The Value of Medicaid: Interpreting Results from the Oregon Health Insurance Experiment

Amy Finkelstein, Nathaniel Hendren, and Erzo F.P. Luttmer*

November 2016

Abstract

We develop a set of frameworks for valuing Medicaid and apply them to welfare analysis of the Oregon Health Insurance Experiment, a Medicaid expansion for low-income, uninsured adults that occurred via random assignment. We estimate that the value of Medicaid to recipients is roughly between one-third and three-quarters of Medicaid's monetary transfers to the external parties who provide partial implicit insurance for the low-income uninsured. Medicaid provides value to recipients through both its expected transfer of resources and its insurance function of moving resources across states of the world. Across approaches, the insurance value to recipients varies considerably, but the transfer value to recipients is stable and always substantial relative to the insurance value. Whether or not the value of Medicaid to recipients exceeds its net (of monetary transfers to external parties) costs depends on the approach used.

1 Introduction

Medicaid is the largest means-tested program in the United States. In 2015, public expenditures on Medicaid were over \$550 billion, compared to about \$70 billion for food stamps (SNAP), \$70 billion for the Earned Income Tax Credit (EITC), \$60 billion for Supplemental Security Income (SSI), and \$30 billion for cash welfare (TANF).[1]

What are the welfare benefits of this large in-kind program? How much is Medicaid valued by recipients? How does this value to recipients compare to the cost of Medicaid or to the monetary transfers Medicaid provides to third parties who, in the absence of Medicaid, implicitly bear some of the costs of covering the low-income uninsured?

[1]See Department of Health and Human Services (2015, 2016)[69, 70], Department of Agriculture (2016)[68], Internal Revenue Service (2015)[71], and Social Security Administration (2016)[72]).

Such empirical welfare questions have received very little attention. Although there is a voluminous academic literature studying the reduced-form impacts of Medicaid on a variety of potentially welfare-relevant outcomes – including health care use, health, financial security, labor supply, and private health insurance coverage – there has been little formal attempt to translate such estimates into statements about welfare. Absent other guidance, academic or public policy analyses often either ignore the value of Medicaid – for example, in the calculation of the poverty line or measurement of income inequality (Gottschalk and Smeeding (1997)[38]) – or makes fairly ad hoc assumptions. For example, the Congressional Budget Office (2012)[67] values Medicaid at the average government expenditure per recipient. In practice, an in-kind benefit like Medicaid may be valued at less, or at more, than expenditures on it (see, e.g., Currie and Gahvari (2008)[23]).

Recently, the 2008 Oregon Health Insurance Experiment provided estimates from a randomized evaluation of the impact of Medicaid coverage for low-income, uninsured adults on a range of potentially welfare-relevant outcomes. The main findings were: In its first one to two years, Medicaid increased health care use across the board – including outpatient care, preventive care, prescription drugs, hospital admissions, and emergency room visits; Medicaid improved self-reported health, and reduced depression, but had no statistically significant impact on mortality or physical health measures; Medicaid reduced the risk of large out-of-pocket medical expenditures; and Medicaid had no economically or statistically significant impact on employment and earnings, or on private health insurance coverage.[2] These results have attracted considerable attention. But in the absence of any formal welfare analysis, it has been left to partisans and media pundits to opine (with varying conclusions) on the welfare implications of these findings.[3]

Can we do better? Empirical welfare analysis is challenging when the good in question – in this case public health insurance for low-income individuals – is not traded in a well-functioning market. This prevents welfare analysis based on estimates of ex-ante willingness to pay derived from contract choices, as is becoming commonplace where private health insurance markets exist (Einav, Finkelstein, and Levin (2010)[28] provide a review). Instead, one encounters the classic problem of valuing goods when prices are not observed (Samuelson (1954)[62]).

In this paper, we develop two main analytical frameworks for empirically estimating the welfare value of Medicaid to recipients, and apply them to the results from the Oregon Health Insurance Experiment. Our first approach, which we refer to as the "complete-information" approach, requires complete specification of a normative utility function and estimates of the causal effect of Medicaid

---

[2]For more detail on these results, as well as on the experiment and affected population, see Finkelstein et al. (2012)[34], Baicker et al. (2013)[10], Taubman et al. (2014)[66], and Baicker et al. (2014)[8].

[3]The results of the Oregon Health Insurance Experiment have received extensive media coverage, but the media drew a wide variety of conclusions as the following two headlines illustrate: "Medicaid Makes 'Big Difference' in Lives, Study Finds" (National Public Radio, 2011, http://www.npr.org/2011/07/07/137658189/medicaid-makes-big-difference-in-lives-study-finds) versus "Spending on Medicaid doesn't actually help the poor" (Washington Post, 2013, http://www.washingtonpost.com/blogs/right-turn/wp/2013/05/02/spending-on-medicaid-doesnt-actually-help-the-poor/). Public policy analyses have drawn similarly disparate conclusions: "Oregon's lesson to the nation: Medicaid Works" (Oregon Center for Public Policy, 2013, http://www.ocpp.org/2013/05/04/blog20130504-oregon-lesson-nation-medicaid-works/) versus "Oregon Medicaid Study Shows Michigan Medicaid Expansion Not Worth the Cost" (MacKinac Center for Public Policy, 2013, http://www.mackinac.org/18605).

on the distribution of all arguments of the utility function. A key advantage of this approach is that it does not require us to model the precise budget set created by Medicaid or impose that individuals optimally consume medical care subject to this budget constraint. However, as the name implies, the information requirements are high; it will fail to accurately measure the value of Medicaid unless the impacts of Medicaid on *all* arguments of the utility function are specified and analyzed. In our application, for example, we specify a utility function over non-health consumption and health, and limit our empirical analysis to estimates of the impact of Medicaid on the distribution of these arguments. In principle, however, the approach requires estimates of the impact of Medicaid on, and the value of, any utility arguments that a creative reader or referee could deem plausibly affected by the program, such as future consumption, marital stability, or outcomes of the recipient's children.

Our second approach, which we refer to as the "optimization" approach, is in the spirit of the "sufficient statistics" approach described by Chetty (2009)[19], and is the mirror image of the complete-information approach in terms of its strengths and weaknesses. By parameterizing the way in which Medicaid affects the individual's budget set, and by assuming that individuals make optimal choices with respect to that budget set, we significantly reduce the implementation requirements. In particular, it suffices to specify the marginal utility function over any single argument. This is because the optimizing individual's first-order condition allows us to value – through the marginal utility of that single argument – marginal impacts of Medicaid on any other potential arguments of the utility function.

We develop two versions of the optimization approach. The "consumption-based optimization approach" values Medicaid's marginal relaxation of the recipient's budget constraint using its co-variance with the marginal utility of consumption; insurance is valuable if it transfers resources from low to high marginal utility of consumption states of the world. The "health-based optimization approach" values a marginal relaxation of the budget constraint using its covariance with the marginal utility of out-of-pocket medical spending; insurance is valuable if it transfers resources from states of the world where the marginal health returns to out-of-pocket spending are low to states where those returns are high. To use these approaches to make inferences about non-marginal changes in an individual's budget set (i.e., covering an uninsured individual with Medicaid), we require an additional statistical assumption that allows us to interpolate between local estimates of the marginal impact of program generosity. This assumption substitutes for the economic assumptions about the utility function in the complete-information approach.

Both approaches provide estimates of the welfare value of Medicaid to recipients, and allow us to decompose this estimate into the value arising from a "transfer component" in which recipients are transferred resources via the free public provision of a good, and the value arising from the (budget-neutral) "pure-insurance component", stemming from Medicaid's ability to move resources across states of the world. Estimating the value of the transfer component to recipients involves a relatively straightforward mapping from empirical quantities; the modeling choices primarily influence our estimates of the insurance component of the value of Medicaid to recipients. We also estimate the impact of Medicaid on government spending and on monetary transfers to providers of

partial, implicit insurance to the "uninsured" (hereafter "external parties"). These estimates provide useful context for interpreting our welfare estimates of the value of Medicaid to recipients.

