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

Web Appendix

Christopher J. Cronin
University of Notre Dame

February 8, 2018

Contents

1 Moral Hazard and Within-Year Decision Making

There are several reasons why one might expect the WYDM and an annual expenditure model to produce different moral hazard effects. One reason is that the WYDM model allows for endogenous health transitions that create dynamic incentives for medical care consumption within the insurance year. As health insurance acquisition leads to greater medical care consumption early in the year due to lower out-of-pocket prices, an individual's health may improve, lessening the need for medical care late in the year. Because this dynamic shift in the incentives for medical care consumption does not exist in an annual expenditure model, one might expect such a model to estimate a larger moral hazard effect.

Another reason the two models may produce different moral hazard estimates is that they make different assumptions about the information available to an individual when medical care decisions are made. For example, the WYDM model assumes that the effective price of medical care is unknown when medical care decisions are made.[1] The competing model assumes that an individual makes an annual expenditure decision while facing a non-linear budget constraint, implicitly assuming that individuals only respond to the effective price of care at the end of the year, which is known with certainty.[2]

The impact of this difference on moral hazard estimates is best understood with a simple example. Assume that the price of a hospital visit can take on two values: p_l, which occurs θ percent of the time, and p_h, which occurs $(1-\theta)$ percent of the time, where $p_h > p_l$.[3] Let Individual U be risk neutral and uncertain about the price of medical care when he makes a medical care decision; he

[1]The *effective* price of a unit of medical care is the *spot* price minus what Keeler et al. (1977) call the *bonus* individuals receive for moving closer to their deductible. There are two reasons that effective prices are unknown in the WYDM model. First, the model assumes that the contemporaneous transaction price (i.e., the sum of insurer and insured's payments) of a unit of medical care is unknown at the time of consumption. Support for this assumption is provided below, but is not a necessary condition for general WYDM models. Second, because future health shocks and medical care preference shocks are unknown, the benefit of moving closer to one's deductible in any given month is unknown.

[2]Bajari et al. (2014) is an exception. The authors assume that an individual selects his total annual medical care expenditure while knowing only the distribution of the proportion of that expenditure he must pay out of pocket; thus, individuals face some uncertainty about their out-of-pocket payments.

[3]The following discussion assumes that no benefit is received from the high price draw that is not also received

represents an individual solving the WYDM model. Let Individual C be identical to Individual U, except that Individual C is certain about his price draw prior to making a medical care decision; he represents an individual solving an annual expenditure model. Consider the consumption decisions of Individuals U and C with and without insurance. If fully insured (i.e., the out-of-pocket price of a visit equals zero), then both individuals will visit the hospital if their willingness to pay for medical care, W, is greater than zero. There are two scenarios to consider when the individuals are uninsured: First, suppose the individuals' willingness to pay for medical care, W, is less than the expected price of medical care, $E[P] = \theta p_l + (1-\theta)p_h$, but greater than p_l. In this scenario, Individual U never visits the hospital when uninsured, as the expected price of medical care is greater than his willingness to pay, while Individual C visits the hospital any time he gets a price draw of p_l. As a result, the moral hazard effect for Individual U would be larger than that for Individual C. On the other hand, suppose $p_h > W > E[P]$.[4] In this scenario, Individual U always visits the hospital when uninsured, while Individual C again only visits the hospital when he gets a price draw of p_l; thus, the moral hazard effect for Individual C would be larger than that for Individual U.

The level of risk aversion in the sample should also impact the moral hazard effects predicted by the models, which can be seen in Figure A1. The figure features a strictly concave utility function for a non-medical good X (with price normalized to one) and the same lottery that is described above. The utility received from the expected value of the price lottery, $U(E[X])$, which is re-interpreted here as a lottery over two levels of non-medical good consumption (i.e., $E[X] = \theta(I - p_l) + (1-\theta)(I - p_h)$, where I is income), is greater than the expected utility of the gamble, $E[U(X)]$, suggesting that the individuals are risk averse. With this utility function, nothing about Individual C's medical care decision-making problem changes because he faces no risk - i.e., in both insured and uninsured states, he makes the same decisions that he made when risk neutral. Risk aversion also has no impact on Individual U's decisions when he is insured. However, consider the behavior of Individual U when he is uninsured. Under risk neutrality, Individual U did not visit

from the low price draw, which is also an assumption of the WYDM model.

[4]The two omitted scenarios are trivial. If $W > p_h$ then both individuals always consume when uninsured and if $W < p_l$ then both individuals never consume when uninsured. In both scenarios, the moral hazard effects for both individuals are identical.

the hospital when $W < E[P]$. In Figure A1, the group of risk neutral individuals not visiting the hospital are represented by area A, where $I - W > E[X]$. Under risk aversion, some individuals who value medical care slightly more than the expected price will still choose not to visit due to their risk aversion (i.e., their fear of a high price draw). The new group of individuals not purchasing care would include those both in area A and area B of Figure A1, as these individuals have income and medical care preferences such that $I - W > CE$, which is the certainty equivalent value of the non-medical care good. In other words, for a fixed W distribution in the population, the moral hazard effect predicted by the WYDM model should be increasing in the level of risk aversion in the population, while risk aversion should have no such effect on the annual expenditure model.

Figure A1: Risk Aversion and Moral Hazard

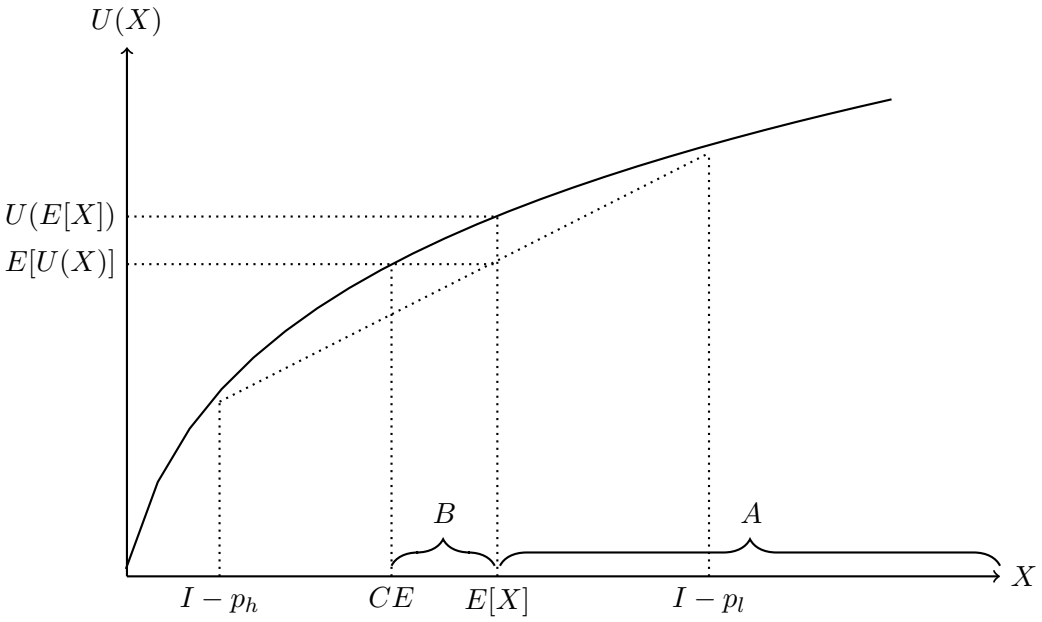

2 Out-of-Pocket Expenditure Equation

Out-of-pocket expenditure resulting from d_t trips to the doctor, h_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^d, p_t^h, 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 = MOX
- 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_1^r = 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 $D * H$ 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.[5]

Out-of-pocket expenditure is calculated as follows:

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

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

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 $= max\{0, D_1 - ADE_t - AHE_t\}$

[5]Order is important because it determines when exactly during the month the deductible or MOX 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 MOX only occur between months, which would simplify the budget constraint; however, it would remove important variation from the data.

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:

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

where

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

$z_t^1 =$ MOX 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:

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

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

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

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:

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

where

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

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

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

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

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

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

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

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

Out-of-pocket expenditure in month t is:

$$O_t = O_t^d + O_t^h + (r_t * p_t^r * C_1^r) \tag{7}$$

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

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

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

3 Data Overview

The empirical exercises included in this research use data from the 1996-1999 cohorts 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, 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.

3.1 Demographic Variables

Individuals are 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.[6] 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.[7] Parents are identified by other persons living in the household 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*).[8] 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 house-

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

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

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

holds. Of those linked, 3 are missing foreign born status and 11 are missing veteran status. All individuals with missing data are assumed to be native born, non-veterans.

3.2 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 an individual receives an offer as they are later removed if the offer is not observed in the MEPS-Insurance Component (IC) Link file. Monthly insurance status (i.e., insured or not) can be determined in the MEPS-Household Component (HC) data file (*insmmyy*). A series of insurance provider type variables in the HC file ensures that an individual is not covered by an additional 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.[9] 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 50% 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.[10] 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

[9] 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.

[10] See Section 4 below for a discussion of the potential bias introduced by this exclusion.

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.[11]

3.2.1 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.
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 any other 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 missing value 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 MOX is coded as zero, meaning comprehensive coverage, but the plan features some cost-sharing, then I recode the MOX to be missing.
6. If hospital and doctor care is free after the deductible is crossed then I set the MOX equal to the stated deductible.

The scope of the missing information problem, after imposing these assumptions, can be observed in Table A1. I list the number of plans with *any* missing characteristic by whether or not

[11] 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.

the plan is held by the employee. Plans are flagged if missing out-of-pocket premium, deductible, hospital coverage, doctor coverage, or MOX. 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.

Table A1: Missing Features in Health Insurance Plans, Logical Imputations

	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 costs an individual $5 out of pocket, then I can impute that he 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:

Table A2: Missing Features in Health Insurance Plans, Expenditure Imputations

	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

3.2.2 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. Table A3 describes the bounds on an acceptable match in each round.

Table A3: Acceptable Match Bounds

	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
MOX	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
- MOX: 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.[12] 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:[13]

- Total premium: $[475, 525]$
- Doctor co-pay: $[5, 15]$
- Doctor co-insurance: $[0, 5]$
- Hospital co-insurance: $[5, 15]$
- Hospital co-pay (day): $[0, 10]$
- Hospital co-pay (stay): $[0, 25]$
- MOX: $[1350, 1650]$
- 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 is found in Table A4.

[12] 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.

[13] 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.

Table A4: Missing Features in Health Insurance Plans, Matching Imputations

	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

3.2.3 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 MOX.[14] The regression includes all plans reporting a MOX and those reporting that the plan has no MOX. These parameters are used to predict the presence of a MOX for plans where it cannot be determined if a MOX exists. I then use an OLS regression of MOX on other plan and company characteristics to predict a MOX for those known to have a one (either through the data or through the previously described prediction), but no recorded value.[15] The same procedure is followed for the deductible; a 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

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

[15] Each regression uses the following set of explanatory variables: total premium, has deductible, deductible, has MOX, MOX, 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., MOX is not included on the right hand side of the MOX regression).

and missing values are replaced with appropriate predictions.

3.2.4 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 MOXs 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.[16] I also assume that prescription drug expenditure is completely unrelated to a plan's deductible and MOX.[17] 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 MOX for 44%. According to the Kaiser Family Foundation's Employer Health Benefits Annual Survey, by 2000 prescription drug expenditure had no relation to the MOX in over 75% of ESHI plans.

[16] 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.

[17] *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 MOX.

3.3 The Estimation Sample

A summary of sample size by inclusion criteria is presented in Table A5. A final summary of chosen and rejected plans is presented in Table A6.

Table A5: Sample Inclusion Criteria

	1996	1997	1998	1999	Total
1996-1999 MEPS Household Component	22601	13683	11137	14178	61599
and single, childless, 19-64 yrs old	4406	2534	2169	2589	11698
and employed in first interview period with offer	1821	923	987	1128	4859
and matches to link file†	749	516	156	688	2109
and no missing interviews	693	472	139	636	1940
and stable insurance status	455	290	98	389	1232

† AHRQ only attempted to collect insurance information for a 25 percent random sample of individuals in 1998, which explains why so many individuals are lost because of no linkage in that year.

Table A6: Insurance Plan Summary

	Held Plans			Rejected Plans		
	Plans	Mean	S.D.	Plans	Mean	S.D.
Premium						
total premium	1119	2057.19	819.09	4164	2207.13	715.25
out-of-pocket premium	1119	343.83	540.31	4164	519.46	610.28
Deductible†						
defined by total expenditure	408	317.98	415.43	829	385.65	567.46
defined by doctor expenditure only	59	191.38	97.04	655	215.69	50.65
defined by hospital expenditure only	28	252.63	218.10	62	150.58	41.24
plan has no deductible	637			2626		
MOX						
MOX	737	1444.98	1037.00	2816	1642.30	1111.48
plan has no MOX	382			1348		
Hospital†						
co-insurance rate	422	17.85	12.77	827	17.17	14.39
co-pay level (per stay)	191	255.04	347.66	750	165.47	209.44
co-pay level (per day)	85	52.28	59.92	123	62.48	88.53
free care past the deductible	431			2583		
Doctor						
co-insurance rate	214	20.81	14.97	775	13.61	12.55
co-pay level	872	10.17	4.47	3219	8.51	4.15
free care past the deductible	46			377		
Network Type						
HMO	1119	0.42		4164	0.51	
FFS	1119	0.47		4164	0.44	
PPO	1119	0.11		4164	0.06	

† These categories are not mutually exclusive. Some plans feature a doctor- and hospital-specific deductible. Moreover, some plans are structured so that an individual pays a daily co-pay, plus a percentage (i.e., co-insurance rate) of the remainder of the bill.

3.4 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.[18] 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.

3.4.1 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 straightforward in the data. The dates of doctor visits and hospital days are consistently reported - 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.

[18] 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 a managed care plan or Medicaid, but was selected as part of a 25% random sample.

