Using Participant Behavior to Measure the Value of Social Programs: The Case of Medicaid*

Trevor S. Gallen†

Winter 2015

Abstract

Social programs frequently have two effects on labor supply: an income effect and a wage effect. Programs produce a wage effect by linking benefits or program premiums to income, generating an implicit marginal tax rate on labor. Programs produce an income effect through the actual cash or in-kind transfer they provide. Conditional on wage and non-wage income elasticities, labor market responses to program structure (or the lack thereof) reveal the value of program participation to beneficiaries. I study a public policy change in Tennessee that disenrolled 12% of its Medicaid population, and use simple calculations to estimate a cash value to beneficiaries of $0.26 cents per dollar spent. Using this same policy change, I estimate a richer model that allows for heterogeneity in family structure, wages, property income, and preferences over healthcare types. This method yields a value of Medicaid of $0.35 per dollar of spending. I find a high variance in the implied distribution of Medicaid's value to beneficiaries, but this high variance can be almost fully explained by the large variation in Medicaid expenditures across recipients.

JEL Classification: H44, H51, I13

---

* I thank Casey Mulligan for numerous helpful conversations and comments. I also thank Gary Becker, John Cochrane, Bruce Meyer, Steven Davis, Clifford Winston, Matt Notowidigdo, Emmanuel Saez, Harald Uhlig, Ben Handel, Yana Gallen, Kris Hult, Ian Fillmore, and Xan Vongsathorn for insightful comments and suggestions. I am grateful to seminar participants at the University of Wisconsin Whitewater and the University of Chicago's Money and Banking Workshop, as well as the Economic Dynamics and Applied Macroeconomics Working Groups at Chicago for useful comments. I acknowledge the University of Chicago Research Computing Center and the Becker Friedman Institute for support of this work.

† Correspondence: University of Chicago Department of Economics, 5757 South University Avenue, Chicago, IL 60637. Email: tgallen@uchicago.edu. Web: home.uchicago.edu/tgallen/

Introduction

Public welfare is a large part of the U.S. economy; in 2013, government transfers comprised more than 17% of personal income. For some social programs with cash benefits and high take-up like Social Security, frictions play a small role, but for programs such as unemployment insurance (Anderson and Meyer, 1997), workers compensation (Ehrenberg, 1988), and Medicaid (Moffitt, 1983), take-up is low and implied frictions are large. This low take-up might be driven by many factors, such as information costs, time costs, or a low value of benefits to recipients. I examine the net value Medicaid's in-kind benefits provide to recipients as revealed by labor-market behavior and take-up behavior.

I identify the value of Medicaid in Tennessee to current beneficiaries using two sources of information. First, Tennessee displays cross-sectional differences in both prices and participation. The degree to which higher prices decrease participation reveals how much people value Medicaid. Second, a policy change in 2005 caused a large disenrollment in Medicaid. Disenrollment removed a large implicit marginal tax rate in the form of Medicaid premiums, resulting in a higher take-home wage, and removed an in-kind transfer, reducing income. However, losing Medicaid makes famiiles feel poorer only insofar as they value Medicaid. Because labor responses to losing Medicaid were small when compared to the counterfactual experience of losing cash income equal to the value of Medicaid expenditures, I conclude that the cash-equivalent value of Medicaid is small.

Tennessee's Medicaid program, called TennCare, provides both of these sources of information through its premium structure and a large disenrollment in the mid-2000's. TennCare was unusual because it was open to adults and children of all income levels and family structure. Rather than exclude enrollees based strictly on income, it instead charged higher premiums for higher-income families. This differs from Medicaid's traditional "notch," which meant that households were disenrolled completely as soon as they passed a specific threshold. However, due to budgetary pressures (Bennett, 2014), Tennessee disenrolled nearly 12% of its Medicaid population in 2005, and in response, labor supply rose (Garthwaite et al., 2013).

Two classic effects explain this change in labor following disenrollment: an income effect and

a wage effect. As households are disenrolled, income is reduced via withdrawal of transfers, households feel poorer and increase labor supply. The wage effect also plays a role in explaining changes to labor supply. Tennessee's relatively open program and premium schedule created a tax on labor markets; for poorer households, access to TennCare had a guaranteed zero out-of-pocket cost, while families above 400% of the poverty line paid more than $10,000 per year in premiums alone. With the 2005 disenrollment, the average family enrolled in Medicaid in 2004 experienced an after-tax wage increase of nearly 6%. The impact this wage increase has on labor depends on the labor supply function.

Considering both income and wage effects, the observable effect of Tennessee disenrollment on labor supply is smaller than expected were Medicaid valued dollar for dollar. Given standard labor-supply elasticities, I find that losing access to Medicaid is equivalent to losing approximately $1266 dollars. Combining this number with federal and state expenditures on Medicaid, the value of Medicaid for an average beneficiary is $0.26 per dollar of spending.

However, averages hide a great deal of heterogeneity in the valuation of Medicaid coverage. To uncover this heterogeneity, I use both cross-sectional variation in the prices that households paid for TennCare and labor choices resulting from disenrollment to estimate the distribution of the value of TennCare. Behavior after disenrollment helps estimate the distribution of the value of Medicaid. Holding constant the effect a change in wages has on labor supply, when a disenrolled household increases labor hours and enrolls in employer-sponsored insurance, it is evidence it valued Medicaid. Similarly, if a disenrolled household does not increase labor hours and remains uninsured, it is evidence it did not value Medicaid. Cross-sectional variation in prices also provides information regarding the distribution of household Medicaid valuations. Households that paid high premiums and still chose to enroll in TennCare reveal a high value of the program.

I find that variation in Medicaid expenditures looks very similar to variation in household values of Medicaid, suggesting a role for homogeneity in valuation per dollar spent, with variation in expenditures taking center stage. Individuals enroll in the same insurance plans, but Medicaid does not spend the same amount on them due to factors such as health status. Although households vary greatly in their valuations, estimated Medicaid spending is extremely skewed,

with the highest 20% of households receiving 81% of spending. Allowing Medicaid expenditures to explain variation in revealed cash values, I find only moderately heterogeneous values of a dollar of Medicaid spending across household percentiles, with a low of $0.18 to a high of $0.43.

In my model, households value consumption, access to healthcare, and leisure. Medicaid has partial substitutes in employer-sponsored insurance (ESI), which agents receive through work and non-group insurance (NGI), and purchase at market prices. The relative utility weight on each type of health insurance differs across households. If the relative weight that agents place on Medicaid is very high, their cash-equivalent value of Medicaid is high. I draw household composition, wages, and property income from the Survey of Income and Program Participation (SIPP) 2004-2008 Tennessee sample. I simulate household's choices (i.e., how much to consume, how much to work, and what type of health insurance to purchase, if any) for a given distribution of preferences over leisure and health insurance. Using the simulated method of moments, I find the distribution of preferences over leisure and health insurance that best matches SIPP and administrative data on labor and health insurance choices by income. I simulate the same households in both 2005 and 2006, where in 2006, many non-poor households no longer had access to TennCare and many purchased ESI or NGI if they wanted health insurance. I rely on both cross-sectional variation in healthcare decisions in the face of different premiums, and panel variation in labor market responses to being disenrolled to estimate my model.

At its core, this paper flips the "Medicaid notch" literature on its head. This literature analyzes the impact of Medicaid's high implicit marginal tax rates on labor supply. Given an elasticity of labor supply, if households value Medicaid, they change their labor supply to enroll in TennCare. The degree to which they do not reduce labor is a measure of how little they value Medicaid. Along with estimating the value of TennCare, I reinterpret results from Dague (2014) and Meyer and Rosenbaum (2001), finding implicit valuations of Medicaid comparable to mine. My cursory estimate of $0.26 per dollar of spending, and heterogeneous agent estimate of $0.35 per dollar, are similar to a $0.28 estimate I extract from Meyer and Rosenbaum (2001), below the $0.48 to $0.68 estimate from Keane and Moffitt (1996) and above the implicit $0.09/dollar spent I extract from in Dague (2014) using a take-up assumption. Further discussion of the implicit valuations I generate are provided below.

On reflection, low revealed valuations of Medicaid in comparison to cost should be obvious for several reasons. First, administrative costs are likely to eat up a significant portion of expenditure. For instance, Mulligan and Philipson (2003) find that for life insurance companies, for every $1 paid in premiums, beneficiaries received $0.70-$0.75 in insurance payments. Second, the alternative to not being enrolled in Medicaid is not drastic for those Cutler and Gruber (1996) describe as "conditionally covered," those who are not enrolled in Medicaid until they become ill or pregnant. Third, the alternative to non-insurance in dire situations is not non-treatment, but medical bankruptcy. For example, Finkelstein and McKnight (2005) find Medicare simply shifted the costs of life-saving treatments from individuals to the government, but did not change whether agents received treatment. Corroborating this evidence, Garthwaite et al. (2014) find hospitals benefit from Medicaid, as it crowds out patients defaulting on medical debt and charitable care provided by the hospital. Finally, in any public program with broad categorical eligibility there will be participants who qualify and enroll but do not value the benefits, bringing down average valuations.

While my paper is agnostic about the normative implications of low-beneficiary valuations for Medicaid, they remain important. Medicaid spending accounts for a sizable portion of transfers to the poor. How Medicaid is valued affects the number of people with consumption below the poverty line. For instance, Meyer and Sullivan (2013) calculate the consumption poverty rate using the fungible value of Medicaid, discussed below. When I run a comparable poverty-rate exercise, instead comparing Medicaid transfers valued 1-to-1 to Medicaid transfers that are valued at $0.30 per dollar, the population lifted out of poverty by Medicaid alone falls by 14 million. How Medicaid is valued is very important, as it makes up a large fraction of consumption expenditures for the poor.

This paper does not answer why Medicaid gives in-kind transfers valued at a fraction of what it spends; many reasons are possible, and a significant literature is devoted to them. One possibility is pure inefficiency. Governments provision goods the private market could better produce, as in Antarctic expeditions discussed by Karpoff (2001). For skeptics, Mulligan and Philipson (2000) may provide a more persuasive answer; those redistributing (the rich, in Mulligan and

Philipson) value some consumption by the poor more than other types of consumption. When government preferences are weighted properly, spending on things the poor do not value but the wealthy do (for the poor) might be efficient. Finally, forced excessive spending on a good might be a transfer to a politically powerful group such as TennCare's Managed Care Organizations, doctors, the pharmaceutical industry, the medical equipment industry, or hospitals, as in Garthwaite et al. (2014). There are many possibilities, not all of which point to inefficiency in a constrained world. Indeed, a pure loss of $0.70 per dollar spent is less likely given the size of the program and the fact that publicly provided Medical care is an outcome reached across many advanced countries. But even in the event of efficiency, my results raise an interesting unanswered question. Since Medicaid's inception, the United States has spent more than $8 trillion on benefits: what party benefitted from the residual (unvalued by recipients) $4.5 to $6 trillion dollars, if not recipients?

Relevant literature

Historically, the Census Bureau has used several methods to value Medicaid. The primary two, discussed by Chiswick (1985), are market-based and recipient-valuation approaches. The market-based approach attempts to find payments made to the same vendors by both Medicaid and market-priced plans to find a comparison base for similar procedures. The recipient-valuation approach looks at overall household spending on a given good in the private market, and compares spending by unsubsidized and subsidized households. The latter is commonly used today as the "fungible value of Medicaid." This approach values Medicaid only if it frees resources. Risk-adjusted Medicaid transfers per enrollee are counted as value only if a recipient's income (without Medicaid) was sufficient to pay for basic food and housing National Research Council (1995) Congressional Budget Office (2010). If an unsubsidized household would not have purchased health insurance, it is attributed zero value. Otherwise, it is typically attributed at average cost, sometimes capped by market value. Since 2012, the CBO simply used average government expenditures of Medicaid. Neither method is satisfactory in understanding the net benefits of Medicaid, with both coming entirely from an expenditures perspective, rather than using revealed preferences.

Literature on Medicaid is large and segmented. Most focuses on estimating whether the large, implicit marginal tax rates Medicaid imposes causes labor-supply shifts. Until 1987, Medicaid and cash welfare had the same cutoff notch, with some households losing a portion of their health insurance and cash welfare if they earned too much. This change produced a cottage industry of papers. Chief among them, Yelowitz (1995) uses a difference-in-difference model to analyze the Medicaid notch around which implicit marginal tax rates are significant. He finds that the Medicaid eligibility notch's level at 80% of its reformed value suppressed labor force participation by 3.3%. Meyer and Rosenbaum (2001), analyzing the same notch in conjunction with changes to the Earned Income Tax Credit and other parts of the tax code, find no impact of the Medicaid notch. Comparing Medicaid's estimated impact on work with an income tax changes, a value of $0.28 per dollar was obtained by Meyer and Rosenbaum.[1]. In my paper, this lack of labor response indicates a large number of low valuations.

