Insurance-Induced Moral Hazard: A Dynamic Model of Within-Year Medical Care Decision Making Under Uncertainty

Web Appendix

Christopher J. Cronin

August 4, 2016

Contents

1 Data Overview

The empirical exercises included in this research use data from the first four cohorts (i.e., 1996-1999) of the Medical Care Expenditure Panel Survey (MEPS), which is collected by the Agency for Healthcare Research and Quality (AHRQ). The within-year decision-making model presented in the paper requires the following information for each individual in the estimation sample: a set of employer-sponsored health insurance (ESHI) plans (i.e., the cost-sharing features of every plan), an insurance decision, a general set of demographic variables, and two-years of medical care consumption decisions, medical care prices, and illness identifiers. Most of these variables cannot be taken directly from the survey, but must be constructed from a set of responses. Furthermore, as is with any large longitudinal data set, many individuals are missing pieces of information because they failed to answer a particular survey question. In some studies this missing variable problem can be handled by simply dropping individuals from analysis, but with an already limited sample, dropping more individuals due to missing variables is not possible here. Therefore, this section serves two purposes. First, it details the construction of variables used in estimation from the raw data. Second, it explains the various data cleaning methods employed to strategically fill missing variable values for the individuals in the estimation sample.

1.1 Demographic Variables

Individuals are first described by the following demographic variables, which are assumed to be exogenous and time-invariant in estimation: sex, age, years of education, income, race (categorized as white or non-white), and MSA (categorized as living in an MSA or not). Sample inclusion requires that individuals are employed and single (i.e., not married and without children). Also, exclusion restrictions are needed to estimate endogenous initial conditions; these are taken from the previous years National Health Interview Survey (NHIS), which is the survey from which MEPS participants are drawn, and include: prior year income, veteran status, and country of birth (categorized as foreign born or not).

Sex, age, race, employment, and MSA status can be taken directly from the data without explanation. Education is taken from the first interview. Later interviews are used if the first interview response is missing.[1] Income is the sum of post-tax income, sale earnings, and tax refund (*ttlpnx, salepnx, refdpnx*). AHRQ provide a cleaned version of these variables with imputations. Individuals self-identify as married (*marry31x*) and I assume separated individuals are not married.[2] Parents are identified by other persons living in the household

[1] Two individuals never report an exact number of years of education; however, they report some college so 14 years of education is assumed.

[2] The marriage condition is meant to eliminate individuals considering family insurance coverage. Separated individuals are considered not married because they are likely to make individual insurance decisions.

who are under 19 if non-student or under 24 if student, who identify them as mother or father (*mopid31x, dapid31x*).

Veteran Status can be extracted directly from NHIS survey data in 1995 and 1996, which corresponds to 1996 and 1997 MEPS participants. In 1997 and 1998, a veteran question is not asked directly but can be inferred from the question "have you ever been honorably discharged" (*miltryds*) and from VA medical care coverage information (*hikindf, hikindg*).[3] Foreign born status can be taken directly from the data in each year. Previous year income is measured on a $1-8$ scale as follows:

1	0-5k
2	5-10k
3	10-15k
4	15-20k
5	20-25k
6	25-35k
7	35-45k
8	45k+

Income is missing for 33% of the NHIS population, so I use a missing indicator in estimation. Furthermore, of the 1,232 individuals in the estimation sample, only 1,179 can be linked to NHIS. Those not matched are new members or previously absent members of NHIS participating households. Of those linked, 3 are missing foreign born status and 11 are missing veteran status. All individuals will missing data are assumed to be native born, non-veterans.

1.2 Insurance Offer Set

A requirement for sample inclusion is that individuals are offered ESHI and that they choose to be continuously insured by an ESHI plan with no switching and no other coverage or uninsured over a 12 month period. Evidence of an insurance offer can be found in both the employment section (*offer31x*) and insurance section (*hpemmyy*) of the survey. When these variables contradict, I assume the individual receives an offer as they are later removed if the offer is not observed in the IC-Link file. Monthly insurance status (i.e., insured or not) can be determined in the main HC data file (*insmmyy*). A series of insurance provider type variables in the HC file ensures that the individual is not covered by an additional

If any of these separated individuals hold family coverage then they are dropped from the sample.

[3] An individual is assumed to be a veteran if miltrydc $= 1$ OR hikindf $= 1$ & age > 18 & male $= 1$ OR hikindg $= 1$ & age > 18 & male $= 1$.

non-ESHI plan. Information in the Person-Round-Plan file identifies individuals switching ESHI plans during the insurance year and those holding multiple ESHI plans.

In the model, an individual selects a health insurance plan from a set of options offered by his employer. Estimation requires data describing the premium and cost-sharing features of both accepted and offered plans. The fact that MEPS collects such data is one of the main reasons that it is used in this research. From 1996-1999 and again in 2001, information about an employee's insurance offer set was gathered from his employer following the first interview.[4] This insurance information is contained in the confidential MEPS IC file. Using a Census Research Data Center (RDC), the IC file was linked with public individual-level data (i.e., MEPS HC file) to form the sample used in estimation. Unfortunately, about 56% of the individuals in the HC file meeting sample inclusion criteria for this research did not have plans included in the IC File and were therefore dropped from the estimation sample.[5] Among the individuals remaining, many have insurance offer sets featuring plans with missing characteristics. At least 1 of the 12 insurance cost-sharing features is missing in 47% of the plans observed in the data. The remainder of this subsection discusses the strategies used to impute these missing values.

Before describing these imputation strategies, it is important to understand that variation in the structure of health insurance plans is the most significant contributor to missing plan features. For example, some plans simply do not have a deductible, which is the same as having a zero deductible. In completing the survey, some firms skipped the deductible question when the plan did not have a deductible, rather than filling in a zero value. As such, the data often cannot distinguish between a skipped question and a plan without a particular cost-sharing feature.[6]

1.2.1 Logical Imputation

I begin this exercise with a series of logical imputations which are informed by the observed features of an insurance plan.

1. If no deductible is listed and it is stated that the plan has no deductible, then a zero deductible is coded. If it is indicated that the plan has a deductible, but no deductible value is recorded, then I note this information and use it later. If the response to the "does this plan have a deductible?" question is missing then the deductible value is left missing until later.

[4] My analysis does not use year 2001 because information on the month a plan-year started, which is critical when taking the data to the model, was not collected.

[5] See Section 4.2 in the paper for a discussion of the potential bias introduced by this exclusion.

[6] There is actually a separate question asking whether or not the plan had a deductible, but many plans are missing this information as well. Therefore, I use the "does this plan have a deductible?" information when possible, but it cannot be consistently relied upon.

2. There are 3 hospital coverage variables: co-insurance, daily co-pay, and per-stay co-pay. If any of these three is given a positive value, then I assume the others (if missing values) do not apply.

3. There are 2 doctor coverage variables: co-insurance and co-pay. If either of these is given a positive value, then I assume the other (if missing values) does not apply.

4. If a hospital co-pay value is recorded, but per-day/per-stay is not stated, then I assume co-pay greater than $100 is per-stay and less than $100 is per-day.

5. If the stop loss is coded as zero, meaning comprehensive coverage, but the plan features some cost-sharing, then I recode the stop loss to be missing.

6. If hospital and doctor care is free after the deductible is crossed then I set the stop loss equal to the stated deductible.

The scope of the missing information problem, after imposing these assumptions, can be observed in the table below. I list the number of plans with *any* missing characteristic by whether or not the plan is held by the employee. Plans are flagged if missing out-of-pocket premium, deductible, hospital coverage, doctor coverage, or stop loss. Total premium and plan type (i.e., HMO, PPO, or FFS) are never missing in the data. Also, note that for many of these missing characteristics I can determine whether or not the plan characteristic exists (e.g., that the plan has a deductible), just not what the true value is.

	Declined Plan	Held Plan	Total
No missing features	2,495	434	2,929
Missing features	1,958	706	2,664
Total	4,453	1,140	5,593

Next, I use observed expenditure information in the data to infer missing cost-sharing features of held plans. MEPS collects not only charge and total expenditure data, but also out-of-pocket and insurer payment for each medical care episode. This information can be used at times to determine deductible levels and doctor/hospital cost-sharing arrangements for held plans only. For example: If a plan begins in January and every trip to the doctor's office over the course of the year cost the individual $5 out-of-pocket, then I can impute that the individual had no deductible and a $5 co-pay at the doctor. For a number of reasons, imputation is rarely this easy; for example, many plans have different cost-sharing characteristics in and out of network, people may not recall the price paid out-of-pocket, and some individuals simply do not pay their medical bills. As such, some subjectivity is necessary to execute this procedure.