Doctor and hospital prices are observed 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.[19] 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 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. This ratio is consistent with the ratio of average prices observed in the data.

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,

[19] 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 allowing for permanent unobserved correlation between the medical care preferences and price draws.

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. See Section 10 below for a discussion of how this assumption is likely to impact my results.

3.4.2 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.[20] By assuming the 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

[20] 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.

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 Table A7) I was able to use average dosage to approximate the number of months.[21] The approximation provided a length for 2,422 out of the 13,592 total fills/refills purchased over the 2 year period.[22]

Table A7: 21 Most Popular Drugs in MEPS

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

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

[22] 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.

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 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, one month and three 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.

3.5 Illness

In addition to reporting 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. An 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.[23] 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 potential for and consequences of misreported illnesses should be considered. First, consider the possibility that there is classical measurement error in the acute and chronic illness variables used in the model. Greene (2012, Ch. 4) shows that for a linear model estimated via ordinary least squares, such measurement error (i) in the dependent variable yields no bias in slope estimates, though there is a precision loss and (ii) in an independent variable causes attenuation bias in slope parameter estimates.[24] Similarly, such measurement error in my illness variables may lead (i) illness production function parameters to be estimated with less precision and (ii) the estimated marginal impact of illness on utility and prices, as well as the impact of lagged illness on future illnesses, to be biased toward zero.

Forecasting how this bias might impact moral hazard estimates in a dynamic model is difficult.

[23] 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.

[24] Note that while these conclusions apply generally to linear models estimated via OLS, they are not guaranteed in the non-linear WYDM model (Schennach, 2004).

In a static or simple two-period model, one could solve for optimal medical care demand using the first order condition; the derivative of medical care demand with respect to insurance/prices would then suggest how the marginal utility of illness impacts moral hazard estimates. In the WYDM model, this first order condition is too difficult to calculate analytically, which is why moral hazard effects are calculated via simulation. That said, the parameters that are potentially biased by the measurement error discussed above represent various *costs* of illness, so correcting the measurement error would presumably increase the cost of illness. As such, one would expect that this measurement error leads health insurance, and the medical care it allows people to consume, to be undervalued, which would lead to an underestimate of moral hazard effects. That said, note that if anything, my estimates are large compared to the existing the literature.

A more concerning possibility is that uninsured or weakly insured individuals, who rarely visit a doctor due to high out-of-pocket prices, might underreport illness.[25] There are two reasons one might suspect such misreporting. First, visiting a doctor could help individuals to recall an illness that took place weeks or even months prior to an interview. Were this the case, uninsured individuals would be less healthy than the data suggest, which could bias moral hazard estimates if price sensitivity varies by health status. According to Zweifel and Manning (2000), findings from the RAND Health Insurance Experiment suggest that the demand for prescription drugs, outpatient visits, preventative care, and emergency room visits respond similarly to changes in prices, while demand for inpatient care is less sensitive to changes in prices. While these findings do not explicitly test for heterogeneity in price sensitivity across illness states, they do suggest that individuals in poor health (i.e., those consuming inpatient hospital visits) are less sensitive to changes in prices (i.e., lower moral hazard estimates) than those in good health. That said, underreporting of illnesses due to recall problems is more likely to occur for individuals in good health, who rarely visit the doctor. For these individuals, there seems to be little evidence of heterogeneity in price sensitivity.[26]

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

[26] Providing empirical evidence against misreporting is difficult. That said, I can report that the share of reported illnesses that are not associated with medical care consumption is decreasing in insurance generosity, for both acute and chronic illnesses. Specifically, the share of reported acute (chronic) illnesses that are not associated with medical care consumption is 48% (30%) for the uninsured, 41% (24%) for the weakly insured (i.e., total premium less than

A second reason one might expect the uninsured to underreport illnesses is if doctor's diagnose illnesses that individuals are not aware of. Such misreporting could bias moral hazard estimates if (in reality) upon receiving health insurance, previously uninsured individuals visit a doctor, causing them to receive a previously unknown diagnosis and, thus, consume even more medical care. This misreporting would cause moral hazard estimates for the uninsured population to be understated.[27] There are two reasons why I do not believe that my moral hazard results are biased by this type of misreporting. First, this argument suggests that my moral hazard estimates are too small; yet again, my estimates are somewhat large compared to the existing literature (see Table 16 in the main manuscript). Second, the model requires only that individuals report the correct number of acute and chronic illnesses, which is easier than reporting the exact illness or intensity. For example, an uninsured individual who never visits a doctor and, therefore, thinks that they have a cold when they actually have strep-throat does not create measurement error in the model. That said, if the same individuals misclassifies the type of illness (e.g., has (chronic) throat cancer, but reports (acute) sore throat) or misses a chronic illness all together (e.g., does not know that they have high blood pressure), then estimates could be biased in the way described above. These examples suggest that misclassification and misreporting are most likely to occur with chronic illnesses that have a long, slow-acting impact on the body, such as hypertension. It is worth noting that Baicker et al. (2013) found no evidence of a significant increase in prevalence or diagnosis of hypertension or high cholesterol levels or on the use of medication for these conditions in the two years following the Oregon Health Insurance Experiment.

Modeling dynamic illness transitions requires that illnesses are observed, and while self-reported illness measures can be biased by endogenous reporting, there is no perfect alternative. A different strategy might be to force survey participants to visit a doctor annually or biannually for a check-up, such as in the Framingham Heart Survey, so that more precise illness information can be collected.

---

the sample median), and 38% (21%) for the strongly insured. While these figures do not rule out measurement error, the fact that uninsured and weakly insured individuals are more likely to report illnesses unassociated with medical visits provides some evidence against the underreporting of illnesses by these individuals.

[27] Moral hazard estimates would also be understated for individuals who are insured in the data. If (in reality) becoming uninsured is associated with not visiting the doctor, causing illnesses to go undiagnosed, then the short-run impact of loosing one's health insurance coverage could be larger than is suggested by the model.

The downside of this survey design is that it naturally exposes participants to information they do not typically receive, which alters decisions and, thus, compromises the external validity of study findings. Ultimately, the objective of this study is to determine how individuals respond to changes in the design and generosity of their health-insurance plans in the short run. Because an individual's illness status both influences and is influenced by medical care decisions, illness transitions must be modeled. Given my focus on short-run changes in medical care consumption, it seems that the most relevant measure of an individual's illness state would include only the illnesses that he is aware of. 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 much more concerning.

3.5.1 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.[28] 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 Table A8).[29]

A few of these changes require explanation.[30] According to the Centers for Disease Control and Prevention (CDC), 17%, 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

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

[29]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.

[30]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.

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 suggests that 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.[31]

Table A8: Recoded ICD-9-CM Codes

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

3.5.2 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

[31] 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 6,208 reported medical conditions are missing an ICD-9-CM code. For these individuals, I look at medical care consumption files in the same periods, which contain an ICD-9-CM code, 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.

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 (ICD-9-CM codes: 747, 753, 755, 757, 758, and 759) are assumed to begin prior to the beginning of the insurance year. 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). 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 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 an individual consumed medical care for an illness in the first interview period that an illness is reported, then the month of consumption, which is observable, 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, just 1,299 illnesses (29%) are in need of a beginning month. For these, a beginning date is drawn uniformly between the first and last month of the interview period 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.[32] The distribution of bound lengths is shown in Table A9.

Table A9: Illness Imputation Bounds

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	

Because some illness beginning and ending dates are imputed, illness counts A_t and C_t are likely measured with error. However, the imputation procedure virtually guarantees *classical* measurement error. Thus, as is discussed in Section 3.5 above, the measurement error in these illness variables is likely to lead (i) all illness production function parameters to be estimated with less precision and (ii) the estimated marginal impact of illness on utility and prices to be biased toward zero.

[32] 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. This strategy would bias the estimates of medical care on illness transitions, as medical care consumers would on average have longer illness periods than non-consumers by construction. With the method used, the productivity/efficacy 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.

3.6 Month-to-Month Variation in State Variables

Table A10: Number of Acute Illnesses from Month t to Month $t+1$

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

Notes: The left-most column labels the number of acute illnesses in period t. The top row labels the number of acute illnesses in period $t+1$. Table cells then contain the percentage of individuals with a particular number of acute illnesses in period t that transition to another number of acute illnesses in period $t+1$. For example, the top left cell says that 'among the individuals with zero acute illnesses in any particular month, 89.4% have zero acute illnesses in the following month.' The right-most column contains the number of person-month observations observed at each illness level.

Table A11: Number of Chronic Illnesses from Month t to Month $t+1$

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

Notes: The left-most column labels the number of chronic illnesses in period t. The top row labels the number of chronic illnesses in period $t+1$. Table cells then contain the percentage of individuals with a particular number of chronic illnesses in period t that transition to another number of chronic illnesses in period $t+1$. For example, the top left cell says that 'among the individuals with zero chronic illnesses in any particular month, 98.5% have zero chronic illnesses in the following month.' The right-most column contains the number of person-month observations observed at each illness level.

Table A12: Changes in Acute Illness Status over Insurance Year

	Intensive Margin		Extensive Margin	
Frequency	Individuals	Percent	Individuals	Percent
0	244	19.8%	431	35.0%
1	153	12.4%	233	18.9%
2	300	24.4%	304	24.7%
3	184	14.9%	139	11.3%
4	150	12.2%	80	6.5%
5	103	8.4%	30	2.4%
6	63	5.1%	9	0.7%
7	24	1.9%	6	0.5%
8	8	0.6%		
9	2	0.2%		
10	1	0.1%		
Total	1232	100%	1232	100%

Notes: This table reports the number of individuals experiencing a particular number of changes in their acute illness state over the course of the insurance year. *Intensive Margin* refers to the number of illness. So, columns 2 reports the total number of individuals experiencing zero to ten changes in their total number of acute illnesses. Column 4 only counts extensive margin changes in illness status - i.e., going from zero to a positive number of acute illnesses or vice versa. Thus, the second entry in column 4, for example, is interpreted as '233 individuals experienced just one transition into or out of having *any* acute illness over the course of the insurance year.'

Table A13: Number of Doctor's Office Visits from Month t to Month $t+1$

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

Notes: The left-most column labels the number of doctor's office visits in period t. The top row labels the number of doctor's office visits in period $t+1$. Table cells then contain the percentage of individuals with a particular number of doctor's office visits in period t that transition to another number of doctor's office visits in period $t+1$. For example, the top left cell says that 'among the individuals with zero doctor's office visits in any particular month, 85.4% have zero doctor's office visits in the following month.' The right-most column contains the number of person-month observations observed at each number of doctor's office visits.

Table A14: Number of Hospital Days from Month t to Month $t+1$

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

Notes: The left-most column labels the number of hospital visits in period t. The top row labels the number of hospital visits in period $t+1$. Table cells then contain the percentage of individuals with a particular number of hospital visits in period t that transition to another number of hospital visits in period $t+1$. For example, the top left cell says that 'among the individuals with zero hospital visits in any particular month, 97.4% have zero hospital visits in the following month.' The right-most column contains the number of person-month observations observed at each number of hospital visits.

Table A15: Prescription Drug Consumption from Month t to Month $t+1$

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

Notes: The left-most column labels whether an individual used prescription drugs in period t. The top row labels whether an individual used prescription drugs in period $t+1$. Table cells then contain the percentage of individuals with a particular prescription drug status in period t that transition to another prescription drug status in period $t+1$. For example, the top left cell says that 'among the individuals not using prescription drugs in any particular month, 89.7% not use any prescription drugs in the following month.' The right-most column contains the number of person-month observations observed not consuming and consuming prescription drugs.

3.7 Variation in Out-of-Pocket Share

The dynamic cost-sharing characteristics of health insurance plans can cause the out-of-pocket price of medical care to change over the course of the year insurance year. That said, 34% of the estimation sample faces a budget constraint where the proportion of medical care prices paid out of pocket is fixed over the course of the year. For this group, the expected out-of-pocket price of a unit of medical care varies over the course of the year only due to changes in illness status. The group is comprised of 113 uninsured individuals, who always pay 100%; 28 fully insured individuals, who always pay 0%; and 278 insured individuals with no dynamic cost-sharing features, who pay 7% of medical care prices out of pocket on average. The rest of the sample faces the possibility of

a change in out-of-pocket prices, of which 30% actually experience a change. The Figure A2 below shows the average share of medical care price paid out of pocket in each insurance month for the 241 individuals who experience a change in this share at some point during the insurance year.