We implement these approaches for welfare analysis of the Medicaid coverage provided by the Oregon Health Insurance Experiment. We use the lottery's random selection as an instrument for Medicaid coverage in order to estimate the impact of Medicaid on the required objects. We use data from study participants to directly measure out-of-pocket medical spending, health care utilization, and health. Our baseline health measure is self-assessed health, which we value using existing estimates of the quality of life years (QALYs) associated with different levels of self-assessed health and an assumed value of a statistical life year (VSLY); we also report estimates based on alternative health measures - such as self-reported physical and mental health, or a depression screen - combined with existing estimates of their associated QALYs. Absent a consumption survey in the Oregon context, we proxy for consumption by the difference between average consumption for a low-income uninsured population and out-of-pocket medical expenditures reported by study participants, subject to a consumption floor. We also implement an alternative version of the consumption-based optimization approach which uses consumption data for a low-income sample in the Consumer Expenditure Survey.

Across the various approaches and specifications, we find that external parties who would otherwise cover some of the health care costs of the low-income uninsured are a major beneficiary of Medicaid. In fact, our baseline estimates indicate that the main beneficiaries of Medicaid are not the recipients themselves, but rather the external parties. Depending on the approach, our baseline estimates indicate that the value of Medicaid to recipients is roughly between one-third and three-quarters of the monetary transfers to these external parties. As a result, if (counterfactually) Medicaid recipients had to pay the government's average cost of Medicaid, we estimate that they would rather be uninsured; specifically, we estimate a welfare benefit to recipients per (gross) dollar of government spending of between $0.2 and $0.5. The large monetary transfers to external parties arise because - in both our data and in other national data sets - the low-income nominally "uninsured" in fact pay only a small share of their medical expenses; as a result, we estimate that 60 cents of every dollar of government spending on Medicaid represents a monetary transfer to external parties.

A distinct, important question is whether the value of Medicaid to recipients exceeds the *net* (of monetary transfers to external parties) resource cost of Medicaid. A priori this is not obvious, and our different approaches reach different conclusions. Because of potential market failures, such as adverse selection, the value of the "pure-insurance" component of Medicaid could exceed the additional resource cost of providing that insurance. However, the value of Medicaid's transfer component to recipients may be less than its net cost if part of the transfer value stems from the moral hazard response (i.e., induced medical spending) to Medicaid. Depending on the approach, we estimate that Medicaid's welfare benefit to recipients per *net* dollar of spending ranges from $0.5 to $1.2; an estimate below $1 suggests that the recipient is not willing to pay the net cost of Medicaid coverage or, in other words, that the insurance value Medicaid provides by moving

resources across states of the world does not exceed its moral hazard costs. We estimate that much of the source of Medicaid's value to recipients comes from the transfer component; depending on the approach used, between 40 and 95 percent of the value of Medicaid to recipients reflects this transfer value, rather than the value of the (budget neutral) insurance product.

Naturally, all of our quantitative results are sensitive to the framework used and to our specific implementation assumptions. We explored sensitivity to a variety of alternative assumptions. Our estimates of the value of the "pure-insurance" component of Medicaid are particularly sensitive, while the value of the "transfer component" to recipients is relatively more stable across approaches and assumptions. However, two primary findings are qualitatively robust across a wide number of alternative specifications: (i) the magnitude of the monetary transfer from Medicaid to external parties is important relative to the value to recipients and (ii) the transfer value to recipients is always substantial relative to its insurance value to recipients. We discuss which modeling assumptions, features of the data, and parameter calibrations are quantitatively most important for the results.

How seriously should our empirical welfare estimates be taken? We leave it to the readers to make up their own minds about the credibility of the welfare estimates. One thing that seems hard to disagree with is that some attempt – or combination of attempts – allows for a more informed posterior of the value of Medicaid than the implicit default of treating the value of Medicaid at zero or simply at gross cost, which occurs in so much existing work. Naturally, our empirical estimates are specific to a particular Medicaid program in Oregon and the people for whom the lottery affected Medicaid coverage. Fortunately, the frameworks we develop can be readily applied to welfare analysis of other public health insurance programs, such as Medicaid coverage for other populations or Medicare coverage.

Importantly, our estimates for the impact on Medicaid beneficiaries only speak to the *recipient's* value of Medicaid. An estimate of the *social* value of Medicaid would need to take account of the social value of any redistribution that occurs through Medicaid. Redistribution generally involves net resource costs that exceed the value to the recipient (Okun 1975 [57]). Accounting for this can be done by weighting the value to recipients by the social marginal utility of income for this group, as in Saez and Stantcheva (2016)[61]. Alternatively, the value to recipients per dollar of net costs can be compared to that of other programs such as the Earned Income Tax Credit that redistribute to a similar group of recipients, as in Hendren (2014)[43].

Our analysis complements other efforts to elicit a value of Medicaid to recipients through quasi-experimental variation in premiums (Dague (2014)[25]) or the extent to which individuals distort their labor earnings in order to become eligible for Medicaid (Gallen (2014)[36], Keane and Moffitt (1998)[47]). These alternative approaches require their own, different sets of assumptions. Interestingly, they yield similar results to our approaches here concerning the relatively low value of Medicaid to recipients relative to its (gross) cost to the government. Yet they do not generally estimate the monetary transfers to external parties or compare recipient value to net costs, or to these monetary transfers.

Our results suggest that a key driving factor behind the relatively low value of Medicaid to recipients compared to gross Medicaid costs is that much of the government spending on Medicaid goes to compensating external parties that would have borne much of the medical costs of the "uninsured" in the absence of formal insurance. This finding complements related empirical work documenting the presence of implicit insurance for the uninsured (Mahoney, 2015)[52] and the role of formal insurance coverage in reducing the provision of uncompensated care by hospitals (Garthwaite et al. (2015)[37] and unpaid medical bills by patients (Dobkin et al., 2016)[26]. However, we know of no prior systematic efforts to estimate and compare the value of Medicaid to recipients and the monetary transfers to external parties in the same context. Given the size of these external monetary transfers relative to Medicaid's value to recipients, our findings suggest that important areas for further work are identifying the ultimate economic incidence and value of these external monetary transfers, and considering the relative efficiency of formal public insurance through Medicaid compared to the previously existing informal insurance system.

The rest of the paper proceeds as follows. Section 2 develops the two theoretical frameworks for welfare analysis. Section 3 describes how we implement these frameworks for welfare analysis of the impact of the Medicaid expansion that occurred via lottery in Oregon. Section 4 presents the results of that welfare analysis. Section 5 provides several benchmarks for interpreting them, and Section 6 explores their sensitivity. The last section concludes.

2 Frameworks for Welfare Analysis

2.1 A simple model of individual utility

Individual welfare is derived from the consumption of non-medical goods and services, c, and from health, h, according to the utility function:

$$u = u\left(c, h\right). \tag{1}$$

We assume health is produced according to:

$$h = \tilde{h}\left(m; \theta\right), \tag{2}$$

where m denotes the consumption of medical care and θ is an underlying state variable for the individual which includes, among other things, medical conditions and other factors affecting health, and the productivity of medical spending. This framework is similar to Cardon and Hendel (2001) [18] who model the value of insurance using a utility function over consumption goods and health, where health is affected by a health shock and medical spending. We normalize the resource costs of m and c to unity so that m represents the true resource cost of medical care. For the sake of brevity, we will refer to m as "medical spending" and c as "consumption."

We conduct our welfare analysis assuming that every potential Medicaid recipient faces the same distribution of θ. Conceptually, we think of our welfare analysis as conducted from behind

the veil of ignorance. Empirically, we will use the cross sectional distribution of outcomes across individuals to capture the different potential states of the world, θ.

We denote the presence of Medicaid by the variable q, with $q = 1$ indicating that the individual is covered by Medicaid ("insured") and $q = 0$ denoting not being covered by Medicaid ("uninsured"). Consumption, medical spending, and health outcomes depend both on Medicaid status, q, and the underlying state of the world, θ; this dependence is denoted by $c(q; \theta)$, $m(q; \theta)$ and $h(q; \theta) \equiv \tilde{h}(m(q; \theta); \theta)$, respectively.[4]

We define $\gamma(1)$ as the value of Medicaid to a recipient, and find $\gamma(1)$ as the implicit solution to:

$$E\left[u\left(c\left(0; \theta\right), h\left(0; \theta\right)\right)\right] = E\left[u\left(c\left(1; \theta\right) - \gamma(1), h\left(1; \theta\right)\right)\right], \tag{3}$$

where the expectations are taken with respect to the possible states of the world, θ. Thus, $\gamma(1)$ is the amount of consumption that the individual would need to give up in the world with Medicaid that would leave her at the same level of expected utility as in the world without Medicaid.[5] Our focus is on empirically estimating $\gamma(1)$.