Three recent papers are salient in the analysis of the effects of Medicaid expansions and contractions. Garthwaite et al. (2013) examine the same natural experiment this paper does. They examine variation generated by Tennessee's rapid withdrawal of Medicaid coverage from 170,000 low-income individuals, mostly childless adults. While the authors main focus interprets labor supply increases through the lens of "job lock," they also find labor-supply distortions consistent with a classic model in which income and substitution effects of Medicaid reduce labor supply. The find withdrawal of Medicaid increased labor supplied by approximately 6.5% among childless adults. While Garthwaite et al. examine disenrollment, Dague et al. (2014) explore the effects of Medicaid expansion using a natural experiment generated by Wisconsin's expansion of public insurance to childless adults. Using both regression discontinuity and propensity-score matching, they find a reduction in labor supply between 0.9% and 10.6%. The paper's broad set of estimates is generated from the inclusion of covariates of age, sex, employment during the prior quarter, and earnings during the prior quarter, and by changing the date of discontinuity. The most compelling regression discontinuity numbers, including covariates and excluding

[1]This estimate is generated using Table IV in Meyer and Rosenbaum (2001) I take the estimated average impact of a $1000 increase in Medicaid spending on employment probability (-0.0096) and divide it by the estimated average impact of a $1000 increase in AFDC and food stamps benefit (conditional on non-work) (-0.0340). Because Meyer and Rosenbaum's regression relates both Medicaid spending and income tax spending to the probability of employment, this estimate also values Medicaid through its impact on work decisions. This method assumes $1000 in welfare benefits paid are equivalent to $1000 in cash: because welfare benefits include food stamps, which may be valued at less than a dollar, it is likely this is a high estimate of the cash-Medicaid exchange rate.

applications from an extension of the application deadline, imply a reduction of 6.32% in the employment rate.

Baicker et al. (2014) study a large, randomized experiment in Oregon. Using administrative data, the paper analyzes results of a randomized experiment, expanding Medicaid coverage to uninsured, low-income adults in 2008. Since many chosen by the experiment to receive Medicaid never followed through, the authors offer two interesting analyses: an intent to treat those who won Oregon's Medicaid lottery, and treatment on the treated analysis on those who receive Medicaid. Their point estimates suggest those who won the Medicaid lottery were only 0.4% less likely to report earnings (interpreted as any work), and those who won the lottery and received Medicaid were 1.6% less likely. However, the confidence intervals on both were wide, with a -1.1% to 0.4% rise for treatment on the treated, and a -4.3% to 1% rise for intent to treat.

The impacts of Medicaid on labor supply have generally been modestly negative, and concentrate on those in poor health (Gruber, 2003)—increasing Medicaid would decrease labor supply among the poor through income and substitution effects away from work. In a second-best world, however, expansions of Medicaid have the capacity to increase labor supply. One explanation, separate from Yelowitz's classic treatment, has been "welfare lock," in which individuals who would otherwise work more in the low-wage labor market do not because they would be unable to find amenable healthcare in the non-group or ESI markets. Removal of large implicit marginal tax rates on some groups via expansion could actually increase labor.

Potential workers in a welfare lock would avoid finding work to keep their healthcare. The opposite might occur when employers are the major source of health insurance; "job lock" occurs when workers keep their jobs only because they are afraid of losing health insurance benefits. This might happen, for example, when a worker would optimally shift to entrepreneurship, but does not in order to keep her ESI. Using the Current Population Survey, Fairlie et al. (2011) examine the impact of healthcare demand on entrepreneurship. Using a difference-in-difference model, examining observations on whether a worker's spouse has ESI, and the Medicare healthcare-availability discontinuity for those without spousal ESI for those aged 65 and older, they find significant "entrepreneurship lock," reducing entrepreneurship among men by about 1%, from

a baseline annual rate of 3%. Another discussion of job lock appears in Gathwaite, Gross, and Nodowidigdo (2013), who argue that the causes of labor changes and increases in ESI are due to job lock. It is difficult to parse job lock and classic substitution between labor and leisure in the TennCare experiment, and I focus on a classical explanation. However, studying job lock among workers with a sudden demand for healthcare shock, Bradley et al. (2013) find that women with breast cancer who depended on own-employment ESI reduced employment by between 8% and 11% less than those who have other heath insurance options while undergoing treatment.

A large body of literature examines the impact of health insurance on labor supplies of various subgroups, much of it summarized by Gruber and Madrian (2002). Gruber and Madrian offer three lessons from the literature: first, health insurance is important for the decision to retire and change jobs. Second, it affects labor supply decisions of secondary earners strongly. Finally, health insurance is not important for labor-market decisions of low-income mothers. The authors argue that overall, the effects of job lock are modest. This observation is key to my paper; the failure of Medicaid to have large labor-market effects reveals low valuations of Medicaid.

Concerning the value of another government-sponsored healthcare system, Finkelstein and McKnight (2005) find that Medicare reduces elderly exposure to out-of-pocket-cost risk substantially. They argue that the compensating variation from the reduction in risk alone is valued between 50% and 75% of Medicare's cost estimates. Since health outcomes did not seem to change, it appears that life-saving treatments were not foregone in the absence of Medicare. Instead, the cost was shifted from individuals to the government. This interpretation will be important to my paper. Medicaid may be crowding out private charity (as in Garthwaite et al. (2014)) or household expenditures (as Medicare does in Finkelstein and McKnight). I do not model ex post and ex ante insurance benefits, I do observe implicit valuations of entire Medicaid contracts, including high valuations that might be due to high transfers, or high value of consumption smoothing.

Similar to this paper, Keane and Moffitt (1996) use the simulated method of moments to produce a structural model of labor supply in a program-participation context. However, they study female head-of-household labor supply in the context of Aid for Dependent Families with Children (AFDC), food stamps, and subsidized housing choices. Examining the large, implicit

marginal tax rates produced by these programs, Keane and Moffitt (1996) find inelastic responses to changes in marginal tax rates, even when wage elasticities are high. The authors do not model Medicaid as a choice. Instead, Medicaid passively conferred benefits when AFDC was chosen. Like the representative household model in this paper, the authors allow for a cash value of Medicaid, unlocked when a household participates in AFDC. Keane and Moffitt estimate a value of $0.48 and $0.64 in return for $1 in Medicaid spending. In contrast to Keane and Moffitt, I focus directly on Medicaid, the decision between various types of health insurance, and many different types of households. I use a natural experiment and premiums to provide variation in valuation, allow for alternative health insurance sources, examine joint household labor-market decisions rather than single choices, and most importantly, estimate and examine the distribution of valuations for Medicaid coverage. I find lower values, around 47% to 63% of Keane and Moffitt's estimates.

Dague (2014), which this paper examines further, analyzes Wisconsin's Medicaid program BadgerCare, which like TennCare, places structure on premiums that enrollees face. BadgerCare features strong increases in monthly premiums for adults and children, going from $0 at 150% of the FPL for both parents and children to $201 before parents are unable to be on BadgerCare after 200% FPL. Child premiums remain $0 until 200% of the FPL, after which they rise to $90.74 per month above 300% of the FPL. Both adults and children display sharp breaks in enrollment when they pass premium thresholds. Dague finds that going from $0 to $10 monthly premium for adults causes a robust drop in the length of enrollment by 1.35 months; with average pooled enrollment months of 10.75, this represents a 13% drop in Medicaid enrollment months due to a $10 monthly fee. When combined with combined with typical estimates of take-up, Dague's implicit distributional results yield an extremely low value for marginal entrants.

TennCare

Medicaid has several "mandatory" populations that states must cover, such as children younger than six below 133% of the FPL, older than six under 100% of the FPL, parents below historic AFDC cutoffs, pregnant women below 133% of the FPL, among others. There are also "optional" populations a state can choose to cover, such as children above 100% of the FPL

or pregnant women above 133% of the FPL. In the baseline reference period of 2005 to 2006, Tennessee covered all mandatory populations along with optional categories such as pregnant women up to 185% of the FPL, individuals dually eligible for Medicaid and Medicare, and several child FPL expansions. Finally, Medicaid has "expansion" or "demonstration" populations. These are covered primarily by Section 1115 Waivers, in which the Federal government allows states to experiment with Medicaid services, eligibility, and delivery systems in a budget-neutral way regarding expected federal contributions. Tennessee's Section 1115 waiver allowed it to expand Medicaid, primarily to the "uninsured," the "uninsurable," and the "medically needy." It also allowed additional, typically elderly, households to be eligible for joint coverage under TennCare and Medicare, even though they were ineligible for traditional Medicaid.

TennCare differs from traditional Medicaid. To expand its program in a plausibly cost-neutral way, Tennessee shifted its entire Medicaid population to Managed Care Organizations (MCOs). These MCOs are in contrast to fee-for-service, in which a state pays physicians for each service provided. MCOs are typically private organizations generating contracts with health providers, paying providers jointly-determined fee schedules, and creating a network of specialists. Although most states use MCOs for at least some of its population, there is frequently a mix between fee-for-service and MCO's for Medicaid populations, depending on health status. Tennessee, like Oregon, Massachusetts, Hawaii, and Washington D.C., are strong outliers concerning the degree to which they use MCOs. As of 2010, low-population and southern states use fee-for-service more heavily, and most other states use MCOs more heavily, but to varying degrees (GAO, 2012).

Since Tennessee uses MCOs more often than other states, there is the question of whether my results are generalizable. I argue yes for two reasons. First, MCOs have become more prevalent over time, and Tennessee's high use of MCOs in 2005 provides insights into current state Medicaid programs (MAPAC, 2011). Second, although Tennessee spent less per enrollee than many other states (at the 10th percentile of spending per enrollee in 2005), it was comparable to other southeastern states, which typically had lower expenditure per enrollee. Figure 2 shows a histogram of total enrollee spending by state for both the nation and southeastern states. Nevertheless, Marton et al. (2014) gives evidence participants may value MCO less than traditional

pay-per-service Medicaid.

TennCare Standard's technical and financial eligibility requirements were simple. Until 2002, to be deemed uninsurable, an individual had to produce a denial letter from a health insurance company, and eventually had to submit to a Medical underwriting process.[2] To be deemed uninsured, an individual had to not have Medical insurance (typically for a year). The "medically needy spend down" category were people with resources who were pushed into poverty when medical expenditures were considered. In 2002, the state attempted to tighten requirements in the face of budgetary pressures, creating a yearly verification process, among other hurdles.

Nevertheless, repeatedly from 1995 until disenrollment, Tennessee's single-audit reports found "weakness in internal control over TennCare eligibility." For example, fake social security numbers made up 1% of all TennCare enrollees. Terminated SSI recipients were not removed from rolls, and could remain on TennCare indefinitely. Even after re-verification procedures in 2005, 22% of enrollee test cases failed to re-verify but were kept on Medicaid rolls. With a deliberately broad program, multiple enrollment avenues, poor internal eligibility control, and lack of re-verification, I treat a failure to be enrolled in Medicaid as a failure in imagination or desire by recipients, rather than the outcome of ineligibility[3]. During the first half of 2005, Tennessee engaged in large disenrollment, ending most optional Medicaid categories for adults. Specifically, it disenrolled the uninsured, uninsurable, and dual-eligible adults, ending their programs. Coverage of children was largely untouched, but the state limited benefits for those not disenrolled from TennCare, increasing copays and limiting benefits for some enrollees.

TennCare's disenrollment targeted only adults. Of those planned to be disenrolled, 121 thousand were uninsured adults, 67 thousand uninsurable adults, and 38 dually enrolled elderly in Medicare and Medicaid (Farrar et al 2007). In Table 1, I display results on the Medicaid populations reported in the closest two comparable reporting periods for Tennessee's administrative data, the Current Population Survey, and the Survey of Income and Program Participation. All three show a dramatic decline in Tennessee's Medicaid Population. My analysis uses SIPP from early

[2]TennCare Rule: 1200-13-14-0.02(3)(c)(3).

[3]I offer robustness checks to eligibility assumptions in the Robustness Appendix.

2005 to early 2006 for a long-difference study of panel-linked individuals. Table 1 suggests that SIPP overstates Medicaid disenrollment. If this is a product of oversampling or poor sampling, this causes no problem in the context of my model's estimation; all observations yield the same information because the preference distribution is homogeneous across people, in the spirit of Becker (1978). Indeed, an oversampling of disenrollees would improve my estimates by allowing better measurement in the affected population. Insofar as there is unmodelled preference heterogeneity by income level (presumably those between 100% and 200% of the FPL), this overweighs my preference distribution toward preferences of the semi-poor, but signs of biases are unclear. Alternatively, the overstatement of disenrollment could be the product of misstatements by interviewees. If interviewees claimed to be disenrolled but maintained enrollment, they presumably did not make changes to their behavior. This is likely to cause me to understate the value of Medicaid. To deal with this, I offer alternative calibrations in the Robustness Appendix.

TennCare disenrollment was a unique experiment because Tennessee offered the same insurance, with no differences in benefits, to all members. Disparities between enrollees were not in benefits, but in how much they had to pay for benefits. For all those below 100% of the poverty line, co-pays and maximum out-of-pocket costs were zero. For all individuals below 100% and children below 200% of the poverty line, premiums were zero. Nevertheless, of the 320 thousand households in Tennessee, representing 770 thousand people who reported being below the poverty line, only 71% reported having at least one person being covered by Medicaid. This lack of take-up is discussed in-depth by (Davidoff et al 2004, Davidoff et al. 2005, Aizer 2003). In Table 2 and Figure 1, I display Tennessee's 2005 premiums by FPL.

TennCare's rules produce cross-sectional contrast in take-up behavior between two groups; a large number of poor households could get Medicaid free but do not, while a large number of wealthier households pay premiums for the same benefits. In addition, disenrollment produced panel differences between groups, since adult disenrollment caused larger proportional losses in income in households that had an adult covered with no children than those that had children. Using SIPP data (discussed in the Data Section below), Appendix Tables A1 and A2 confirm that the primary role of TennCare is to provide health insurance to poor members of the population; those with no property income have heads with TennCare coverage three times higher

than heads of families with no property income. TennCare coverage also focuses on children, and the average age of those with TennCare is lower than the average age of those with ESI, NGI, or no health insurance, as in Appendix Table A3. Finally, Appendix Table A4 displays changes to TennCare coverage for adults (those over 18 years old) and children (those 18 and under) between 2005 and 2006. Adults were disenrolled heavily, but children were not. In fact, children were slightly more likely to be on TennCare in 2006, while adults were 32% less likely to be on TennCare in the same year, relative to 2005.