Figure A2: Average Share of Medical Care Price Paid Out of Pocket

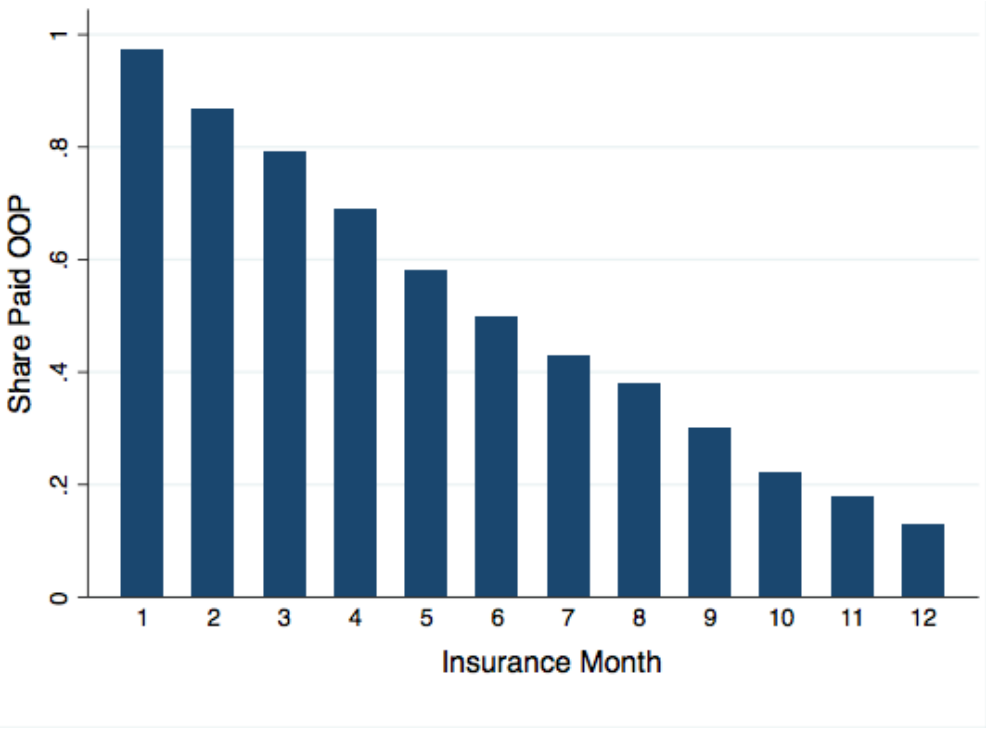

3.8 Medical Care Price and Demand Correlation

In this section, I present results from several regressions of monthly medical care demand on prices. I measure monthly medical care demand using any consumption and logged total expenditure. I focus on doctor and hospital demand only, as prescription drug cost-sharing is not observed in the data (see Section 3.2.4 above). I measure two types of medical care prices: marginal/spot and expected end-of-year. Ellis (1986) argues that rational, forward-looking individuals should respond only to expected end-of-year prices; however, several empirical studies have found individuals to be more responsive to spot prices (e.g., Einav et al., 2015; Abaluck et al., 2015; Dalton et al., 2017). Rather than taking a position on how forward-looking individuals are, I examine responsiveness to both prices.

Marginal prices are measured as the share of first dollar medical care expenses paid out of pocket in each month of the year. For simplicity, doctor and hospital deductibles are added together, while

doctor and hospital co-insurance rates are averaged, so individuals face a single out-of-pocket share across medical care types.

Expected end-of-year prices are more difficult to measure. Using a sample of new employees entering a firm over a calendar year, Aron-Dine et al. (2015) measure expected end-of-year prices for a plan-month as one minus the fraction of new employees joining the plan in the month that eventually cross their deductible. I generally observe just one person per plan and view few new employees, ruling out this strategy. Instead, I calculate the share of last dollar medical care expenses paid out of pocket at the end of the year for each person given their insurance holdings. Again, I aggregate across doctor and hospital insurance characteristics to create a single share. I then regress the observed end-of-year, out-of-pocket share on monthly insurance characteristics (e.g., distance to deductible, distance to MOX, co-insurance rate, month of insurance year) and interactions.[33] I use these estimates to predict a monthly, expected end-of-year out-of-pocket share for each individual, conditional on their observed insurance holdings and past medical care consumption.

Summary statistics of the model's key variables are presented in Table A16. Note that 419 individuals have no dynamic cost-sharing features and, therefore, never experience a change in expected end-of-year prices. Moreover, only 241 individuals cross their deductible or MOX over the course of the year, generating a change in their marginal price. The correlation between marginal and expected end-of-year prices is 0.96.

Table A16: Monthly Demand and Price Variables

	Mean	Median	Std. Dev.	Min	Max
Any Consumption	0.22	0.00	0.41	0	1
Total Expenditure	73.39	0.00	738.05	0	51,002
Marginal Price	0.43	0.12	0.45	0	1
Expected EOY Price	0.37	0.12	0.38	0	1

Variables are calculated from 14,784 person-month observations.

Regression results are found in Table A17. Each table entry measures the correlation between medical care demand and price, conditional on controls, from a different regression. Importantly, prices are endogenous in all specifications. The purpose of this exercise is to show that across many

[33]The results from this procedure are intuitive. Distance to deductible, distance to MOX, and the plan's co-insurance rate are all positively associated with the end-of-year share owed out of pocket. The R^2 from the regression is 0.79.

specifications, prices and demand are negatively correlated and that this relationship, with few exceptions, is statistically significantly different from zero. The full structural model is needed to control for the endogeneity of prices.

Column 1 presents price sensitivity estimates for four models that do not include individual-level fixed effects; all uncover negative and significant correlation between prices and demand. Including individual fixed effects reduces the negative correlation (Column 2). Note that the fixed effects likely eliminate much of the bias in the parameter due to endogenous insurance selection, but not all, as early year demand still affects later year prices for those facing non-linear budget constraints. In the structural model, controlling for permanent unobserved heterogeneity approximates an individual-level fixed effect. Moreover, note that standard errors increase substantially with the inclusion of an individual fixed effect, as only individuals experiencing a change in both prices and demand at some point over the year contribute to fixed effect estimates.[34]

Table A17: Linear Monthly Medical Care Demand Regressions

	(1)		(2)	
	Coef.	S.E.	Coef.	S.E.
Any Consumption				
Marginal Price	-0.118	0.012	-0.068	0.020
Expected EOY Price	-0.133	0.015	-0.035	0.029
Log Expenditure				
Marginal Price	-0.644	0.061	-0.506	0.100
Expected EOY Price	-0.702	0.076	-0.383	0.139
Individual FE			X	

The table presents the price coefficients from several regressions of medical care demand on prices, sex, race, education, age, MSA, and income. All regressions contain 14,784 observations. Standard errors are clustered at the individual level. When the dependent variable is *Any Consumption* (e.g., the top left cell), the parameter is interpreted as the (percentage point) decrease in the likelihood of visiting the doctor or hospital in a month when the co-insurance rate is increased from zero (i.e., free care) to one (i.e., no insurance). When the dependent variable is *Log Expenditure* (e.g., the bottom left cell), the parameter (times 100) is interpreted as the percent decrease in monthly expenditure when the co-insurance rate is increased from zero to one.

[34]When marginal prices are used in estimation, only the 241 individuals crossing their deductible or MOX contribute to identification. When Expected EOY prices are used, only the 813 individuals with dynamic cost-sharing features contribute to estimation.

4 Representativeness of the data

Sample selection due to employee/employer non-response might matter for the paper's substantive results. Table A18 compares four subsamples from of the nationally representative MEPS data.

- Sample A: Employed individuals between the ages of 24 and 64 who receive a health insurance offer from their employer.
- Sample B: sample A individuals, further restricted to be single and childless.
- Sample C: sample B individuals, further restricted to those that have an insurance alternative set match in the link file
- Estimation Sample: sample C individuals, further restricted to exclude any individual holding multiple plans, switching plans or insurance status, and/or holding a plan with characteristics that are inconsistent with medical expenditure data.

Sample A is the most policy relevant, while sample B is a nationally representative benchmark for the estimation sample. Compaing sample C and the estimation sample allows one to distinguish between employer/employee non-response and other sources of endogenous sample selection.

A key simulation in the paper compares average spending without insurance coverage against average spending when individuals are able to select their preferred insurance among the alternatives that their employer provides (i.e., the main moral hazard estimate of 53.1%). Ultimately, this figure is a function of my particular sample's (i) health insurance alternative sets and (ii) estimated price sensitivity.[35] Thus, if employer/employee response rates are to impact the moral hazard estimate, they must affect either (i) the representativeness of the health insurance plans offered to my estimation sample or (ii) the representativeness of individuals in the estimation sample as it relates to their level of price sensitivity. I address each below.

[35] Recall that moral hazard effects are calculated via simulation. The model is first simulated assuming everyone is uninsured; then a second time, forcing everyone to select their preferred plan (i.e., choosing to be uninsured is not an option). Because individuals are allowed to select their preferred plan in the second simulation, moral hazard effects are larger with a more generous alternative set than a less generous one. That said, the moral hazard estimate which compares total annual medical expenditure without insurance to that with full coverage (i.e., column 2 of Table 14) *should not* be biased by a non-representative health insurance offer set in the estimation sample.

Table A18: Representativeness of the Sample

	Sample A		Sample B		Sample C		Est. Sample	
	Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.
Demographic								
age	39.9	10.8	36.9	12.2	38.3	12.1	39.5	11.7
education (highest grade completed)	13.25	2.66	13.4	2.5	13.4	2.5	13.6	2.5
income (in 1996 dollars)	33,873.5	23,817.4	30,985.6	21,432.6	32,354.6	22,330.1	33,716.0	21,944.7
male	0.54		0.53		0.51		0.49	
lives in a MSA	0.80		0.84		0.83		0.83	
Hispanic	0.16		0.14		0.12		0.12	
black	0.13		0.16		0.15		0.14	
region - northeast	0.18		0.19		0.18		0.19	
region - midwest	0.23		0.24		0.25		0.25	
region - south	0.36		0.35		0.35		0.34	
region - west	0.23		0.22		0.22		0.22	
Employment								
federal employee	0.04		0.03		0.07		0.09	
state/local government employee	0.16		0.14		0.18		0.19	
blue collar	0.40		0.40		0.36		0.35	
employer benefit - paid vacation	0.88		0.86		0.87		0.90	
employer benefit - retirement plan	0.70		0.65		0.70		0.75	
industry - manufacturing	0.10		0.11		0.12		0.11	
industry - sales	0.17		0.17		0.15		0.15	
industry - repair services	0.10		0.12		0.10		0.10	
industry - professional services	0.16		0.15		0.18		0.18	
number of employees	177.6	190.3	170.8	188.5	186.6	192.1	192.9	193.8
Health								
health status - excellent	0.34		0.34		0.34		0.34	
health status - very good	0.35		0.35		0.36		0.35	
health status - good	0.24		0.23		0.23		0.24	
health status - fair/poor	0.07		0.08		0.07		0.07	
total annual MC exp.†	1,597.5	5,001.6	1,490.4	4,033.5	1,657.3	4,043.0	1635.3	3063.0
ever insured in first survey year	0.96		0.95		0.96		0.91	
holder of ESHI for 12 consecutive mons	0.77		0.81		0.85		0.91	
Sample Size	17,527		4,859		2,109		1,232	

† This expenditure level was take directly from the MEPS data and includes types of medical care spending not included in this analysis (e.g., dental and eye care, home healthcare, medical equipment, etc.).

Notes: Sample A contains employed individuals between the ages of 19 and 64 who receive a health insurance offer from their employer. Sample B contains all Sample A individuals, further restricted to be single and childless. Sample C contains Sample B individuals, further restricted to those that have an insurance alternative set match in the link file.

4.1 Representativeness of the Insurance Offer Set

Here, I establish that individuals in the estimation sample receive a similar ESHI offer set to individuals in samples A and B, where the main differences between the offer sets are due to the fact that the estimation sample has a larger proportion of government employees that tend to receive generous ESHI offer sets.

The share of public/private employees across samples is documented in Table A18 above. Samples A and B have a smaller number of state and federal employees than the estimation sample. The main difference is in the number of federal employees, which increases roughly three-fold in the estimation sample. The difference exists because no federal employees are excluded due to employer non-response. The abundance of government employees in the estimation sample should increase the average size of the insurance offer set. According to my data, government employees are offered an average of almost 10 plans, while the average private sector employee is offered 2 plans. The implications for the estimated moral hazard effects depend on whether plans selected by government employees are more or less generous on average. The data suggest that the average government plan is more generous. For $2,000 in total annual medical spending, the average government-sponsored plan observed in the data covers 89% of spending, while the average plan sponsored by a private firm covers 83%. The *most generous* plan offered to government and private sector employees cover 94% and 84%, respectively, on average.

I also consider the representativeness of the ESHI offer set faced by private sector employees in the estimation sample. Ideally, I would compare the offer sets of those in samples A and B to those in the estimation sample; however, the offer sets for samples A and B are not observable. Thus, I attempt to compare both firm and insurance characteristics of the estimation sample to nationally representative data outside the MEPS.[36] First, using BLS data, I show that firms in the estimation sample are slightly larger than would be expected if randomly sampled from the population in 1999.

[36] I am not able to determine whether differences between the estimation sample and the nationally representative results are due to the selection of single, childless people into firms with unique health insurance offerings or employer/employee non-response in the construction of the estimation sample.

Table A19: Distribution of Employee Firm Size

Firm Size	BLS - 1999	Estimation Sample
1-50	26.2%	19%
50-250	18.8%	17%
250-1000	14.2%	14%
1000+	37.3%	50%

This finding is unsurprising, as one might expect larger firms with designated HR departments to be most likely to respond to the MEPS survey. With respect my moral hazard estimates, this finding should again lead my estimates to be slightly too large, as large employers offer larger, more generous insurance alternatives sets than small employers. Second, I compare the average number of plans and premiums offered to estimation sample individuals against the nationally representative offerings summarized in the Kaiser Family Foundation 1999 Employer Health Benefit Survey.[37] The offerings are quite similar. Among the employers of individuals contained in the estimation sample, 41% offer only one plan, 17% offer two plans, and 42% offer three or more plans. According to the Employer Health Benefit Survey, of employer nationwide 35% offer only one plan, 15% offer two plans, and 50% offer three or more plans. Moreover, the average plan offered to an individual in the estimation sample had a total premium of $181, while the average plan nationwide had a premium of $177 (in 1996 dollars).

In summary, individuals in the estimation sample are more likely to work for the government or large firm than the average individual in the population. As a result, the average insurance alternative set observed in the estimation sample is likely to be more generous than the average alternative set in the population. Because moral hazard effects are calculated via simulation, where individuals are allowed to select their preferred plan in the *insured* state, my main estimates may be larger than would be calculated using a nationally representative sample. That said, when moral hazard effects are calculated for a fixed insurance plan, such as in column 2 of Table 14, no such bias should exist.

[37] The 1999 Employer Health Benefit Survey interviewed 2,694 randomly selected public and private employers in the United States; thus, the full estimation sample is used in comparison, not just private sector employees.

4.2 Representativeness of Individuals

The estimation sample is non-randomly selected. Above, I discuss how this selection could lead to insurance alternative sets that are not representative of those seen in the population at large, which could impact the estimation of moral hazard effects. A more fundamental problem would be that non-random selection leads the underlying price sensitivity of individuals in the estimation sample to be different from the population at large, as this would cause the structural parameters of the model to be estimated with bias, affecting all model simulations.

