B: The Role of Health Shocks						
[1] i.i.d. choice specific shocks only	0.0353	0.0520	0.0622	0.0696	0.0761	0.0832
[2] i.i.d. and health shocks only	0.0409	0.0642	0.0793	0.0905	0.1002	0.1101
[3] Incremental contribution of health shocks (%)	12.20	13.36	12.98	12.49	12.03	11.48
[4] All but health shocks	0.0386	0.0701	0.0957	0.1178	0.1386	0.1617
[5] All shocks included	0.0459	0.0913	0.1317	0.1674	0.2003	0.2343
[6] Incremental contribution of health shocks (%)	15.90	23.22	27.33	29.63	30.80	30.99
Panel C: The Role of Bequest Motive Shocks						
[1] i.i.d. choice specific shocks only	0.0353	0.0520	0.0622	0.0696	0.0761	0.0832
[2] i.i.d. and bequest motive shocks only	0.0352	0.0579	0.0739	0.0866	0.0980	0.1113
[3] Incremental contribution of bequest shocks (%)	-0.22	6.46	8.88	10.16	10.93	11.99
[4] All but bequest shocks	0.0444	0.0759	0.0997	0.1187	0.1357	0.1526
[5] All shocks included	0.0459	0.0913	0.1317	0.1674	0.2003	0.2343
[6] Incremental contribution of bequest shocks (%)	3.27	16.87	24.30	29.09	32.25	34.87
Panel D: Contributions of i.i.d. Choice Specific Shocks						
[1] Lower bound (%)	71.24	40.96	23.39	12.13	5.09	0.51
[2] Upper bound (%)	81.26	69.99	66.74	65.35	64.55	63.64

6.2 Disentangling the Contribution of Income, Health and Bequest Motive Shocks to Lapsations

In this section, we present a series of counterfactual simulations aimed at disentangling the contributions of income, health and bequest shocks, including both observed and unobserved components, to the lapsation of life insurance policies observed in the data. We present our results in four panels in Table 12. There are two sub-panels in Panels A-C. Let us first discuss Panel A, which illustrates the contribution of income shocks to the lapsation of 1994 policyholders. In the shaded sub-panel, we use as baseline the model’s prediction of the cumulative lapsation rates of 1994 policyholders when only i.i.d choice specific shocks are present (the first row), and examine how the addition of income shocks increases the model’s predicted lapsation rates (the second row).²⁴ Notice that the addition of income shocks to the i.i.d choice specific shocks leads to more lapsation. The incremental contribution of income shocks accounts for about 6.75% of the total lapsation predicted by the model when all shocks are included in 1996.²⁵ The incremental contributions of income shocks over time are reported in the third row. It reveals that the importance of income shocks are increasing over time in explaining lapsation. By 2006, income shocks alone were able to explain about 12.89% of the predicted lapsation.

The bottom, un-shaded, sub-panel in Panel A uses a different baseline. The baseline is instead the model’s prediction of lapsation rates when *all but* incomes shocks are included (reported in the fourth row of Panel A). This baseline prediction is compared to the predicted lapsation rates when all shocks are included (fifth row of Panel A). The difference is attributed to the incremental contribution of income shocks (sixth row of Panel A). Using this baseline, we see that the contribution of income shocks to lapsation is also increasing over time, from 9.59% in 1996 to 33.63% in 2006.²⁶

Panel B of Table 12 carries analogous calculations to illustrate the contribution of health shocks to the lapsations of 1994 life insurance policyholders. The shaded sub-panel shows that if we use the predicted lapsation rates with only i.i.d choice specific shocks as the baseline, the incremental contribution from adding health shocks using this baseline is more or less stable over time, staying

²⁴Note that the first row numbers in Panel A of Table 12 are identical to the numbers in the first row of Panel D of Table 11.