Income and Substitution Effects

There were two major effects of TennCare's disenrollment on labor supply: an income effect and wage effect. This section discusses the basic labor supply implications of TennCare's disenrollment in a simple model, before main effects get absorbed in a structural model. Losing access to Medicaid means losing access to an in-kind transfer, making households poorer. This pure income effect increased labor supply. In addition, when a household was disenrolled, its reward to work increased. Sliding-scale premiums acted as an implicit marginal tax rate; when purchasing Medicaid was no longer an option, the household effectively had a higher wage. The wage effect itself may be decomposed into a a positive substitution effect and a negative income effect on work, if leisure is a normal good.

As Figure 1 shows, the change in the implicit marginal tax rates from disenrollment were not small. For a single-parent household enrolled in TennCare with one child moving from 250% to 350% of the FPL by earning $12830 more, yearly premiums increase by $4500, from $3000 per year to $7500 per year. Linearizing the tax schedule, this household faced a 35% implicit marginal tax rate from premiums alone. The average household in Tennessee in 2005 did not face quite so high a rate. To calculate the average, I use a binning method similar to Mulligan (2013). First, I linearize the step function of premiums regarding income by family size by Medicaid enrollees per family. I then weight the relevant tax rates for each type of family, averaging to find an average implicit marginal tax rate of nearly 12% for families on Medicaid[4].

[4]See Data Appendix for details.

Removing an additional tax of 10% from a family with an ordinary tax rate of 20% results in a 14% increase in the reward to work.

While a change in exposure to premiums alone offers one source of identification for the value of Medicaid, it is not the only source. Medicaid is an in-kind transfer whose loss can be valued in cash. A given cash value of Medicaid has an associated labor-supply change from disenrollment, determined by the income elasticity of labor supply. According to the Centers for Medicaid and Medicare Services (CMMS), the average TennCare state and federal expenditures in 2005 were $4,966. The average pre-tax family income of an adult on Medicaid in 2005 was $31,723. Medicaid is large: if TennCare coverage was valued 1-to-1, a randomly disenrolled individual with a 20% tax rate would experience a nearly 14% drop in post-tax income (exclusive of wage and premium effects), holding behavior constant. This decreases to a 3.6% drop in income if Medicaid is valued at only $0.26 per dollar spent.

One additional income effect occurred even for households remaining on TennCare. Since Tennessee also cut benefits of those still on TennCare (such as pharmaceutical drug benefits), real Medicaid expenditures per patient fell. Real spending per TennCare enrollee fell 13% between 2005 and 2006, with a 10% drop in nominal spending and 3% in inflation rate.[5]

TennCare's disenrollment generated separate wage and income effects, and equation 1 offers a simple way of summarizing a single agent's response to those effects. The change in labor from a wage change comes from the percentage change in wage multiplied by the Marshallian (or uncompensated) wage elasticity of labor supply, and the change in labor from an in-kind income change comes from the percentage change in income generated by the loss of in-kind transfer income, multiplied by income elasticity.

$$\Delta \log L = \epsilon^M \Delta \log w + \xi \eta \Delta \log H \tag{1}$$

where ϵ^M is the Marshallian (or uncompensated) income elasticity of labor supply, η is the income elasticity of labor supply, Δ denotes change in labor L, wage w, and change in healthcare

[5]Although one-third of TennCare disenrollees had higher expenditures than average, it does not account for much of the change, which was driven primarily by lower spending on remaining enrollees.

transfer H. ξ, the crucial quantity to estimate using this equation, translates changes in healthcare transfers to cash income. By measuring all quantities and using parameters for which there is strong evidence in the literature, I estimate ξ as the residual.

Changes in labor, wages, and healthcare between 2005 and 2006 are easy to measure, and much literature is devoted to measuring labor elasticities. Table 3 displays the baseline and 2005-2006 increase in hours and drop in Medicaid to be used in equation 1. Alternative measurements are discussed in Appendix Table A5.

For any representative agent model with a standard budget constraint and utility over consumption and leisure, the Slutsky equation shown in equation 2 holds. Rather than describe a representative agent's preferences, implicitly defining the relevant elasticities, I work directly with elasticities.

$$\epsilon^M = \epsilon^H + s_L \eta \tag{2}$$

where ϵ^H is the Hicksian (or compensated) elasticity of labor supply, and s_L is the share of labor income in total income. Fundamentally, the Slutsky equation decomposes the total effect that a change in wages has on labor supply into a substitution and income effect, and I use it to calculate income elasticity, conditional on Marshallian, Hicksian, and the share of wages in income. The use of the Slutsky equation offers crucial discipline, as it restricts the joint range of elasticities I use in calibration.

Rather than taking a firm stand on elasticities, I calibrate to a range of elasticities based on the literature. For example, balanced growth preferences often demand low Marshallian elasticity to accord with small changes in aggregate per-capita working hours for prime-age individuals, even while real wages change by two orders of magnitude, as in Ramey and Francis (2009). Both macroeconomic and microeconomic estimates agree on a value of around 0.5 to 0.59 for the Hicksian elasticity of aggregate hours, when considering frictions (Chetty et al., 2011). Marshallian and Hicksian elasticities, combined with the Slutsky equation and a share of wages in income, imply a residual income elasticity. I find the income elasticity associated with a range of Marshallian and Hicksian elasticities. Figure 3 depicts income elasticities commensurate with the Slutsky equation, and some are large. My baseline estimate is -0.71.

Although consistent with most microeconomic foundations of macroeconomic models, the assumption that the Slutsky equation holds is not an innocuous assumption. My estimates accord with macroeconomic evidence of balanced growth preferences and Chetty et al. (2011), but the literature on income elasticities suggests values lower than those I use. Comparing players of differently sized lottery winnings, Imbens et al. (2001) find an income elasticity of -0.11, though they find evidence that the size of the income change matters. Using variances of inheritance size, Holtz-Eaken et al. (1993) find an income elasticity near zero. My typical value, though higher than some recent literature, is of smaller magnitude than Pencavel (1987)'s income elasticity of (negative) one, or Hausman (1988)'s similar value, and is closer to the relatively high values implied in Kimball and Shapiro (2008)

A baseline income elasticity of -0.71 is thus large. Contrary to intuition, this causes a high estimate of ξ, the cash value of Medicaid. Since the Slutsky equation holds for regular preferences and budget constraints, a reduction in income elasticity can only come from either a reduction in the Hicksian elasticity (for which there appears to be agreement between sources of variation) or in the Marshallian elasticity. Increasing the Marshallian elasticity lowers my estimates since disenrollment had a large wage effect. Consequently, I consider my high-income elasticity to be conservative for its purpose of estimating ξ since it yields the highest value for ξ. With this in mind, I present results for a series of income elasticities, including some consistent with modern literature. These results are shown in Table 4.

Of course, it could be that micro and macro estimates of the Hicksian elasticity of labor supply are incorrect. I could instead set the Marshallian elasticity to zero and alter the Hicksian elasticity, thereby changing the income elasticity. This ends my estimation's robustness for a simple reason. A Marshallian elasticity of zero implies that permanent wage changes to have no effect on labor supply in the long run. As the Hicksian elasticity moves towards zero, my estimated income elasticity falls and income changes have no effect on labor as well: in order to explain any movements, there must have been large changes in Mediciad value. Table 5 offers this alternative inference exercise.

Since the exercise embedded in equation 1 with consistent elasticities generated by equation 2 is equivalent to a representative agent model with preferences summarized by my stated elasticities, I analyze a homogenous group for which these effects are unlikely to be lost to noise. Literature on Medicaid often concerns itself with information frictions or stigma costs (Moffitt, 1983). To difference-out these concerns, I examine the change in labor supply for households on Medicaid, comparing changes in labor during the disenrollment year to changes in other years. Therefore, Δs denote a double difference: the change between 2005 and 2006 in comparison to the change between two non-disenrollment years, 2004 and 2005. By doing so, I compare households that already overcame information frictions and shame hurdles, and my results are robust to misstatement of (1) by a constant. To compare like households in this example, I use panel data in SIPP to measure the change in labor between 2005 and 2006 for those enrolled in Medicaid in 2005, beyond the change in labor between 2004 and 2005 for those enrolled in Medicaid in 2006. Ordinarily, about 14% of Medicaid enrollees in a given year transition out of Medicaid by the following year. Between 2005 and 2006, this reached 60%, reflecting disenrollment and high turnover. Measured labor supply rose by 2.6 percentage points more between 2005 and 2006. The representative drop in income generated by disenrollment between 2005 and 2006 was $\xi \cdot \$2326$, in comparison to average medicaid family income of a little more than \$16,000. Using the linearized premium schedule described previously, the increase in the reward to work (increase in take-home wages) due to disenrollment was 5%.

With a Marshallian elasticity of zero and a Hicksian elasticity of 0.5, I estimate the cash conversion rate, ξ, to be 0.26, as shown in Table 4 and 5. This is a robust result; if the Marshallian elasticity were 0.16 (at the high end of estimates found in the literature), I would instead attribute more of labor's change to the wage change, and the estimate for ξ remain at 0.26, increasing only slightly. If SIPP misrepresented labor's change by 50%, and labor actually increased by 3.9%, my estimate for ξ would be 0.38.

Figure 4 depicts my cash equivalent exchange rate by measured change in labor supplied for several values of Marshallian and Hicksian elasticities. For reasonable Marshallian elasticities, estimates are similar because increasing the Marshallian elasticity decreases the income elasticity, which decreases the estimate of ξ, but also increases the labor response to a wage change,

increasing the estimate of ξ. The Hicksian elasticity is much more important: setting it low while the Marshellian elasticity is zero means that there are no wage effects and only small income effects: consequently, to explain any change in labor there must have been a large change in income. Nevertheless, for a reasonable range around Chetty's Micro-Macro consensus Hicksian elasticity, values do not range much below 0.2 or above 0.5.

Because this exercise concerns itself with labor's response conditional on an income change, my choice of period is a natural concern. I prefer comparing 2005-2006 with the 2004-2005 period using SIPP, because Tennessee's labor market was healthy in both the "pre" and the "post" periods, but there was relatively little net disenrollment in Medicaid between 2004 and 2005, while there was a large disenrollment between 2005 and 2006. However, it may be the case that the 2005-2006 time period does not give enough time for the labor market to reach equilibrium.

In 2005's healthy economy, the national job monthly finding rate at a monthly level was around 40% (Shimer 2012). With any reasonable job finding rate and employment exit probability in Tennessee, the labor market would converge to equilibrium in a matter of months following the shock generated by disenrollment. Analyzing the change in Medicaid take-up over a longer time horizon will do little to account for frictions in job search and will pick up more secular changes, swamping out disenrollment's effects and adding noise to my estimate of the cash-equivalent value of Medicaid.

Unfortunately, whether or not it is desirable, two-year difference-in-difference comparisons using the 2004 SIPP are not possible because SIPP is only four years long and disenrollment occurs in the middle of the second year. Instead, I use an alternative Medicaid comparison group to double-difference. Rather than using 2004-2005 Tennessee as a baseline compared to 2005-2006 Tennessee, I use the 2005-2007 Census Bureau's East South Central Division[6] excluding Tennessee as a baseline and compare it to Tennessee over the same two years. State fixed effects in Medicaid recipient behavior swamp variation in Medicaid, making the pure cross sectional comparison very noisy: I find that while the rest of Tennessee's division has a much smaller reduction in Medicaid enrollees, these other states had larger increases in measured aggregate

[6]The ESC Division includes Kentucky, Mississippi, and Alabama.

hours between 2005-2007, implying negative Medicaid valuations[7]. This "backwards" comparison also true of the 2005-2006 cross-state comparison.

While cross-state comparisons may not contain useful information about the level of ξ, they do provide information about the robustness of my timeframe choice. While the estimate of the value of Medicaid is negative using 2005-2006, it is more negative using the 2005-2007 comparison group. If both negative values were caused primarily by a constant state fixed effect, then the reduction in the value of Medicaid between the 2005-2006 and 2005-2007 calculations suggest that my 2005-2006 timeframe estimates a conservatively high value for Medicaid. Like my choice of state experiment, I conclude that my choice of timeframe appears to yield a higher value of Medicaid than other potentially reasonable alternatives.

A Bayesian version of this estimation procedure accounts for uncertainty in elasticities and labor data. I assume uncertainty about Marshallian and Hicksian elasticities and the change in labor supply in equations 3-5. Uncertainty about elasticities comes from the literature, and uncertainty about the change in labor ΔL is a half percentage point standard error.

$$\epsilon^M \sim \text{Uniform}(0, 0.16) \tag{3}$$

$$\epsilon^H \sim \text{Uniform}(0.4, 0.7) \tag{4}$$

$$\Delta L \sim \mathcal{N}(0.026, 0.005) \tag{5}$$

With a uniform prior for ξ between zero and one, Figure 5 depicts the posterior distribution of ξ. The modal value is 0.22, with a 25th and 75th percentile of 0.19 and 0.29. Since my baseline estimate of 0.26 uses a Marshallian elasticity of zero, and higher Marshallian elasticities decrease the estimate of ξ slightly, my baseline estimate is slightly higher than the mean or modal estimates. Table 6 gives percentiles of the posterior distribution. The message from this exercise is that the robustness checks in Figure 4 are robust even when allowing the Marshallian, Hicksian, and measured labor change to vary: reasonable uncertainty about labor response or

[7]Medicaid's two-year change in hours for families with Medicaid patients in Tennessee in March 2005 was 1.01 hours per year, while the decline in Medicaid was -0.23 contracts. The two-year change in hours for other states was an increase of 0.74 hours, while the decline in Medicaid was -0.30.