Unfortunately, one cannot test whether the underlying price sensitivity among estimation sample individuals differs from that of the general population. Given that the primary reason for non-random selection is due to employer non-response to a survey, it does not seem that one should *expect* estimation sample individuals to be uniquely price sensitive/insensitive; however, it is possible. One way to evaluate the likelihood that estimation sample individuals are common in their price sensitivity is to compare their observable characteristics with those of the nationally representative samples A and B, then ask if observable differences in characteristics are likely to reflect differences in underlying price sensitivities.

Table A18 compares individual-level characteristics across four samples. Starting in the first row and moving down, one can see that the estimation sample contains a lower proportion of males than samples A and B. Some of this difference is likely associated with the estimation sample also having a larger proportion of government employees (discussed above) and, therefore, a lower proportion of blue collar workers. Half of the difference in the proportion of males across samples is reflected in sample C, meaning that males are either less likely to report their employer's information or their employers are less likely to respond. The estimation sample also contains fewer Hispanics, which (i) comparing samples A and B is somewhat due to the fact that Hispanics are more likely to be married and have children and (ii) comparing samples B and C, could be somewhat due to MEPS containing non-citizens, who do not wish to report the identify of their employers. I regard both the difference in males and Hispanics across the samples as relatively small and, thus, unlikely to have a large impact on overall price sensitivity in the model. Even if the differences were large, I'm not aware of any research suggesting men and/or Hispanics are more/less price sensitive. In terms of health and medical care expenditures, the estimation sample is similar to sample A.

In summary, there are few observable differences in the individuals that comprise samples A and B, which are nationally representative, and the estimation sample. Moreover, none of these differences are associated with more/less price sensitivity in an obvious way. As such, despite non-random selection into the estimation sample, there is little reason to believe that model parameters are biased due to sample selection.

5 Estimation Procedure

I estimate the model's parameters, Ω, using a nested fixed point algorithm (Rust, 1987). The inner algorithm solves the DPP for a given set of parameters and for each unobserved type. The outer algorithm uses the resulting probabilities and densities to calculate the likelihood function, $L(\Omega)$, and attempts to improve the likelihood value using a BHHH gradient method (Berndt et al., 1974).

An individual contributes to the likelihood function the product of his observed illness state, medical care price, medical care choice, and insurance choice probabilities. Illness state probabilities take on closed forms due to the logit assumptions in Section 3 of the paper. I assume that doctor's prices are drawn from a Singh Maddala distribution; hospital and prescription drug prices are distributed log-normal.[38] I allow the scaling parameter of these distributions to vary by state variables.

I assume that ϵ_t^{dhr} and ϕ_y^j follow a Type 1 Extreme Value (T1EV) distributions, which simplifies estimation in two ways. First, when ϵ_t^{dhr} is T1EV, Equation 8 in the paper becomes

$$V(\mathbf{S_t}|\mu^k, I_y^j) = EC + ln\left(\sum_{d=0}^{D}\sum_{h=0}^{H}\sum_{r=0}^{1} exp\left(\overline{V}_{dhr}(\mathbf{S_t}|\mu^k, I_y^j)\right)\right) \qquad \forall t \tag{10}$$

where EC is Euler's Constant. The assumption simplifies solution to the optimization problem, as calculation/simulation of a $[(D+1)*(H+1)*2]-1$ dimensional integral would be required if a continuous distribution without a closed form (e.g., normal) were chosen. Second, because choice

[38] For each medical care price, I tested a total of eleven parametric distributions for goodness-of-fit following the methodology of Jones et al. (2014). Details are provided in Section 7 below.

errors are additive and distributed T1EV, choice probabilities have the following closed forms.

$$P(m_t^{dhr} = 1|\mathbf{S_t}, \mu^k, I_y^j) = \frac{exp\left(\overline{V}_{dhr}(\mathbf{S_t}|\mu^k, I_y^j)\right)}{\sum_{d'=0}^{D}\sum_{h'=0}^{H}\sum_{r'=0}^{1} exp\left(\overline{V}_{d'h'r'}(\mathbf{S_t}|\mu^k, I_y^j)\right)} \quad \forall t, \forall dhr \tag{11}$$

$$P(I_y^j = 1|\mathbf{S_0}, \mu^k) = \frac{exp\left(\overline{Q}_j(\mathbf{S_0}, \mu^k)\right)}{\sum_{j'=0}^{J^i} exp\left(\overline{Q}_{j'}(\mathbf{S_0}, \mu^k)\right)} \quad \forall y, \forall j. \tag{12}$$

The likelihood contribution for individual i in month t conditional on μ^k and I_y^j is

$$\begin{aligned} L_{it}(\Omega|\mu^k, I_y^j) = &\prod_{a=0}^{4}\left[\pi_t^a(\cdot|\mu^k)\right]^{\mathbb{1}_{A_t=a}} \prod_{c=0}^{2}\left[\gamma_t^c(\cdot|\mu^k)\right]^{\mathbb{1}_{C_t^+=c}} \\ &\prod_{d=0}^{D}\prod_{h=0}^{H}\prod_{r=0}^{1}\Big[f^d(p_t^d|\mu^k)^{[1-m_t^{0hr}]} f^h(p_t^h|\mu^k)^{[1-m_t^{d0r}]} f^r(p_t^r|\mu^k)^{[1-m_t^{dh0}]} \\ &P(m_t^{dhr} = 1|\mathbf{S_t}, I_y^j, \mu^k)\Big]^{m_t^{dhr}} . \end{aligned} \tag{13}$$

The first row contains the illness state contribution for month t. Price densities comprise the second row. The $[1 - m_t^{0hr}]$ exponent, for example, ensures that the price of a doctor visit in month t contributes to the likelihood function only if an individual actually visits the doctor, which is the only time that I observe this price. The third row contains month t medical care choice probabilities.

An individual's health and illness states entering the year are a function of prior medical care consumption and related behaviors, as well as his personal preferences for good health and medical care, making these initial conditions endogenous. To control for this endogeneity, $L_{i1}(\Omega|\mu^k, I_y^j)$ appears as above with the first row replaced by

$$\prod_{s=0}^{4}\left[\tilde{\eta}^s(\cdot|\mu^k)^{\mathbb{1}_{HS=s}}\right] \prod_{a=0}^{4}\left[\tilde{\pi}_1^a(\cdot, s|\mu^k)\right]^{\mathbb{1}_{A_1=a}} \prod_{c=0}^{4}\left[\tilde{\gamma}_1^c(\cdot, s|\mu^k)\right]^{\mathbb{1}_{C_1=c}} . \tag{14}$$

Initial probabilities $(\tilde{\eta}, \tilde{\pi}, \tilde{\gamma})$ are separately estimated from transition probabilities, with exclusion restrictions, and are allowed to vary by modeled permanent unobserved heterogeneity.[39] An individual is allowed to have between zero and four initial chronic illnesses. I also control for the en-

[39]Exclusion restrictions are income last year, missing last year's income, veteran status, and foreign birth status.

dogeneity of initial self-reported health (excellent, very good, good, fair, poor), which is allowed to impact illness transition probabilities over the course of the year, using ordered logit probabilities.[40]

The total likelihood contribution for individual i conditional on μ^k is then written

$$L_i(\Omega|\mu^k) = \prod_{j=1}^{J^i} \left[P(I_{iy}^j = 1|\mathbf{S_0}, \mu^k) \prod_{t=1}^{T} L_{it}(\Omega|\mu^k, I_y^j) \right]^{I_{iy}^j}. \tag{15}$$

The contribution of individual i unconditional on unobserved type is $L_i(\Omega) = \sum_{k=1}^{K} \theta^k L_i(\Omega|\mu^k)$.

6 Insurance Alternative Set Generosity as an Instrument

The structural model uses a combination of (i) variation in employer provided insurance alternative sets and (ii) joint modeling of insurance and medical care demand with permanent unobserved heterogeneity to separately identify moral hazard effects from selection (see Section 5.3 of the paper for details). The first strategy argues that if health insurance alternative sets are exogenously determined, then the generosity of an individual's alternative set serves as an instrument for his selected plan. In other words, the generosity of the alternative set affects the generosity of the held plan, but only affects medical care demand through it's impact on the held plan. An additional control for endogenous selection is used (i.e., (ii) above) because all insurance alternative sets may not be exogenously determined, as some individuals may select their employer due in part to the insurance benefits offered. That said, if the exogeneity condition holds for many individuals, the identification strategy should help in reducing *some* selection bias.

In this section, I use a series of reduced-form models to show how instrumenting for annual medical care prices (i.e., measures of held insurance plan generosity), with measures of the generosity of one's health insurance alternative set, can reduce bias in moral hazard estimates. I present results for two measures of medical care demand: (i) logged total annual medical care expenditure and (ii)

[40] A model that is more consistent with Grossman (1972) might allow for self-reported health status, or *health stock*, to update in each month after medical care is consumed. Unfortunately, health status is reported only once per interview period; a maximum of three times over the course of a health insurance year. As such, I allow initial health status to impact illness transition probabilities and rely on the sum of acute and chronic illnesses to capture monthly changes in an individual's health stock.

a binary measure of any expenditure over the course of the year. Both measures focus on doctor and hospital expenditure only, as insurance cost-sharing for prescription drugs is not observed in the data.

Medical care demand is regressed on three measures of annual medical care prices: (i) the expected end-of-year share of last dollar medical care expenditure paid out of pocket, (ii) the actual/observed end-of-year share of last dollar medical care expenditure paid out of pocket by an individual, and (iii) the mean marginal out-of-pocket share paid by an individual over the course of the year.[41] Aron-Dine et al. (2013) provide motivation for consumer responsiveness to each of these three prices.

To measure the generosity of each individual's health insurance alternative set, I begin by simulating $2,000 in total annual medical care expenditures by each individual under each of their available plans. I assume that this expenditure is split evenly between hospital and doctor's office care. From this simulation, I create two instruments: (i) the minimum out-of-pocket dollars required from an individual's available plans and (ii) the average out-of-pocket dollars required from an individual's available plans. I also tried specifications including quadratics of these instruments and total number of offered plans as an additional instruments. These additional specifications weakened the instrument set.

Table A20 contains the results of this analysis. Note that all models include controls for sex, race, education, age, MSA, and income. My primary findings are presented in the first two rows, where medical care demand is measured as logged total annual medical care expenditure. Using OLS, a 10 percentage point increase in the co-insurance rate is found to increase total annual expenditures by between 11.6 and 25.7 percent, depending on the price measure used. This specification does not control for the endogeneity of prices, which is primarily due to insurance selection. Using 2SLS to control for this endogeneity, I find that a 10 percentage point increase in the co-insurance rate increases total annual expenditure by between 8.33 and 10.2 percent. Using each of the three price

[41] I describe how monthly, expected end-of-year shares are calculated in Section 3.8 above. The annual measure used here is simply the expected end-of-year share calculated in the first month of the year, when health insurance is selected. As is also done in the previous section, when calculating out-of-pocket shares, doctor and hospital deductibles are added together, while doctor and hospital co-insurance rates are averaged, so individuals face a single out-of-pocket share across medical care types.

Table A20: OLS and 2SLS Estimates of Price Sensitivity

	Expected OOP EOY Share		Actual OOP EOY Share		Mean Marginal Share	
Dependent Variable	Est.	S.E.	Est.	S.E.	Est.	S.E.
Log Total Annual Expenditure						
OLS	-1.160	0.259	-2.568	0.188	-1.844	0.194
2SLS	-1.022	0.548	-1.006	0.511	-0.833	0.428
Hausman Test	0.107		10.656		6.973	
(P-value)	(0.744)		(0.001)		(0.008)	
Hansen J-Stat	2.830		2.749		2.687	
(P-Value)	(0.243)		(0.253)		(0.261)	
Any Annual Exp.						
OLS	-0.198	0.039	-0.283	0.032	-0.225	0.032
2SLS	-0.342	0.083	-0.330	0.077	-0.273	0.064
Hausman Test	3.851		0.444		0.667	
(P-value)	(0.050)		(0.505)		(0.414)	
Hansen J-Stat	0.953		0.802		0.722	
(P-Value)	(0.621)		(0.670)		(0.697)	
First Stage F-Stat	55.48		43.65		61.19	
(P-value)	(0.000)		(0.000)		(0.000)	

† First Stage F- and J-Stats correspond to 2SLS regressions only. All regressions include 1,232 observations. The following additional controls are included in all specifications: sex, race, education, age, MSA, and income.

measures, the joint F-test suggests that the instrument set is strong and the Hansen J-test suggests that the instruments are exogenous. The Hausman test shows that the OLS and 2SLS estimates are significantly different from one another when the actual end-of-year price and mean marginal price is used, suggesting that the price measures are in fact endogenous.

These results provide clear evidence that the generosity of one's health insurance alternative set can be used as an instrument for the generosity of one's held plan. Here, moral hazard effects were reduced by controlling for selection bias, suggesting that individuals adversely select their health insurance plan among their available alternatives based in part on their expected future expenditure. Note that the Hansen test also suggests that the generosity of one's health insurance alternative set is exogenously determined, which is relevant for the structural model; though, this is a weak specification test.

In the second set of rows, I estimate how prices affect whether an individual has any annual expenditures. The main take-away from this analysis is that correcting for selection bias consistently *increases* moral hazard estimates, but the OLS and 2SLS estimates are significantly different from one another in just one specification. These results suggest that adverse selection does not occur

on the extensive margin - i.e., people do not select into health insurance based on an expectation of any vs. no consumption; rather, they select in based on expectations of high vs. low expenditure. This finding could also mean that extensive margin selection occurs in the selection of an employer, but this seems unlikely.[42]

7 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, he 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 Figure A3 below), as potentially large price draws motivate both the purchase of health insurance and provide a disincentive for frequent medical care consumption.