We emphasize that $\gamma(1)$ measures welfare from the perspective of the individual recipient. A social welfare perspective would also account for the fact that Medicaid benefits a low-income group. Saez and Stantcheva (2016)[61] show that in general this can be accomplished by scaling the individual valuation by a social marginal welfare weight, or the social marginal utility of income.

2.2 Complete-information approach

In the complete-information approach, we specify the normative utility function over all its arguments and require that we can observe all the arguments of this utility function both with insurance and without insurance. It is then straightforward to solve equation (3) for $\gamma(1)$.

We assume that the utility function takes the following form:

Assumption 1. *Full utility specification for the complete-information approach.*

The utility function has the following form:

$$u(c, h) = \frac{c^{1-\sigma}}{1-\sigma} + \tilde{\phi} h,$$

where σ denotes the coefficient of relative risk aversion and $\phi = \tilde{\phi}/E[c^{-\sigma}]$ denotes the marginal value of health in units of consumption.

---

[4]We assume that q affects health only through its effect on medical spending. This rules out an impact of insurance, q, on non-medical health investments as in Ehrlich and Becker (1972)[27].

[5]Note that $\gamma(1)$ is measured in terms of consumption rather than income, and is therefore not necessarily interpretable as "willingness to pay". However, if we also assume (a) individual optimization and (b) an income elasticity of demand for h of zero when individuals face a zero price for medical care (as is the case at $q = 1$ in our baseline specification), then $\gamma(1)$ is interpretable as "willingness to pay". Specifically, $\gamma(1)$ corresponds to the compensating variation for gaining Medicaid from the perspective of the uninsured and the equivalent variation for losing Medicaid from the perspective of the insured. Because of the well-known transitivity property of equivalent variation, it can then be compared to other policies targeted to the insured.

Utility has two additive components: a standard CRRA function in consumption c with a coefficient of relative risk aversion of σ, and a linear term in h.

With this assumption, equation (3) becomes, for $q = 1$:

$$E\left[\frac{c\left(0;\theta\right)^{1-\sigma}}{1-\sigma} + \tilde{\phi}h\left(0;\theta\right)\right] = E\left[\frac{\left(c\left(1;\theta\right) - \gamma(1)\right)^{1-\sigma}}{1-\sigma} + \tilde{\phi}h\left(1;\theta\right)\right]. \tag{4}$$

We use equation (4) to solve for $\gamma(1)$. This requires observing the distributions of consumption and expected health outcomes that occur if the individual were on Medicaid ($c\left(1;\theta\right)$ and $E[h\left(1;\theta\right)]$) and if he were not ($c\left(0;\theta\right)$ and $E[h\left(0;\theta\right)]$). One of these is naturally counterfactual. We are therefore in the familiar territory of estimating the distribution of "potential outcomes" under treatment and control (e.g., Angrist and Pischke (2009) [5]).[6]

We can decompose $\gamma(1)$ into two economically distinct components: the value of Medicaid to recipients comes from both average increases in resources for the individual and from a (budget-neutral) better allocation of those resources across states of the world. We refer to these throughout as, respectively, the "transfer component" and the "pure-insurance component" of the value of Medicaid to recipients.

The transfer term, denoted by T, measures the value to Medicaid recipients of receiving the expected consumption and medical spending under Medicaid rather than receiving the expected consumption and medical spending in the uninsured state; in other words, it represents the value of Medicaid to recipients arising from its role as a transfer program. The transfer term is given by the solution to the equation:

$$\frac{E\left[c\left(0;\theta\right)\right]^{1-\sigma}}{1-\sigma} + \tilde{\phi}E\left[\tilde{h}\left(E[m(0;\theta)];\theta\right)\right] = \frac{\left(E\left[c\left(1;\theta\right)\right] - T\right)^{1-\sigma}}{1-\sigma} + \tilde{\phi}E\left[\tilde{h}\left(E[m(1;\theta)];\theta\right)\right]. \tag{5}$$

Approximating the health improvement $E\left[\tilde{h}\left(E[m(1;\theta)];\theta\right) - \tilde{h}\left(E[m(0;\theta)];\theta\right)\right]$ by $E\left[\frac{d\tilde{h}}{dm}\right]E\left[m(1;\theta) - m(0;\theta)\right]$, we implement the calculation of T as the solution to:

$$\frac{E\left[c\left(0;\theta\right)\right]^{1-\sigma}}{1-\sigma} - \frac{\left(E\left[c\left(1;\theta\right)\right] - T\right)^{1-\sigma}}{1-\sigma} = \tilde{\phi}E\left[\frac{d\tilde{h}}{dm}\right]E\left[m(1;\theta) - m(0;\theta)\right].$$

Evaluating this equation requires an estimate of $E\left[\frac{d\tilde{h}}{dm}\right]$, the slope of the health production function between $m(1;\theta)$ and $m(0;\theta)$, averaged over all states of the world. We estimate $\frac{d\tilde{h}}{dm}$ using an approach described in Section 4.2.2 below. This expression shows that Medicaid spending that increases consumption (c) increases T dollar-for-dollar; however, increases in medical spending (m)

---

[6]Our particular specification of the utility function affects the set of potential outcomes we need to estimate. The additivity of utility from consumption and health allows us to estimate the marginal consumption and marginal health distributions under each insurance status. With complementarities, such as estimated in Finkelstein et al. (2013) [32], we would need to estimate the causal effect of insurance on joint distributions. The linearity assumption in h allows us to restrict our health estimation to average health under each insurance status. Because we allow for curvature in utility over consumption – to reflect the fact that individuals are risk averse – we must estimate the distribution of consumption under each insurance status.

due to Medicaid may increase T by more or less than a dollar depending on the health returns to medical spending as described by the health production function, $\tilde{h}(m;\theta)$.[7]

The pure-insurance term, denoted by I, is given by:

$$I = \gamma(1) - T. \qquad (6)$$

The pure-insurance term measures the value of Medicaid that results from the (budget-neutral) reallocation of a given amount of resources across different states of the world. The pure-insurance value will be positive if Medicaid moves resources towards states of the world with a higher marginal utility of consumption and a higher health return to medical spending.

2.3 Optimization approaches

In the optimization approaches, we reduce the implementation requirements of the complete-information approach through two additional economic assumptions: We assume that Medicaid only affects individuals through its impact on their budget constraint, and we assume individual optimizing behavior. These two assumptions allow us to replace the full specification of the utility function (Assumption 1) by a partial specification of the utility function.

Assumption 2. *(Program structure) We model the Medicaid program q as affecting the individual solely through its impact on the out-of-pocket price for medical care $p(q)$.*

Importantly, this assumption rules out other ways in which Medicaid might affect c or h, such as through direct effects on provider behavior (e.g., an effect of Medicaid on a provider's willingness to treat a patient or how the provider treats that patient).

For implementation purposes, we assume the out-of-pocket price of medical care $p(q)$ is constant in m although, in principle, one could extend the analysis by allowing for a nonlinear price schedule. We denote out-of-pocket spending on medical care by:

$$x(q,m) \equiv p(q)m. \qquad (7)$$

We allow for implicit insurance for the uninsured by not requiring that those without Medicaid pay all their medical expenses out of pocket (i.e., we do not impose that $p(0) = 1$).

Assumption 3. *Individuals choose m and c optimally, subject to their budget constraint.*
Individuals solve:

$$\max_{c,m} u\left(c, \tilde{h}(m;\theta)\right) \; subject \; to \; c = y(\theta) - x(q,m) \quad \forall m, q, \theta.$$

We let $y(\theta)$ denote (potentially state-contingent) resources.

---

[7]By the standard logic of moral hazard, if consumers optimally choose m, they would value the increase in health arising from the Medicaid-induced medical spending at less than its cost, since they chose not to purchase that medical spending at an unsubsidized price. Note, however, that we have not (yet) imposed consumer optimization.

The assumption that the choices of c and m are individually optimal is a nontrivial assumption in the context of health care where decisions are often taken jointly with other agents (e.g., doctors) who may have different objectives (Arrow (1963)[6]) and where the complex nature of the decision problem may generate individually suboptimal decisions (Baicker, Mullainathan, and Schwartzstein (2015)[9]). In particular, Baicker, Mullainathan, and Schwartzstein (2015)[9] highlight certain types of care - including preventive care - as examples of care that individuals undervalue. The fact that Medicaid increases use of preventive care (Finkelstein et al. (2012)[34]) could call into question the assumption of the optimization approach that individuals equalize marginal utilities across health and consumption.