Marshallian and Hicksian elasticities, does not affect estimates of ξ drastically.

This exercise highlights the main effects of TennCare disenrollment on labor supply. First, there were large changes in the reward to work from a reduction in the implicit marginal tax on labor of disenrolled by 10%. Second, there were large income effects since a cash-equivalent loss of \$2075 per disenrolled individual motivated them to work. In the next section, I introduce an explicit model of healthcare, labor, and consumption choice. This model includes agents heterogeneous in their value of healthcare in general, and of Medicaid specifically. Unlike this exercise, the utility function in the main model differentiates consumption of goods and healthcare. This allows me to model the choice set facing families in Tennessee in 2005 better. Nonetheless some of the primary drivers of identification come from the changes in labor behavior described above.

Heterogeneous Household Model

I use part of the general equilibrium model of Gallen (2014)'s analysis of the Affordable Care Act to give more structure to the relationship between Medicaid and labor supply. This model introduces household heterogeneity in preferences, and differentiates consumption of healthcare services and other market goods since health insurance might not be equivalent to other forms of consumption. Medicaid has partial substitutes in employer-sponsored insurance (ESI), which agents receive through work and non-group insurance (NGI), which agents can purchase at market prices. I allow multiple substitutes in healthcare production, itself additively separable from consumption. The model also allows heterogeneity in household demographics and nonlinearities in the budget constraint. The tradeoff between the previous simple exercise and this more complex heterogeneous model is clear; to take the model closer to data, the source of results becomes more opaque. Nevertheless, the fundamental income and wage effects introduced previously influence the decisions of every household.

Households have utility over consumption, a composite healthcare good, and leisure. The household has constant Frisch elasticity of labor supply utility over household consumption (C_j), composite family healthcare (H_j), and labor hours L_j, which takes the following form for household j with number of non-elderly members N_j and number of non-elderly adults $1 + N_{s,j}$:

$$U(C, L, H) = N_j \log\left(\frac{C_j}{N_j}\right) + \kappa N_j \log\left(\frac{H_j}{N_j}\right) - \sum_{i=1}^{N_{s,j}} \psi_j \frac{\epsilon}{1+\epsilon}(L_{i,j} - \nu)^{\frac{1+\epsilon}{\epsilon}} \tag{6}$$

where κ denotes the relative utility from healthcare, ψ denotes the disutility from labor, ν is the fixed cost of working, and ϵ is the constant Frisch elasticity of labor supply.

The household gets utility over composite healthcare H, which is generated by the CES combination of various types of healthcare across the household:

$$H_i = \left(\omega + a_{E,i}\left(\sum_{j=1}^{N_j} H_{E,j}\right)^{\frac{\sigma-1}{\sigma}} + a_{N,i}\left(\sum_{j=1}^{N_j} H_{N,j}\right)^{\frac{\sigma-1}{\sigma}} + a_{G,i}\left(\sum_{j=1}^{N_j} H_{G,j}\right)^{\frac{\sigma}{\sigma-1}}\right)^{\frac{\sigma}{\sigma-1}} \tag{7}$$

where $H_{E,j}$ indicate whether family member j has health insurance through one of the family's employers. Similarly, $H_{N,j}$ indicates whether the agent owns contracts for non-group insurance, and $H_{G,j}$ indicates whether agent j has Medicaid. ω is the baseline level of Medical care for any family; it plays a dual function of allowing healthcare to be an inferior or luxury good. The elasticity of substation between healthcare types is governed by σ. As σ moves toward infinity, ESI, NGI, and Medicaid become perfect substitutes.

The budget constraint states that consumption is equal to after-tax income minus non-group healthcare expenditures and Medicaid premiums. After-tax income is generated by the sum of labor income wL and property income I with expenditures on employer-sponsored insurance removed before proportional income tax τ_j:

$$C_j = (1 - \tau_j)\left(\sum_{i=1}^{1+S_{N,j}} w_{i,j}L_{i,j} + I_j - P_E H_{E,j}\right) - P_N H_{N,j} - P_G H_{G,j} \tag{8}$$

where $w_{i,j}$ is the wage of household member i in household j, P_E is the cost of ESI, P_N is the cost of NGI, and P_G is the income-contingent cost of Medicaid. I omit denoting income dependence in $P_G(Inc_j)$ for legibility.

I allow households to "purchase" government insurance directly if they are below 200% of the federal poverty line. Additionally, they may purchase H_G for children at any point on the FPL,

though they may have to pay premiums. I allow for free gaming of TennCare, as discussed in the section of this paper describing TennCare, so households may sign up and pay premiums without paying an enrollment cost other than stated premiums. Allowing free gaming might ignore some costs. For example, some individuals had to be uninsured for more than a year to be classified "uninsured." Others had to submit private insurance rejections or submit to an underwriting process by TennCare to be defined "uninsurable." Alternative specifications include small "inconvenience" fees. Since I measure the net benefit of Medicaid to households, this does not move distributions; while Dague (2014) demonstrates a large number of individuals accruing a low net benefit, adding small costs increases the mean pre-inconvenience estimate for Medicaid while leaving the post-inconvenience estimate unchanged. I do not move distributions; the only way inconveniences deter individuals from Medicaid is if they were implausibly large, or if Medicaid had a value lower than application costs. This paper focuses on the latter.

Preference parameters over healthcare types a_E, a_N, a_G are independent Pearson type I distributions. These distributions are four-parameter generalizations of the beta distribution, taking in a mean, variance, skew, and kurtosis as parameters. I set the excess kurtosis to zero, and estimate the first three moments to fit the data. For $j \in \{E, N, G\}$, I estimate parameters μ_j, σ_j^2, and γ_j for the Pearson type I distribution, denoted Beta*[8]:

$$a_j \sim \text{Beta}^* \left(\mu_j, \sigma_j^2, \gamma_j^2, 0\right) \tag{9}$$

Finally, I assume preference parameter ψ, the disutility of labor, is also generated by a Pearson type I distribution with zero excess kurtosis:

$$\psi \sim \text{Beta}^* \mathcal{N} \left(\mu_\psi, \sigma_\psi^2, \gamma_\psi^2, 0\right) \tag{10}$$

This heterogeneous model has several advantages and disadvantages. Heterogeneity in ψ allows for observation of households with similar wages but different hours and healthcare choices.

[8]In my Robustness Appendix, I change the assumption of independence between preferences over healthcare types by allowing utility weight for healthcare κ to vary, changing the interpretation of a_E, a_N, and a_G to be the value of that specific healthcare, above and beyond the value reflected by κ. This reduces my estimated values of Medicaid by increasing the value of each person's next best opportunity. For instance, people who really like Medicaid are likely to enjoy it because they value all kinds of healthcare: they therefore do not have to be paid as much to drop Medicaid. In this sense, my assumption is conservative with regards to my relatively low estimated value of Medicaid.

Heterogeneity in healthcare production weights a helps explain combinations of ESI, NGI, and government healthcare among otherwise identical households. a weights, along with the structure imposed by Medicaid premiums, contain a distribution of welfare implications for Medicaid expansions, and help explain why average and marginal Medicaid take-up are different. Wage, income, and household count heterogeneity are crucial to generate a distribution of marginal individuals.

Modeling the value of a "contract" rather than expenditures directly, I disconnect the average spending of Medicaid from the value a household places on a Medicaid contract. With a program like Medicaid, some households are transferred more than others are. This distribution of spending contributes to a distribution of a_G values. Nevertheless, both the distribution and average valuations are comparable and informative; when I find the average Medicaid dollar is valued at \$0.35, it must be that somewhere in the distribution dollars are not being valued by recipients at what they cost.

Understanding Identification

Before estimating the model, I discuss identification and offer a simple example, estimating a comparable parameter using results from Dague (2014). My model has much in common with discrete-choice product differentiation problems from McFadden (1974), or revealed values of patents generated by the binary choice of renewal from Pakes (1986). However, my model adds a twist; not all households face the same prices, and may alter decisions to change their prices endogenously. For example, in traditional Medicaid notch literature, we might imagine a Medicaid household that chose to work an additional hour, being disqualified for Medicaid, and losing \$5000 of Medicaid expenditures for \$10 in pre-tax income. This indicates a very low value of Medicaid; at equilibrium, it is costless to motivate a household to work a little less since it is indifferent between paid labor and leisure. If the same household purchases employer-sponsored insurance, it is further evidence that while it values healthcare, it does not value the healthcare Medicaid offers.

Generally, I observe only discrete health insurance choices; if a household chooses labor hours

that place it above the poverty line and the household purchases ESI rather than dropping labor hours and accessing Medicaid, it must be that its valuation of Medicaid plus leisure relative to ESI and labor is less than some threshold of indifference between the two options. If Medicaid's take-up among those who would have to pay very little for it was 60%, and was 20% among those who have to pay $2000 for it, it suggests that 40% of the population has an extremely low valuation, and 20% have a valuation higher than $2000. Unfortunately, I never gain information about tails beyond the maximum amount a family pays for Medicaid. I infer the shape of the tail beyond that last point from distributional assumptions. Consequently, I pay careful attention to robustness checks. Even without reliable top bounds, much valuable, public-policy information can be gleaned. A simple example using Dague's (2014) results offers further intuition into identification of my large, heterogeneous agent model.

A simple example: extracting results from Dague (2014)

Dague (2014) uses regression discontinuity to analyze the effects of discrete premium hikes in Wisconsin's BadgerCare program on the length of Medicaid enrollment. A significant result in Table 1.4 from Dague interests me; she finds that increasing the monthly premium from $0 to $10 decreases the average Medicaid spell length by 1.35 months. With an average enrollment spell of 10.69 months for those below 150% of the FPL, this represents a 12.6% decline in Medicaid contracts due to a $10 monthly premium. Another way of putting this is that at any given point, the 12.6% of people who would have been on Medicaid are unwilling to pay $10 per month to stay on Medicaid. As Dague discusses, a common estimate for take-up among zero premium individuals is 62% (Sommers et al., 2012). I treat these two observations as censored moments, given information about the underlying distribution of Medicaid valuations. I interpret Dague's results as suggesting 12.6% of Medicaid enrollees who value net benefits above $0 do not value them above $10 a month ($120 a year). Similarly, I interpret the low take-up estimates that Sommers et al. (2012) gives as suggesting 38% of individuals do not value the net benefits of Medicaid above $0. Assuming a normal distribution of net benefits, I express the two equation, two unknown system in equation 11, letting Φ be the CDF of net benefits and μ and σ^2 be its mean and variance:

$$\begin{pmatrix} \Phi(0|\mu,\sigma^2) \\ \frac{\Phi(10|\mu,\sigma^2)-\Phi(0|\mu,\sigma^2)}{\Phi(0|\mu,\sigma^2)} \end{pmatrix} = \begin{pmatrix} 0.38 \\ 0.126 \end{pmatrix} \tag{11}$$

The likelihood in equation (11) assumes net valuations below zero exist; having net valuations at or below zero is the only way to reconcile low take-up among people knowledgeable about the program. However, when analyzing the value of Medicaid to beneficiaries, I use only positive values of those who sign up.

As displayed in Table 7, I find the average cash value of Medicaid among beneficiaries as $551. Since Dague observes data from 2008 through 2010 but my Medicaid expenditure data include 2009, I conservatively assume the average of 2008 and 2009 spending ($5855 and $7584, respectively) to generate a mean cash exchange value of $0.08 per dollar spent. Clearly, this exercise depends on functional form, and several additional functional forms are given in Table 7.

Nevertheless, I draw three lessons from my identification exercise with Dague's results. First, an increase of $10 in monthly premiums for adults results in a dramatic decline in total exposures to Medicaid in Wisconsin–more than 12%– suggesting a large mass of low valuations. Second, it provides a simple example to the idea behind my cross-sectional variation premium's identification: a series of censored valuations. This exercise's vulnerability to types of distributions leads me to use a more flexible, four-parameter distribution, which I could not do with only the two significant discontinuity results for adults in Dague (2014).

Data

I estimate the heterogeneous model using moments from SIPP. I focus on panel waves 4 and 5 (March, 2005), and 7 and 8 (March, 2006), in Tennessee. These data provide household, family, and individual responses to questions about income, healthcare coverage, demographics, and labor hours for 2,499 individuals in approximately 1,000 families. Important for the study of the effects of TennCare disenrollment, these participants were surveyed in 2005 and re-surveyed in 2006.

Using this panel variation provides an advantage over other methods, such as analyzing syn-

thetic cohorts in the CPS. Summary statistics on the population used in this paper are provided in Table A6. The average number of people in a family was 2.42, shrinking slightly in 2006 to 2.37. The average age in the entire sample, including children and the elderly, was 36.33. Wages among working-age individuals aged 25 to 55 was 15.23 dollars per hour when imputing missing data for non-working adults using a Heckman selection procedure, using own children under 18 as a selection instrument. Of this prime working-age population, 82.02% were employed, consistent with BLS estimates for this age group. The median household had nearly no property income and less than half the population had education beyond high school. These statistics changed little between 2005 and 2006.