I also use this step to exclude some individuals and plans from the final sample. First, I eliminate individuals with large amounts of medical care spending information that cannot possibly be generated under their current coverage. Second, I remove redundant plans, or plans within a specific individuals offer set that are exactly the same according to the model's characterization of insurance or differ only by start month. After removing these redundant plans and making the imputations described above, the number of plans with missing characteristics is as follows:

	Declined Plan	Held Plan	Total
No missing features	2,398	485	2,883
Missing features	1,847	641	2,488
Total	4,245	1,126	5,371

1.2.2 Matching Method

After the logical edits described above, I turn to more traditional imputation techniques. I first use a matching procedure that fills in missing insurance characteristics with those of plans matching on observable characteristics. To execute this procedure, I first order all 7,028 plans in the original IC-Link file by total premium and network type. I then categorize each plan by its missing characteristics (e.g., missing nothing, missing deductible only, missing deductible and doctor coverage, etc.). Beginning with the plans missing the fewest characteristics, I fill in missing information using values from similar plans where the variable is observed. There are 6 rounds of matching, where the level of similarity required for a match is relaxed further in each round. The table below describes the bounds on an acceptable match in each round.

	Round					
	1	2	3	4	5	6
Total premium	0	0	25	25	50	50
Total deductible	0	0	25	25	50	50
Doctor deductible	0	0	25	25	50	50
Hospital deductible	0	0	25	25	50	50
Doctor co-pay	0	0	5	5	10	10
Doctor co-insurance	0	0	5	5	10	10
Hospital co-insurance	0	0	5	5	10	10
Hospital daily co-pay	0	0	10	10	20	20
Hospital stay co-pay	0	0	25	25	50	50
Stop loss	0	0	150	150	300	300
Plan type	0	0	0	0	0	0

In order to match, every characteristic of the matched plan must fall within the appropriate characteristic bounds centered at the observed characteristics of the plan with some piece of missing information. For example: assume plan A is missing all deductible information, but has values for all other plan characteristics. Further, assume these characteristics are as follows:

- Total premium: 500
- Doctor co-pay: 10
- Doctor co-insurance: none
- Hospital co-insurance: 10
- Hospital co-pay (day): none
- Hospital co-pay (stay): none
- Stop loss: 1500
- Plan type: HMO

In the first two search rounds, plan A would only be matched if it finds a plan with the exact same characteristics described above.[7] If a match is found then deductible, hospital deductible, and doctor deductible information from the matched plan are recorded for plan A. In the third and fourth rounds, plan A is matched if it can find a plan with characteristics fitting in the following bounds:[8]

- Total premium: $[475, 525]$
- Doctor co-pay: $[5, 15]$
- Doctor co-insurance: $[-5, 5]$
- Hospital co-insurance: $[5, 15]$
- Hospital co-pay (day): $[-10, 10]$

[7] Only plans with no missing information are available for match. After the first round of matching, some plans previously ineligible for match become eligible, which is why there are two rounds of matching for each set of bounds.

[8] There are several notes to be made here. First, in every round, the plan types must match exactly. Second, the out-of-pocket premium is not used as a criteria for matching because it varies across employers, meaning has little to do with the commonality of plans in the marketplace. Third, in many instances, it can be determined with certainty that a plan with a missing deductible does in fact have a non-zero deductible. In these cases, the match plan must have a non-zero deductible, in addition to the restrictions described above.

- Hospital co-pay (stay): $[-25, 25]$
- Stop loss: $[1350, 1650]$
- Plan type: HMO

In each round all plans with missing information look for a match before moving on to the next round. A summary of missing plans after the matching procedure can be seen below.

	Declined Plan	Held Plan	Total
No missing features	3,973	924	4,897
Missing features	272	202	474
Total	4,245	1,126	5,371

1.2.3 Regression Method

After using this matching technique, there are still 474 plans with missing information. However, the existence of missing information here does not necessarily imply that the observable characteristics of these plans are outside the bounds of normalcy. Because the matching procedure requires similarity in every observable characteristic, it is possible that a plan missing only one characteristic, with all other observed characteristics at the sample mean, may not find a match simply because that combination of features is not seen in another plan.

I turn to regression methods to determine missing values for the remaining plans. To reduce the number of regressions required, I assume that all remaining plans are defined by the most popular insurance structure, meaning doctor and hospital specific deductibles, hospital co-pay, and doctor co-insurance are ruled out. I begin by estimating whether or not a plan has a stop loss.[9] The regression includes all plans reporting a stop loss and those reporting that the plan has no stop loss. These parameters are used to predict the presence of a stop loss for plans where it cannot be determined if a stop loss exists. I then use an OLS regression of stop loss on other plan and company characteristics to predict a stop loss for those known to have a one (either through the data or through the previously described prediction), but no recorded value.[10] The same procedure is followed for the deductible; a

[9] For this regression and all other dichotomous outcomes I estimate logit parameters.

[10] Each regression uses the following set of explanatory variables: total premium, has deductible, deductible, has stop loss, stop loss, hospital cost-sharing generosity, doctor's office cost-sharing generosity, firm size, number of firm enrollees, federal plan indicator, year dummies, HMO, and PPO. I also use a number of missing variable dummies that change across equations as information is filled in. Some variables must be removed from this set so that they are not on the right and left hand side of the regression (e.g., stop loss is not included on the right hand side of the stop loss regression).

logit for any deductible and an OLS regression to predict the level. Hospital co-insurance, doctor's office co-pay, and out-of-pocket premium equations are all estimating using Tobit models and missing values are replaced with appropriate predictions.

1.2.4 Prescription Drug Cost-Sharing

The employer questionnaire used to gather insurance information asks whether each plan "covers" outpatient prescriptions (99% of all plans in the sample do), but does not ask the co-insurance or co-pay level, whether there is a separate deductible for prescription drugs, or whether cumulative deductibles and stop losses apply to prescription drugs expenditure. Therefore, I assume that HMO, PPO, and FFS plans have a fixed prescription drug co-insurance rate of 13%, 17%, and 19% respectively. These rates are consistent with the average rates in the 1996 MEPS Abstraction file.[11] I also assume that prescription drug expenditure is completely unrelated to a plan's deductible and stop loss.[12] This assumption is also informed by the 1996 MEPS Abstraction file, which finds that prescription drug expenditure had no relation to a deductible for 77% of held plans and had no relation to the stop loss for 44%. According to the Kaiser Family Foundation's Employer Health Benefits Annual Survey, by 2000 prescription drug expenditure had no relation to the stop loss in over 75% of ESHI plans.

1.3 Medical Care Consumption and Pricing

Survey participants report all medical care consumption that takes place during an interview period. To the best of their ability, individuals provide the date and location of consumption, illness treated, procedure performed, price charged, and the cost-sharing arrangement with their insurance company. Given the length of time covered by an interview period and the number of details associated with any consumption episode, the ability of individuals to recall this information is questionable. To counteract the recall problem, survey administrators contacted all reported medical facilities to verify as much information as possible.[13]

[11] Co-pays are a much more popular form of cost-sharing for prescription drugs than co-insurance rates (68% vs. 32% in 1996). However, co-pays make the number and timing of refills a relevant factor in analysis, which I would like to abstract from. Further, 80% of all ESHI plans feature multi-tier prescription drug coverage in the form of co-pays by 2000 (Kaiser EHBAS). Thus, correctly implementing co-pays for prescription drug coverage would require both a quantity and quality decision by individuals.

[12] *Unrelated* means that the insurer and insured share the total cost of prescription drugs from the first day of an insurance year to the last, irrespective of accumulated expenditure. Also, out-of-pocket prescription drug expenditure does not contribute to the accumulated expenditure relevant for the cumulative deductible or stop loss.

[13] Contact with medical providers was made via telephone interview and mailed survey materials. Providers were contacted for every consumption episode, except office-based visits. AHRQ only contacted office providers if (1) the patient received Medicaid, (2) the household held a managed care plan and was selected as part of a 75% random sample, or (3) the household did not hold an managed care plan or Medicare, but

For each medical care type, AHRQ provides both imputed (File 1) and non-imputed (File 2) data, where the former contains episode level consumption variables that combine individual and provider responses. File 1 is used in my analysis. Note that contacting providers increases the reliability of all variables associated with particular visit, but individuals are relied upon entirely to report visits. If they fail to report a particular visit, no medical provider is contacted to verify their absence. The information below describes how I categorize reported medical care consumption into specific types, how consumption units are priced, and how prescription drug consumption dates are determined.

1.3.1 Doctor and Hospital

MEPS classifies all medical care consumption into inpatient, outpatient, ER, office-based, and prescription drug. For the purposes of this research, inpatient, outpatient, and ER visits are all classified as hospital care. A single trip to the ER or single outpatient visit is equivalent to one day of an inpatient stay, in terms of consumption decision and price draw. An inpatient stay of 5 days constitutes 5 separate decisions to consume medical care. Identifying doctor's office visits is straight-forward in the data. The dates of doctor visits and hospital days are consistently reported in the data (99.9% of visits are accompanied by a date). For the few doctor visits and hospital days that are missing a date, a month is selected at random from the appropriate interview period.

Doctor and hospital prices are observed in the data only when an individual consumes care. Therefore, an individual only receives a price contribution to his likelihood function in months when he chooses to consume.[14] The medical care price information needed for the likelihood contribution is taken directly from the data. The total price paid for medical care, which is the sum of the insurer's payment and insured's payment, is used as the total price. If an individual consumes multiple units of the same type of medical care in a month then the average price paid for that month is used, as each individual can only have one price contribution per month.