²⁵That is, $(0.0384 - 0.0353) / 0.0459 \approx 6.75\%$ where 0.0459 is lapsation rates predicted by the model when all shocks are included (reported in the fifth row of the table, as well as Panel B of Table 10). The other percentages are calculated analogously.

²⁶There are two other possible counterfactual baselines that we do not report. We could have used “i.i.d choice specific shock and health shocks” as baseline and contrast it with “i.i.d choice specific shock, health shocks and income shocks” (which is the same as “all but bequest motive shocks”). Alternatively, we could have used “i.i.d choice specific shock and bequest motive shocks” as the baseline and contrast it with “i.i.d choice specific shock, bequest motive shocks and income shocks” (which is the same as “all but health shocks”). Note that the information needed to carry out these calculations is presented in other rows in Table 12. For space reasons, we do not present these calculations separately.

at about 12% throughout the years. If we use the predicted lapsation rates when all but health shocks are included as the baseline, the incremental contribution from adding health shocks goes from 15.90% in 1996 to 31% in 2006.²⁷

Panel C of Table 12 shows the contribution of bequest motive shocks to life insurance lapsation. As in Panels A and B, the top shaded sub-panel calculates the incremental contribution of bequest motive shocks using the predicted lapsation with only i.i.d choice specific shocks as the baseline, and the bottom sub-panel uses the predicted lapsation with all shocks except for bequest shocks as the baseline. We find that the importance of bequest motives in explaining lapsation increases over time. In 1996 the bequest motive shocks explain between -0.22% to 3.27% of the lapsation; but by 2006, it explains between 11.99% to 34.87% of the lapsation.

Panel D of Table 12 bounds the contributions of i.i.d choice specific shocks in explaining the lapsations. The lower bounds are calculated as the residuals after subtracting the upper bound contributions from income, health and bequest motive shocks.²⁸ Panel D reveals that lapsation of life insurance policies are largely driven by i.i.d choice specific shocks for younger individuals, but for surviving policyholders an ever larger fraction of lapsations is explained by either income, or health, or bequest motive shocks. By 2006, between more than 1/3 to almost 100 percent of the lapsations are driven by one of these shocks.

To summarize, the simulations reported in Table 12 indicate that when individuals are young, most of the life insurance policy lapses are driven by i.i.d choice specific shocks, and the rest is explained, in descending order of importance, by health shocks and income shocks; the bequest motive shocks only account for a very minor fraction of lapsation. However, as policyholders get older, the importance of the i.i.d choice specific shocks declines dramatically, and the three shocks eventually account for about the same fraction of lapsation (ranging from about 12% to around 30%). If anything, the bequest motive shocks are more important in explaining lapsation among older individuals than the income and health shocks.

Policy Implications. Our findings on the relative contributions of income, health, bequest motive, and i.i.d. choice shocks to the lapsation rates of 1994 policyholders have implications regarding the possible consumer welfare effects of emerging life settlement markets. As we mentioned in the Introduction, theoretical studies by Daily, Hendel and Lizzeri (2008), Fang and Kung (2010a) and Fang and Kung (2010b) show that the reasons for lapsation will determine whether life settle-

²⁷It is worthwhile pointing out that the discrepancy in the results for different baselines is due to the nonlinearities in our model.

²⁸For example, we obtain 71.24% lower bound number for year 1996 from $1 - 9.59\% - 15.90\% - 3.27\%$ where 9.59%, 15.90% and 3.27% are respectively the upperbound contributions of income, health and bequest motive shocks reported in Panels A to C. The other bounds are calculated analogously.

ment markets can improve consumer welfare. Specifically, [Fang and Kung \(2010b\)](#) show that a key determinant for whether consumers may benefit from life settlement markets is whether lapsation is driven by factors that are positively correlated with the marginal utility of income. Thus, to the extent we found in [Section 5.3](#) that decreases in log income lead to a higher probability of lapsation, and decreases in log income certainly lead to higher marginal utility of income, the fraction of lapsation that can be attributed to changes in income, both observed and unobserved, should be a potential source of welfare gain for consumers when life settlement markets are introduced. On the other hand, lapsation driven by i.i.d choice specific shocks are not positively correlated to the marginal utility of income, and thus lapsation driven by i.i.d. choice specific shocks can lead to a reduction in consumer welfare. Analogously, to the extent that health shocks and bequest motive shocks are orthogonal to income shocks, and are thus not necessarily positively correlated to the marginal utility of income, we suspect that the fractions of lapsation attributable to these two shocks are likely sources for consumer welfare reduction when life settlement is introduced.