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 controlling for (i) permanent unobserved heterogeneity and (ii) endogenous selection into medical care consumption in the full model would not cause one distributional assumption to dominate another.

The eleven distributions I test are shown in Table A21 below. Jones et al. (2014) test the same distributions 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, which is also characterized by long, thin right tails. All other distributions considered are nested and limiting

[42]Not shown are results measuring medical care demand using total annual doctor and hospital visits. Find significant reductions in moral hazard effects across specifications for doctor visit demand. For hospital visit demand, standard errors are large, so no significant differences exits between OLS and 2SLS specifications.

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 Table A21.[43]

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 A21 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.

Table A22 compares various percentile values of the observed distribution with their simulated counterparts using the distributional assumptions discussed above. The simulated prescription drug and doctors office distributions each fit the observed data well. The simulated hospital distribution

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

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. Hospital prices are similarly truncated when simulating the structural model.

Figure A3: Empirical Price Distributions

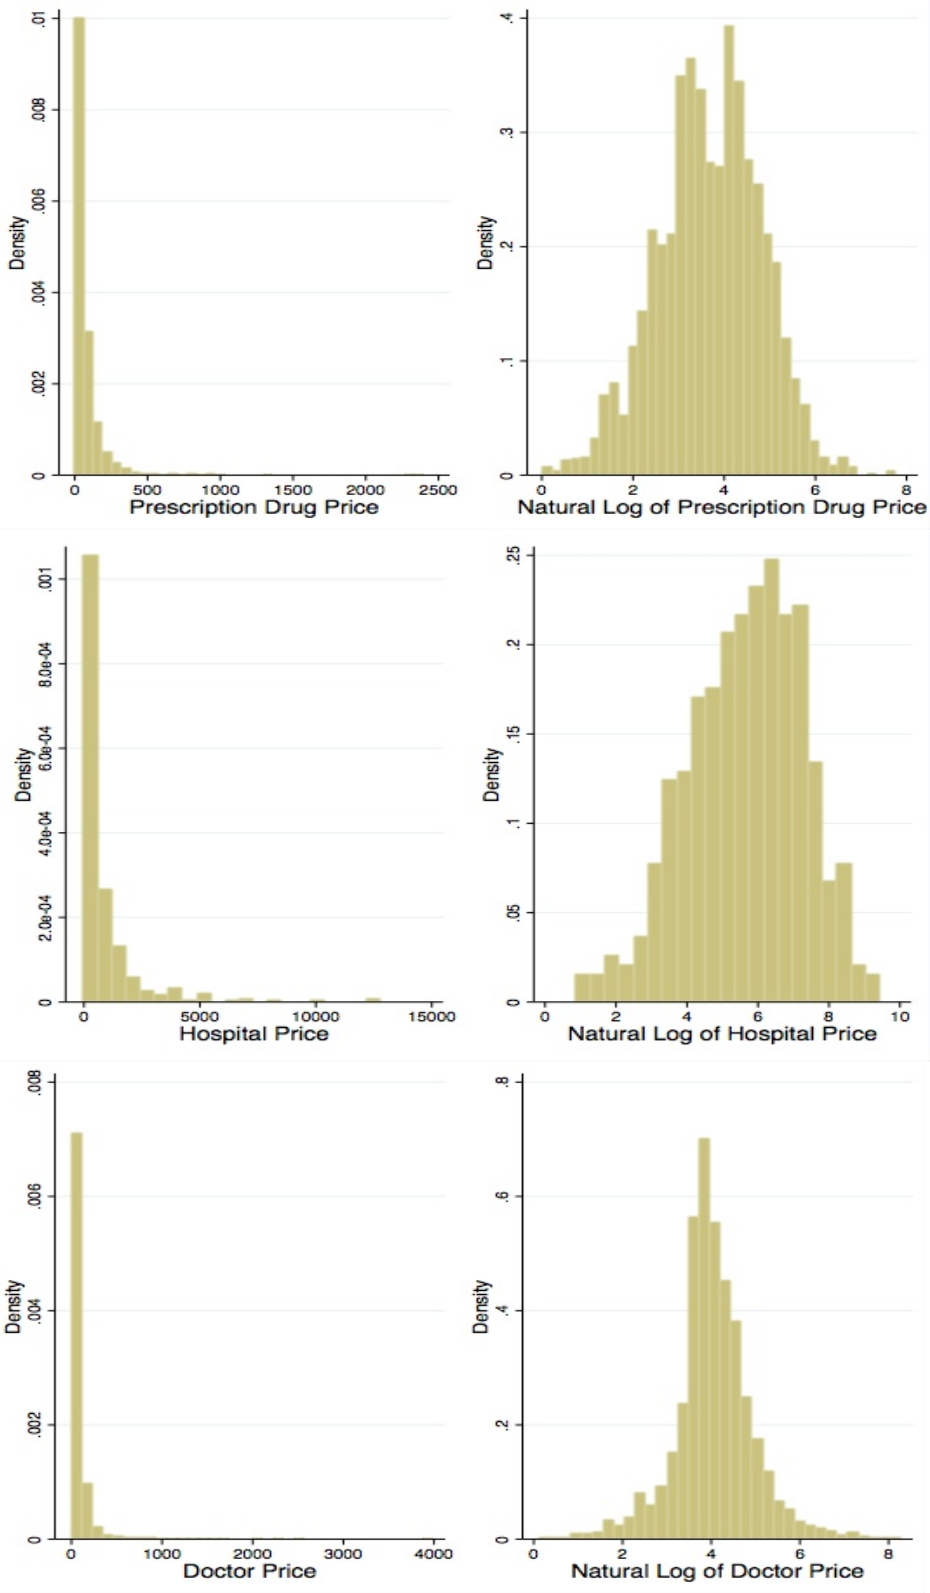

Table A21: Goodness of Fit Across Parametric Distributions

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

Table A22: Price Simulations for Selected Distributions

	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

8 Additional Structural Parameters: Price, Initial Condition, and Closing Function

Table A23: Structural Price Parameter Estimates

	Doctor Price		Hospital Price		Prescription Price	
	Est.	S.E.	Est.	S.E.	Est.	S.E.
constant	46.0687	6.1075	4.8073	0.4075	2.7601	0.0669
male	0.6465	1.3201	0.2742	0.1164	0.1557	0.0148
non-white (black or Hispanic)	0.9701	1.7763	0.1114	0.1387	-0.0900	0.0161
education (highest grade completed)	-0.0634	0.2940	0.0368	0.0230	0.0001	0.0032
age	0.1573	0.0687	0.0116	0.0054	0.0033	0.0007
lives in a MSA	12.3417	1.5972	-0.2760	0.1441	0.1272	0.0200
income (in 1996 dollars)	0.0934	0.0379	0.0096	0.0040	0.0027	0.0003
initial health	-0.1577	0.7375	0.0174	0.0582	0.0899	0.0077
HMO	-3.2647	1.4617	0.2219	0.1433	-0.1424	0.0140
PPO	-3.4918	2.4293	-0.2162	0.1868	0.0852	0.0296
no insurance	-15.3438	2.2738	-2.4222	0.2072	-0.4426	0.0307
acute illnesses	-4.6980	1.4736	0.0810	0.1398	0.1036	0.0180
acute illnesses (squared)	1.2520	0.3657	0.0057	0.0347	-0.0071	0.0049
chronic illnesses	-1.8965	1.1936	-0.0358	0.1096	0.3636	0.0142
chronic illnesses (squared)	0.4138	0.2183	0.0006	0.0151	-0.0177	0.0024
shape 1†	2.1544	0.0391	1.1566	0.0385	0.8033	0.0041
shape 2†	0.9031	0.0359	*	*	*	*

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.

Table A24: Initial Condition Parameter Estimates

	Health Status		Acute Illness		Chronic Illness	
	Est.	S.E.	Est.	S.E.	Est.	S.E.
Initial Condition						
constant	2.1265	0.4270	-1.5873	0.6021	-5.2351	0.7405
male	-0.1440	0.1339	-0.5475	0.1530	-0.3923	0.1797
non-white (black or Hispanic)	0.2043	0.1356	-0.4217	0.1754	-0.3111	0.2061
education (highest grade completed)	-0.1005	0.0269	0.0390	0.0315	0.1512	0.0410
age	0.0204	0.0056	0.0051	0.0147	0.0568	0.0169
lives in a MSA	-0.1345	0.1629	0.0895	0.2026	-0.3269	0.2185
income (in 1996 dollars)	-0.0056	0.0041	-0.0005	0.0049	-0.0042	0.0050
March/April (indicator)	*	*	-0.0777	0.2817	*	*
May/June (indicator)	*	*	0.3950	0.2912	*	*
July/August (indicator)	*	*	0.3756	0.2082	*	*
September/October (indicator)	*	*	0.1512	0.2692	*	*
November/December (indicator)	*	*	0.0232	0.3829	*	*
initial health	*	*	0.2107	0.1605	0.5382	0.1918
initial health*age	*	*	0.0010	0.0064	0.0032	0.0072
last year income	-0.0748	0.0404	-0.0557	0.0494	-0.0275	0.0567
last year income missing	-0.2294	0.2686	-0.4274	0.3344	-0.5290	0.3826
veteran (indicator)	0.1353	0.2168	0.0150	0.2621	-0.1838	0.3026
foreign born (indicator)	-0.2105	0.2212	0.2807	0.2486	-0.1373	0.2836
threshold 1	1.5161	0.0736	1.6354	0.1066	1.7092	0.1230
threshold 2	3.3160	0.1437	3.0146	0.1997	3.0549	0.1970
threshold 3	6.4692	0.6487	4.1915	0.3699	4.3233	0.3311

Table A25: Closing Function Parameter Estimates

	Estimate	SE
Closing Function		
doctor visits	-0.0741	0.0970
doctor visits (squared)	0.0102	0.0153
hospital days	-0.0723	0.3118
hospital days (squared)	-0.0594	0.0922
Rx consumption	-0.3432	0.1669
acute illness†	-105.9790	15.0042
acute illness (squared)	3.3967	1.7980
acute illness*age	0.6927	0.1733
chronic illness†	-884.2178	348.2993
chronic illness (squared)‡	-26.8637	14.8387
chronic illness*age	-0.8261	2.7786

† Parameters are scaled by a factor of 100 in estimation to improve numerical optimization.
‡ Parameter is scaled by a factor of 10 in estimation to improve numerical optimization.

9 Analysis of Model Fit

To assess the model's ability to explain unique features of the data, I use the model to simulate a new dataset and compare key moments of 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. I then forward simulate annual insurance decisions and monthly medical care decisions, illness transitions, and medical care prices for each of the 1,232 observed individuals, using observed insurance offer sets, demographics, and initial conditions. The resulting dataset contains 1,478,400 $[= 1,232 * 12 * 100]$ person-month observations.

Table A38 summarizes the observed and simulated data. The model is able match illness and price moments well, which is unsurprising, given that these moments are fairly stable over the insurance year. Medical care consumption patterns are more difficult to explain, as consumption varies substantially by illness state and the annual distribution is non-normal. The first row of the third section in Table A38 shows that the likelihood of medical care consumption in any given month in the simulated data matches that of the observed data well. More importantly, the next three rows show that the model is capturing the relationship between illness and consumption, suggesting that much of the variation in simulated medical care consumption is being created by illness and not just the logit structure. Furthermore, mean and median annual medical care consumption counts in the simulated data are close to those in the observed data, though the model underpredicts prescription drug consumption by a small amount.

While the model matches mean annual medical care consumption for all three types of care, the variances of the consumption distributions are larger in the observed data. The source of this discrepancy is better understood by looking at the entire distribution of annual medical care consumption in Table A39. Notice that for each type of medical care consumption, at least 28% of the observed population consumes zero units in a year; however, 14.12% and 27.52% of the observed population consume more than ten doctor visits and prescription drug months in a year, respectively. As the table reveals, the model does a fairly good job of capturing the fact that many individuals consume no medical care in a calendar year, while many others consume large amounts.[44] That said, the model tends to underpredict both the share of the population with zero

[44]The DFM improves the fit of all three consumption distributions. As is explained in Section 11 below, the model

Table A26: Observed and Simulated Outcomes

	Observed			Simulated		
	Mean	Median	S.D.	Mean	Median	S.D.
Illness						
% entered year with chronic illness	38.23			37.42		
% with chronic illness by year's end	48.54			47.15		
ave. number of chronic illnesses†	1.78	1.00	1.05	1.71	1.00	1.05
% with ≥ 1 acute illness in year	81.41			85.29		
ave. total months with acute illness	5.37	4.00	4.69	5.27	4.00	4.19
ave. number of acute illnesses†	1.52	1.00	0.81	1.45	1.00	0.74
Medical Care Prices‡						
doctor visit price	90.61	55.00	159.18	88.17	55.60	191.66
hospital day price	828.48	329.60	1435.91	807.03	355.30	1637.26
prescription drug price	74.37	41.93	113.63	75.25	40.36	109.03
Medical Care Consumption						
% consuming (any month)	43.53			42.62		
% consuming when $A_{it} \neq 0$	68.06			64.82		
% consuming when $C_{it} \neq 0$	72.80			66.27		
% consuming when $A_{it} = C_{it} = 0$	8.69			13.94		
annual doctor visits	4.89	2.00	8.69	4.86	2.00	6.62
annual hospital days	0.65	0.00	2.00	0.65	0.00	1.51
months with Rx cons.	4.45	2.00	4.96	4.28	2.00	4.66
Medical Care Expenditure						
total annual medical care exp.	1473.13	358.34	3656.88	1267.70	406.25	3062.88
Held Insurance Types						
% No Insurance	9.17			11.34		
% HMO	37.99			37.46		
% PPO	10.14			12.46		
% FFS	42.69			38.74		

† Conditional on any chronic/acute illness.
‡ 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.

medical care consumption and the maximum amount of consumption for each type of care.[45]

Figure A6 shows both observed and simulated annual expenditure distributions on a log scale. Total annual medical care expenditure is a good indicator of overall model fit, as it encompasses the fit of all three consumption and price distributions. The figure reveals an underprediction of zero medical care expenditure (i.e., 17% of the simulated sample vs. 22% of the observed sample),

that allows for four unobserved types generates significantly more zero consumers and right-tail consumers than the model with only one unobserved type. Adding a fifth mass point did not further improve the fit of these distributions, though it increased estimation time substantially.