One benefit of using SIPP is the high reporting rate. Analyzing a host of transfer programs, Meyer et al. (2009) document the relative performance of the Panel Survey of Income Dynamics, SIPP, CPS, ACS, and the Consumer Expenditure Survey, finding that for most programs, SIPP far outperforms other surveys in program participation report rates when compared to administrative data.

Although SIPP is better than other standard datasets at estimating a number of statistics, it undercounts income and Medicaid participation, among other issues (Meyer et al., 2009; Card et al., 2004). For example, SIPP understates income (Czajka and Denmead, 2008). Since I am not running regressions on a micro-dataset, but instead calibrating moments, I "fix" my measurements. For example, although I use SIPP to calibrate the relative proportion of people on Medicaid by income, I use administrative data to produce the aggregate number of people, and force the sum over income to equal administrative counts, assuming a constant rate of undercounting across income levels.

Analyzing only people matched in both waves, healthcare coverage choices in 2005 and 2006 are shown in Table A7. In both years, approximately 60% of the population had ESI, and about 10% had NGI. The proportion of the population with TennCare fell by almost 20%—from 22% in 2005 to 18% in 2006. Using the panel aspect of the data, it is possible to examine how individuals transitioned from TennCare to other types of health insurance coverage (or no coverage). Table 8 shows the transition matrix for healthcare coverage. Fewer than 50% of those previously

on TennCare stayed on TennCare (some transitioned into TennCare). Only 10% transitioned to ESI, and fewer than 5% to NGI. Most of those who no longer had TennCare in 2006 were uncovered.

I use moments generated from families in SIPP to match my model of family decision-making. Following a given family from 2005 to 2006, I use the panel aspect of SIPP, which allows me to both eliminate compositional noise from changing samples and eliminate the assumption-laden constriction of a synthetic panel. Additionally, SIPP's monthly question about Medicaid coverage provides greater detail than the CPS's March Supplement question, which only refers to coverage "last year." The next section describes the estimation in detail.

Estimation

I estimate the model described by equations (6)-(10). Specifically, I take the distribution of $a_{E,i}$, $a_{N,i}$, $a_{G,i}$, and ψ_i to each be three-parameter beta distributions, defined by some nonnegative interval. Given household composition, ϵ, σ, and prices, I simulate the distribution of preference parameters and find the maximizing choice for any household.

I produce a distribution of nearly 1,000 households by extracting all sub-families defined in SIPP, including sub-family wages, property income, SIPP-defined federal poverty level, number of children, and the presence of a spouse. To populate my simulation, I sample 20,000 of these households, with replacement. I use a two-step Heckman selection procedure to fill in blank wages, using children under 5 as a participation instrument, with education, number of own children in the household, marital status, and age as controls. The method of drawing samples from SIPP has the advantage of being distribution-free (except for non-observed wages) with regards to the cross-correlations between drawn values, which is particularly convenient insofar as the marginal distributions of spousal wages, or wages by spousal presence, are not independent.

Conditional on a set of distributional parameters and household composition, I simulate the joint labor and healthcare choices a family makes to maximize utility. For any family, all healthcare

variables are limited to be a discrete variable between zero and the number of people in the family, $1 + N_{s,i} + N_{c,i}$ (representing one insurance contract possible per person). Labor for an adult might take discrete hour values between zero and 50 hours. Period length is calibrated to a year.

The Centers for Medicaid and Medicare Services in 2009 provided data on prices. Included in the analysis are all state and federal spending categories. For Tennessee, this yields $4331 in 2005, and $4477 in 2006 (both in 2005 dollars). Prices for employer-sponsored insurance and non-group insurance are generated from the Medical Expenditure Panel Survey, supplemented by measurements of full out-of-pocket costs by Gabel et al. (2003), reflecting the full cost of ESI and NGI. Finally, for government healthcare prices P_G, which depend on income, I input the premiums from Table 2 directly.

In the simulated method of moments framework discussed by McFadden (1989), McFadden and Rudd (1994), and Gouriroux and Monfort (1997), I target average labor, ESI, NGI, and Medicaid take-up and purchase rates, and the same rates by FPL. In addition to this cross-sectional variation, I use the experimental variation from the panel, targeting labor and healthcare responses to disenrollment. These target moments are listed in Table 9.

Conditional on the distribution of preference parameters, and given a distribution of wages and household composition drawn directly from SIPP, I can simulate a host of families with preferences, wages, and other attributes according to my parameters and data. Given household and preferences, I numerically maximize utility, and generate any simulated statistics, comparing them to cross-sectional and panel data. I then choose parameters to make these simulated moments match data. I simulate 10,000 families and allow them choices over labor, spousal labor (if present), ESI, NGI, and Medicaid, subject to the heterogeneous agent model outlined in equations 6-8. I assume households face a discrete choice: health contracts can not be purchased partially, but only in integer numbers, up to the maximum number of household members. Similarly, labor can be chosen in increments of five hours, from zero to fifty possible hours. ESI can only be purchased if someone in the household works. The set of discrete choices for a given household is given by the Cartesian product of all feasible combinations of choices in equation (12)-16).

$$L_j \in \{0, 5, ..., 50\} \tag{12}$$

$$L_{S,j} \in \begin{cases} \{0, 5, ..., 50\} & \text{if } N_{S,j} = 1 \\ 0 & \text{otherwise} \end{cases} \tag{13}$$

$$H_{E,j} \in \begin{cases} \{0, 1, ..., 1 + N_{S,j} + N_{c,j}\} & \text{if } \sum_{i=1}^{1+N_{s,j}} L_{i,j} > 0 \\ 0 & \text{otherwise} \end{cases} \tag{14}$$

$$H_{N,j} \in \{0, 1, ..., 1 + N_{S,j} + N_{c,j}\} \tag{15}$$

$$H_{G,j} \in \begin{cases} \{0, 1, ..., 1 + N_{S,j} + N_{c,j}\} & \text{if Percent of FPL } < 2 \\ \{0, ..., N_{c,j}\} & \text{otherwise} \end{cases} \tag{16}$$

Table 9 presents simulated moments and their corresponding targets. I estimate the "baseline" value of healthcare, ω, the weight on healthcare κ, the fixed cost of working ν, and the twelve distributional parameters. These fifteen moments are matched to twenty-one statistical moments generated from SIPP. Most moments match adequately to represent Tennessee in 2005, and the change in Tennessee's aggregate variables from 2005 to 2006. Generally, the model captures the important cross-sectional and panel variation I target. Two moments are poorly matched: the change of the poor population on Medicaid, and the hours of work by income level.

My estimation fails to achieve the change in poor population because in reality, the poor (who should not have been impacted by the disenrollment much) appear to have been impacted. This might have been, for example, because if a parent could not sign up, he/she also withdrew his/her children. In my model, there is little incentive to do this, but fixed costs of a family signing up might drive this result. The reason my does not capture the degree to which hours are responsible for income differences is likely due to a lack of demographics adjustment for ψ, which is common in labor preferences like these (Kimball and Shapiro, 2008). I maintain the same distribution for ψ no matter the number of potential workers in a household, whether a head of household, a spouse, or both work. This reflects an assumption that it is a deep structural parameter, rather than a reduced form parameter. Although this offers intuitive appeal, in which humans have the same distribution of distaste for leisure no matter their martial ar-

rangements, I sacrifice fit.

Results

How much do people who receive Medicaid value it?

After estimation, I can simulate the population and calculate the compensating variation necessary to make someone voluntarily drop Medicaid take-up to zero. For those individuals whom I estimate are at the margin of enrolling in Medicaid or not, the value is zero. While they might value healthcare and Medicaid, their next best alternative (which might include ESI or NGI) is nearly equivalent. This is an important distinction. I find the value of the existence of Medicaid in a world that holds the existence and prices of ESI and NGI constant. The compensating variation is how much a family would have to be paid to disenroll the entire family from government healthcare. This paper examines the net benefit of Medicaid to those currently enrolled. There might be unmeasured positive value to those Cutler and Gruber (1996) call conditionally covered, who are signed up for Medicaid by a hospital when they go in for Medicaid treatment such as a delivery but are not otherwise enrolled; I only estimate the value to those currently enrolled.

It is possible to find this break-even point for each of the nearly 20% of households in my simulated population on Medicaid. The distribution of compensating variation is displayed in Figure 6. It is highly skewed, with a median valuation of $373 and a mean valuation of $1764. This is the mean of compensating variation, rather than the value. Since some households pay for Medicaid, including the out-of-pocket expenses and premiums increases the skew slightly. Table 10 shows a summary of this compensating variation.

A clear lesson from Table 10 is that consumer surplus from Medicaid is large, though not as large as costs. Among those who receive it, the average recipient would be willing to pay $1764 to keep Medicaid rather than be disenrolled. With a mean expenditure of $4966 in 2005 and $4331 in 2006, Medicaid is valued, on average, at about $0.35 per dollar spent. Of course, this calculation depends on the mean expenditure on Medicaid. It is well known that the distribution of Medicaid payments is not singular, but is instead highly skewed. I examine the distribution

of Medicaid expenditures separately.

The skew of compensating variation also helps elucidate why Medicaid take-up is so low and variable across states (Sommers et al., 2012), and why enrollment decreases so much when a $10 monthly fee is introduced Dague (2014). For some proportion of households, it is treated as nearly valueless. This is not an artifact of distributional choice; with a guaranteed zero out-of-pocket cost in Tennessee for low-income children, but tens of thousands not enrolled, the value of the program must be low for some households. The presence of near-zero valuations are a product of the data and revealed preferences, not structure.

Another study on Medicaid, Finkelstein et al. (2012), helps us understand why. For those included in Oregon's Medicaid, the chances of a doctor's office visit in the last 12 months were about 69%, the probability of using prescription drugs was nearly 57%, and other measurements of utilization were lower. This suggests that even among those on Medicaid, the probability of using it is often less than 40%. With such low utilization observed, Medicaid might have some catastrophic insurance value, but some households should treat it as valueless because they never use it.

Because I do not allow households the option of valuing future Medicaid enrollment, my analysis only gives the current value of benefits to beneficiaries. If a household dodges premiums by not enrolling until it has a need for Medicaid, I miss the option value of Medicaid for these households since I analyze only the proportion of the population that is enrolled in a given year, rather than over its lifetime. Similarly, I may overstate the value of Medicaid for those who stay on Medicaid simply because they may use it later, and don't want to deal with re-verification frictions.

Understanding the source of low utilization is important for designing Medicaid policy. I reveals that households value Medicaid coverage disparately. However, measured by expenditures, households also receive different amounts of health services through Medicaid. The next section provides evidence consistent with a homogeneous valuation of about thirty-five cents per dollar of spending, and that large differences in valuation are driven by expenditures, rather than by

differences in per-dollar valuation. This suggests the problem of take-up is likely to be one of spending, rather than heterogeneous valuation.

The difference between my cursory calculations and my more realistic heterogeneous model should not be surprising. First, cursory calculations accord with heterogeneous agent calculations when higher Marshallian elasticities and lower income elasticities are considered, as shown in Table 4. The heterogeneous agent model examines the consequences of a nonlinear tax schedule directly, which has the capacity to be more distortionary than its linearized counterpart does. I use both cross-sectional and panel data to estimate the heterogeneous agent model rather than pure time-series variation. If the time-series variation I use is less representative of the long-run effects than cross-sectional variation, my simple time-series analysis might overstate the value of Medicaid. Similarly, if there are cross-sectional frictions that I fail to model in my heterogeneous agent model, it might understate the value of Medicaid.

Differential Medicaid Spending

The distribution of Medicaid spending is highly skewed (NIHCM (2012)). An estimate of the cumulative distribution of spending, replicated from Figure 1 in NIHCM (2012), is depicted in Figure 7. It is comparable to the distribution of Illinois Medicaid spending (Pollack 2013). If ex ante all households knew how much Medicaid expenditure they were going to receive, the distribution of a_G's might be driven by the distribution of expenditures. In this section, I examine this proposition and find that although a significant portion of the a_G distribution can be explained by the distribution of costs, there remains a significant deficit in the value of Medicaid expenditures.

In Table 11, I take the distribution of cumulative spending by percentile that NIHCM (2012) produces and transform it into average spending by percentile. As Table 11 shows, the average spending varies dramatically by percentile; the top 5% of spenders make up nearly 50% of spending. To allow cost data to explain my distribution of valuations (which did not use cost data, but did use unrelated, imputed subsidized premium data by household), I juxtapose those who valued Medicaid the least with those who were spent on the least. Lining up these categories, I

find a degree of uniformity within the population, especially as Medicaid expenditures begin to grow. For example, the 10% of people between the 60th and 70th percentiles only get 4.8% of spending, or $2384/person. According to my model, they value it at $638, yielding a value of $0.26 per person, not dissimilar to the average valuation. The last column of Table 11 displays the "exchange rate" of Medicaid per dollar of consumption. It varies, especially for households that receive very little in Medicaid expenditures. Among the 50% of Medicaid households that receive more than 95% of spending, exchange rates are much more homogenous, ranging between 0.18 and 0.43. This method is only suggestive: by lining up value and expenditures, I give the distribution of expenditures its best shot at explaining values.