Thought experiment: marginal expansion in Medicaid. To make further progress valuing Medicaid – and to invoke the envelope theorem, which applies given Assumption 3 – it is useful to consider the thought experiment of a "marginal" expansion in Medicaid and thus consider $q \in [0, 1]$. In this thought experiment, q indexes a linear coinsurance term between no Medicaid ($q = 0$) and "full" Medicaid ($q = 1$), so that we can define $p(q) \equiv qp(1) + (1 - q)p(0)$. Out-of-pocket spending in equation (7) is now:

$$x(q, m) = qp(1)m + (1 - q)p(0)m. \tag{8}$$

A marginal expansion of Medicaid (i.e., a marginal increase in q), relaxes the individual's budget constraint by $-\frac{\partial x}{\partial q}$:

$$-\frac{\partial x(q, m(q; \theta))}{\partial q} = (p(0) - p(1))m(q; \theta). \tag{9}$$

The marginal relaxation of the budget constraint is thus the marginal reduction in out-of-pocket spending at the current level of m. It therefore depends on medical spending at q, $m(q; \theta)$, and the price reduction from moving from no insurance to Medicaid, $(p(0) - p(1))$. Note that $-\frac{\partial x}{\partial q}$ is a program parameter that holds behavior constant (i.e., it is calculated as a partial derivative, holding m constant).

We define $\gamma(q)$ – in parallel fashion to $\gamma(1)$ in equation (3) – as the amount of consumption the individual would need to give up in a world with q insurance that would leave her at the same level of expected utility as with $q = 0$:

$$E\left[u\left(c\left(0; \theta\right), h\left(0; \theta\right)\right)\right] = E\left[u\left(c\left(q; \theta\right) - \gamma(q), h\left(q; \theta\right)\right)\right]. \tag{10}$$

2.3.1 Consumption-based optimization approach

If individuals choose c and m to optimize their utility function subject to their budget constraint (Assumptions 2 and 3), the marginal welfare impact of insurance on recipients $\frac{d\gamma}{dq}$ follows from applying the envelope theorem to equation (10):

$$\frac{d\gamma}{dq} = E\left[\frac{u_c}{E[u_c]}\left((p(0) - p(1))m(q; \theta)\right)\right], \tag{11}$$

where u_c denotes the partial derivative of utility with respect to consumption. Appendix A.1 provides the derivation. Due to the envelope theorem, the optimization approaches do not require us to estimate how the individual allocates the marginal relaxation of the budget constraint between increased consumption and health. Intuitively, because the individual chooses consumption and health optimally (Assumption 3), a marginal reallocation between consumption and health has no first-order effect on the individual's welfare.

The representation in equation (11), which we call the *"consumption-based optimization approach,"* uses the marginal utility of consumption to place a value on the relaxation of the budget constraint in each state of the world. In particular, $\frac{u_c}{E[u_c]}$ measures the value of money in the current state of the world relative to its average value, and $((p(0) - p(1))m(q; \theta))$ measures how much a marginal expansion in Medicaid relaxes the individual's budget constraint in the current state of the world. A marginal increase in Medicaid benefits delivers greater value if it moves more resources into states of the world, θ, with a higher marginal utility of consumption (e.g., states of the world with larger medical bills, and thus lower consumption). As we discuss in Appendix A.1, nothing in this approach precludes individuals from being at a corner with respect to their choice of medical spending.[8]

We can decompose the marginal value of Medicaid to recipients in equation (11) into a transfer term (T) and a pure-insurance term (I). The decomposition is:

$$\frac{d\gamma(q)}{dq} = \underbrace{(p(0) - p(1))E[m(q; \theta)]}_{\text{Transfer Term}} + \underbrace{Cov\left[\frac{u_c}{E[u_c]}, (p(0) - p(1))m(q; \theta)\right]}_{\substack{\text{Pure-Insurance Term} \\ \text{(consumption valuation)}}}. \qquad (12)$$

Although implemented differently, the transfer and pure-insurance term are conceptually the same as in the complete-information approach above. The transfer term measures the recipients' valuation of the expected transfer of resources from the rest of the economy to them; under our assumption of consumer optimization, this value cannot exceed the cost of the transfer, and will be driven below cost by any moral hazard response to insurance. In other words, if Medicaid, by subsidizing the price of medical care p increases medical spending, this increased medical spending will be valued at less than its cost. The "pure-insurance" term measures the benefit of a budget-neutral reallocation of resources (i.e., relaxing or tightening the recipient's budget constraint) across different states of the world, θ.[9] The movement of these resources is valued using the marginal utility of consumption in each state. The pure-insurance term will be positive for risk-averse individuals

---

[8]This is because an individual values the mechanical increase in consumption from Medicaid according to the marginal utility of consumption, regardless of the extent to which he or she has options to substitute increases in other goods, such as health, for this increase.

[9]This is analogous to moving resources across people in the optimal tax formulas, where the welfare impact of increasing the marginal tax rate on earnings financed by a decrease in the intercept of the tax schedule is given by the covariance between earnings and the social marginal utility of consumption (see, e.g., Piketty and Saez (2013)[59] equation (3)).

as long as Medicaid re-allocates resources to states of the world with higher marginal utilities of consumption.

We arrive at a non-marginal estimate of the total welfare impact of Medicaid, $\gamma(1)$, by integrating with respect to q:

$$
\begin{aligned}
\gamma(1) &= \int_0^1 \frac{d\gamma(q)}{dq} dq \\
&= \underbrace{(p(0)-p(1)) \int_0^1 E[m(q;\theta)] dq}_{\text{Transfer Term}} + \underbrace{\int_0^1 Cov\left(\frac{u_c}{E[u_c]}, (p(0)-p(1))m(q;\theta)\right) dq}_{\substack{\text{Pure-Insurance Term} \\ \text{(consumption valuation)}}} \quad (13)
\end{aligned}
$$

which follows from the fact that $\gamma(0) = 0$, by definition.

Implementation

We estimate the transfer term and pure-insurance term separately, and then combine them.

Pure-insurance term.   Evaluation of the pure-insurance term in equation (12) requires that we specify the utility function over the consumption argument. We assume the utility function takes the following form:

Assumption 4. *Partial utility specification for the consumption-based optimization approach.*
The utility function takes the following form:

$$
u(c,h) = \frac{c^{1-\sigma}}{1-\sigma} + v(h)
$$

where σ denotes the coefficient of relative risk aversion and $v(.)$ is the subutility function for health, which can be left unspecified.

With this assumption, the pure-insurance term in equation (12) can be re-written as:

$$
Cov\left(\frac{c(q;\theta)^{-\sigma}}{E[c(q;\theta)^{-\sigma}]}, (p(0)-p(1))m(q;\theta)\right). \quad (14)
$$

Interpolation.   We can use the above equations to calculate the marginal value of the first and last units of insurance ($\frac{d\gamma(0)}{dq}$ and $\frac{d\gamma(1)}{dq}$ respectively). However, we do not observe $q \in (0,1)$ and therefore do not observe m for these intermediate values. Moreover, with only a partial specification of the utility function, we cannot derive how an optimizing individual would vary m for non-observed values of q. Therefore, we require an additional assumption to obtain an estimate of $\gamma(1)$ in the optimization approaches. For our baseline implementation, we make the following statistical assumption (we explore sensitivity to other approaches below):

Assumption 5. *(Linear Approximation) The integral expression for $\gamma(1)$ in equation (13) is well approximated by:*

$$\gamma(1) \approx \frac{1}{2}\left[\frac{d\gamma(0)}{dq} + \frac{d\gamma(1)}{dq}\right].$$

Assumption 5 allows us to use estimates of $\frac{d\gamma}{dq}(0)$ and $\frac{d\gamma}{dq}(1)$ to form estimates of $\gamma(1)$. This approximation is illustrated by Figure 1. The solid line shows that the value of a marginal expansion of Medicaid coverage, $\frac{d\gamma}{dq}(q)$, may be a nonlinear function of the degree of Medicaid coverage, q. The area under this curve is the true value, $\gamma(1)$, of obtaining Medicaid coverage (for the hypothetical $\frac{d\gamma}{dq}(q)$ curve we drew). The dashed line shows our linear approximation of $\frac{d\gamma}{dq}(q)$ and the resulting area under the dashed line represents our estimate of $\gamma(1)$.