There are two issues with the medical care pricing data that must be addressed. The first issue is bundle pricing. At times, rather than pricing individual doctor visits or hospital days, a price is set for a fixed number of visits or for general treatment of an ailment. This pricing strategy is rarely seen in the data (283 of the 13,819 visits reported in the data are priced as part of a bundle) but must be dealt with. I assume that all visits in the bundle are

was selected as part of a 25% random sample.

[14]This data feature is similar to the labor literature, where wages are only observed for individuals who choose to work. In the labor literature, an employment decision is included in estimation to control for the endogeneity of wages (i.e., those who are likely to earn high wages are also likely to work). The model presented in this paper controls for the endogeneity of prices (i.e., those who are likely to receive low price draws should also be likely to consume) by modeling a consumption decision and be allowing for permanent unobserved correlation between the medical care preferences and price draws.

chosen independently. If all visits in the bundle are of the same type (i.e., doctor or hospital) then I assume that the price is spread out equally over each visit. If the consumption bundle consists of both hospital days and doctor visits, then an 8/1 price ratio is assumed between hospital and doctor prices.[15]

The second issue is with observed medical care prices for the uninsured. It should be the case that the list (charged) price is equal to the transaction (paid) price for all uninsured individuals, but this is not the case. (See footnote 2 in the paper for terminology.) At times, the uninsured negotiate directly with their doctors for lower payments or simply fail to pay a medical bill, which is observed as zero payment in the data. These payments are an issue because I use the observed amount paid for medical care as the total price in the model, meaning (without correction) the uninsured face lower average medical care prices. Left unchanged, these low prices would incorrectly incentivize zero coverage in simulation. In theory, the model should include negotiation costs, medical care quality, and individual credit-scores, which would capture the costs associated with the low payments of the uninsured; however, this would significantly complicate the model. Instead, I use the list price as the total price for the uninsured in estimation.

1.3.2 Prescription Drug

I assume that an individual decides whether or not to consume any prescription drugs each month, not the number of prescriptions to fill. In addition to making estimation more tractable, this assumption acknowledges the doctor's role in the prescription drug decision. Frequently, doctors prescribed multiple drugs such that consuming two, as opposed to one, prescription in no way reflects a marginal decision by an individual.

The dates of prescription drug consumption are much more difficult to extract from the data than the dates of doctor and hospital visits. The model requires that for each month I observe whether or not any prescription drugs are consumed and total drug spending. The goal is then to take from the data the beginning and ending consumption months for each prescription, as well as the total amount spent on each prescription over all consumption months. This total is divided evenly over the consumption months. The prescription price each month used in the model is then the sum of all prescription expenditures within the month.

The first month that a particular prescription is filled can be found in the data for most drugs (73%). For those missing a date, I match the prescriptions to illnesses and medical care consumption in the same interview period using ICD-9-CM codes.[16] By assuming the

[15] This ratio is consistent with the ratio of average prices observed in the data.

[16] Every illness and medical care consumption entry in the data file contains an ICD-9-CM condition code. These codes are used throughout the data cleaning process to match consumption to illness and vice versa.

prescription drug is first taken in the month that the illness began or in the same month as a doctor or hospital visit, I can impute some prescription drug beginning dates (85% have a beginning date after this assumption). For the remaining prescriptions, I use the number of refills in each interview period and the beginning and ending months of each interview period to randomize a starting date between reasonable bounds.

The last month that a drug is taken is more difficult to determine. The interview period containing each refill of a prescription can be found in the data; however, only the initial fills are accompanied by an exact date. Thus, I know the date that each individual began medication, the number of refills, and the interview periods the refills were purchased in. Using this information and the quantity of each prescription, which is usually the number of pills and strength of each pill, I approximate the number of months that each refill covers. For the 21 most popular drugs in the data set (see below) I was able to use average dosage to approximate the number of months.[17] The approximation provided a length for 2,422 out of the 13,592 total fills/refills purchased over the 2 year period.[18]

For less popular drugs, I use two methods for determining the last month that a drug is taken. First, for drugs in which every refill is the same quantity and consumption spans multiple interview periods, I use the average prescription length in periods prior to the last to calculate an average refill length. I then apply the average length to refills in the last interview period, which provides a length for 4234 additional fills/refills. For single prescription fills and for prescriptions with varying refill pill counts I was advised by a medical professional to make the following assumptions:

1. For prescriptions treating acute illnesses, assume fills/refills are prescribed for one month.

2. For prescriptions treating chronic illnesses with less than 90 pills, assume fills/refills are prescribed for one month.

3. For prescriptions treating chronic illnesses with more than 90 pills or not in a pill form, assume fills/refills are written for three months.

The first assumption is based on the fact that doctors rarely prescribe multi-month prescriptions for acute illnesses. Such illnesses are curable; thus, if the illness continues

[17] Average dosage is taken from "Mosby's: Pharmacology in Nursing." All assumptions were verified by a physician.

[18] This procedure (and others in this section) was done for all fills/refills over the 2 years that an individual was interviewed, even though I only use 1 year of information in estimation. I begin with data covering two years for several reasons. Most importantly, consumption of a prescription in the second year, which began in the first year, tells me that the ending consumption date for that drug runs past the end of the year, providing a solid ending date in the model. Also, because individual's enter the estimation sample in the month that their held plan begins, imputation is easier when all prescriptions over the two-year period are analyzed.

Drug	**Average Dosage**
Birth Control	Monthly packages of 21, 28, and 30 pills
Zoloft	50 - 150 mg/day
Zocor	5 - 10 mg/day
Zestril	10 - 40 mg/day
Vasotec	5 - 40 mg/day
Synthroid	0.125 - 0.5 mg/day
Prozac	10 - 40 mg/day
Premarin	0.3 - 1.25 mg/day
Paxil	20 - 50 mg/day
Naproxen	250 - 500 mg/day
Lotensin	5 - 40 mg/day
Glucophage	500 - 800 mg/2-3 times per day
Flonase	2 sprays/day
Claritin	10 mg/day
Atenolol	10 - 100 mg/day
Amoxicillin	500 - 800 mg/2-3 times per day
Toprol	25 - 100 mg/day (max 400)
Ranitidine	150 mg/2 times per day
Provera	5 - 10 mg/day
Norvasc	5 - 10 mg/day
Ibuprofen	300 - 800 mg/3-4 times per day

for more than a month after a prescription has been taken, it is likely that the illness has been misdiagnosed, meaning other treatments need to be explored. The second and third assumptions result from several common prescribing practices. First, many insurance companies set the maximum prescription size that they will cover for particular drugs at 30 or 90 days. As such, 1 month and 3 month prescriptions have become standard prescription lengths for physicians. Second, doctors rarely prescribe less than one pill per day because keeping up with the medication schedule is difficult for patients. Third, longer prescriptions are often prescribed for patients with chronic illnesses because the drug is meant to control the illness, not cure it. Prescriptions are then written for longer periods of time with fewer check-ups needed between refills.

These assumptions provide a length for every fill/refill in the sample. These lengths are then added together to make a total length for each prescription drug consumption episode. To validate this technique, I review the ending month assigned to each drug to ensure that it falls in the correct interview period, which is observed. The test results in a 90% success rate. For those predicted to continue their prescription consumption beyond their last known interview period where consumption was reported (i.e., the 10% that failed the test above), interview beginning and ending months and number of refills in the final

period are used to randomly impute a more realistic ending month.

1.4 Illness

In addition to reporting all medical care consumption in each interview, individuals are asked to report all "health problems (experienced during the current interview period) including physical conditions, accidents, or injuries that affect any part of the body as well as mental or emotional health conditions, such as feeling sad, blue, or anxious about something." Participants are told explicitly to include ailments even if they did not seek professional medical care. The individual's description of the illness is recorded as verbatim text, which is later coded to 5-digit ICD-9-CM codes by professional coders. If an illness is identified during the interview as a "priority condition" then an expanded set of questions is asked, including the exact date the illness began.[19] In all interviews after the first, participants are reminded of illnesses reported in prior interview periods and are asked if the illness has "bothered" them since the beginning of the current interview period. Therefore, one can determine the interview period in which the illness began and ended for all non-priority conditions. For priority conditions, one can determine the exact month that the illness began and the interview period in which it ended.

Before moving on to describe how illnesses are classified and how more precise illness dates are determined, the likelihood of unreported illnesses should be considered. In addition to standard recall problems, these data do not contain undiagnosed illnesses. Individuals cannot know about the presence of some illnesses, especially certain chronic illnesses (e.g., hypertension, breast cancer, heart disease, etc.), without the illness being identified by a doctor. The model assumes that illnesses arrive as random (though endogenous) shocks, which do not require a doctor to be identified. Empirically, the missingness of undiagnosed illnesses may be problematic if those rarely seeking medical care, and therefore not being diagnosed, appear healthier than they actually are.[20] If the objective of this research were to study the role of medical care consumption in determining health over the life course, then this issue of undiagnosed illness would be especially concerning. Data that required participants to take a physical examination, which would reduce the probability of undiagnosed illness, may be more valuable for such an objective. However, the focus of this research is the measurement of moral hazard within a single health insurance year. Endogenous health outcomes are modeled to correct potential bias in moral hazard estimates, but correctly predicting long-term health outcomes is not among the principle objectives. Therefore, for the purposes of this research, the most valuable measure of illness is the measure that an

[19] Certain conditions were designated as priority conditions by AHRQ because their prevalence, expense, or relevance to policy called for further inquiry. Some examples are cancer, diabetes, HIV/AIDS, hypertension, arthritis, stomach ulcers, and back problems. A full list can be found on the MEPS website.