[45]The maximum number of observed annual doctor and hospital visits is 103 and 37, respectively. The corresponding maximums in the simulated data are 67 and 22. Furthermore, 22.7% of the observed population consumes prescription drugs in every month while only 13.2% of the simulated population does.

an overprediction of the median expenditure, and an underprediction of the mean expenditure (see Table A38). These differences in the expenditure distribution can be explained by the differences in medical care consumption discussed above. Furthermore, Table A38 reveals that the model underpredicts insurance possession by 2.3 percentage points, which makes sense given that the simulated expenditure distribution has a mean and standard deviation that are smaller than what is observed in the data. In other words, the model slightly underpredicts insurance possession because individuals are not worried enough about being in the right tail of the expenditure distribution.

Table A27: Observed and Simulated Annual Consumption

	Doctor visits		Hospital days		Prescriptions	
	Obs.	Sim.	Obs.	Sim.	Obs.	Sim.
0	28.17	24.22	75.81	71.93	35.47	31.14
1	16.15	16.04	12.09	15.43	13.31	16.22
2	12.66	11.47	4.87	4.75	5.60	7.95
3-5	18.43	19.34	4.79	5.88	10.47	9.67
6-9	10.46	12.29	1.46	1.60	7.63	10.47
10+	14.12	16.67	0.97	0.41	27.52	24.55

Notes: The table reports annual consumption levels for the observed and simulated samples. The values are the percentage of the samples consuming at each level. For example, the top left entry states that 28.17% of the observed sample visited a doctor's office zero times during the insurance year. The prescription drug levels are measured in consumption months.

10 Price Sensitivity across Types of Medical Care

To my knowledge, there are no papers in the structural literature that attempt to separate price sensitivity for hospital, doctor, and prescription drug care. Using RAND HIE data, Manning et al. (1987) finds that price sensitivity for inpatient hospital care is on par with outpatient care for the 0-25 co-insurance range, but somewhat smaller for the 25-95 co-insurance range. Using Oregon HIE data, Finkelstein et al. (2012) find a 30% increase in hospitalization and 20% increase in hospital days in response to Medicaid acquisition, which reflects greater price sensitivity than that for prescription drugs but less sensitivity than that for outpatient doctor office visits. As such, my results suggest greater sensitivity to hospital prices, relative to other forms of medical care, than the existing literature. This result is partially explained by the fact that I categorize inpatient (11%), outpatient (62%), and ER (27%) visits as hospital visits, while the existing literature typically

Figure A4: Annual Medical Care Expenditure

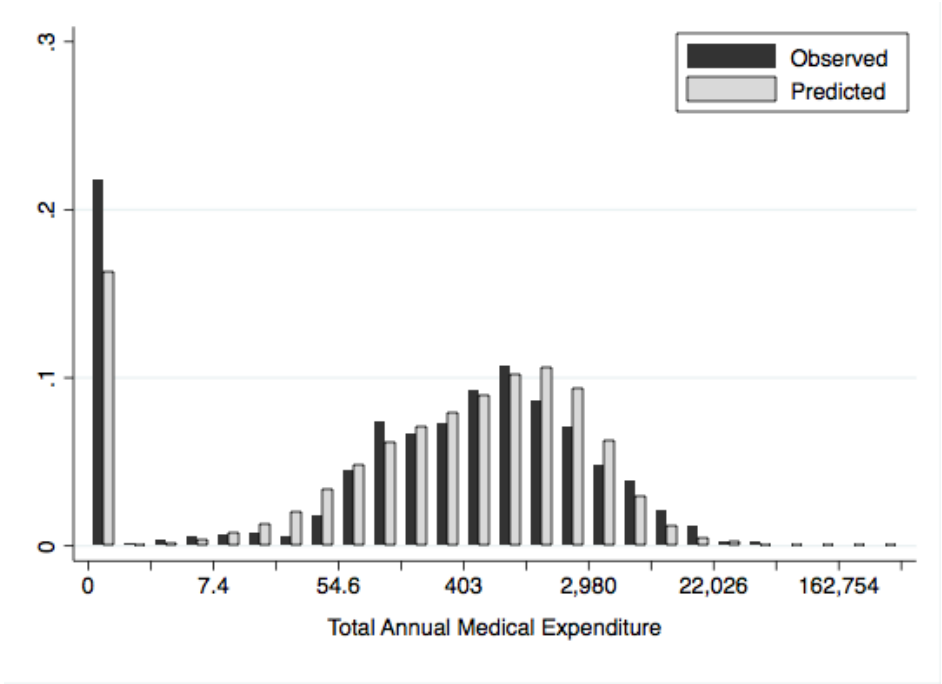

Notes: This figure compares observed and simulated annual expenditure distributions. The horizontal axis is log-scaled, so that markers are spaced $exp(1.0)$ apart.

isolates inpatient care; patients are likely to display much more price sensitivity for outpatient and ER visits. Greater sensitivity to hospital prices is also likely to result from several differences between the underlying data generating process and the WYDM model, which I discuss below. I argue that these differences have a minimal impact on the model's structural parameters, which determine the full distribution of moral hazard effects. However, it is easy to see how these failures overstate how undesirable hospital care is to uninsured patients, which affects the decomposition of moral hazard effects by doctor, hospital, and prescription drug care.

First, in the US, hospitals serve as a form of insurance for the uninsured. The Emergency Medical Treatment and Active Labor Act, passed in 1986, requires all Medicare accepting hospitals to stabilize patients facing an emergency (loosely defined), regardless of their ability to pay. As such, many individuals go to the hospital when sick, even if they know that they will be unable to pay the resulting bill. When hospitals ultimately bill these individuals, they can (i) pay the bill, (ii) negotiate for a lower price, or (iii) simply refuse to pay the bill. When (iii) occurs, hospitals often report failure of payment to a credit agency, harming the individual's credit score, and have the option of selling the debt to a collection agency. Options (ii) and (iii) are not allowed in the

model - I assume that all uninsured individuals pay the full charge price for care (see Section 3.4.1 above).[46] This decision likely causes an underestimation of patient price sensitivity, meaning my moral hazard estimates are too small. Relative to the insured, many uninsured patients are observed go to the hospital despite the (modeled) high prices, specifically because they know that they will not ultimately pay their full bill. An alternative strategy would be to code the observed price paid as the charged price, which would increase estimated price sensitivity.

Second, hospital emergency rooms frequently serve as urgent treatment centers, particular for the uninsured, who often do not have a primary care provider (PCP). In addition to the fact that the emergency room must treat an individual independent of their ability to pay, emergency rooms are first-come, first-serve and do not take appointments; care is received within a few hours. Because doctor's offices take appointments, PCPs often cannot see patients for several days after being contacted. If an individual is not a previous patient of a PCP, the wait can be weeks or even months. As such, there may be a waiting cost associated with going to a doctor's office, particularly for uninsured individuals, that does not exist at the emergency room. I do not model the potential waiting cost of going to the doctor; thus, uninsured individuals are again observed visiting the hospital, in relatively high proportions, despite high prices, which likely causes my estimates of price sensitivity to be understated. That said, note that if this cost is regarded by individuals as being permanent over the course of the year, then the model should capture it through preferences for medical care and the permanent unobserved heterogeneity; however, if regarded as time varying, then the model might miss this cost.

Third, the model assumes that all treatment is elective, though some certainly is not. I think of this issue as an omission of *illness severity* in the model, which prevents the impact of treatment on future health from varying by illness severity. The omission of illness severity is best understood

[46] Negotiating lower prices is popular in the data. I observe the uninsured take a total of 41 visits to the hospital, though some visits are for multiple days. 15% of the visits result in the uninsured individual paying the full charged amount. The average amount paid is half the charged amount. Only two individuals do not pay their bill at all. When I assume that uninsured patients pay the observed *paid price* for a hospital visit, rather than the observed *charged price*, average total annual expenditure on hospital visits drops from $1108 to $521, but this is due almost entirely to large reductions in the highest prices. In the footnote of Table 3, I mention that removing the largest hospital expenditure by an uninsured individual decreases the mean from $1108 to $647.

with an example: imagine someone in the data were hit by a bus. In the data, one would observe a large increase in medical care consumption, regardless of the prices being charged, that the model can not justify with the (average) improvements in illness probabilities estimated in the illness production function. Such behavior again causes the WYDM model to understate consumer price sensitivity, at least for the class of treatments which are elective. Again, some of this behavior is sure to be explained by permanent unobserved heterogeneity, but not all.[47]

These three differences between the WYDM model and the true data generating process, if anything, should lead the moral hazard effect to be *underestimated*. There are several reasons that this threat is not overly concerning. First, my estimates are a bit larger than those found in the existing literature. Second, all three differences should be most relevant for uninsured individuals, who make up just 9% of the sample. Moreover, because many insured individuals hold plans with deductibles, the model does not rely on uninsured individuals alone to identify responsiveness to full medical care prices. Third, the three threats are most relevant for hospital visits, which are rare; moreover, the parameters most important for measuring price sensitivity are the CRRA parameters, $\{\omega_{00}, \ldots, \omega_{03}\}$, which are not specific to any particular type of medical care. As a result, the impact that unmodeled hospital costs/benefits can have on these parameters is limited. That said, these factors do influence the interpretation of the estimated moral hazard effect. Specifically, 53.1% is the percentage of mean (insured) annual medical care expenditure that occurs only due to insurance possession *according to the model*. In other words, 53.1% is the moral hazard effect in a world where the three discrepancies between the model and the real world discussed above do not exist.

While these three issues should not have a large impact on my overall moral hazard estimate, they are likely to impact the decomposition of the moral hazard effect into doctor, hospital, and prescription drug care. Note that all three issues rationalize why individuals go to the hospital, even when prices are high and hospital visit productivity is relatively low. Without these incentives in the model, the simulation suggests that uninsured individuals virtually never go to the

[47] The model captures just one year of behavior. If an individual is hit by a bus early in the year, he is likely to consume lots of medical care independent of prices for the full year. The model would then suggest that this individual *likes* medical care, by placing more population weight on an unobserved type with strong permanent preferences for medical care. However, if such a severe illness/injury did not produce consecutive months of treatment, then the model will understate the importance of visiting the hospital in the event of such serious illnesses.

hospital. However, as long as the model captures the underlying level of price sensitivity correctly, total expenditure given current prices should be correct. In otherwords, while underpredicting hospital consumption when uninsured, the model should compensate by overpredicting doctor and prescription drug use when uninsured.[48]

11 The Role of Permanent Unobserved Heterogeneity

Permanent unobserved heterogeneity plays a number of roles in the model. Virtually all of these roles are related to the fact that illness *severity* cannot be perfectly observed in the data. I observe illnesses (i.e., ICD-9 codes) that can be categorized as acute or chronic. I measure an individual's *wellness* by the number of illnesses they possess. This measure of wellness is clearly imperfect, as one intense illness can cause much more harm/disutility and motivate much more medical care consumption than many minor illnesses. Thus, some relevant measure of wellness is likely unobservable in the model. There are several ways in which this unobserved wellness could bias some parameter estimates (e.g., the impact of medical treatment on illness) and generate heterogeneity in others (e.g., the marginal utility of a doctor's office visit), which I attempt to capture by modeling permanent unobserved heterogeneity. I discuss how permanent unobserved heterogeneity affects the model's parameters and substantive conclusions below. Note that I have re-estimated the structural parameters of the WYDM model with one and two unobserved types to aid in this discussion.[49]

[48] There is evidence of this overprediction in Tables 3 and 15. In Table 15 average annual expenditure when uninsured is \$340.96 and \$253.00 for doctor and prescription drug consumption, respectively. In Table 3, observed average annual expenditure when uninsured is \$289.67 and \$111.48 for doctor and prescription drug consumption, respectively.

[49] Note that the first unobserved type is always normalized to have no effect on the endogenous variables of the model; thus, the model with one unobserved type does not allow for any correlation between the models unobservables (i.e., the standard conditional independence assumption). The model with two types allows for such correlation.

11.1 Effect on Model Parameters and Simulations

Tables A28 - A32 below contain re-estimated model parameters.[50] One can think of the permanent unobserved heterogeneity as allowing for heterogeneity in some parameters, while correcting bias in others. With respect to prices, illnesses, and initial conditions, the permanent unobserved heterogeneity allows for variation in the respective constant terms. In the contemporaneous utility function, the permanent unobserved heterogeneity (i) captures wide variation in preferences for medical care and (ii) allows this variation in preferences to be related to the unexplained likelihood that an individual becomes ill or receives a high price draw. For example, in Tables A28 and A31, one can see that in moving from one (i.e., no permanent unobserved heterogeneity) to two unobserved types, the model begins to identify a sub-population of individuals who's negative tastes for prescription drugs are not as strong. Identifying this group is important for correctly predicting the distribution of prescription drug consumption over the health insurance year, which can be seen in Table A33. With only one unobserved type, the model cannot explain why the two most popular types of prescription drug consumers are those that never consume (35.5%) and those that virtually always consume (27.5%). Without unobserved heterogeneity, the model predicts that 11.8% of the population never consumes and 12.1% consumes in 10 or more months, both far from the observed percentages. By allowing for just two unobserved types, the model is able increase these percentages to 24.7% and 20.2%, respectively. Moving to four unobserved types brings these percentages even closer to the observed levels. Capturing this preference heterogeneity has a similar effect on the simulated distribution of doctors office visits, which is also characterized by extreme consumers; some consuming a lot and some little. Taken together, because permanent unobserved heterogeneity improves the fit of the medical care consumption distributions, it also significantly improves the fit of the overall medical care expenditure distribution, which can be seen in Figure A5. Without permanent unobserved heterogeneity, the model predicts too few zero and extreme spenders and too many moderate spenders.