Transfer term. Evaluation of the transfer term in equation (12) does not require any assumptions about the utility function. However, integration in equation (13) to obtain an estimate of the transfer term requires that we know the path of $m(q;\theta)$ for interior values of q, which are not directly observed. We therefore use the above Assumption 5 to integrate between our estimate of the transfer term at $q = 0$ and at $q = 1$.

We can obtain lower and upper bounds for the transfer term without such integration. Under the natural assumption that average medical spending under partial insurance lies between average medical spending under full insurance and average medical spending under no insurance (i.e., $E[m(0;\theta)] \leq E[m(q;\theta)] \leq E[m(1;\theta)])$[10], we obtain lower and upper bounds for the transfer value of Medicaid as Medicaid's impact on the out-of-pocket price of medical care, $p(0) - p(1)$, times medical spending at, respectively, the uninsured and insured levels:

$$[p(0) - p(1)]E[m(0;\theta)] \leq (p(0) - p(1))\int_0^1 \frac{E[m(q;\theta)]}{dq}dq \leq [p(0) - p(1)]E[m(1;\theta)]. \qquad (15)$$

2.3.2 Health-based optimization approach

The consumption-based optimization approach values Medicaid by how it relaxes the budget constraint in states of the world with different marginal utilities of consumption. Here, we show that one can alternatively value Medicaid by how it relaxes the budget constraint in states of the world with different marginal utilities of out-of-pocket spending on health. This requires a stronger assumption than Assumption 3, which states that individuals optimize; we now require that individual choices satisfy a first-order condition:

Assumption 6. *The individual's choices of m and c are at an interior optimum and hence satisfy the first-order condition:*

$$u_c(c,h)\,p(q) = u_h(c,h)\frac{d\tilde{h}(m;\theta)}{dm} \quad \forall m,q,\theta. \qquad (16)$$

_____

[10]A downward-sloping demand function for m would be sufficient for this assumption to hold.

The left-hand side of equation (16) is the marginal cost of medical spending in terms of forgone consumption. The right-hand side of equation (16) is the marginal benefit of additional medical spending, which equals the marginal utility of health $u_h(c, h)$, multiplied by the increase in health provided by additional medical spending, $\frac{d\tilde{h}}{dm}$.

With this assumption, we can use use equation (16) to replace the marginal utility of consumption, u_c in equation (11) with a term depending on the marginal utility of health, u_h, yielding:

$$\frac{d\gamma}{dq} = E\left[\left(\frac{u_h}{E[u_c]}\frac{d\tilde{h}(m;\theta)}{dm}\frac{1}{p(q)}\right)((p(0) - (p(1))m(q;\theta))\right]. \tag{17}$$

We refer to equation (17) as the *"health-based optimization approach."* Analogous to the consumption-based optimization approach, the first term between parentheses measures the value of money in the current state of the world relative to its average value, and the second term between parentheses measures by how much Medicaid relaxes the individual's budget constraint in the current state of the world. From the health-based optimization approach's perspective, the program delivers greater value if it moves more resources to states of the world with a greater return to out-of-pocket spending (i.e., states of the world where the return to out-of-pocket spending is higher because the individual has chosen to forgo valuable medical treatment due to underinsurance).

However, unlike the consumption-based optimization approach, the health-based optimization approach will be biased upward if individuals are at a corner solution in medical spending, so that they are not indifferent between an additional \$1 of medical spending and an additional \$1 of consumption.[11] In other words, Assumption (6) is stronger than Assumption (3) because it requires that individuals' optimization leads them to an interior solution in m.

As was the case with the consumption-based optimization approach, the marginal value of Medicaid to recipients in equation (17) can be decomposed into a transfer term and a pure-insurance term:

$$\frac{d\gamma(q)}{dq} = \underbrace{(p(0) - p(1))E[m(q;\theta)]}_{\text{Transfer Term}} + \underbrace{Cov\left(\frac{u_h}{E[u_c]}\frac{d\tilde{h}(m;\theta)}{dm}\frac{1}{p(q)}, (p(0) - (p(1))m(q;\theta)\right)}_{\substack{\text{Pure-Insurance Term} \\ \text{(health valuation)}}}. \tag{18}$$

---

[11]If the individual is at a corner solution with respect to medical spending, then the first term between parentheses in equation (17) is less than the true value that the individual puts on money in that state of the world (i.e., $\left(\frac{u_h}{E[u_c]}\frac{d\tilde{h}(m;\theta)}{dm}\frac{1}{p(q)}\right) < \frac{u_c}{E[u_c]}$), generating upward bias in the covariance term in equation (18) below because $(p(0) - p(1))m$ is below its mean at the corner solution $m = 0$. The transfer term in equation (18) is not affected by corner solutions because the transfer term does not depend on utility and, hence, is not biased when our estimate of the value that the individual puts on money in an particular state of the world is biased.

Implementation

Since evaluation of the transfer term does not require any assumptions about utility, it is exactly the same as in the consumption-based optimization approach. However, evaluation of the pure-insurance term will once again require a partial specification of the utility function. This time, the partial specification is over health rather than consumption:

Assumption 7. *Partial utility specification for the health-based optimization approach.*
   The utility function takes the following form:

$$u(c, h) = \tilde{\phi} h + \tilde{v}(c),$$

where $\tilde{v}(.)$ is the subutility function for consumption, which can be left unspecified.

Given Assumption 7, the pure-insurance term in the health-based optimization approach in equation (18) can be written as:

$$Cov\left(\frac{d\tilde{h}(m; \theta)}{dm}\frac{\phi}{p(q)}, (p(0) - (p(1))m(q; \theta)\right). \tag{19}$$

The term $\phi \equiv \frac{\tilde{\phi}}{E[\tilde{v}'(c)]}$ is, as in the complete-information approach, the marginal value of health in units of consumption. As before, we require an additional (statistical or economic) assumption to obtain an estimate of $\gamma(1)$ in the optimization approaches from $\frac{d\gamma(0)}{dq}$ and $\frac{d\gamma(1)}{dq}$, and in our baseline implementation we make the same statistical assumption as in the consumption-based optimization approach (see Assumption 5). Implementation of equation (19) requires that we estimate the marginal health return to medical spending, $\frac{d\tilde{h}}{dm}$. We describe the estimation of $\frac{d\tilde{h}}{dm}$ in Section 4.2.2 below.

2.3.3   Comment: Endless Arguments

The option of using either a health-based optimization approach (equation 18) or a consumption-based optimization approach (equation 12) to value a marginal expansion of Medicaid is an example of the multiplicity of representations that are a distinguishing feature of "sufficient statistics" approaches (Chetty (2009)[19]). The logic of the "pure-insurance" term is also highly related to the broad insights from the the asset-pricing literature – where the introduction of new financial assets can be valued using their covariance with the marginal utility of income, which itself can have multiple representations, such as in the classic consumption CAPM (see, e.g., Cochrane (2005)[20]). The pure-insurance term plays a key role in overcoming the requirement in the complete-information approach of having to specify a utility function over all variables on which Medicaid has an impact.

   Relatedly, a key distinction between the complete-information and the optimization approaches comes from the fact that the optimization approach allows one to consider marginal utility with respect to *one* argument of the utility function. Combined with additive separability assumptions

(i.e., Assumptions 4 and 7), we can value Medicaid without knowledge of the marginal valuation of other arguments in the utility function.

The complete-information approach, by contrast, requires adding up the impact of Medicaid on *all* arguments of the utility function. In the above model, we assumed the only arguments were consumption and health. If we were to allow other potentially utility-relevant factors that might be conjectured to be impacted by health insurance (such as leisure, future consumption, or children's outcomes), we would also need to estimate the impact of the program on these arguments, and value these changes by the marginal utilities of these goods across states of the world. As a result, there is a potential methodological bias to the complete-information approach; one can keep positing potential arguments that Medicaid affects if one is not yet satisfied by the welfare estimates.