[20] I thank Jessica Vistnes and Steven Hill of AHRQ for their thoughts and suggestions on this topic.

individual is aware of. In the short-run, individuals are most influenced by the illnesses that they actually know they have.

1.4.1 Classification of ICD-9-CM Codes

To estimate the model, all illnesses must be categorized as acute or chronic. I classify illnesses using the Chronic Condition Indicator (CCI) program, which identifies chronic illnesses by 5-digit ICD-9-CM codes.[21] For the purposes of this research, a chronic illness differs from an acute illness in that it is assumed to never fully subside. The program defines a chronic illness as "a condition that lasts 12 months or longer and meets one or both of the following tests: (a) it places limitations on self-care, independent living, and social interactions; (b) it results in the need for ongoing intervention with medical products, services, and special equipment." I must then reclassify some of the illnesses observed in the data (see below).[22]

A few of these changes require explanation.[23] According to the Centers for Disease Control and Prevention (CDC), 16%, of individuals 14-49 have some form of the Herpes Simplex virus (Genital Herpes), yet most do not know they have it. While it satisfies the definition of chronic illness imposed by this research in that the disease cannot be fully eliminated, it behaves like an acute illness. Infected individuals may experience outbreaks a few times a year, but potentially not at all. During these episodes individuals use medical care to treat the symptoms, but outside these outbreaks live normal lives. Furthermore, the rare occurrence of this illness in the data suggest that the individuals only report the virus during an outbreak, meaning it is also reported as an acute illness. Carpel Tunnel syndrome is coded as an acute illness because treatment (splint, physical therapy, or surgery) can restore the wrist to full health. The other respiratory disease recode applied to only one individual, who reports having had this disease for 15 years.[24]

[21] The Chronic Condition Indicator was developed as part of the Healthcare Cost and Utilization Project (HCUP), which is sponsored by AHRQ.

[22] This list is not an exhaustive list of all ICD-9-CM codes that should be reclassified given the difference in chronic illness definitions. The illnesses listed here are only the *observed* illnesses that need to be changed.

[23] None of these assumptions are likely to have a large impact in estimation. Each of these changes affects fewer than 25 individuals, except Acute Reaction to Stress and Chronic Sinusitis, which still affect fewer than 75 individuals.

[24] Other notes: For the purposes of this research, near/far sightedness and astigmatism are not considered illnesses though they appear frequently in the condition files. Only 198 of the 6208 reported medical conditions are missing an ICD-9-CM code. For these individuals, I look at medical care consumption in the same periods to try to infer what the illness was. If there is no medical care consumption in the period and the illness is present in every period over the two year span, I assume that it is a chronic illness. Otherwise, I assume that it is an acute illness.

ICD-9-CM	Illness	Changed to
054	Herpes Simplex	Acute
239	Unspecified cancer	Chronic
308	Acute reaction to stress	Acute
309	Adjustment reaction	Acute
354	Carpel Tunnel	Acute
360	Disorder of the globe	Acute
436	Acute Cerebrovascular Disease	Acute
473	Chronic Sinusitis	Acute
474	Chronic disease of the tonsils and adenoids	Acute
519	Other respiratory disease	Chronic
562	Diverticula of the intestines	Acute
625	Premenstrual Syndrome	Acute
730	Unspecified Osteomyelitis	Acute
V10	Past cancer	Chronic

1.4.2 Illness Beginning and Ending Dates

The model requires that the data reveal the total number of acute and chronic illnesses an individual has in each of the 12 months spanning the insurance year. To determine this information, I need to know the beginning and ending month of each reported illness. For every illness reported in the data, it can be determined in which *interview period* the illness began and ended, but not necessarily the *month*. Thus, in what follows, I describe the procedures used to impute the beginning and ending month for each illness. I begin with a few simplifying assumptions. First, all congenial diseases observed in the data are assumed to begin prior to the beginning of the insurance year.[25] Second, I assume that all illnesses reported in consecutive interview periods represent one continuous illness. Furthermore, throughout the file there are examples of the same illness being reported multiple times by the same individual with gaps in reporting (e.g., had illness during 1st and 3rd interview). At times, it is reasonable to assume that the multiple records describe one continuous illness. Most of the time, it is unlikely that the illness lasted 6 or 7 months, so it is best assumed that the illness occurred on two separate occasions (e.g., a common cold). I resolve these issues case by case. In total, there are 414 instances where such a *gap rule* is needed. Rules are generally made at the illness level (e.g., any gap for a common cold implies two illnesses), for which there are 100. After these edits, I observe 5,586 independent illnesses for the estimation sample over the 2 year period of study. I add to these illnesses 37 illness records that are necessary to match medical care consumption information (e.g., I observe that an individual takes insulin in every month but has no record of diabetes). I also remove any

[25] I observe the following ICD-9-CM codes: 747, 753, 755, 757, 758, 759.

illness known to begin after the end of the insurance year. This step reduces the number of illnesses to 4,482.

Of these illnesses, 29% are classified by AHRQ as "priority conditions" (mostly chronic and important acute illnesses), so the month that the illness began is known. A beginning date must be imputed for the remaining illnesses. To aid in imputation, I first assume that if the individual consumed medical care for an illness in the first interview period that an illness is reported, then the month of consumption is the month that the illness began. This assumption is equivalent to assuming that people do not wait more than a month to go to the doctor if they are going to go at all. After this assumption, 1,299 illnesses are in need of a beginning month. For these, a beginning data is drawn uniformly between the first and last month of the interview where the illness was first reported. For the first interview period, I account for the fact that the illness may have begun prior to the start of the year. This randomization procedure results in 302 illnesses having start dates that are past the end of the insurance year. They are dropped and 4,180 illnesses remain.

Illness ending months are more difficult to extract because they are never observed explicitly. I only observe the last interview period in which an illness was reported. Fortunately, chronic illnesses are assumed to never end, so only acute illnesses need an approximated ending month. Further, many acute illnesses begin during the 12 months of the insurance year, but end in an interview period after the close of the insurance year, meaning no ending month is required for the model. Thus, only 2,592 illnesses need an ending date. Randomization is not necessary for 343 illnesses because they start and end in the same interview period and the beginning date is the last month of the period. There are then only 2,249 illnesses that need an ending month to be imputed. For these, I draw an ending month at random from the months included in the interview period that the illness reportedly ended in. If an illness ends in the same interview period that it begins, the lower bound is set to the beginning month.[26] The distribution of bound lengths is shown below:

[26]The lower bound could be increased, which would decrease the draw range, by setting it equal to the month of last corresponding consumption in the last interview period that the illness was reported. I did not use this strategy because it almost certainly would have lead estimates to reflect that medical care increases the length of illness, as those consuming medical care would on average have longer illness periods than non-consumers by construction. With the method used, the productivity of medical care is identified by the data alone. If medical care helps an individual avoid illness in the following interview period, then it is productive. If it does not, then medical care may be harmful.

Bound Length	Frequency	Cumulative Percentage
1	355	15.8%
2	437	35.2%
3	391	52.6%
4	368	69.0%
5	371	85.5%
6	197	94.2%
7	101	98.7%
8	29	100.0%
Total	2249	

1.5 Month-to-Month Variation in State Variables

Table 1: Number of Acute Illnesses from Month t to Month $t+1$

			Transition Probability in $t+1$						
		Obs	0	1	2	3	4	5	6+
t	0	8,669	89.4	9.5	0.9	0.1	0.0	0.0	0.0
	1	3,918	20.8	69.1	8.9	1.2	0.1	0.0	0.0
	2	1,461	5.8	23.4	60.5	8.8	1.4	0.0	0.0
	3	473	1.7	7.6	27.5	52.9	9.5	0.9	0.0
	4	177	0.0	1.7	10.7	22.0	53.1	10.7	1.7
	5	53	0.0	1.9	3.8	7.6	30.2	39.6	17.0
	6+	33	0.0	3.0	0.0	0.0	12.1	24.2	60.1

Table 2: Number of Chronic Illnesses from Month t to Month $t+1$

		Transition Probability in $t+1$						
	Obs	0	1	2	3	4	5+
0	8,309	98.5	1.4	0.1	0.0	0.0	0.0
1	3,481	0.0	97.4	2.4	0.2	0.0	0.0
2	1,686	0.0	0.0	96.9	2.9	0.2	0.1
t 3	846	0.0	0.0	0.0	96.9	2.8	0.2
4	195	0.0	0.0	0.0	0.0	95.9	4.1
5+	267	0.0	0.0	0.0	0.0	0.0	100.0