Figure 7 shows the inverse CDF of log spending from the data alongside my model's theoretical discount. That is, the distributions are comparable when I multiply the distribution of valuations by a number between three and five. Figure 7 brings an important question to the fore: when allowing for heterogeneity, there is a difference between the average value of a dollar of Medicaid and the valuation of a dollar of Medicaid to an average family. A distributional fit for both is depicted in Figure 7. The fit is good; the model captures a large amount of variation in Medicaid valuation. Outside data shows that it captures the distribution of Medicaid payments accurately, suggesting the population could be homogenous in its valuation of Medicaid. In addition, the right tail of valuations is only identified by structural assumptions, but looks quite similar to the tail of expenditures, differing only by a constant.

Conclusion

I find that disparate models yield comparable results. A simple, cursory calculation and a microeconomic model of heterogeneous households yield the similar qualitative conclusion: the in-kind transfer of Medicaid is valued at approximately $0.35 per dollar spent, far below a natural benchmark of $1. Both robustness checks in the representative agent model and using cost distributions for heterogeneous agents urge caution in my point estimate. Ranges from $0.14 to $0.40 per dollar are not contradicted by reasonable alternative methods, data, or parameter assumptions.

Key to this result is the recognition that low take-up and failure to reduce hours in response to the Medicaid notch dramatically is evidence that it is not very valuable to households per dollar spent. Although it is possible that information or stigma costs are underlying causes behind my preference parameters, they represent a real phenomenon in household utility, relevant for public policy analysis. Medicaid's compensating variation is not high. My results are in accord with two simple observations. If Medicaid was valuable, take-up would be higher. Similarly, if Medicaid was valuable, few people would earn small amounts above the Medicaid notch, wherever it were for their state.

Although average results are similar, heterogeneity reveals that it is highly unlikely that Medicaid is efficiently enrolling households. Marginal households (those near indifference to Medicaid coverage) have very low valuations, while inter-marginal households have dramatically higher valuations. Although this is clear from lack of take-up among households with zero out-of-pocket costs and take-up among households paying thousands of dollars for Medicaid, my model uses structure to quantify the distribution of valuations. There is large variation in Medicaid valuations, but there is also large variation in costs. I find that nearly all variation in Medicaid valuations can be explained by heterogeneous Medicaid payments. Doing so leaves a homogenous valuation of Medicaid spending in terms of cash: approximately \$0.26-\$0.35 per dollar.

References

Anderson, Patricia M. and Bruce D. Meyer, "Unemployment Insurance Takeup Rates and the After-Tax Value of Benefits ," 1997, *112* (3), 913–937.

Baicker, Katherine, Amy Finkelstein, Jae Song, and Sarah Taubman, "The Impact of Medicaid on Labor Force Activity and Program Participation: Evidence from the Oregon Health Insurance Experiment," *NBER Working Paper*, 2014, (19547), 1–13.

Becker, Gary S., *The Economic Approach to Human Behavior*, University of Chicago Press, 1978.

Bennett, Christina Juris, *TennCare, one state's experiment with Medicaid expansion*, Nashville, Tennessee: Vanderbilt University Press, 2014.

Bradley, Cathy J., David Neumark, and Scott Barkowski, "Does employer-provided health insurance constrain labor supply adjustments to health shocks? New evidence on women diagnosed with breast cancer," *Journal of Health Economics*, 2013, *32* (5), 833 – 849.

Card, David, Andrew K.G. Hildreth, and Lara D. Shore-Sheppard, "The Measurement of Medicaid Coverage in the SIPP: Evidence From a Comparison of Matched Records," *Journal of Business and Economic Statistics*, 2004, *22*, 410–420.

Chetty, Raj, Adam Guren, Day Manoli, and Andrea Weber, "Are Micro and Macro Labor Supply Elasticities Consistent? A Review of Evidence on the Intensive and Extensive Margins," *American Economic Review: Papers and Proceedings*, 2011, *101* (2), 471–475.

Chiswick, Barry R., *The evaluation of Census Bureau procedures for the measurement of non cash benefits and the incidence of poverty*, Washington D.C.: U.S. Department of Commerce,

Congressional Budget Office, "The Distribution of Household Income and Federal Taxes," 2010.

Cutler, David M and Jonathan Gruber, "The Effect of Medicaid Expansions on Public Insurance, Private Insurance, and Redistribution ," 1996, *The American Economic Review* (86), 2.

Czajka, John. L. and Gabrielle Denmead, "Income Data for Policy Analysis: A Comparative Assessment of Eight Surveys. Final Report.," *Mathematica Policy Research*, 2008.

Dague, Laura, "The Effect of Medicaid premiums on enrollment: A regression discontinuity approach," *Journal of Health Economics*, 2014, *37*, 1–12.

_ , Thomas DeLeire, and Lindsey Leininger, "The Effect of Public Insurance Coverage for Childless Adults on Labor Supply," *Working Paper*, 2014, pp. 1–32.

Ehrenberg, Ronald G., "Workers' Compensation, Wages, and the Risk of Injury," in Jr. John F. Burton, ed., *How taxes effect economic activity*, Ithaca:: ILR Press, 1988.

Fairlie, Robert W., Kanika Kapur, and Susan Gates, "Is employer-based health insurance a barrier to entrepreneurship?," *American Economic Review*, 2011, *30* (1), 146–162.

Finkelstein, Amy and Robin McKnight, "What Did Medicare Do (And Was It Worth It)?," *NBER Working Paper*, 2005, *11609*, 1–38.

_ , Sarah Taubman, Bill Wright, Mira Bernstein, Jonathan Gruber, Joseph P. Newhouse, Heidi Allen, Katherine Baicker, and The Oregon Health Study Group, "The Oregon Health Insurance Experiment: Evidence from the First Year," *Quarterly Journal of Economics*, 2012, *127* (3), 1057–1106.

Gabel, Jon, Roland McDevitt, and Laura Gando, "Insurance Market Participation Under Symmetric Information," *Unpublished mimeo*, 2003.

Gallen, Trevor S., "The Labor Market Effects of the Affordable Care Act," *Working Paper*, 2014.

Garthwaite, Craig, Tal Gross, and Matthew J. Notowidigo, "Public Health Insurance, Labor Supply, and Employment Lock," *NBER Working Paper*, 2013, pp. 1–67.

_ , _ , and _ , "Hospitals as Insurers of Last Resort ," 2014, *Working Paper*, 1–57.

Gouriroux, Christian and Alain Monfort, *Simulation-Based Econometric Methods*, Oxford University Press, 1997.

Government Accountability Office (United States) (GAO)

Government Accountability Office (United States) (GAO), "Medicaid: States Use of Managed Care ," 2012, *GAO-12-872R.*

Gruber, Jonathan, *Medicaid*, Vol. Means Tested Transfer Programs, University of Chicago Press,

_ **and Brigitte C. Madrian**, "Health Insurance, Labor Supply, and Job Mobility: A Critical Review of the Literature," *NBER Working Paper*, 2002, (8817), 1–58.

Hausman, Jerry, "Labor supply," in Henry J. Aaron and Joseph A. Pechman, eds., *New Perspectives in Workers Compensation*, Ithaca:: The Brookings Institution, 1988, pp. 27–72.

Holtz-Eaken, Douglas, David Joulfaian, and Harvey S. Rosen, "The Carnegie Conjecture: Some Empirical Evidence," *The Quarterly Journal of Economics*, 1993, *108* (2), 413–435.

Imbens, Guido W., Donald B. Rubin, and Bruce I. Sacerdote, "Estimating the Effect of Unearned Income on Labor Earnings, Savings and Consumption: Evidence from a Survey of Lottery Players," *The American Economic Review*, 2001, *91* (4), 778–794.

Karpoff, Jonathan M., "Public versus Private Initiative in Arctic Exploration: The Effects of Incentives and Organizational Structure," *Journal of Political Economy*, 2001, *109* (1), 38–78.

Keane, Michael and Robert Moffitt, "A structural model of multiple welfare program participation and labor supply," 1996, (1080-96).

Kimball, Miles S. and Matthew D. Shapiro, "Labor Supply: Are the Income and Substitution Effects Both Large or Both Small?," 2008, (14208).

Marton, James, Aaron Yelowitz, and Jeffery C. Talbert, "A tale of two cities? The heterogeneous impact of medicaid managed care ," 2014, *36*, 47–68.

McFadden, Daniel, "Conditional Logit Analysis of Qualitative Choice Behavior," in P. Zarembka, ed., *Frontiers in Econometrics*, New York: Academic Press, 1974, pp. 105–142.

_ , "A Method of Simulated Moments for Estimation of Discrete Response Models without Numerical Integration," *Econometrica*, 1989, *57* (5), 995–1026.

_ **and Paul A. Rudd**, "Estimation by Simulation," *Review of Economics and Statistics*, 1994, *57* (4), 591–608.

Medicaid and CHIP Payment and Access Commission (MAPAC)

Medicaid and CHIP Payment and Access Commission (MAPAC), "Report to the Congress: The Evolution of Managed Care in Medicaid ," 2011.

Meyer, Bruce D. and Dan T. Rosenbaum, "Welfare, the Earned Income Tax Credit, and the Labor Supply of Single Mothers," *Quarterly Journal of Economics*, 2001, *116* (3), 1063–1114.

_ **and James X. Sullivan**, "Winning the War: Poverty from the Great Society to the Great Recession," *NBER Working Paper*, 2013, *18718.*

_ **, Wallace K. C. Mok, and James X. Sullivan**, "The Under-Reporting of Transfers in Household Surveys: Its Nature and Consequences," *NBER Working Paper*, 2009, *15181*, 1–94.

Moffitt, Robert, "An Economic Model of Welfare Stigma," *The American Economic Review*, 1983, *73* (5), 1023–1035.

Mulligan, Casey, "Average Marginal Labor Income Tax Rates under the Affordable Care Act," *NBER Working Paper*, 2013, *19365.*

_ **and Tomas Philipson**, "Merit Motives and Government Intervention: Public Finance in Reverse," *NBER Working Paper*, 2000, *7698.*

_ **and** _ , "Insurance Market Participation Under Symmetric Information," *Unpublished mimeo*, 2003.

National Research Council, *Measuring Poverty: A New Approach*, Washington DC: The National Academies Press, 1995.

National Institute for Health Care Management Foundation (NIHCM)

National Institute for Health Care Management Foundation (NIHCM), "The Concentration of Health Care Spending," 2012.

Pakes, Ariel S, "Patents as Options: Some Estimates of the Value of Holding European Patent Stocks," *Econometrica, Econometric Society*, July 1986, *54* (4), 755–84.

Pencavel, John, "Labor supply of men: A survey," in O. Ashenfelter and R. Layard, eds., *Handbook of Labor Economics*, 1 ed., Vol. 1, Elsevier, 1987, chapter 01, pp. 3–102.

Ramey, Valerie A. and Neville Francis, "A Century of Work and Leisure," *American Economic Journal: Macroeconomics*, 2009, *1* (2), 189–224.

Sommers, Benjamin D., Meredith Roberts, Katherine Swartz, and Arnold M. Epstein, "Reasons for the wide variation in Medicaid participation rates among states hold lessons for coverage expansion in 2014," *Health Affairs*, 2012, *31* (5), 909–919.

Yelowitz, Aaron S., "The Medicaid Notch, Labor Supply, and Welfare Participation: Evidence from Eligibility Expansions," *Quarterly Journal of Economics*, 1995, *110* (4), 909–939.

Tables

Table 1: Summary Statistics for Tennessee

Source	March 2005	March 2006	Percent Change
	Medicaid Coverage		
Administrative Data	1,358,164	1,215,467	-10.5%
CPS	961,440	855,469	-10.0%
SIPP	1,324,889	1,044,420	-21.2%
	Employment		
CES (employment)	2,718,400	2,765,300	1.7%
LAUS (employment)	2,928,004	2,997,779	2.4%
QECW (employment)	2,655,631	2,712,948	2.2%
CPS (current employment status)	2,689,610	2,769,195	3.0%
CPS (any employment in year)	3,031,009	3,088,733	1.9%
CPS (yearly hours/2000)	2,822,555	2,879,721	2.0%
SIPP (linked only, any hours, normalized)	2,993,209	3,046,816	1.8%
	Wages		
QECW (weekly wage/40)	16.28	16.50	1.38%
	Property Income		
BEA: State Personal Income(annual per cap.)	$31,282.82	$32,651.21	4.37%
BEA: State Nonwage Income (annual per cap.)	$9,524.54	$10,009.39	5.09%
BEA: State Transfer Income (annual per cap.)	$5,466.24	$5,564.56	1.80%
	Inflation		
CPI (national)	187.1	193.1	3.2%
	Spending		
Medicaid Spending per enrollee	$4966	$4477	-9.83%
	Population		
Census (all ages)	5,964,308	6,056,196	1.5%
Census (18-64)	3,773,510	3,826,118	1.4%
Census (18+)	4,520,309	4,592,977	1.6%
LAU Noninstitutional	4,586,532	4,670,063	1.8%
CPS (all ages)	5,857,114	5,867,031	.2%
CPS (18-64)	3,676,413	3,684,484	.2%
SIPP (not linked, all ages)	5903724	5,878,453	-.4%
SIPP (not linked, 18-64)	3,735,685	3,730,953	-.1%

Table 1: This table displays different sources for Medicaid Coverage, Employment, income, wages, and other variables of interest using administrative, CPS, and SIPP data between March 2005 and March 2006. CPS Medicaid estimates use the Minnesota Population center's health insurance weight (hinswt), while SIPP variables use standard individual population weights. QCEW variables are quarterly measurements. Because SIPP employment counts only look at linked individuals, I reweight the linked population to sum to the full SIPP population.