2.4 Gross and net costs

We benchmark our welfare estimates, $\gamma(1)$, against Medicaid costs. We consider only medical expenditures when estimating program costs. This abstracts from any potential administrative costs associated with Medicaid. It also abstracts from any labor supply responses to Medicaid which may impose fiscal externalities on the government via their impact on tax revenue.[12] Under these assumptions, the average cost to the government per recipient, which we denote by G, is given by:

$$G = E\left[m\left(1;\theta\right) - x(1,\theta)\right]. \tag{20}$$

This gross cost per recipient, G, is higher than the net cost to society; some component of public Medicaid spending replaces costs previously borne by external parties (non-recipients).

Medicaid's net cost per recipient, which we denote by C, is given by:

$$C = E\left[m\left(1;\theta\right) - m\left(0;\theta\right)\right] + E\left[x\left(0, m\left(0;\theta\right)\right) - x\left(1, m\left(1;\theta\right)\right)\right]. \tag{21}$$

Net cost per recipient consists of the average increase in medical spending induced by Medicaid, $m\left(1;\theta\right) - m\left(0;\theta\right)$, plus the average decrease in out-of-pocket spending due to Medicaid, $x\left(0, m\left(0;\theta\right)\right) - x\left(1, m\left(1;\theta\right)\right)$.

We decompose gross costs to the government, G, into net costs, C, and monetary transfers to external parties:

$$G = C + N.$$

We denote by N the monetary transfers by Medicaid from the government to providers of implicit insurance for the uninsured. The monetary transfers to external parties are given by the amount of medical spending that went unpaid by the uninsured:

$$N = E\left[m\left(0;\theta\right)\right] - E\left[x\left(0, m(0;\theta)\right)\right]. \tag{22}$$

---

[12]In general such fiscal externalities should be included in program costs, however, in the context of the Oregon Health Insurance Experiment, there is no evidence that Medicaid affected labor market activities (Baicker et al. (2014)[8]).

In other words, N denotes monetary transfers to the providers of implicit insurance who, in the absence of Medicaid, would have paid for medical spending that was not covered by the out-of-pocket spending of uninsured individuals.

3 Application: the Oregon Health Insurance Experiment

We apply these approaches to welfare analysis of the Medicaid expansion that occurred in Oregon in 2008 via a lottery. The lotteried program, called OHP Standard, covers low-income (below 100 percent of the federal poverty line), uninsured adults (aged 19-64) who are not categorically eligible for OHP Plus, Oregon's traditional Medicaid program.[13] OHP Standard provides comprehensive medical benefits with no patient cost-sharing and low monthly premiums (\$0 to \$20, based on income). We focus on the welfare effects of Medicaid coverage after approximately one year.[14]

3.1 Empirical framework

In early 2008, the state opened a waiting list for the previously closed OHP Standard. It randomly selected approximately 30,000 of the 75,000 people on the waiting list to have the opportunity – for themselves and any household members – to apply for OHP Standard. Following the approach of previous work on the Oregon experiment, we use random assignment by the lottery as an instrument for Medicaid; more details on our estimation strategy and implementation can be found in Appendix A.2.1.

When analyzing the mean impact of Medicaid on an individual outcome y_i (such as medical spending m_i, out-of-pocket spending x_i, or health h_i), we estimate equations of the following form:

$$y_i = \alpha_0 + \alpha_1 Medicaid_i + \epsilon_i, \tag{23}$$

where $Medicaid$ is an indicator variable for whether the individual is covered by Medicaid at any point in the study period. We estimate equation (23) by two-stage least squares, using the following first-stage equation:

$$Medicaid_i = \beta_0 + \beta_1 Lottery_i + \nu_i, \tag{24}$$

in which the excluded instrument is the variable "$Lottery$" which is an indicator variable for whether the individual was selected by the lottery.[15] Previous work has used the the lottery as an instrument for Medicaid to examine the impact of Medicaid on health care utilization, financial well-being, labor market outcomes, health, and private insurance coverage (Finkelstein et al. (2012)[34], Baicker et

---

[13]Eligibility for OHP Plus requires both income below a threshold and that the individual be in a covered category, which includes, for example, children, those on TANF, and those on SSI.

[14]Throughout, we use the term "Medicaid" to refer to coverage by either OHP Standard or OHP Plus. In practice, the increase in Medicaid coverage due to the lottery comes entirely from an increase in coverage by OHP Standard (Finkelstein et al. (2012) [34]).

[15]Appendix A.2 describes how we estimate the quantile effects of Medicaid.

al. (2013)[10], Baicker et al. (2014)[8], and Taubman et al. (2014)[66]). Finkelstein et al. (2012)[34] provides supporting evidence on the assumptions required to use the lottery as an instrument for Medicaid coverage.

Winning the lottery increased the probability of being on Medicaid at any time during the subsequent year by about 25 percentage points. This "first-stage" effect of lottery selection on Medicaid coverage is below one because many lottery winners either did not apply for Medicaid or were deemed ineligible. As a result, all of our estimates of the impact of Medicaid are local average treatment effects (LATEs) of Medicaid for the compliers - i.e., those who are covered by Medicaid if and only if they win the lottery (see, e.g., Angrist and Pischke (2009) [5]). Thus in our application, "the insured" ($q = 1$) are treatment compliers and "the uninsured" ($q = 0$) are control compliers. Differences between treatment and control compliers reflect the impact of Medicaid (i.e., α_1) in the IV estimation of equation (23). As we discuss in more detail in the Conclusion, the impact of Medicaid - and our estimates of the value of Medicaid - could well differ in other populations.

In our results below, we report characteristics and outcomes separately for treatment and control compliers. Estimation of these objects is standard (see, e.g., Abadie (2002)[1]; Abadie (2003)[2]; and Angrist and Pischke (2009)[5]). For example, uninsured individuals who lost the lottery include both control compliers and never-takers; since uninsured individuals who won the lottery provide estimates of characteristics of never-takers and the first stage gives an estimate of the share of individuals who are compliers, we can combine these estimates to infer the characteristics of control compliers. Likewise, insured individuals who won the lottery include both treatment compliers and always-takers; we can use insured individuals who lost the lottery to estimate characteristics of always-takers and thus identify the characteristics of treatment compliers.

3.2 Data and summary statistics

The data from the Oregon Health Insurance Experiment were previously analyzed by Finkelstein et al. (2012)[34] and are publicly available at www.nber.org/oregon. Data on Medicaid coverage (q) are taken from state administrative records; all other data elements from the Oregon Health Insurance Experiment in the baseline analyses are derived from information supplied by approximately 15,500 respondents to mail surveys sent about one year after the lottery to individuals who signed up for the lottery.

Table 1 presents descriptive statistics on the data from the Oregon Health Insurance Experiment. The first column reports results for the full study population. Columns 2 and 3 report results for the treatment compliers and control compliers respectively. Panel A presents demographic information. The population is 60 percent female and 83 percent white; about one-third are between the ages of 50-64. The demographic characteristics are balanced between treatment and control compliers (p-value = 0.12).

We focus in this subsection on outcome variables that are relatively easy to measure: medical spending (m), out-of-pocket medical spending (x), and out-of-pocket prices (p). The next two subsections describe how we estimate the impact of Medicaid on the monetized value of health (h)

and on consumption (c); neither the value of health nor consumption are directly measured in our data, so estimation of each requires additional assumptions. Panel B presents summary statistics on key outcome measures in the Oregon data.

Medical spending m. Survey responses provide measures of utilization of prescription drugs, outpatient visits, ER visits, and inpatient hospital visits. To turn these quantity estimates into estimates of total annual medical spending, Finkelstein et al. (2012)[34] annualized the quantity reports and summed them up, weighting each type of use by its average cost (expenditures) among the low-income publicly insured non-elderly adults in the 2002-2007 (pooled) Medical Expenditure Survey (MEPS). Importantly, the MEPS data on expenditures reflect actual payments (i.e., transacted prices) rather than contract or list prices (MEPS (2013), page C-107)[53]).

We estimate that Medicaid increases total medical spending by about \$900. On average, annual medical spending is about \$2,700 for control compliers ($q = 0$) and about \$3,600 for treatment compliers ($q = 1$).

Out-of-pocket spending x. We measure annual out-of-pocket spending for the uninsured ($q = 0$) based on self-reported out-of-pocket medical expenditures in the last six months, multiplied by two.[16] Average annual out-of-pocket medical expenditures for control compliers is $E\left[(x(0, m(0, \theta))\right] = \$569$.