Table 3: Number of Doctor's Office Visits from Month t to Month $t+1$

		Transition Probability in $t+1$						
	Obs	0	1	2	3	4	5	6+
0	11,863	85.4	10.2	2.4	0.9	0.3	0.1	0.3
1	1,763	61.6	24.7	9.0	3.0	1.0	0.5	0.3
2	566	44.9	26.5	15.2	7.4	2.5	1.8	1.8
t 3	252	25.0	25.4	16.7	12.3	7.9	5.1	7.5
4	125	15.2	20.8	15.2	13.6	19.2	6.4	9.6
5	63	11.1	17.5	19.1	12.7	12.7	4.8	22.2
6+	104	7.7	8.7	11.5	17.3	14.4	6.7	33.7

Table 4: Number of Hospital Days from Month t to Month $t+1$

		Transition Probability in $t+1$						
	Obs	0	1	2	3	4	5+	
0	14,342	97.4	2.0	0.2	0.1	0.1	0.2	
1	306	81.7	12.4	4.6	1.0	0.7	0.7	
2	51	62.8	13.7	15.7	4.0	2.0	2.0	
t 3	26	80.8	11.5	0.0	0.0	3.9	3.9	
4	20	65.0	5.0	5.0	10.0	10.0	5.0	
5+	39	66.7	5.1	0.0	7.7	2.6	18.0	

Table 5: Prescription Drug Consumption from Month t to Month $t+1$

		Trans. Prob. in $t+1$	
	Obs	0	1
t 0	9,711	89.7	10.3
1	5,073	11.5	88.5

1.6 Out-of-Pocket Share

Figure 1: Average Share of Medical Care Price Paid Out-of-Pocket

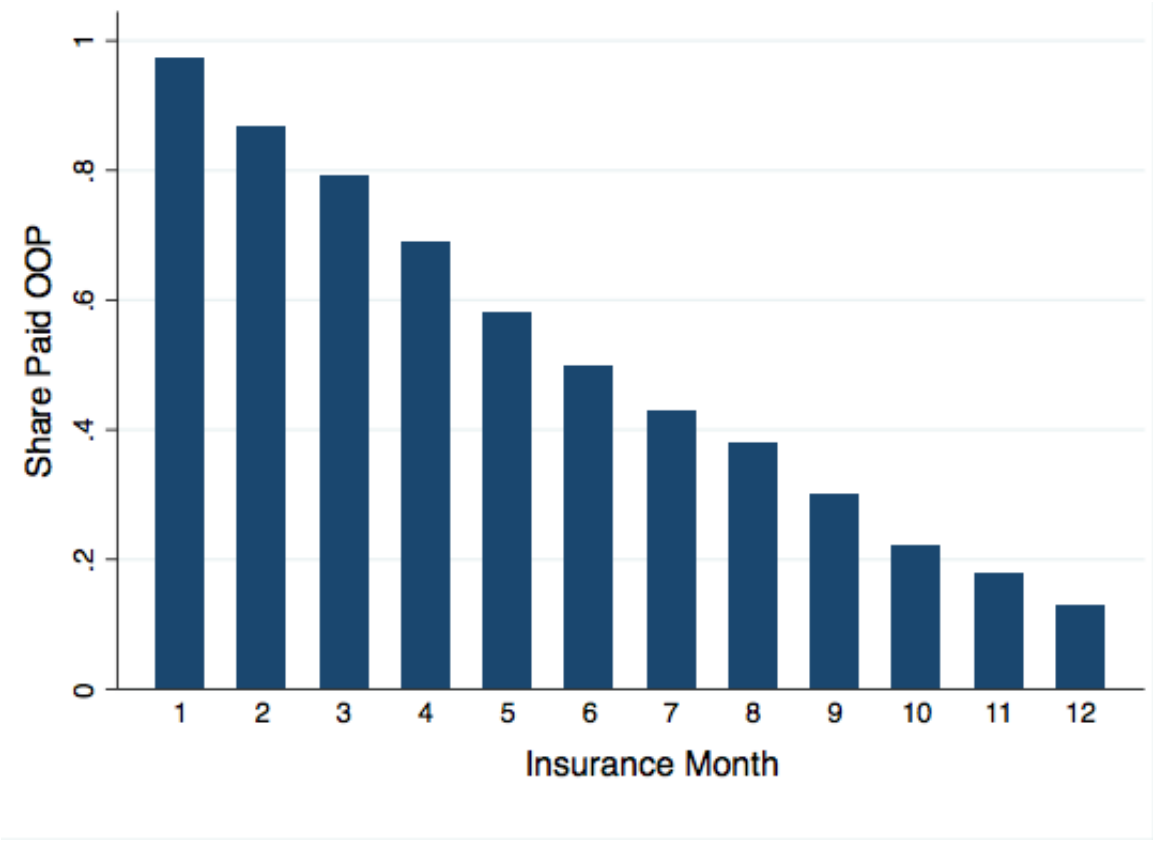

Notes: This figure shows the average share of medical care price paid out of pocket in each insurance month for the 215 individuals who experience a change in this share at some point during the insurance year.

2 Selection of Parametric Price Distributions

The model assumes that an individual makes medical care consumption decisions facing unknown prices; thus, in order to solve the dynamic optimization problem, the individual must integrate over the three price distributions. Because decisions are made based on the full price distribution rather than a single price draw, it is important that the parametric distribution chosen to model each price closely matches the empirical distribution observed in the data. It is particularly important that the long, thin right tail of the price distributions is captured (see figures below), as potentially large price draws motivate both the purchase of health insurance and provide a disincentive for frequent medical care consumption.