In addition to allowing for heterogeneity in preferences for medical care, modeling permanent unobserved heterogeneity reduces bias in other parameters, by capturing correlation in the models unobservables. For example, Table A28 reveals that as unobserved types are added, individuals

[50] Price, closing function, and initial condition parameters are available upon request.

dislike acute illnesses more. This preference parameter is identified by differences in medical care consumption across illness states - i.e., the fact that individuals with more acute illnesses consume more medical care (and take on the costs associated with consuming medical care) than individuals with no/fewer acute illnesses suggests a distaste for acute illness. Table A32 reveals 49% of the sample is described by types 2, 3, and 4; types that tend to have low preferences for doctor and hospital care but face a high likelihood of illness. Once the model can condition on this permanent correlation between medical care preferences and illness, the positive time-varying relationship between acute illness occurrence and medical care consumption becomes stronger, which is reflected by an increase in the estimated disutility for acute illness.[51]

11.2 Effect on Moral Hazard

The primary challenge in estimating moral hazard effects is the endogenous selection of health insurance - i.e., the tendency of individuals who expect to consume lots of medical care to purchase the best health insurance. Presumably, the two most likely reasons that one might expect to consume lots of medical care are: (1) a strong preference for medical care consumption and (2) personal knowledge of an existing or impending illness. If individuals with strong preferences for medical care, severe illnesses, or impending illnesses select into generous plans and then spend a lot, it will seem as if insurance led to this spending. Likewise, if healthy individuals with weak preferences for medical care select into a non-generous plan and do not spend, it will seem as if a lack of insurance caused their lack of spending. To prevent assigning endogenous selection to moral hazard effects, researchers typically control for observables correlated with high future expenditure, such as age and observed illness, when modeling insurance decisions; however, individuals may still select insurance based on their own preferences for medical care, unobserved current illnesses, and

[51] Another example: Allowing for permanent unobserved heterogeneity has an impact on the effectiveness of prescription drugs. Table A29 suggests that prescription drugs become a more effective treatment for acute illnesses with additional unobserved types. This finding likely suggests that those regularly selecting into using prescription drugs have more severe unobserved acute illnesses, which bias the effectiveness of the drugs towards zero. Table A30 suggests that prescription drugs become a *less* effective treatment for chronic illnesses with additional unobserved types, which likely suggests that a continuous drug regimen is important to preventing the development of chronic illnesses, not just recent prescription drug use.

anticipated future illnesses.

As is discussed above, permanent unobserved heterogeneity plays a critical role in capturing both variation in medical care preferences and permanent unobserved illness severity - the exact characteristics upon which we generally worry that individuals adversely select health insurance. Moreover, as is evident in Equation 14 in the paper, the model assumes that individuals make insurance decisions conditional on their unobserved type; thus, allowing for permanent unobserved heterogeneity is critical to controlling for adverse selection in the model. Such is evident in Table A34, which shows the impact that permanent unobserved heterogeneity has on moral hazard estimates. First, note that as I move from 1 to 2 to 4 unobserved types, the percent of (insured) mean annual medical care expenditure explained by insurance possession (i.e., the moral hazard effect) falls from 57.4% (s.d. 3.1) to 56.1% (s.d. 4.2) to 53.1% (s.d. 4.0). Second, the insurance-induced spending distribution shifts significantly. With only one unobserved type, roughly 95% of the population responds to insurance acquisition by consuming more medical care. With additional types, more and more individuals do not respond at all, while the far right tail of the distribution thickens. These changes are consistent with Table A33 and Figure A5, which highlight the impact that preference variation has on medical care consumption. The unobserved types identify sub-populations that are unlikely to use medical care, even when ill (i.e., types 2 and 3, but type 4 in particular). As such these individuals are less responsive to changes in the out-of-pocket cost of medical care that is associated with insurance acquisition.

Table A28: WYDM Model: Utility Parameters with Various Mass Points

	1 Mass Point		2 Mass Points		4 Mass Points	
	Est.	S.E.	Est.	S.E.	Est.	S.E.
RA constant	0.0746	0.0062	0.0639	0.0056	0.0594	0.0066
RA age	0.0011	0.0004	0.0013	0.0003	0.0010	0.0004
RA non-white	0.0034	0.0097	0.0035	0.0109	0.0040	0.0114
RA male	0.0440	0.0152	0.0484	0.0189	0.0213	0.0139
acute illnesses	-11.0128	2.0317	-15.2972	2.7167	-28.5934	4.4581
acute illnesses (squared)	-0.7202	0.3778	0.2392	0.4940	2.7975	0.7919
acute illnesses*age	0.1816	0.0333	0.2077	0.0388	0.1894	0.0370
chronic illnesses	-5.9103	3.0156	-14.5791	4.9399	-3.5427	5.208
chronic illnesses (squared)	-0.7400	0.2066	0.0845	0.3409	-0.7278	0.6208
chronic illnesses*age	0.2709	0.0801	0.2791	0.1287	-0.1531	0.1761
doctor visits	-2.5541	0.0908	-2.5540	0.1291	-1.5748	0.1832
doctor visits (squared)	0.2117	0.0152	0.2030	0.0184	-0.1772	0.0276
doctor visits*age	-0.0054	0.0011	-0.0090	0.0015	-0.0008	0.0012
doctor visits*male	-0.1102	0.0193	-0.1383	0.0251	0.0408	0.0203
hospital days	-5.6667	0.3694	-5.5670	0.3703	-5.4660	0.4032
hospital days (squared)	0.8668	0.0922	0.8625	0.0941	0.8503	0.0986
hospital days*age	0.0103	0.0034	0.0094	0.0039	0.0082	0.0034
hospital days*male	0.1992	0.0591	0.1957	0.0648	0.1998	0.0586
any Rx consumption	-3.7111	0.3007	-5.4695	0.4523	-5.8783	0.5979
any Rx consumption*age	-0.0625	0.0079	-0.0491	0.0104	0.0258	0.0068
any Rx consumption*male	-1.7109	0.1167	-1.7069	0.1565	-0.3121	0.1361
any consumption	-1.3351	0.0666	-1.3298	0.0720	-2.1247	0.1151
any consumption*male	-0.4851	0.0600	-0.5293	0.0628	-0.4811	0.0753
negative consumption†	0.1000		0.1000		0.1000	
discount factor†	0.9960		0.9960		0.9960	

† These parameters are not estimated. Note that β is set to 0.996, instead of the traditional 0.95, because decisions are made monthly in the model. $0.996^{12} \approx 0.95$. The parameter ω_{70} measures the disutility of each dollar of negative consumption (i.e., outspending monthly income, requiring an individual to use savings or to borrow) and is not estimated due to weak identification. Other model parameters do not seem to be sensitive to different fixed values for this parameter.

Table A29: WYDM Model: Acute Illness Parameters with Various Mass Points

	1 Mass Point		2 Mass Points		4 Mass Points	
	Est.	S.E.	Est.	S.E.	Est.	S.E.
constant	-1.6263	0.1082	-1.9042	0.1088	-2.1088	0.0940
male	-0.3808	0.0377	-0.3550	0.0415	-0.3296	0.0374
non-white (black or Hispanic)	-0.0712	0.0358	-0.0960	0.0396	-0.0824	0.0373
education (highest grade completed)	-0.0184	0.0056	-0.0079	0.0059	0.0016	0.0053
age	-0.0032	0.0019	-0.0008	0.0020	-0.0030	0.0018
lives in a MSA	-0.0111	0.0477	-0.0102	0.0488	-0.0016	0.0465
income (in 1996 dollars)	0.0055	0.0010	0.0051	0.0149	0.0045	0.0011
initial health	-0.0049	0.0168	0.0150	0.0172	0.0081	0.0156
acute illness	3.4285	0.0473	3.3408	0.0501	3.2414	0.0560
acute illness (squared)	-0.2528	0.0118	-0.2309	0.0140	-0.1965	0.0160
consecutive mons. with acute illness†	-0.0149	0.0128	0.0197	0.0143	0.0179	0.0131
consecutive mons.* initial month	0.0420	0.0120	0.0224	0.0140	0.0115	0.0122
chronic illness	0.3254	0.0328	0.2467	0.0335	0.1515	0.0374
chronic illness (squared)	-0.0464	0.0064	-0.0366	0.0061	-0.0201	0.0088
doctor visits	-0.0510	0.0064	-0.0516	0.0070	-0.0481	0.0071
doctor visits (squared)	0.0067	0.0009	0.0065	0.0010	0.0053	0.0009
hospital days	-0.1090	0.0181	-0.1063	0.0187	-0.0946	0.0181
hospital days (squared)	0.0221	0.0039	0.0219	0.0043	0.0193	0.0040
Rx consumption	-0.1048	0.0130	-0.1389	0.0182	-0.1475	0.0212
acute illness*doctor visits	-0.0067	0.0010	-0.0063	0.0010	-0.0058	0.0010
acute illness*hospital days	-0.0012	0.0026	-0.0016	0.0026	-0.0016	0.0021
acute illness*Rx consumption	-0.0027	0.0031	0.0032	0.0041	-0.0041	0.0040
chronic illness*doctor visits	0.0012	0.0006	0.0028	0.0007	0.0013	0.0006
chronic illness*hospital days	-0.0035	0.0014	-0.0033	0.0015	-0.0017	0.0014
chronic illness*Rx consumption	0.0252	0.0047	0.0154	0.0049	0.0063	0.0036
threshold 2	3.1819	0.0404	3.2651	0.0411	3.2557	0.0415
threshold 3	5.9939	0.0696	5.9962	0.1287	6.0154	0.0727
threshold 4	8.1308	0.1244	8.1858	0.1308	8.2134	0.1261

Month indicators are included in the regression but are not reported here.

† I control for the number of consecutive months with *any* acute illness leading up to the current month to capture the impact of a more severe acute illnesses on illness transition probabilities. Due to the nature of the data, this count begins in the first month of the insurance year, meaning an individual entering the third month of the year with two consecutive acute illness months may actually have had an acute illness even longer. To control for this measurement error, I interact consecutive months with an acute illness with an indicator that equals one if an individual has had an acute illness in every month since the first month of the year.

Table A30: WYDM Model: Chronic Illness Parameters with Various Mass Points

	1 Mass Point		2 Mass Points		4 Mass Points	
	Est.	S.E.	Est.	S.E.	Est.	S.E.
constant	-5.1434	0.0912	-5.5784	0.1308	-5.0130	0.1723
male	0.2663	0.0171	0.2458	0.0255	-0.0182	0.0281
non-white (black or Hispanic)	-0.0157	0.0044	-0.0211	0.0081	-0.0173	0.0130
education (highest grade completed)	0.0077	0.0011	0.0315	0.0037	0.0188	0.0044
age	0.0390	0.0020	0.0376	0.0028	0.0047	0.0044
lives in a MSA	0.0352	0.0059	0.0293	0.0101	-0.0326	0.0177
income (in 1996 dollars)	0.0002	0.0001	-0.0002	0.0002	0.0002	0.0004
initial health	0.0148	0.0022	0.0386	0.0052	0.0546	0.0114
acute illness	0.0447	0.0146	0.0502	0.0205	-0.2161	0.0445
acute illness (squared)	-0.0023	0.0028	-0.0010	0.0002	0.0561	0.0112
consecutive mons. with acute illness	-0.0209	0.0031	-0.0163	0.0037	-0.0143	0.0067
consecutive mons.*initial period	0.0168	0.0023	0.0018	0.0023	-0.0002	0.0044
chronic illness	0.3689	0.0213	0.3014	0.0258	0.3153	0.0444
chronic illness (squared)	-0.0243	0.0029	-0.0176	0.0029	-0.0137	0.0062
doctor visits	-0.0729	0.0089	-0.0594	0.0093	-0.0346	0.0076
doctor visits (squared)	0.0058	0.0009	0.0042	0.0009	0.0028	0.0007
hospital days	-0.0283	0.0190	-0.0221	0.0137	-0.0392	0.0129
hospital days (squared)	0.0064	0.0044	0.0042	0.0031	0.0072	0.0031
Rx consumption	-0.4607	0.0467	-0.3065	0.0418	-0.1869	0.0378
threshold 2	2.5771	0.2405	2.6267	0.2451	2.8027	0.2648

Month indicators are included in the regression but are not reported here.