Our baseline analysis assumes that the insured have zero out-of-pocket spending (i.e., $x(1, m(1; \theta)) = 0$), since Medicaid in Oregon has zero out-of-pocket cost sharing, no or minimal premiums, and comprehensive benefits.[17] We explore sensitivity below to using the self-reported out-of-pocket spending for the insured for $x(1, m(1; \theta))$; naturally, this reduces our estimate of the value of Medicaid to recipients.

Out-of-pocket prices p. The optimization approaches require that we define the out-of-pocket price of medical care with Medicaid, $p(1)$, and without Medicaid, $p(0)$. Our baseline analysis assumes $p(1) = 0$; i.e., those with Medicaid pay nothing out of pocket towards medical spending. This is consistent with our baseline assumption that $x(1, m) = 0$, and as noted we examine sensitivity to this assumption below.

We measure $p(0)$ as the ratio of mean out-of-pocket spending to mean total spending for control compliers ($q = 0$), i.e., $\frac{E[x(0, m(0; \theta))]}{E[m(0; \theta)]}$. We estimate $p(0) = 0.21$. In other words, we estimate that

---

[16]To be consistent with our treatment of out-of-pocket spending when we use it to estimate consumption (discussed below in subsection 3.4), we impose the same two adjustments here. First, we fit a log normal distribution on the out-of-pocket spending distribution. Then, we impose a per capita consumption floor by capping out-of-pocket spending so that per capita consumption never falls below the floor; this cap binds for less than 0.3 percent of control compliers.

[17]This assumes that the uninsured report their out-of-pocket spending without error but that the insured (some of whom report positive out-of-pocket spending in the data) do not. This is consistent with a model of reporting bias in which individuals are responding to the survey with their typical out-of-pocket spending, not the precise spending they have incurred since enrolling in Medicaid. In this instance, there would be little bias in the reported spending for those who are not enrolled in Medicaid (since nothing changed), but the spending for those recently enrolled due to the lottery would be dramatically overstated because of recall bias.

the uninsured pay only about \$0.2 on the dollar for their medical spending, with the remainder of the uninsured's expenses being paid by external parties. This will have important implications for our welfare results below. It is therefore important to note that our estimate that the uninsured pay relatively little of their medical expenses out of pocket is not an artifact of our setting or of our data.[18]

3.3 Measuring requisite health (h) inputs

3.3.1 Approaches to valuing changes in health (h)

Both the complete-information approach and the health-based optimization approach require that we measure and value the impact of Medicaid on health. We have several measures of health in the Oregon data. For our baseline analysis, we use the widely-used five-point self-assessed health question that asks "In general, would you say your health is:" and gives the following response options: "Excellent, Very Good, Good, Fair, Poor." In prior work on the Oregon Health Insurance Experiment, Medicaid coverage was estimated to improve this measure of self-reported health, in addition to two other measures: the Patient Health Questionnaire (PHQ) depression screen, and the Short Form health questionnaire's (SF-8) measures of recent physical and mental health problems (Finkelstein et al. (2012)[34]; Baicker et al. (2013)[10]). We conduct sensitivity analysis below to using these other measures of health.

A key challenge for welfare analysis is how to *value* changes in a given measure of health. When the health measure is mortality, the standard approach is to use estimates of the value of a statistical life year (VSLY), which has been estimated using various approaches. There is no evidence of an impact of Medicaid on mortality in the Oregon Health Insurance Experiment, although the confidence intervals are, not surprisingly, wide (Finkelstein et al. (2012)[34]). For non-mortality health measures, such as are observed in the Oregon Health Insurance Experiment, the standard approach involves two steps: first map these health measures into a cardinal utility scale, expressed in terms of quality-adjusted life years (QALYs), and then scale it by an estimate of the VSLY to arrive at a monetary value.[19] We discuss these two steps below. Appendix A.4 provides more detail on the first of these steps: how health measures are mapped to QALYs in general, and on the specific mappings we use in our baseline and sensitivity analyses.

Mapping measured health to QALY units. We map our baseline self-assessed health measure into QALYs using the mapping estimated in Van Doorslaer and Jones (2003)[73]. They apply the

---

[18]The Kaiser Commission on Medicaid and the Uninsured estimates that the average uninsured person in the U.S. paid \$500 out of pocket but incurred total medical expenses of \$2443 (Coughlin et al. (2014)[21], Figure 1), suggesting that on average the uninsured in the U.S. pay only 20% of their total medical expenses. Likewise, Hadley et al. (2008)[41] estimate that the uninsured pay only 35% of their medical costs out of pocket. To verify this is also true when focusing on low-income populations in the U.S. as a whole, we analyzed out-of-pocket spending using the Medical Expenditure Panel Survey (MEPS) from 2009-2011. We estimate that uninsured adults aged 19-64 below 100 percent of the federal poverty line pay about \$0.33 out of pocket for every dollar of their medical expenses.

[19]By definition, a year lived in perfect health yields one QALY and a year in which one is dead yields zero QALYs. A year lived in an intermediate health state yields a QALY between zero and one.

"standard gamble approach," which is one of the two principal methods used to translate self-reported health into QALYs. Specifically, they ask respondents their preferences over hypothetical outcomes in order to elicit a probability v such that respondent reports being indifferent between living in a particular health state and facing a gamble consisting of living in perfect health with probability v and being dead with probability $1 - v$. One year lived in this particular health state is assigned a QALY of v.

Panel B of Table 1 shows results for our baseline self-assessed health measure, reported in QALY units. Treatment compliers are less likely to respond than control compliers that they are in poor or fair health and more likely to describe their health as good, very good, or excellent. Weighting the effect of Medicaid on each health state by the associated QALY of that health state, our estimates indicate that Medicaid increases health by by 0.05 QALYs.

Choosing a VSLY. The value of a statistical life year is defined as the value of a statistical life (VSL) divided by the remaining life expectancy. The value of a statistical life is based on individuals' (hypothetical) willingness to pay out of their own income for a small change in the probability of death. Estimation of the VSL is challenging, but there exists a large literature, reviewed by Viscusi (1993)[74] and Cropper, Hammitt, and Robinson (2011)[22], that uses various approaches to do so.[20] We take as a "consensus" estimate from this literature Cutler's (2004) [24] choice of $100,000 for the VSLY for the general US population.

An added challenge in our context is how to adjust this VSLY to our low-income population. Our primary goal is to estimate the *individual's* willingness to pay (or, more precisely, "willingness to give up consumption") for Medicaid out of their own income. Naturally, lower-income people will be willing to pay less (i.e., forego less consumption) to obtain an additional statistical life year. To be clear, this does not mean that society's willingness to pay for an additional statistical life year is lower for lower-income populations. Rather, society may scale up the individuals' willingness to pay by a social welfare weight (as in Saez and Stantcheva (2016)[61]) to arrive at a social willingness to pay, or adjust for the cost of redistributing to low-income populations (as in Hendren (2014)[43]). But, for the purpose of estimating the individual's willingness to pay γ, we must scale the VSLY estimate from Cutler (2004) [24] to reflect the lower incomes in our population.

We therefore scale the VLSY by the ratio of the marginal utility of consumption for our population and the general population. With CRRA utility over consumption (see Assumption 1), our baseline assumption (see below) of a coefficient of relative risk aversion $\sigma = 3$, and per capita consumption for our population that is about 40 percent of the general population's[21], this implies

---

[20]A common approach is to estimate a hedonic model on choices made in labor markets or product markets that involve a tradeoff between money and a risk of death (e.g., Aldy and Viscusi (2008)[3]). Methods relying on stated-preference approaches are also widely used (e.g., Lindhjem et al. (2011)[51]). Another method is to infer the VSL from politicians' choices that tradeoff risks of death and monetary benefits among constituents (e.g., decisions on speed limits as used by Ashenfelter and Greenstone (2004)[7]).

[21]Specifically, using data from the Consumer Expenditure Survey, we estimate that median per capita consumption for families below the federal poverty line and headed by an uninsured adult is $7,234, compared to $19,364 in the general US population. Details of the sample definition are identical to those listed in subsection A.5.1 except that the sample is not limited to singles.

a VSLY (ϕ) of $0.4^3 \cdot \$100,000 \approx \$5,000$. This is of course substantially below the VSLY for the general population; we report sensitivity to alternative assumptions below.