Our finding that i.i.d choice specific shocks explain the bulk of the policy lapsation when individuals are relatively young (when they are in their early 60s) suggests that life settlement is likely to lead to a welfare loss for relatively young policyholders; but may lead to welfare gains for older policyholders in their early seventies, as changes in income become a more important source for lapsation. We should emphasize, however, that these implications from our analysis are only suggestive; a more definite study of the welfare effect of the life settlement market would require that we estimate a fully structural model of the behavior of both the consumers who choose life insurance policies and the life insurance companies who offer such policies.

6.3 Discussion

We now discuss two important issues. The first issue is about the identification of the three components of the serially-correlated unobservable state variables intended to capture income, health and bequest motive shocks. To the best of our knowledge, this is the first paper that allows for more than one unobservable state variable. So a natural question is whether the distributions of such unobservable state variables can be separately identified. This is obviously an important question to be addressed in future research. For now, we would first like to emphasize that in this paper, we tried to anchor the interpretation of these three shocks by restricting that each has the same effect on behavior as their respective observable counterparts (see [Section 5.1](#)).

To give further evidence that patterns in the data justify the inclusion of multiple dimensions of serially correlated unobserved state variables, we also estimated a series of alternative models where we include only a subset of the three shocks. In the same spirit of [Heckman and Singer \(1984\)](#) for the case of unobserved types, we ask whether the inclusion of additional unobserved

Table 13: Log Likelihoods of Various Specifications of Unobservable State Variables

Specification	Serially Correlated Unobservables?			Log Likelihood
	Income?	Health?	Bequest Motive?	
All	Yes	Yes	Yes	-9,164.13
1	No	Yes	Yes	-9,246.65
2	Yes	No	Yes	-9,233.80
3	Yes	Yes	No	-9,219.66
4	No	No	Yes	-9,274.13
5	No	Yes	No	-9,288.69
6	Yes	No	No	-9,285.87
None	No	No	No	-9,338.59

state variables increases the log-likelihood of the estimated model. In Table 13, we report the log-likelihood of the estimated models with various specifications of the unobservable shocks. In particular, the specification labeled “All” corresponds to the model estimated in Section 5 where we include all three unobservable state variables, and the specification labeled “None” corresponds to the model estimated in Section 4 where we do not include any unobservable state variables. In specifications labeled 1 to 6, various combinations of the three unobservable state variables are included in the estimation. From the last column in Table 13, we can see that the inclusion of the additional unobservable state variables significantly increases the log-likelihood of the models. For example, in specifications 4-6, we estimated models with only one of the unobservable state variables respectively. The log-likelihoods of these models improve over specification “None”. Similarly, in specifications 1-3, we estimated models with two of the three unobservable state variables; and again, the log-likelihoods of these models improve over specifications with only one of the unobservables. Finally, the log-likelihood of the model with all three shocks is higher than specification 1-3. The results in Table 13 show that the data indeed seems to be more consistent with a model using all three serially-correlated unobserved state variables.²⁹

The second issue is regarding our finding that the importance of the unobserved state variables in explaining lapsation increases over time. The concern is whether this is a mechanical result due to the way we simulate the unobservable state variables using SMC. In particular, recall that the initial distribution of the unobservable state variables are assumed to have smaller support than in later periods (see Sections 5.1 and 5.2). While this is a possibility, we would like to make two