Figure 2: Empirical Price Distributions

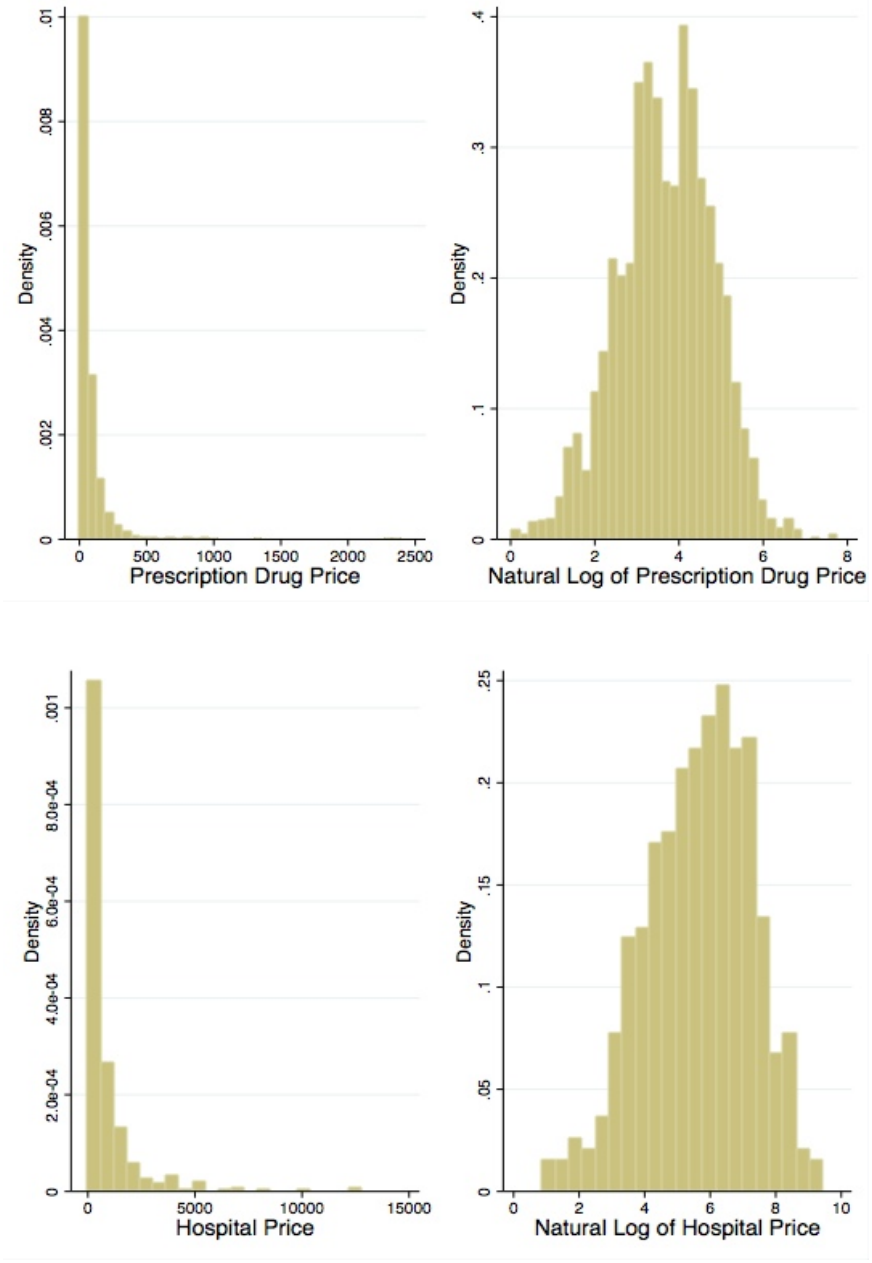

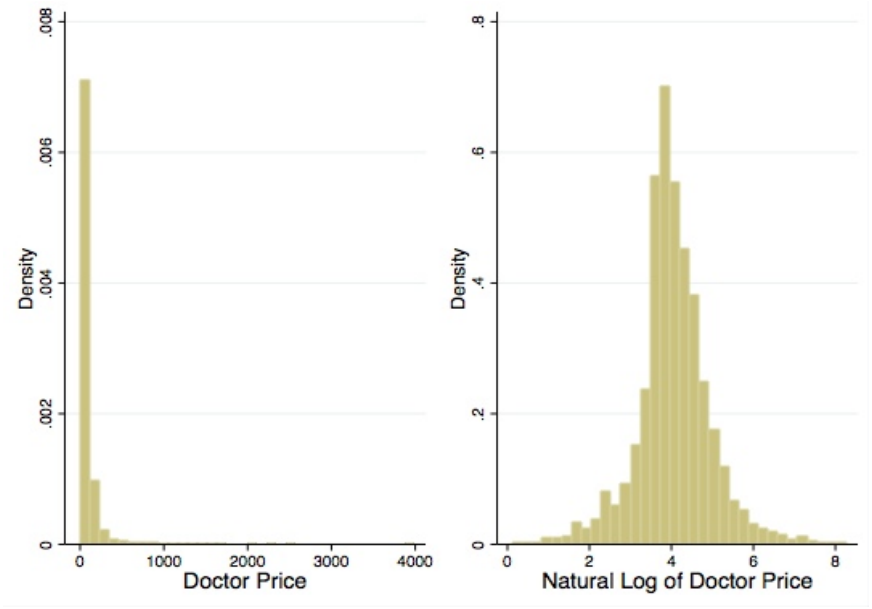

For each type of medical care, I test the fit of eleven commonly used continuous distributions to the empirical price distribution. The task of selecting appropriate distributions is complicated by the structural model. Estimating the model is time-intensive, so testing various distributional assumptions for each medical care price within the full model is not feasible. Instead, I perform the distributional tests outside the model, implicitly assuming that (1) the unobserved correlation between endogenous variables (i.e., μ) and (2) control for endogenous selection into treatment in the full model would not cause one distributional assumption to dominate another.

The eleven distributions I test are shown in the table below. Jones et al. (2014) test the same distributions, plus the generalized gamma distribution, in their analysis of English hospital inpatient cost data. The first distribution, generalized beta of the second kind (GB2), is a four parameter distribution developed and studied extensively by McDonald (1984). The distribution is frequently used to model income distributions, which are also characterized by long, thin right tails. All other distributions considered are nested and limiting cases of the GB2 distribution. Each distribution is characterized by one scale parameter and between zero (exponential) and three (GB2) shape parameters. As in the full model, I allow the scale parameter to vary by observable covariates (see Table 9 in the paper for a list of covariates).

The Singh-Maddala, Dagum, Beta of the Second Kind, Lomax, and Fisk distributions are each nested within the GB2. As such, WALD and/or Likelihood Ratio tests can be used to determine goodness of fit for these distributions relative to the GB2. Akaike Information Criteria (AIC) and Bayesian Information Criteria (BIC) allow all eleven models to

be compared. These statistics are presented in the table below.[27]

Distribution	**Rx Price**		**Log Hospital Price**		**Doctor Price**	
	AIC	**BIC**	**AIC**	**BIC**	**AIC**	**BIC**
Generalized Beta of the Second Kind	57016.33	57155.18	1805.26	1892.64	33334.57	33461.80
Singh-Maddala	57067.20	57199.44	1811.59	1894.82	33461.00	33582.16
Dagum	57120.65	57252.89	1805.36	1888.58	33470.77	33591.94
Beta of the Second Kind	57032.74	57164.98	1882.48	1965.70	33637.00	33758.16
Lomax	57347.75	57473.37	2628.36	2707.43	34518.68	34633.79
Fisk	57152.48	57278.10	1890.84	1969.90	33469.39	33584.50
Gamma	57468.9	57594.56	1880.47	1959.53	34654.04	34769.15
Log-Normal	56933.81	57059.44	1942.36	2021.42	33778.93	33894.04
Weibull	57490.60	57616.23	1809.05	1888.11	34771.34	34886.45
Exponential	57489.03	57608.04	2626.36	2701.26	34770.40	34879.45
Normal	67127.40	67253.02	1808.84	1887.90	41027.58	41142.69

The AIC and BIC criteria clearly identify the log-normal distribution as the best fit for the empirical prescription drug price distribution. I experienced several problems with non-concave regions when the eleven parametric distributions were fit to the empirical hospital price distribution; thus, analysis was conducted on log-hospital prices. Using log prices, the AIC identifies the GB2 distribution as the best fit, followed by the normal distribution; however, the BIC identifies the normal distribution as the best fit, followed by the Weibull and Dagum. Though the AIC and BIC do not identify the same distribution as the most preferred, these results suggest that the normal distribution fits well relative to the alternatives. Note also that fitting the normal distribution to the empirical log-price distribution is equivalent to fitting the log-normal distribution to the unlogged empirical price distribution. The latter is done in the paper. Finally, the last two columns of the table identify the GB2 distribution as the best fit for the empirical doctor price distribution, followed by the Singh-Maddala distribution. Note, however, that the GB2 distribution does not have a CDF with a closed form, meaning sampling from this distribution requires advanced sampling methods. These sampling methods, such as the Accept-Reject Method, can be prohibitively time-consuming in estimation; thus, I used the Singh-Maddala distribution in estimation, which does have a CDF with a closed form.

The next table compares various percentile values of the observed distribution with their simulated counterparts using the distributional assumptions discussed above. The

[27] WALD and Likelihood Ratio results are available upon request. Evaluation of goodness of fit using these tests is consistent with AIC and BIC results.

simulated prescription drug and doctors office distributions each fit the observed data well. The simulated hospital distribution fits the observed distribution fairly well, except for in the right tail, where it is heavier. The simulated hospital price distribution reported below has been truncated at the maximum observed hospital price in the data. Without this truncation, the simulated hospital data have a maximum of $43,381.58, mean of $1,151.18, and standard deviation of $3,362.99.

	Rx Distribution		**Hospital Distribution**		**Doctor Distribution**	
	Observed	Simulated	Observed	Simulated	Observed	Simulated
mean	74.06	75.75	828.43	1000.47	90.56	88.73
s.d.	112.91	108.75	1437.44	2077.06	159.23	172.95
min	1.03	0.43	2.41	2.48	1.15	1.16
1%	2.57	2.79	4.84	3.86	5.00	8.50
5%	5.69	5.97	18.58	19.71	12.62	16.45
10%	9.33	9.13	33.00	34.48	23.17	21.99
25%	19.51	18.80	94.16	85.48	38.00	35.05
50%	41.62	40.67	329.60	249.25	55.00	56.88
75%	88.06	87.47	937.64	878.97	91.17	96.41
90%	169.63	169.56	1942.06	2521.07	154.77	163.79
95%	323.79	258.58	3554.83	3554.83	244.70	233.40
99%	467.73	559.66	7180.39	12834.00	739.50	549.27
max	2399.62	1368.19	12835.9	12834.00	3995.00	4836.05

3 Out-of-Pocket Expenditure Equation

Out-of-pocket expenditure resulting from v_t trips to the doctor, s_t hospital days, and consuming prescription drugs r_t in month t can be calculated as a function of accumulated out-of-pocket doctor's office expenditure entering the month, ADE_t; accumulated out-of-pocket hospital expenditure entering the month, AHE_t; the month t total price of a doctor's office visit, hospital days, and prescription drug consumption, p_t^v, p_t^s, and p_t^r, respectively; and the cost-sharing features of one's insurance plan. The cost-sharing features that may impact the out-of-pocket price of care are as follows:

- D_1^d = doctor's office specific deductible
- D_1^h = hospital specific deductible
- D_2 = stop loss
- C_1^d = doctor's office co-insurance rate

- C_2^d = doctor's office co-pay level
- C_1^h = hospital co-insurance rate
- C_2^{h1} = hospital co-pay level defined by day
- C_2^{h2} = hospital co-pay level defined by stay
- C_3 = prescription drug co-insurance level

To ensure that the out-of-pocket price of an additional unit of care is captured regardless of the unique combination of insurance features a plan contains, I categorize each plan into one of four general types and allow different pieces of the out-of-pocket function below to change depending on the plan type. These four types are as follows:

- Type 1: composite deductible (or no deductible at all) and hospital co-pay defined by day (or hospital co-pay does not at all exist).
- Type 2: composite deductible (or no deductible at all) and hospital co-pay defined by stay.
- Type 3: separate deductible(s) and hospital co-pay defined by day (or hospital co-pay does not at all exist).
- Type 4: separate deductible(s) and hospital co-pay defined by stay.