Table 2: TennCare Premiums by FPL Level

Income level	Single	Family
$< 100\%$ FPL	\$0	\$0
100%-150% FPL	\$240	\$480
150%-200%	\$420	\$840
200%-250%	\$1,200	\$3,000
250%-300%	\$1,800	\$4,500
300%-350%	\$2,400	\$6,000
350%-400%	\$3,000	\$7,500
400%-500%	\$4,200	\$10,500
500%-600%	\$5,400	\$13,500
600%+	\$6,600	\$16,500

Table 2: This table depicts TennCare premiums in 2005 and 2006, by Federal Poverty Level (FPL) as described by the TennCare Bureau of Regulations. They are also graphically depicted in Figure 1. TennCare's premiums are of critical importance in both my simple example and my more detailed heterogeneous agent model. In the simple example, I linearize the premium schedule in this table for several different types of households, generating the average marginal tax rate by weighting each tax rate by the frequency of households experiencing it. For example, a household of one at the poverty line (\$9570 in 2005) experiences tax rate estimated by long-differencing premiums of 11.5% (the maximum premium, 6600, divided by an income increase of $6 \cdot 9570$), not far from the 12% a linear regression estimates. In the heterogeneous example, households directly face the cliffs in this premium schedule. The removal of an 11.5% tax rate due to disenrollment may cause large changes in labor supply.

Table 3: Diff-in-diff: Behavior of households on Medicaid Last Year

Data	2005 (pre-disenrollment)	2006 (disenrollment)
Δ Hours	1.07	1.76
Δ Medicaid coverage	-0.14	-0.60

Table 3: This table displays the central difference-in-difference (year-on-year, first quarter) data for this paper's "representative agent" results from SIPP. The pre-disenrollment baseline is measured by last year's change in Medicaid enrollment and hours for a family with Medicaid enrollees. The -0.46 drop in Medicaid coverage is valued at 14% of the family's income, when valued dollar for dollar, while while an increase in hours by 0.69 is a 2.6% increase in hours.

Table 4: Simple Estimates of Medicaid's Value

Marshallian elasticity	Income elasticity	ξ	Note
-0.10	-0.86	0.25	Low Marshellian
0.00	-0.71	0.26	Preferred (BGP)
0.16	-0.49	0.26	High Marshellian
0.39	-0.16	0.30	Largest Modern Income

Table 4: This table displays the key estimate of the cash value of \$1 in Medicaid (ξ), along with several inferential analyses. Specifically, it holds the Hicksian elasticity ϵ^H at 0.5, and adjusts the Marshallian elasticity, calculating the income elasticity via the Slutsky equation. My core result is in the second row, with my preferred Hicksian and Marshallian elasticities. Other rows offer different Marshallian elasticities, with some change in estimated cash value of Medicaid.

Table 5: Alternative Estimates of Medicaid's value

Hicksian elasticity	Income elasticity	ξ	Note
0.11	-0.16	1.16	Low Hicksian
0.30	-0.42	0.42	60% average Hicksian
0.50	-0.71	0.26	Preferred (Micro-Macro Bounds)
0.7	-1.00	0.18	Cross-country macro variation

Table 5: Like Table 4, this table displays estimates of the cash value of \$1 in Medicaid (ξ). However, it abandons the assumption that we know the Hicksian elasticity of labor supply. By shutting down both sources of explanation (wage and income) with a zero Marshallian and low Hicksian elasticity, I naturally estimate high values of Medicaid. Along with Table 4, this table indicates that either an elastic response to wages or an elastic response to income is enough to calculate low values of Medicaid.

Table 6: Percentiles of ξ Distribution

Percentile	ξ value
5%	0.14
10%	0.15
25%	0.19
50%	0.23
Mean	0.25
Mode	0.22
75%	0.29
90%	0.34
95%	0.38

Table 6: This table depicts percentiles of the estimated posterior Bayesian distribution of ξ, the cash exchange rate of Medicaid benefits, after taking uncertainty about the prior distribution of ξ, the change in labor, and the Marshallian and Hicksian elasticity into account. The estimates are fairly robust: from a place of reasonable agnosticism about ξ, a simple model of Tennessee's episode would place our estimates between \$0.14 and \$0.38 per dollar.

Table 7: Implied ξ values of Dague 2014

Functional Form	Implied yearly cash value	Implied cash value per dollar
Normal Distribution	\$551	0.08
Gumbel Distribution	\$458	0.07
Triangular Distribution	\$540	0.08
Student's t with 2 d.o.f	\$758	0.11
Exponential distribution with 0 mass point	\$196	0.03

Table 7: This table shows estimates of the cash value of Medicaid and implied exchange rate of cash to Medicaid using results from Dague (2014) and Sommers et al. (2012). Because functional form is important, several different functional forms are given along with the mean value of Medicaid to beneficiaries divided by \$6720, the average Wisconsin medicaid spending in the period. This gives the estimated cash exchange rate for Medicaid implied by Dague (2014)'s discontinuity results. The implied cash value is generated by truncating distributions are truncated at zero. The reason these distributions are chosen rather than other possibilities is that we observe a number of individuals with net cash value of Medicaid below zero. Consequently, distributional families defined only on the positive number line cannot meet my first moment unless they have a mass point at or below zero. I do estimate an exponential distribution with a mass point at zero.

Table 8: Healthcare coverage transitions

	2006			
2005	ESI	TennCare	NGI	None
ESI	66.02	1.77	5.10	27.11
TennCare	9.99	49.46	4.17	36.37
NGI	19.23	1.24	47.85	31.68
None	17.13	8.58	8.35	65.94

Table 8: All statistics are weighted using SIPP provided population weights

Table 9: Data and Simulated Moments

Moment	Simulation	Target
Average hours/population	24.98	21.12
Prop. of tot pop on Medicaid	0.25	0.23
Prop. of Medicaid beneficiaries that are poor	0.47	0.43
Prop. of Medicaid beneficiaries that are semi-poor	0.43	0.46
Prop. of Medicaid beneficiaries that are rich	0.08	0.11
Proportion of poor that are medicaid beneficiaries	0.57	0.58
Proportion of semi-poor that are medicaid beneficiaries	0.25	0.22
Proportion of rich that are medicaid beneficiaries	0.04	0.05
Proportion of total population on ESI	0.60	0.56
Proportion of total population on NGI	0.08	0.10
Proportion of total population employed	0.53	0.49
Change of total population on Medicaid	-0.02	-0.02
Change of poor population on Medicaid	0.004	-0.04
Change of semi-poor population on Medicaid	-0.04	-0.06
Change of rich population on Medicaid	-0.01	-0.01
Change in total labor provided	0.017	0.017
Proportion of Medicaid-removed going to ESI	0.49	0.44
Proportion of Medicaid-removed going to NGI	0.05	0.06
Hours of poor conditional on working	9.5	15.04
Hours of semi-poor conditional on working	32.4	28.39
Hours of rich conditional on working	42.1	35.48

Table 9: This table depicts matched data and simulated moments. Coverage is measured in average contracts per family. "Poor" families are those earning below 100% of the federal poverty line. "Semi-Poor" families are those earning between 100% and 300% of the FPL. "Rich" families are those earning above 300% of the FPL.

Table 10: Distribution of Medicaid Compensating Variation

Concept	Compensating Variation
Minimum	<$0.01
1st percentile	<$2.2
10th Percentile	<$23
25th Percentile	$93
50th Percentile	$373
Mean	$1764
75th Percentile	$1558
90th Percentile	$4904
99th Percentile	$19008
Maximum	$65407

Table 10: This table depicts the distribution of compensating variation for a Medicaid contract. Compensating variation answers how much a household must be paid to make them indifferent to Medicaid disenrollment, while the full value adds premiums and co-pays (if paid) to compensating variation. There are many households that barely value Medicaid, and receive it only because it has essentially no out of pocket cost for households who only cover children. .

Table 11: Attributing Valuation Distribution to Cost Distribution

Percentile	Proportion of Population	Proportion of Spending	Average Spending	Average Value	Medicaid Exchange Rate
0-15.4%	15.4%	0%	$0	<$0.01	N/A
15.4%-20%	4.6%	0.1%	$108	$42	0.39
20%-30%	10%	0.3%	$149	$64	0.43
30%-40%	10%	0.9%	$447	$128	0.29
40%-50%	10%	1.6%	$795	$222	0.28
50%-60%	10%	2.7%	$1341	$373	0.28
60%-70%	10%	4.8%	$2384	$638	0.27
70%-80%	10%	8.4%	$4171	$1140	0.27
80%-90%	10%	16%	$7946	$2295	0.29
90%-95%	5%	15.7%	$15593	$4904	0.31
95%-99%	4%	27.7%	$34390	$8015	0.23
99%-100%	1%	21.8%	$108260	$19008	0.18

Table 11: This table takes the the 12 percentile categories of Medical spending that NIHCM (2012) use and produces an estimate of the distribution of Medicaid spending by group. I then "line up" low valuations with low percentiles in the distribution of Medicaid spending, to get an idea of how well the distribution of valuations is explained by the distribution of spending. While average spending data is generated from separate data as average value data, the two are surprisingly comparable: typically, valuation is about 20% the value of spending. The last column displays the "Medicaid exchange rate" by percentile, which ranges from 0.18 to 0.43.

Figures

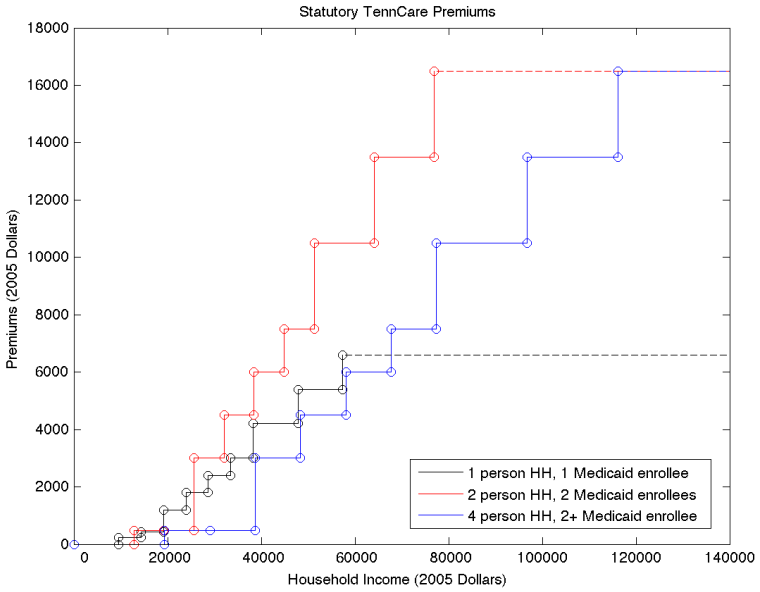

Figure 1: This figure depicts statutory TennCare premiums by income level for three types of example families. The three families are 1) A single individual on Mediciad 2) A two-person household on a family plan 3) A four-person household on a family plan (two or more members on Medicaid). For each type of family, premiums are zero until they hit the poverty line: \$9,570, \$12,830, and \$16,090 respectively. Upon reaching 100% of the FPL,premiums rise at an (approximate) 14%, 26%, and 18% implicit tax rate until \$55,860, \$74,940, and \$113,100, after which the implicit tax rate is zero again.

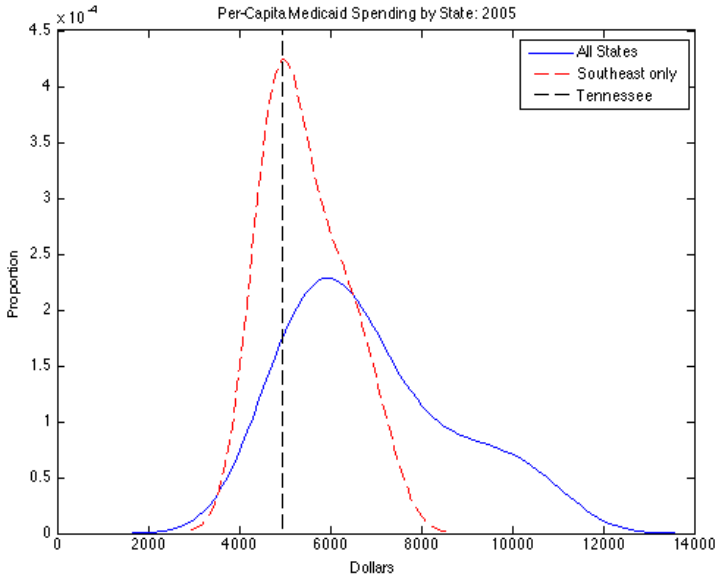

Figure 2: This figure depicts state spending per Medicaid enrollee in 2005 from the yearly Centers for Medicaid and Medicare Services data. The blue line is a kernel density estimation of per capita spending for all states. The dashed red line is the density of all Southeastern states (Alabama, Arkansas, Florida, Georgia, Kentucky, Louisiana, Mississippi, North Carolina, South Carolina, Tennessee, Virginia, and West Virginia). The dashed black line denotes Tennessee's 2005 per capita Medicaid spending. While at the 10th percentile in the nation it is nearer to the mean of Southeastern states, which have lower average spending more generally. This suggests differences in the cost of living are driving differences in spending per enrollee in Tennessee relative to the nation, rather than MCO usage.