Table A31: Unobserved Heterogeneity Parameter Estimates - 2 Types

	Type 1	Type 2	
		Est.	S.E.
Mass Point Location			
acute illness prob.	0.0000	0.2556	0.0472
chronic illness prob.	0.0000	0.1623	0.0286
doctor visit pref. (linear)	0.0000	-0.3152	0.0505
doctor visit pref. (squared)	0.0000	0.0396	0.0067
hospital day pref. (linear)	0.0000	-0.4414	0.1321
hospital day pref. (squared)	0.0000	0.0661	0.0374
any Rx pref.	0.0000	2.5870	0.1760
any consumption pref.	0.0000	0.3287	0.1212
doctor visit price dist.	0.0000	4.4706	1.2863
hospital day price dist.	0.0000	-0.0423	0.1072
monthly Rx price dist.	0.0000	0.9271	0.0172
initial acute illness prob.	0.0000	0.8501	0.1536
initial chronic illness prob.	0.0000	1.5867	0.1646
initial health status prob.	0.0000	0.1500	0.1376
Type Probabilities†			
parameter estimate	0.0000	-0.7510	0.0756
type probability	67.9	32.1	

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

Table A32: Unobserved Heterogeneity Parameter Estimates - 4 Types

	Type 1	Type 2		Type 3		Type 4	
		Est.	S.E.	Est.	S.E.	Est.	S.E.
Mass Point Location							
acute illness prob.	0.0	0.569	0.059	0.556	0.066	0.573	0.064
chronic illness prob.	0.0	0.985	0.099	0.978	0.101	1.109	0.106
doctor visit pref. (linear)	0.0	-1.599	0.144	-1.636	0.150	-1.336	0.169
doctor visit pref. (squared)	0.0	0.423	0.025	0.397	0.028	0.412	0.025
hospital day pref. (linear)	0.0	-1.536	0.274	-1.543	0.297	-1.352	0.323
hospital day pref. (squared)	0.0	0.243	0.082	0.230	0.097	0.282	0.086
any Rx pref.	0.0	0.788	0.549	0.759	0.563	-3.788	0.652
any consumption pref.	0.0	1.213	0.227	0.832	0.275	1.631	0.146
doctor visit price dist.	0.0	-3.161	2.402	-2.816	2.729	-13.005	2.330
hospital day price dist.	0.0	-0.120	0.198	-0.110	0.213	-0.445	0.180
monthly Rx price dist.	0.0	0.770	0.032	-0.747	0.033	-0.090	0.040
initial acute illness prob.	0.0	1.360	0.195	1.155	0.234	0.899	0.232
initial chronic illness prob.	0.0	2.300	0.222	2.102	0.254	1.460	0.251
initial health status prob.	0.0	0.387	0.173	0.385	0.208	0.192	0.212
Type Probabilities†							
parameter estimate	0.0	-0.961	0.094	-1.363	0.106	-1.129	0.122
type probability	51.0	19.5		13.0		16.5	

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

Table A33: Observed and Simulated Annual Consumption

	Doctor visits				Hospital days				Prescriptions			
	Obs.	1 Type	2 Types	4 Types	Obs.	1 Type	2 Types	4 Types	Obs.	1 Type	2 Types	4 Types
0	28.2	17.2	19.7	24.2	75.8	71.4	67.0	71.93	35.5	11.8	24.7	31.14
1	16.2	16.4	15.4	16.0	12.1	17.3	18.5	15.43	13.3	17.0	15.0	16.22
2	12.7	13.1	11.8	11.5	4.8	5.0	5.0	4.75	5.6	11.8	11.6	7.95
3-5	18.4	23.9	23.1	19.3	4.8	5.0	6.0	5.88	10.5	26.7	15.5	9.67
6-9	10.4	21.0	16.2	12.3	1.5	1.1	0.4	1.60	7.6	21.7	15.0	10.47
10+	14.1	13.4	13.8	16.7	1.0	0.2	0.1	0.41	27.5	12.1	20.2	24.55

Notes: The table reports annual consumption levels for the observed and simulated samples. The values are the percentage of the samples consuming at each level. For example, the top left entry states that 28.2% 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 A5: Annual Medical Care Expenditure - 1, 2, and 4 Mass Points

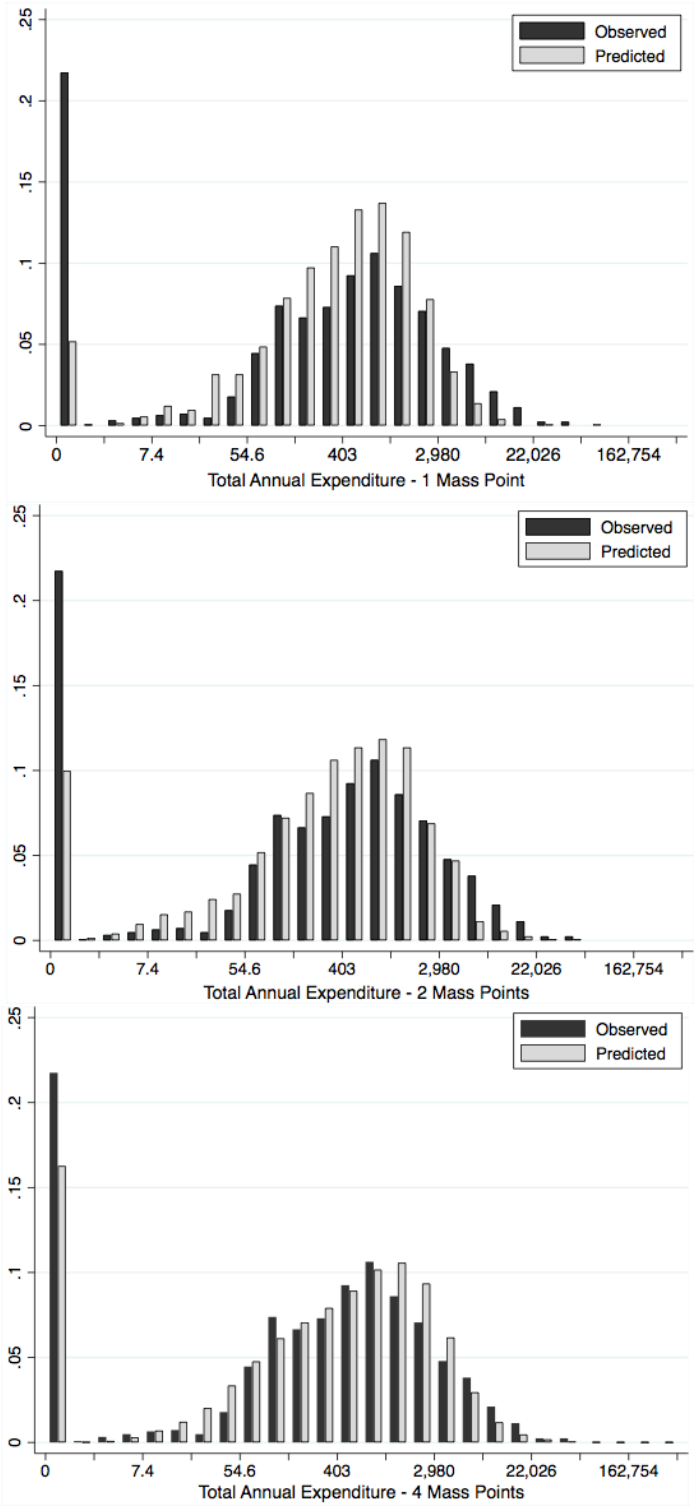

Notes: These figure compare observed and simulated annual expenditure distributions. The horizontal axis is log-scaled, so that markers are spaced $exp(1.0)$ apart. The (top, middle, bottom) figure reports simulated expenditure from a model with (1, 2, 4) unobserved mass points

Table A34: Predicted Increase in Annual Expenditure

Percentile	1 Mass Point	2 Mass Points	4 Mass Points
5	0	0	0
10	3	0	0
15	12	3	0
20	23	11	0
25	35	20	0
30	50	31	0
35	68	46	0
40	90	65	0
45	118	90	0
50	154	121	0
55	198	160	17
60	254	212	83
65	326	275	164
70	418	361	285
75	538	476	462
80	702	642	722
85	948	896	1,143
90	1,381	1,358	1,904
95	2,451	2,502	3,651
99	7,666	8,146	10,906
Mean	658	661	772
Mean Uninsured Exp.	488	511	663
Mean Insured Exp.	1,146	1,172	1,435

Notes: This table displays various percentiles of the distribution of the predicted increase in mean annual medical care that results from insurance acquisition. All three columns compare a counterfactual where an individual is forced into his most preferred plan among those offered to him by his employer to a counterfactual where individuals are forced to be uninsured.

12 Annual Expenditure Model

This section describes the parameter estimates and model fit analysis of the representative annual expenditure (RAE) model.[52] 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

[52]The RAE model is constructed by imposing several assumptions on the WYDM model. One such assumption is that an individual knows all future price and preference shocks, which significantly complicates the solution to the DPP as the state space in any given month contains these future shocks. To simplify this solution, I assume that each individual receives only one price draw for each type of medical care and one preference draw for each medical care consumption bundle, rather than a new set of draws in each month of the year. These draws then become regressors in the interpolation procedure described in Section 5.1 of the paper, so that in any month, t, the maximal expected future value of consumption in month, $t+1$ (see Equation 8 of the paper), is a function of *known* future shocks.

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.

12.1 Parameter Estimates

Table A35 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. Arcidiacono et al. (2007) describe a similar tension between forward-looking and myopic versions of their dynamic model of smoking and drinking decisions.

The parameters in Table A36 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 data a bit better and estimation was significantly faster than in the within-year decision-making model.

Price parameters can be observed in Table A37. 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 A35: 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 A36: 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 A37: 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.

12.2 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 A38), 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 A39). 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 A38: 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 A39: 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 A6: Annual Medical Care Expenditure

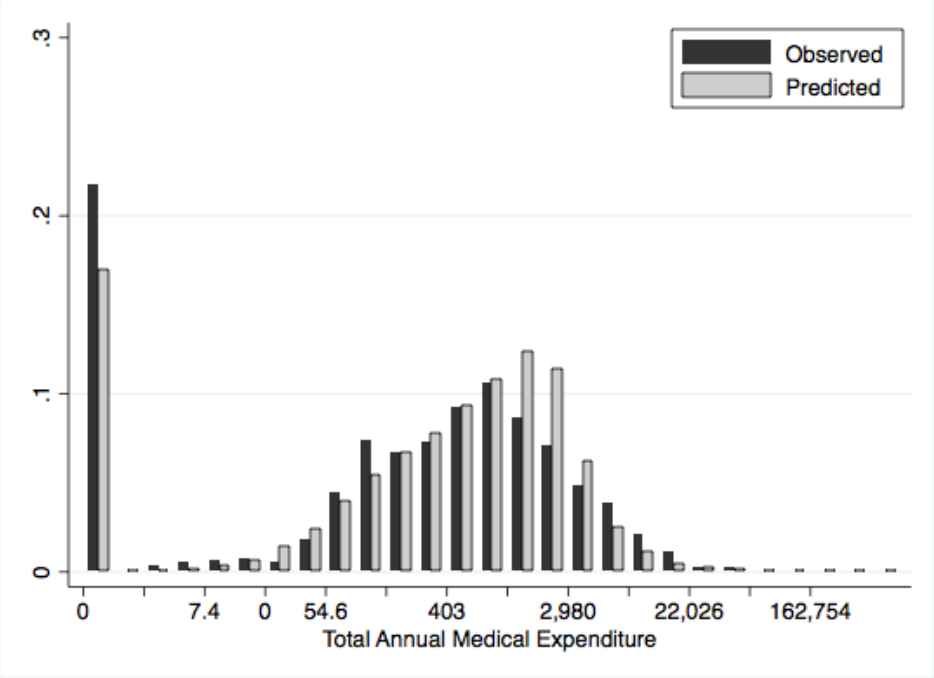

References

Abaluck, J., J. Gruber, and A. Swanson (2015): "Prescription drug utilization under Medicare Part D: A dynamic perspective," *NBER Working Paper*, 20976.

Arcidiacono, P., H. Sieg, and F. Sloan (2007): "Living rationally under the volcano? An empirical analysis of heavy drinking and smoking," *International Economic Review*, 48, 37–65.

Aron-Dine, A., L. Einav, and A. Finkelstein (2013): "The RAND Health Insurance Experiment, Three Decades Later," *Journal of Economic Perspectives*, 27, 197–222.

Aron-Dine, A., L. Einav, A. Finkelstein, and M. Cullen (2015): "Moral Hazard in Health Insurance: Do Dynamic Incentives Matter?" *Review of Economics and Statistics*, 97, 725–741.

Baicker, K., S. Taubman, H. Allen, M. Bernstein, J. Gruber, J. P. Newhouse, E. Schneider, B. Wright, A. Zaslavsky, A. Finkelstein, and the Oregon Health Study Group (2013): "The Oregon Experiment - Medicaid's Effects on Clinical Outcomes," *New England Journal of Medicine*, 368.

Bajari, P., H. Hong, A. Khwaja, and C. Marsh (2014): "Moral Hazard, Adverse Selection, and Health Expenditures: A Semiparametric Analysis," *RAND Journal of Economics*, 45, 747–763.

Berndt, E., B. Hall, R. Hall, and J. Hausman (1974): "Estimation and Inherence in Nonlinear Structural Models," *Annals of Economic and Social Measurement*, 3, 103–116.

Dalton, C. M., G. Gowrisankaran, and R. Town (2017): "Salience, Myopia, and Complex Dynamic Incentives: Evidence from Medicare Part D," *NBER Working Paper*, 21104.

Einav, L., A. Finkelstein, and P. Schrimpf (2015): "The response of drug expenditure to nonlinear contract design: evidence from medicare part D," *The Quarterly Journal of Economics*, 130, 841–899.

Ellis, R. P. (1986): "Rational Behavior in the Presence of Coverage Ceilings and Deductibles," *RAND Journal of Economics*, 17, 158–175.

Finkelstein, A., S. Taubman, B. Wright, M. Bernstein, J. Gruber, J. Newhouse, H. Allen, K. Baicker, , and the Oregon Health Study Group (2012): "The Oregon Health Insurance Experiment: Evidence from the First Year," *The Quarterly Journal of Economics*, 127.

Greene, W. H. (2012): *Econometric analysis*, Pearson Education Limited.

Grossman, M. (1972): "On the Concept of Health Capital and the Demand for Health," *Journal of Political Economy*, 80, 223–255.

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

Keeler, E. B., J. P. Newhouse, and C. E. Phelps (1977): "Deductibles and the Demand for Medical Care Services: The Theory of a Consumer Facing a Variable Price Schedule under Uncertainty," *Econometrica*, 45, 641–656.

Manning, W. G., J. P. Newhouse, N. Duan, E. B. Keeler, and A. Leibowitz (1987): "Health Insurance and the Demand for Medical Care: Evidence from a Randomized Experiment," *The American Economic Review*, 77, 251–277.

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

Rust, J. (1987): "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," *Econometrica*, 55, pp. 999–1033.

Schennach, S. M. (2004): "Estimation of nonlinear models with measurement error," *Econometrica*, 72, 33–75.

Zweifel, P. and W. G. Manning (2000): "Moral hazard and consumer incentives in health care," *Handbook of Health Economics*, 1, 409–459.