The strategy for calculating out-of-pocket expenditure is to separately determine the amount of pre-deductible and post-deductible expenditure first for doctor visits and then for hospital days. Doctor and hospital expenditure must be derived separately because the cost-sharing features of nearly every plan require individuals to pay a different proportion of the total cost for each service. Two assumptions are needed to calculate out-of-pocket expenditure each month. First, it is assumed that the total price charged for a doctor visit and the total price charged for a hospital day are fixed for an individual within any month. This assumption makes the problem empirically tractable. As mentioned, an individual solves his optimization problem by integrating over three distributions of possible prices. If this price were allowed to vary for each trip to the doctor/hospital in a month, then a $V * S$ dimensional integral would need to be solved in every month. Further, given that the factors influencing the total prices one faces in a month are not changing, it seems reasonable to assume that the total price does not change. Second, the equation assumes that all doctor visits within a month are made before any hospital decisions take place. While this assumption is clearly not ideal, it is necessary if one hopes to avoid modeling the specific order in which an individual decides to visit the doctor and hospital, which would cause the

size of the alternative set to explode.[28]

Out-of-pocket expenditure is calculated as follows:

Pre-deductible spending on doctor's office care is:

$$o_t^{1d}(v_t, p_t^v, \Psi_t) = \begin{cases} p_t^v * v_t & \text{if } \ 0 \leq p_t^v * v_t \leq x_t^1 \\ x_t^1 & \text{if } \ 0 \leq x_t^1 < p_t^v * v_t \end{cases} \tag{1}$$

where

If Type 1: $x_t^1 =$ composite deductible remaining $= max\{0, D_1 - ADE_t - AHE_t\}$

If Type 2: $x_t^1 =$ composite deductible remaining $= max\{0, D_1 - ADE_t - AHE_t\}$

If Type 3: $x_t^1 =$ doctor deductible remaining $= max\{0, D_1^d - ADE_t\}$

If Type 4: $x_t^1 =$ doctor deductible remaining $= max\{0, D_1^d - ADE_t\}$

Regardless of type, post-deductible spending on doctor visits is:

$$o_t^{2d}(v_t, p_t^v, \Psi_t) = \begin{cases} over_t^d & \text{if } \ over_t^d \leq z_t^1 \\ over_t^d - z_t^1 & \text{if } \ over_t^d > z_t^1 \end{cases} \tag{2}$$

where

$over_t^d = max\left\{0, int\left\{\frac{(p_t^v * v_t) - x_t^1}{p_t^v}\right\} * C_2^d\right\} + max\left\{0, (p_t^v * v_t) - x_t^1 - int\left\{\frac{(p_t^v * v_t) - x_t^1}{p_t^v}\right\} C_2^d\right\} * C_1^d$

$z_t^1 =$ stop loss remaining $= max\left\{0, D_2 - ADE_t - AHE_t - o_t^{1d}\right\}$

Total amount spent out-of-pocket on doctor visits in month t is:

$$O_t^d = o_t^{1d} + o_t^{2d} \tag{3}$$

[28] Order is important because it determines when exactly during the month the deductible or stop loss is passed. Alternatives to this assumption are: (1) Force an individual to select not only the number of visits but the order as well; however, this is likely to be empirically infeasible. (2) Stipulate that non-linear changes in medical care prices that happen when one crosses the deductible or stop loss only occur between months, which would greatly simplify the budget constraint; however, it would remove important variation from the data.

Conditional on month t doctor spending, pre-deductible spending on hospital care is:

$$o_t^{1h}(s_t, p_t^s, \Psi_t) = \begin{cases} p_t^s * s_t & \text{if } \ 0 \leq p_t^s * s_t \leq x_t^2 \\ x_t^2 & \text{if } \ 0 \leq x_t^2 < p_t^s * s_t \end{cases} \tag{4}$$

where

If Type 1: $x_t^2 =$ composite deductible remaining $= max\{0, D_1 - ADE_t - AHE_t - O_t^d\}$

If Type 2: $x_t^2 =$ composite deductible remaining $= max\{0, D_1 - ADE_t - AHE_t - O_t^d\}$

If Type 3: $x_t^2 =$ hospital deductible remaining $= max\{0, D_1^h - AHE_t\}$

If Type 4: $x_t^2 =$ hospital deductible remaining $= max\{0, D_1^h - AHE_t\}$

Post-deductible spending on hospital care is:

$$o_t^{2h}(s_t, p_t^s, \Psi_t) = \begin{cases} over_t^h & \text{if } \ over_t^h \leq z_t^2 \\ over_t^h - z_t^2 & \text{if } \ over_t^h > z_t^2 \end{cases} \tag{5}$$

where

$z_t^2 =$ stop loss remaining $= max\{0, D_2 - ADE_t - AHE_t - O_t^d - o_t^{1h}\}$

$A = int\left(\frac{(p_t^s * s_t) - x_t^2}{p_t^s}\right)$, where $int(w)$ rounds w to the nearest integer

If Type 1: $over_t^h = max\left\{0, A * C_2^{h1}\right\} + max\left\{0, (p_t^s * s_t) - x_t^2 - A * C_2^{h1}\right\} * C_1^h$

If Type 2: $over_t^h = max\left\{0, min\{(p_t^s * s_t) - x_t^2, C_2^{h2}\}\right\} + max\left\{0, (p_t^s * s_t) - x_t^2 - C_2^{h2}\right\} * C_1^h$

If Type 3: $over_t^h = max\left\{0, A * C_2^{h1}\right\} + max\left\{0, (p_t^s * s_t) - x_t^2 - A * C_2^{h1}\right\} * C_1^h$

If Type 4: $over_t^h = max\left\{0, min\{(p_t^s * s_t) - x_t^2, C_2^{h2}\}\right\} + max\left\{0, (p_t^s * s_t) - x_t^2 - C_2^{h2}\right\} * C_1^h$

The total amount spent out-of-pocket on hospital care in month t is:

$$O_t^h = o_t^1 h + o_t^2 h \tag{6}$$

Out-of-pocket expenditure in month t is:

$$O_t = O_t^d + O_t^h + (r_t * p_t^r * C_3) \tag{7}$$

Accumulated doctor and hospital expenditure entering month $t+1$ is:

$$ADE_{t+1} = \begin{cases} ADE_t + O_t^d & \text{if } t > 0 \\ 0 & \text{if } t = 0 \end{cases} \tag{8}$$

$$AHE_{t+1} = \begin{cases} AHE_t + O_t^h & \text{if } t > 0 \\ 0 & \text{if } t = 0 \end{cases} \tag{9}$$

4 Annual Expenditure Model

This section describes the parameter estimates and model fit analysis of the representative annual expenditure (RAE) model. The RAE model is estimated using a nested fixed point algorithm, where an individual's dynamic programming problem is solved via backwards recursion in the inner loop and a BHHH gradient method is used to improve the likelihood function in the outer algorithm. Model fit is analyzed by comparing key moments in the observed and simulated data. The simulated data are constructed by sampling from the joint error distribution and permanent unobserved heterogeneity distribution 100 times for each individual.

4.1 Parameter Estimates

Table 6 reports estimated preference parameters. Note first that utility does not vary with acute and chronic illness in the RAE model as illness is exogenous, meaning the parameters on these variables (i.e., $\omega_{10}, ..., \omega_{22}$ in the within-year model) are not identified. The main difference in parameter estimates between the two models is that preferences for medical care consumption are greater in the RAE model. This result occurs almost by design. In the within-year decision-making model, people may consume medical care either because they like it (i.e., a preference motive) or because it improves health and they like health (i.e., an investment motive). In the RAE model, because illness is exogenously determined the only mechanism through which the model can rationalize positive medical care consumption in the data is through preferences, meaning preferences must be more positive.

The parameters in Table 7 seem to capture unobserved heterogeneity that is similar to what was found by the within-year model - the most representative unobserved group is Type 1; relative to Type 1 individuals, the other types have lower preferences for doctor and hospital visits, mostly higher preferences for prescription drugs, and pay lower prices at the doctor's office and hospital. I did allow for one additional mass point in the unobserved heterogeneity distribution in the RAE model simply because it fit the model a bit better and estimation was significantly faster than in the within-year decision-making model.

Price parameters can be observed in Table 8. Most parameters have the same signs that were observed in the within-year decision-making model, but magnitudes vary. Most notably, living in an MSA has a significantly smaller impact on price in the RAE model for both doctor visits and prescription drug consumption, while the sign of its effect flips from negative to positive for hospital visit prices. Probably the most noticable difference between the model parameters is the impact that being uninsured has on prices, which flips from negative in the within-year model to positive in the RAE model for both doctor and hospital visits.

Table 6: Preference Parameter Estimates

	Parameter	Estimate	SE
Utility Function			
RA constant	ω_{00}	0.0897	0.0068
RA age	ω_{01}	0.0014	0.0002
RA non-white (black or Hispanic)	ω_{02}	0.0304	0.0066
RA male	ω_{03}	-0.0279	0.0058
doctor visits	ω_{30}	0.7204	0.2183
doctor visits (squared)	ω_{31}	-0.7162	0.0942
doctor visits*age	ω_{32}	0.0000	0.0004
doctor visits*male	ω_{33}	-0.0331	0.0109
hospital days	ω_{40}	-2.8035	0.2005
hospital days (squared)	ω_{41}	0.3252	0.0512
hospital days*age	ω_{42}	0.0092	0.0022
hospital days*male	ω_{43}	0.1252	0.0496
any Rx consumption	ω_{50}	-1.1415	0.1261
any Rx consumption*age	ω_{51}	0.0203	0.0025
any Rx consumption*male	ω_{52}	-0.7768	0.0717
any consumption	ω_{60}	-2.8536	0.1564
any consumption*male	ω_{61}	-0.7185	0.0654
negative consumption†	ω_{70}	0.1000	
Other			
discount factor†	β	0.9960	
log-likelihood value‡	$L(\Omega)$	-47732.18507	

† Not currently estimated. Note that β is set to 0.996, instead of the traditional 0.95, because this is a monthly model. $0.996^{12} \approx 0.95$

‡ The log-likelihood value with only one unobserved mass point is -53052.5436.