²⁹We also conducted formal likelihood ratio tests of the specifications 1-6 and “None” against specification “All.” The tests are all in favor of specification “All”.

counter-arguments. First, even though the unobservable state variables in the earlier periods have smaller support (in fact, just one point support in the initial period), these points in the support were chosen to best fit the data; thus there is no *a priori* reason that the unobservable state variables in the early periods should have less impact just because they have a small support. Second, if the simulated unobservable state variables were pure noise that the individuals do not take into account, then the unobservables' importance in explaining the observed lapsations should not have changed over time at all, that is, the importance should be about zero in all periods. Thus the fact that the importance of the unobservable state variables were found to be increasing over time is an indication that these simulated unobservable state variables are capturing something informative.

7 Conclusion

In this paper, we empirically investigate the contributions of income, health and bequest motive shocks to life insurance lapsation. We present a dynamic discrete choice model of life insurance decisions allowing for serially correlated unobservable state variables. The model is designed to deal with the data reality where researchers only observe whether an individual has made a new life insurance decision (i.e., purchased a new policy, or added to/changed an existing policy) but do not observe the actual policy choice or the choice set from which the new policy is selected. The semi-structural dynamic discrete choice model allows us to bypass these data limitations. We empirically implement the model using the limited life insurance holding information from the Health and Retirement Study (HRS) data.

We deal with serially correlated unobserved state variables using posterior distributions of the unobservables simulated from Sequential Monte Carlo (SMC) methods. Relative to the few existing papers in the economics literature that have used similar SMC methods, our paper is the first to incorporate multi-dimensional serially correlated unobserved state variables. In order to give the three unobservable state variables in our empirical model their desired interpretations as unobserved income, health and bequest motive shocks, this paper proposes two channels through which we can anchor these unobservables to their related observable variables. We present our estimation results separately for dynamic models without and with serially correlated unobservable state variables. We illustrate the importance of serially correlated unobservables by showing the features in the data that a dynamic model without unobservable state variables is unable to capture, and how with unobservable state variables the model fit dramatically improves (c.f. Tables 8 and 10). We also show that in the model with unobserved state variables, the contribution of the shocks to unobservables is much larger than the contribution of the shocks to observed state

variables (Table 11).

Our estimates for the model with serially correlated unobservable state variables are sensible and yield implications about individuals' life insurance decisions consistent with the both intuition and existing empirical results. In a series of counterfactual simulations reported in Table 12, we find that a large fraction of life insurance lapsation is driven by i.i.d choice specific shocks, particularly when policyholders are relatively young. But as the remaining policyholders get older, the role of such i.i.d. shocks gets less important, and more of their lapsation is driven either by income, health or bequest motive shocks. Income and health shocks are relatively more important than bequest motive shocks in explaining lapsation when policyholders are young, but as they age, the bequest motive shocks play a more important role.

Our empirical findings have important implications regarding the effect of the life settlement industry on consumer welfare. As shown in theoretical analysis in [Daily, Hendel and Lizzeri \(2008\)](#) and [Fang and Kung \(2010a,b\)](#), the theoretical predictions about the effect of life settlement on consumer welfare crucially depend on why life insurance policyholders lapse their policies. If bequest motive shocks are the reason for lapsation, then the life settlement industry is shown to reduce consumer welfare in equilibrium; but if income shocks are the reason for lapsation, then life settlements may increase consumer welfare. To the extent that we find both income shocks and bequest motive shocks play important roles in explaining life insurance lapsations, particularly among the elderly population targeted by the life settlement industry, our research suggests that the effect of life settlement on consumer welfare is ambiguous. Unfortunately, our "semi-structural", partial equilibrium model of life insurance decisions (which is necessitated by data limitations) is not suitable for a quantitative general equilibrium evaluation of the welfare impacts of introducing a life settlement market. This is an important, but challenging, area for future research.

References

American Council of Life Insurers, "2009 Life Insurers Fact Book," 2009.