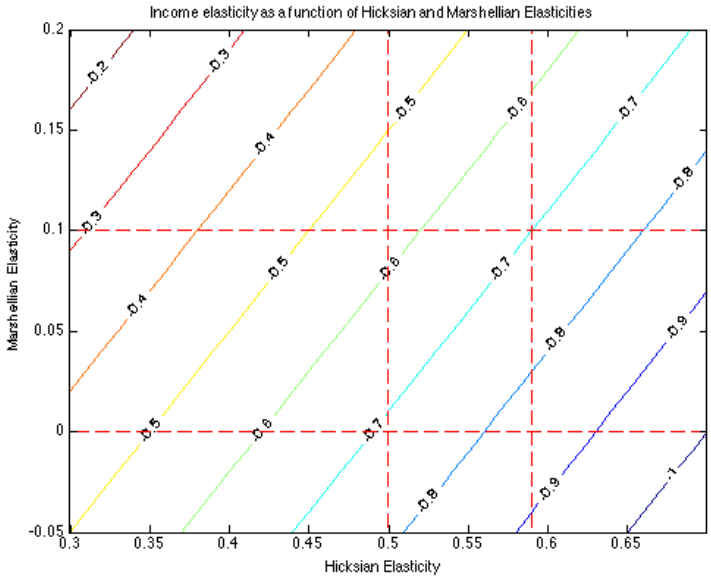

Figure 3: This figure depicts the income elasticity (which is negative) calculated from the Slutsky equation with respect to various Marshallian and Hicksian wage elasticities of labor supply, given a wage share of 0.7. Contours show combinations that give the same income elasticity. The dashed vertical red lines depict two reasonable bounds for Hicksian elasticities, while the dashed horizontal lines do the same for Marshallian elasticities.

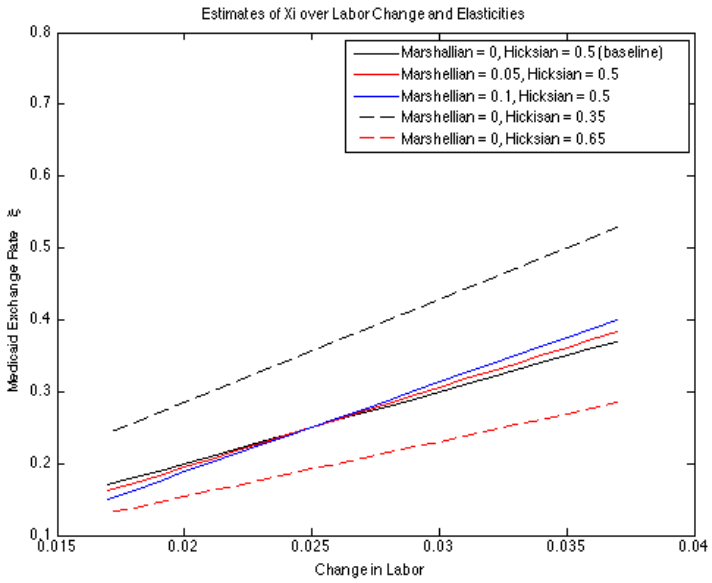

Figure 4: This figure summarizes the two most important dimensions of robustness: the degree to which labor changed, and the Hicksian elasticity. The Marshallian elasticity is unimportant for my purposes: holding the Hicksian elasticity constant, if it is raised, the income elasticity falls but so does the labor change needed to be explained by income effects. The labor change's importance matters, and is the core input into this exercise. The Hicksian elasticity matters, especially when wage effects are turned off, because it directly determines the income elasticity: for a given change in labor, as it gets smaller, the implied change in income must have been larger. Beyond simple robustness, this figure gives intuition for the relatively sharp bounds in Figure 5.

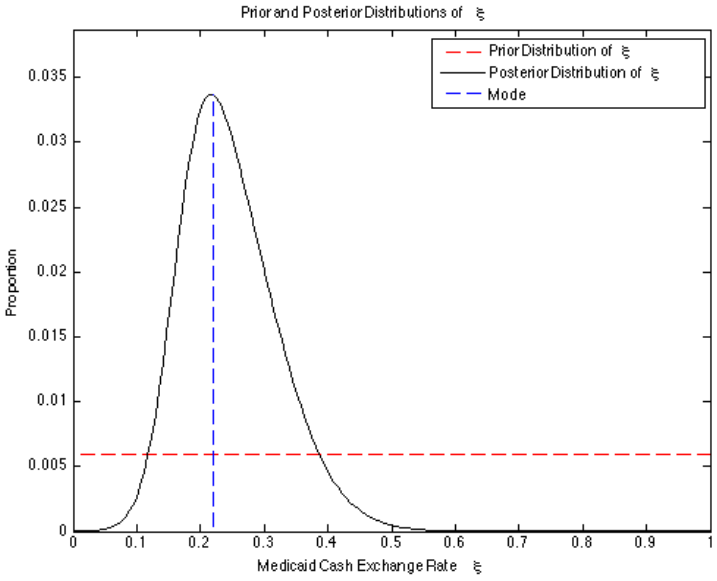

Figure 5: This figure depicts the the prior and posterior distributions of the cash exchange rate ξ given uncertainty about elasticities and change in labor supply for a simple model. It is calculated using equations (3)-(5) and Bayes rule. The modal estimate is lower than my baseline estimate because my baseline case uses a Marshallian elasticity of zero, while Bayes rule is typically integrating over higher values for the Marshallian elasticity. The reason this figure cannot produce Medicaid exchange rates much above 0.5 is depicted graphically in Figure 4.

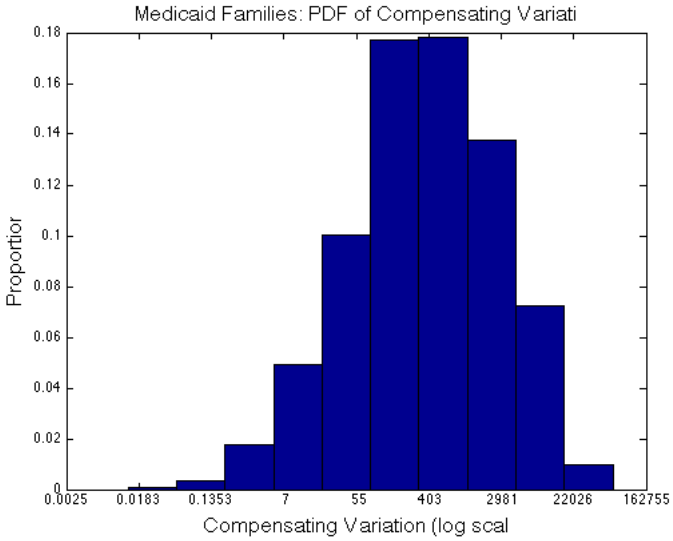

Figure 6: This figure displays the distribution of compensating variation (in logs) in the population of Medicaid recipients. It answers the question "how much would a Medicaid patient be willing to accept in cash to disenroll their entire family from Medicaid?" The top panel displays the probability density function of log-EV. Values displayed on the x-axis are in dollars, though the x-axis is transformed by taking the natural log of EV.

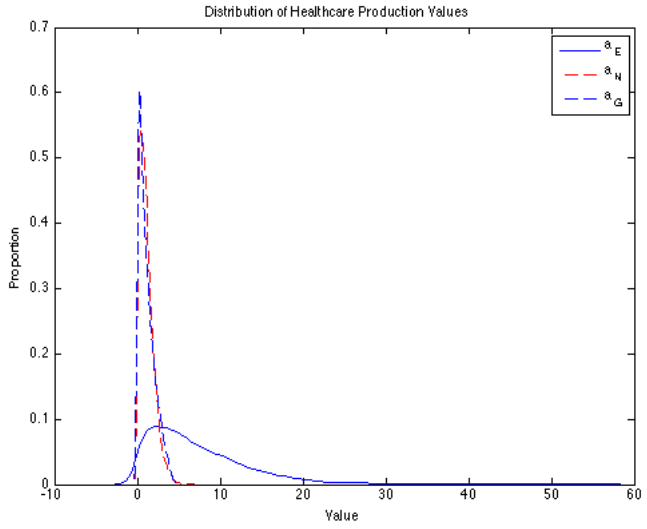

Figure 7: This figure displays the distributions of a_E, a_N, and a_G. It makes clear that while there are few people who have low valuations of ESI, many have low valuations of NGI and government healthcare. It additionally indicates that NGI and ESI are fairly comparable in their distributions of valuation.

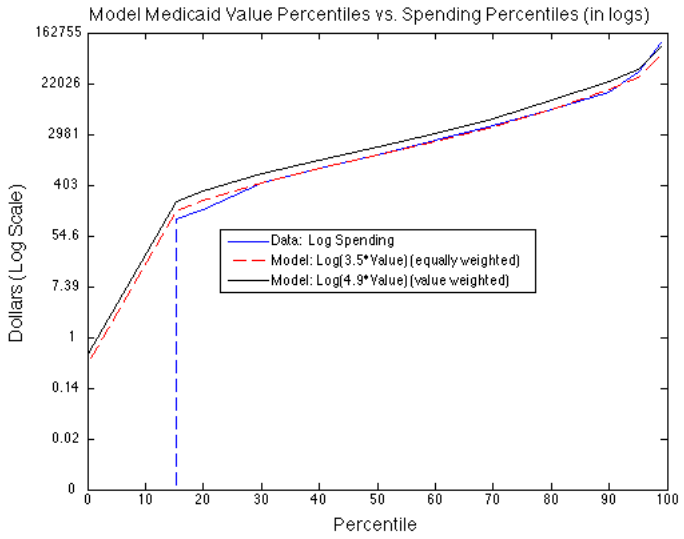

Figure 8: This figure displays the cumulative compensating variation valuation for Medicaid, along with the cumulative distribution of payments generated from Tennessee's Medicaid spending and NIHCM (2012). Importantly, the "model" results were generated without any use of expenditure data. The blue values are generated by applying NIHCM (2012) to Tennessee's spending. The red and black values are generated by the model using a simple count of "contracts" and Tennessee's FPL-subsidized premiums. The red values are shifted by a factor of 3.5, a "best-fit" to the expenditure line equally weighting individuals. The black values are shifted by a factor of 4.9, reflecting a value-weighted fitting. Both CDFs line up, suggesting the distribution of valuations and distributions of expenditures differ only by a constant, suggesting that the distribution of valuations may be explained nearly entirely by the distribution of costs.

Appendix Tables

Table A1: Healthcare coverage by property income (2005)

	No property income	Positive property income
ESI	0.37	0.70
NGI	0.07	0.10
TennCare	0.38	0.12
None	0.19	0.08

Table A1: This table depicts the frequency of healthcare coverage by family property income. This table reinforces the wage-based difference in healthcare coverage targeted in Table 9. All statistics are weighted using SIPP provided population weights.

Table A2: Medicaid Coverage by Federal Poverty Line

FPL	2005	2006	%Δ
$< 100\%$	61.3%	58.6%	-4.4%
$> 100\%$ & $< 300\%$	25.5%	17.2%	-32.5%
$> 300\%$	5.9%	4.5%	-23.7%

Table A2: This table shows Medicaid coverage by FPL in 2005 and 2006, showing that the wealthier groups showed large declines in Medicaid participation, while those below 100% of the FPL saw a much smaller decline. All statistics are weighted using SIPP provided population weights.

Table A3: Mean age by healthcare coverage (2005)

ESI	37.06
NGI	51.56
TennCare	30.99
None	31.61

Table A3: This figure depicts the average age of those covered by various healthcare types. While my model does not take into account age differences, this table notes that TennCare participants are typically younger than ESI participants, who in turn are younger than those that use non-group insurance. All statistics are weighted using SIPP provided population weights.

Table A4: TennCare coverage change, adults vs. children

	2005	2006
Adult	17.67%	12.07%
Child	34.27%	36.52%

Table A4: All statistics are weighted using SIPP provided population weights.

Table A5: Diff-in-diff: Behavior of households on Medicaid: 2005-2007

Data	Not Tennessee	Tennessee
Change in Hours for households	0.74	1.01
Change in Medicaid coverage	-.30	-.23

Table A5: This table depicts an alternative difference-in-difference model for the representative agent model. Rather than differencing changes in time on changes in time, this table depicts changes in time for different states. Comparing Tennessee to the rest of the East South Central Region (KY, MS, AL), it had a high increase in hours even as it had a lower decrease in Medicaid contracts over the course of two years. Over two years, Tennessee's change in labor hours was more mild than that of surrounding states, while their Medicaid disenrollment was larger.

Table A6: SIPP sample summary statistics

	March, 2005			March, 2006		
	mean	median	std. dev	mean	median	std. dev
Family size	2.42	2	1.30	2	2.37	1.28
Age	36.33	36	21.90	36.95	36	22.01
Wage ($)	15.23	14.15	4.03	15.55	13.92	4.78
Weekly hours worked	33.18	40	19.57	40	32.21	19.46
Proportion employed	82.02	-	-	79.08	-	-
HH property income ($)	96.26	2	443.38	118.58	1	634.97
Education (%)						
Less than high school		24.48			21.24	
High school diploma		26.06			28.46	
Some college		29.58			31.61	
College diploma or more		19.88			18.69	
N		2499			2328	

Table A6: All statistics are weighted using SIPP provided population weights. Wages, hours, and labor market participation are computed only for individuals age 25 to 55. Wages are imputed for those who do not work using a Heckman selection procedure described in the text.

Table A7: Healthcare coverage choices

	2005	2006
ESI	0.60	0.57
NGI	0.10	0.12
TennCare	0.22	0.18
None	0.08	0.13
N	2499	2328

Table A7: All statistics are weighted using SIPP provided population weights.