Table 7: Permanent Unobserved Heterogeneity Parameter Estimates

	Param.	Type 1	Type 2		Type 3		Type 4		Type 5	
			est.	s.e.	est.	s.e.	est.	s.e.	est.	s.e.
Mass Point Location										
doctor visit preference (linear)	μ_{3a}	0.0	-1.689	0.218	-1.956	0.219	-1.720	0.222	-0.994	0.247
doctor visit preference (squared)	μ_{3b}	0.0	0.772	0.094	0.785	0.093	0.783	0.094	0.499	0.102
hospital day preference (linear)	μ_{4a}	0.0	-0.864	0.222	-0.932	0.240	-1.371	0.293	-0.109	0.243
hospital day preference (squared)	μ_{4b}	0.0	0.131	0.059	0.150	0.062	0.272	0.073	0.047	0.061
any Rx preference	μ_5	0.0	5.270	0.300	3.872	0.198	0.484	0.130	1.545	0.145
any consumption preference	μ_6	0.0	2.366	0.338	2.567	0.235	2.593	0.173	1.404	0.199
doctor visit price distribution	μ_7	0.0	-2.070	2.645	-3.478	2.978	-8.215	2.524	-2.343	2.933
hospital day price distribution	μ_8	0.0	-0.369	0.348	-0.977	0.357	0.607	0.433	-0.201	0.326
monthly Rx price distribution	μ_9	0.0	0.538	0.053	-0.917	0.053	0.404	0.057	-0.683	0.056
Type Probabilities†										
parameter estimate	θ	0.0	-0.874	0.087	-1.398	0.106	-1.211	0.115	-1.145	0.119
type probability		43.9	18.3		10.8		13.1		14.0	

† Probabilities are derived from estimation of parameter θ^k such that $Pr(\mu^k) = \frac{exp(\theta^k)}{\sum_{k'=1}^{4} exp(\theta^{k'})}$.

Table 8: Structural Price Parameter Estimates

	Doctor Price		Hospital Price		Prescription Price	
	est.	s.e.	est.	s.e.	est.	s.e.
constant	46.9878	5.7496	5.6409	0.6983	2.9301	0.0779
male	0.6260	1.2046	-0.0499	0.1998	0.1173	0.0135
non-white (black or Hispanic)	-0.8887	1.7526	0.0928	0.2179	-0.0729	0.0155
education (highest grade completed)	-0.2124	0.2743	-0.0442	0.0381	0.0084	0.0030
age	0.1399	0.0640	0.0254	0.0091	0.0027	0.0006
lives in a MSA	9.7953	1.5129	0.2321	0.2112	-0.0216	0.0173
income (in 1996 dollars)	0.0324	0.0312	0.0068	0.0061	0.0045	0.0003
initial health	0.1987	0.6573	0.0888	0.0900	0.1429	0.0072
HMO	-3.5723	1.2022	-0.0204	0.2035	-0.0234	0.0123
PPO	-5.0467	1.8801	-0.1811	0.3563	0.0872	0.0254
no insurance	40.7014	3.9699	0.7040	0.3152	-0.1435	0.0277
acute illnesses	-2.5374	1.2450	-0.0725	0.2128	0.0947	0.0137
acute illnesses (squared)	0.7118	0.2776	0.0609	0.0548	-0.0088	0.0032
chronic illnesses	-2.5370	1.0971	-0.0874	0.1564	0.3565	0.0112
chronic illnesses (squared)	0.3927	0.1728	-0.0024	0.0216	-0.0246	0.0015
shape 1†	2.4087	0.0437	1.2578	0.0686	0.84.02	0.0039
shape 2†	0.7753	0.0368	*	*	*	*

Month indicators are included in all regressions but are not reported here.
† The hospital and prescription drug price data are fit to log-normal distributions. The estimated *shape* parameters of these distributions correspond to the standard error of the distribution. The doctor price data is fit to a Singh-Maddala distribution, which has two shape parameters.

4.2 Model Fit

In large part, the RAE model fits the data in a way that is similar to the within-year decision-making model. Both prices and consumption fit fairly well (Table 9), though it is again difficult to explain large amounts of consumption both at zero and in the far right tail of the distribution in a model where all medical care must be rationalized (Table 10). The most significant deficiency of the RAE is its inability to explain such a small number of uninsured individuals. The problem is intuitive. In the RAE model, insurance reduces the out-of-pocket price of medical care for the insured; however, because individual's can observe prices, they can also protect themselves against high price draws by simply choosing not to seek care when prices are high. As there is no health in the model, the only consequence to not seeking care is preference-based, which is (1) mild, compared to the consumption sacrificed with a high price draw and (2) contemporaneous. In summary, insurance provides individuals little protection the RAE model because prices are known, so it is difficult for

the model to explain why individuals carry it.

Table 9: Observed and Simulated Outcomes

	Observed			Simulated		
	mean	median	s.d.	mean	median	s.d.
Medical Care Prices†						
doctor visit price	90.61	55.00	159.18	95.34	61.59	157.75
hospital day price	828.48	329.60	1435.91	800.58	312.71	1631.47
prescription drug price	74.37	41.93	113.63	73.75	39.73	113.26
Medical Care Consumption						
annual doctor visits	4.89	2.00	8.69	4.92	3.00	5.62
annual hospital days	0.65	0.00	2.00	0.63	0.00	1.37
months with Rx cons.	4.45	2.00	4.45	4.46	2.00	4.77
Medical Care Expenditure						
total annual medical care expenditure	1473.13	358.34	3656.88	1338.87	512.07	3664.40
Held Insurance Types						
% No Insurance	9.17			23.92		
% HMO	37.99			31.12		
% PPO	10.14			08.97		
% FFS	42.69			35.97		

† Prices are only observed when an individual consumes medical care; thus, the simulated mean and standard deviation are calculated only from individuals consuming care in simulation.

Table 10: Observed and Simulated Annual Consumption

	Doctor visits		Hospital days		Prescriptions	
	obs.	sim.	obs.	sim.	obs.	sim.
0	28.17	22.17	75.81	70.53	35.47	31.13
1	16.15	14.91	12.09	16.42	13.31	14.58
2	12.66	10.59	4.87	5.47	5.60	7.13
3-5	18.43	18.70	4.79	5.92	10.47	11.27
6-9	10.46	15.15	1.46	1.45	7.63	8.66
10+	14.12	18.15	0.97	0.20	27.52	27.05

The table reports annual consumption levels for the observed and simulated samples. The values are the proportion of the samples consuming at each level. For example, the top left entry states that 28.22% of the observed sample visited a doctor's office zero times during the insurance year. The prescription drug levels are measured in consumption months.

Figure 3: Annual Medical Care Expenditure

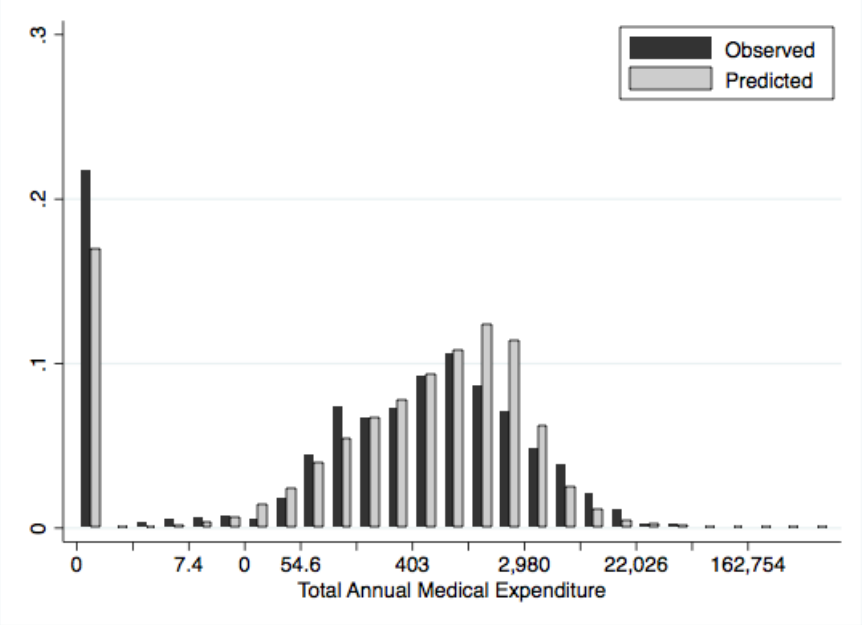

References

Jones, A. M., J. Lomas, and N. Rice (2014): "Applying Beta-Type Size Distributions to Healthcare Cost Regressions," *Journal of Applied Econometrics*, 649–670.

McDonald, J. B. (1984): "Some generalized functions for the size distribution of income," *Econometrica*, 647–663.