Blevins, Jason, "Sequential Monte Carlo Methods for Estimating Dynamic Microeconomic Models," Working Paper, Ohio State University, 2011.

Cawley, John and Thomas Philipson, "An Empirical Examination of Information Barriers to Trade in Insurance," *American Economic Review*, 1999, Vol. 89, 827–846.

Chandik, Mark, "The Growing Life Settlement Industry: Is Anyone Watching Out for Consumers?," Testimony presented at the California Senate Bank-

- ing, Finance and Insurance Committee on Life Settlements, 2008, available at http://www.sen.ca.gov/ftp/sen/committee/standing/banking/info_hearings/backgrounds/20-08_life_settlement_background.
- Daily, Glenn, Igal Hendel, and Alessandro Lizzeri, "Does the Secondary Life Insurance Market Threaten Dynamic Insurance?," *American Economic Review Papers and Proceedings*, 2008, 98 (2), 151–156.
- Doucet, Arnaud, Nando de Freitas, and Neil Gordon, *Sequential Monte Carlo methods in practice*, New York: Springer, 2001.
- Fang, Hanming and Edward Kung, "How does the life settlements market affect the primary market for life insurance?," *Working paper*, 2010a.
- and —, "The Welfare Effect of Life Settlement Market: The Case of Income Shocks," Working Paper, University of Pennsylvania and Duke University, 2010b.
- Fernandez-Villaverde, Jesus and Juan F. Rubio-Ramirez, "Estimating Macroeconomic Models: A Likelihood Approach," *Review of Economic Studies*, 2007, 74, 1059–1087.
- Fox, Jeremy T, "Semiparametric Estimation of Multinomial Discrete-Choice Models Using a Subset of Choices," *Rand Journal of Economics*, 2007, Vol. 38, No. 4, 1002–1019.
- Gilbert, Jersey and Ellen Schultz, *Consumer Reports Life Insurance Handbook*, Consumer Reports Books: Yonkers, NY, 1994.
- Gordon, N.J., D.J. Salmond, and A.F.M. Smith, "Novel Approach to Nonlinear/Non-Gaussian Bayesian State Estimation," 1993, 140 (2), 107–113.
- He, Daifeng, "The Life Insurance Market: Asymmetric Information Revisited," forthcoming, *Journal of Public Economics*, 2009.
- , "Is Dynamic Selection Present in the Life Insurance Market?," Working paper, College of William and Mary, 2010.
- Heckman, James and Burton Singer, "A Method for Minimizing the Impact of Distributional Assumptions in Econometric Models for Duration Data," *Econometrica*, 1984, 52 (2), 271–320.
- Hendel, Igal and Alessandro Lizzeri, "The Role of Commitment in Dynamic Contracts: Evidence from Life Insurance," *Quarterly Journal of Economics*, 2003, 118 (1), 299–327.

- Hong, Han, Ronald A. Gallant, and Ahmed Khwaja, “Estimating a Dynamic Oligopolistic Game with Serially Correlated Unobserved Production Costs,” *Working paper*, 2008.
- Hu, Yingyao and Matthew Shum, “Nonparametric Identification of Dynamic Models with Unobserved State Variables,” *mimeo*, 2009.
- Kasahara, Hiroyuki and Katsumi Shimotsu, “Nonparametric Identification of Finite Mixture Models of Dynamic Discrete Choice,” *Econometrica*, 2009, *Vol. 77, No. 1*, 135–175.
- Keane, Michael P. and Kenneth I. Wolpin, “The Solution and Estimation of Discrete Choice dynamic programming models by simulation and interpolation: Monte Carlo evidence,” *The Review of Economics and Statistics*, 1994, *Vol. 76*, 648–672.
- Kung, Edward, “Efficient Mortgage Design in an Equilibrium Model of Housing and Mortgage Markets,” *Working paper*, Duke University, 2012.
- McFadden, Daniel, “Modelling the Choice of Residential Location,” in A. Karlqvist, L. Lundqvist, F. Snickars, and J. Weibull, eds., *Spatial Interaction Theory and Planning Models*, 7 North Holland: Amsterdam, 1978, *Vol. I*, 75–96.
- Norets, Audriy, “Inference in Dynamic Discrete Choice Models with Serially Correlated Unobserved State Variables,” *Econometrica*, 2009, *Vol. 77, No. 5*, 1665–1682.
- Olsson, J. and T. Rydén, “Asymptotic Properties of Particle Filter-Based Maximum Likelihood Estimators for State Space Models,” *Stochastic Processes and their Applications*, 2008, *118 (4)*, 649–680.
- Rust, John, “Structural Estimation of Markov Decision Processes (1994) in R. Engle and D. McFadden (eds.) *Handbook of Econometrics Volume 4*, 30823139, North Holland.” in R. Engle and D. McFadden (eds.) *Handbook of Econometrics*, Volume 4, North Holland, 1994, pp. 3082–3139.

A Appendix

In this appendix, we include two tables that provide the summary statistics for the key state variables conditional on having life insurance (Table A.1) and conditional on not having life insurance (Table A.2).

Table A.1: Summary Statistics of State Variables Conditional on Having Life Insurance.

Variable Description	Wave					
	1996	1998	2000	2002	2004	2006
Age of respondent	61.071 (4.3397)	62.989 (4.3234)	64.839 (4.3184)	66.825 (4.2841)	68.766 (4.3120)	70.688 (4.2611)
Household income (\$1,000s)	65.895 (74.367)	66.099 (76.907)	67.793 (83.742)	65.583 (74.072)	67.665 (79.049)	68.957 (83.542)
Log household income	10.689 (1.1552)	10.700 (1.0806)	10.714 (1.0907)	10.690 (1.1304)	10.754 (0.9468)	10.778 (0.8574)
Whether ever diagnosed with high blood pressure	0.3978 (0.4895)	0.4356 (0.4959)	0.4794 (0.4997)	0.5210 (0.4997)	0.5695 (0.4953)	0.6144 (0.4869)
Whether ever diagnosed with diabetes	0.1326 (0.3392)	0.1475 (0.3547)	0.1616 (0.3682)	0.1937 (0.3953)	0.2164 (0.4119)	0.2388 (0.4265)
Whether ever diagnosed with cancer	0.0591 (0.2359)	0.0792 (0.2701)	0.1010 (0.3014)	0.1345 (0.3412)	0.1633 (0.3697)	0.1915 (0.3936)
Whether ever diagnosed with lung disease	0.0687 (0.2530)	0.0748 (0.2632)	0.0773 (0.2671)	0.0865 (0.2811)	0.1034 (0.3045)	0.1139 (0.3178)
Whether ever diagnosed with heart disease	0.1846 (0.3880)	0.2107 (0.4079)	0.2338 (0.4233)	0.2655 (0.4417)	0.3051 (0.4606)	0.3433 (0.4749)
Whether ever diagnosed with stroke	0.0431 (0.2030)	0.0518 (0.2218)	0.0590 (0.2357)	0.0644 (0.2456)	0.0759 (0.2650)	0.0910 (0.2877)
Whether ever diagnosed with psychological problem	0.0496 (0.2171)	0.0632 (0.2433)	0.0670 (0.2500)	0.0774 (0.2673)	0.0860 (0.2804)	0.0920 (0.2892)
Whether ever diagnosed with arthritis	0.3896 (0.4877)	0.4392 (0.4964)	0.4774 (0.4996)	0.5357 (0.4988)	0.5810 (0.4935)	0.6234 (0.4847)
Sum of above conditions	1.3250 (1.2062)	1.5020 (1.2606)	1.6565 (1.2996)	1.8785 (1.3325)	2.1006 (1.3765)	2.3085 (1.3949)
Whether married	0.8718 (0.3343)	0.8638 (0.3430)	0.8594 (0.3477)	0.8591 (0.3480)	0.8618 (0.3451)	0.8463 (0.3608)
Whether has children	0.9460 (0.2261)	0.9460 (0.2261)	0.9426 (0.2327)	0.9455 (0.2270)	0.9437 (0.2305)	0.9368 (0.2434)
Age of youngest child	25.094 (9.819)	27.104 (9.713)	29.126 (9.620)	31.207 (9.688)	33.342 (9.547)	35.129 (9.545)
# of live respondents	2,927	2,739	2,524	2,313	2,187	2,011

Note: Standard deviations are in parenthesis.

Table A.2: Summary Statistics of State Variables Conditional on Not Having Life Insurance.

Variable Description	Wave					
	1996	1998	2000	2002	2004	2006
Age of respondent	61.328 (4.4532)	63.313 (4.4617)	65.363 (4.2256)	66.968 (4.2963)	68.885 (4.1783)	70.634 (4.2181)
Household income (\$1,000s)	38.518 (65.056)	37.115 (40.092)	51.930 (103.065)	53.221 (86.519)	60.548 (117.678)	57.684 (66.599)
Log household income	9.787 (1.9069)	9.904 (1.6405)	10.161 (1.5933)	10.273 (1.4337)	10.417 (1.3169)	10.536 (1.0812)
Whether ever diagnosed with high blood pressure	0.4369 (0.4966)	0.4471 (0.4977)	0.4617 (0.4990)	0.5009 (0.5005)	0.5633 (0.4964)	0.6392 (0.4807)
Whether ever diagnosed with diabetes	0.2273 (0.4196)	0.2313 (0.4221)	0.2460 (0.4311)	0.2597 (0.4389)	0.2722 (0.4455)	0.3022 (0.4596)
Whether ever diagnosed with cancer	0.0657 (0.2480)	0.0925 (0.2901)	0.0988 (0.2987)	0.1039 (0.3054)	0.1418 (0.3492)	0.1722 (0.3779)
Whether ever diagnosed with lung disease	0.0884 (0.2842)	0.0947 (0.2931)	0.0887 (0.2846)	0.0946 (0.2930)	0.1285 (0.3350)	0.1282 (0.3346)
Whether ever diagnosed with heart disease	0.2323 (0.4228)	0.2379 (0.4263)	0.2419 (0.4287)	0.2746 (0.4467)	0.3043 (0.4606)	0.3443 (0.4756)
Whether ever diagnosed with stroke	0.0960 (0.2949)	0.0947 (0.2931)	0.0968 (0.2959)	0.1002 (0.3005)	0.1115 (0.3151)	0.1264 (0.3326)
Whether ever diagnosed with psychological problem	0.1364 (0.3436)	0.1211 (0.3267)	0.1391 (0.3464)	0.1299 (0.3365)	0.1229 (0.3286)	0.1300 (0.3367)
Whether ever diagnosed with arthritis	0.3965 (0.4898)	0.4868 (0.5004)	0.5121 (0.5004)	0.5121 (0.5003)	0.5709 (0.4954)	0.6099 (0.4882)
Sum of above conditions	1.6793 (1.4358)	1.8062 (1.4366)	1.8851 (1.3487)	1.9759 (1.3556)	2.2155 (1.4626)	2.4524 (1.5067)
Whether married	0.6919 (0.4623)	0.7181 (0.4504)	0.7681 (0.4224)	0.7829 (0.4126)	0.7694 (0.4216)	0.7821 (0.4132)
Whether has children	0.9116 (0.2842)	0.9053 (0.2931)	0.9234 (0.2662)	0.9091 (0.2877)	0.9074 (0.2902)	0.9212 (0.2696)
Age of youngest child	22.902 (11.510)	25.057 (11.559)	26.923 (11.362)	29.542 (10.362)	31.242 (10.527)	34.027 (10.292)
# of live respondents	397	456	498	541	530	547

Note: Standard deviations are in parenthesis.