3.3.2 Estimating the (heterogeneous) impact of medical spending on health

The health-based optimization approach also requires that we estimate (and value using the above approach) the health return to medical spending, $\frac{d\tilde{h}}{dm}$ (see equation (19)). Estimating the health production function is notoriously challenging (see, e.g., Almond et al. (2010)[4] for one approach). In our case, the challenges are compounded by the fact that we must estimate *heterogeneity* in these returns across the values of the (endogenous) choice of medical spending m.

To estimate the health returns to medical spending, we use the lottery as an instrument for medical spending. This assumes that Medicaid affects health only via an impact on medical spending. To estimate the returns to medical spending at separate values of the (endogenous) choice of m, we assume that heterogeneity in m can be proxied using a set of observable variables θ^K, and assume that the health production function is constant for all m conditional on θ^K.[22] We use measures of baseline medical and financial status for θ^K. Appendix A.6 provides more detail on our implementation of this approach and the resulting estimates. With these estimates, we can calculate the slope of the health production function conditional on θ^K , $E\left[\frac{\partial \tilde{h}}{\partial m}|\theta^K\right]$. The estimates are fairly imprecise but are suggestive of larger health returns for more financially constrained individuals.

3.4 Measuring requisite consumption (c) inputs

Both the complete-information approach and the consumption-based optimization approach require that we measure consumption. Specifically, the complete-information approach requires that we estimate the impact of Medicaid on the distribution of consumption, while the consumption-based optimization approach requires that for the "pure-insurance term" we estimate the joint distribution of consumption and out-of-pocket spending for the uninsured.[23]

The difficulty in obtaining high-quality consumption data is a pervasive problem for empirical research on a wide array of topics. Ours is no exception. Consumption data are not available for participants in the Oregon study.

We take two different approaches to measuring consumption. The consumption proxy approach uses data from the Oregon Health Insurance Experiment to estimate the required consumption metrics for both the complete-information and the consumption-based optimization approach. The CEX consumption measurement approach uses national data from the Consumer Expenditure

---

[22]This approach omits any value of insurance within each value of θ^K, and thus likely understates the true value of insurance. However, it provides a parsimonious methodology for implementing the health-based optimization approach with our data. Its empirical limitations also highlight the importance of further work aimed at identifying not only the average return on medical spending, but also its heterogeneity.

[23]Equation (14) suggests that we need to estimate the joint distribution of $c(0; \theta)$ and $(p(0) - p(1))m(0; \theta)$ at $q = 0$. Since $p(1) = 0$ by assumption, this reduces to the joint distribution of consumption c and out-of-pocket spending $x(0, m(0; \theta)) = p(0)m(0; \theta)$. We need to estimate this joint distribution only for the uninsured (so for $q = 0$) because our assumption that Medicaid provides full insurance (i.e., $p(1) = 0$) implies that the marginal value of additional insurance for the fully insured (so for $q = 1$) is zero.

Survey (CEX) to directly estimate the pure-insurance term for the uninsured in the consumption-based optimization approach. We describe each in turn.

3.4.1 Consumption proxy approach for the complete-information and consumption-based optimization approaches

We proxy for non-medical per capita consumption c using the individual's out-of-pocket medical spending, x, combined with average values of non-medical expenditure and out-of-pocket medical expenditure. Letting $\bar{c}$ denote the average non-medical expenditure for the population, we define the consumption proxy as:

$$c = \bar{c} - (x - \bar{x})/n, \tag{25}$$

where n denotes family size and $\bar{x}$ denotes average per capita out-of-pocket medical spending among control compliers; average family size among compliers is about 2.9 (see Table 1). Our approach accounts for within-family resource sharing by assuming that consumption is shared equally within the family, i.e., the impact of a given amount of out-of-pocket medical spending on non-medical consumption is shared equally within families.[24] This seems a reasonable assumption given the joint nature of many components of consumption; however, in the sensitivity analysis below, we also report results in which we assume the other extreme: that the out-of-pocket spending shock is borne entirely by the individual with the spending.

This consumption proxy approach makes several simplifying assumptions. First, it assumes that the only channel by which Medicaid affects consumption is by reducing out-of-pocket spending; it rules out Medicaid affecting consumption by changing income, which is a reasonable assumption in our context.[25] Second, it assumes that, absent out-of-pocket spending, per capita consumption would be the same for all individuals in the Oregon study. This is an assumption made for convenience and unlikely to be literally true. However, it may not be a terrible approximation of reality, since heterogeneity in non-medical consumption (in the absence of out-of-pocket spending) may be limited in our relatively narrowly defined population: uninsured adults 19-64 below the federal poverty line in Oregon. Finally, it does not allow for the possibility of any intertemporal consumption smoothing through borrowing or saving. Such opportunities are likely limited in our low-income study population but presumably not zero; by not allowing for this possibility, we likely bias upward our estimate of $\gamma(1)$. This is the prime motivation for our alternative measurement approach for consumption that we discuss in subsection 3.4.2 below.

---

[24]This same logic implies that the welfare benefits from Medicaid are also shared among family members. This is captured in the optimization approach by equation (13); this equation values any dollar flowing to the family by the marginal utility of consumption of the individual irrespective of whether dollar is used to benefit the individual or other family members. However for the complete-information approach, it requires that we we replace $\gamma(1)$ by $\gamma(1)/n$ when estimating equation (3).

[25]Prior analysis of the Oregon Health Insurance Experiment showed no evidence of a direct impact of Medicaid on income (Finkelstein et al. (2012)[34], Baicker et al. (2014)[8]).

Implementation. We estimate c in equation (25) in several steps. We use the Oregon survey data to measure x (as described above), and also family size n. We estimate $\bar{c}$ as mean per capita non-medical consumption in a population that has similar characteristics as participants in the Oregon study, but for which consumption data are available, namely families that live below the federal poverty line, have an uninsured household head, and are in the Consumer Expenditure Survey (CEX).[26]

We estimate the impact of Medicaid on the distribution of out-of-pocket spending x. To do so, we make the parametric assumption that out-of-pocket spending is a mixture of a mass point at zero and a log-normal spending distribution and then estimate the distribution of out-of-pocket spending x for control compliers using standard, parametric quantile IV techniques. Appendix A.2 describes these techniques in more detail, and also reports that the parametric model fits the data quite well.

Because there is unavoidable measurement error in our approach to estimating c, and because welfare estimates are naturally sensitive to c at low values, we follow the standard procedure for ruling out implausibly low values of c (e.g., Brown and Finkelstein (2008)[14], Hoynes and Luttmer (2011)[45]) by imposing an annual consumption floor. Our baseline analysis imposes a consumption floor of \$1,977, which corresponds to the 1st percentile of non-medical consumption for our uninsured sample in the CEX. We impose the consumption floor by capping the out-of-pocket spending drawn from the fitted log-normal distribution at $\bar{x} + n(\bar{c} - c_{floor})$, where $\bar{x}$ is average per capita out-of-pocket medical spending as in equation (25). Our baseline consumption floor binds for less than 0.3 percent of control compliers. In the sensitivity analysis below, we explore sensitivity to the assumed value of the consumption floor. Finally, we map the fitted, capped out-of-pocket spending distribution to consumption using equation (25).

Figure 2 shows the resultant distributions of consumption for control compliers ($q = 0$) and treatment compliers ($q = 1$). Average non-medical consumption for control compliers ($q = 0$) is \$9,214 with a standard deviation of \$1,089. For treatment compliers ($q = 1$), consumption is simply average non-medical consumption for the insured (\$9,505), since by assumption $x(1, m) = 0$.[27] The difference between the two lines in the figure shows the increase in consumption due to Medicaid for the compliers.

3.4.2 Consumer Expenditure Survey approach to estimating the "pure-insurance term" for the consumption-based optimization approach

A concern with our consumption proxy approach is that it assumes that changes in out-of-pocket spending x translate one for one into changes in consumption if the individual is above the con-

---

[26]Details of the sample definition are identical to those listed in subsection A.5.1 except that the sample is not limited to singles.

[27]Average non-medical consumption for the low-income uninsured (i.e., $\bar{c}$) is \$9,214 in the CEX. To account for the fact that non-medical consumption for the uninsured is presumably lowered due to out-of-pocket medical costs which are 0 for the insured, we assume that average non-medical consumption for the insured is $\bar{c} + (\bar{x})/n$ (see equation (25)) where $\bar{x}$ denotes average out-of-pocket spending for the uninsured.

sumption floor. If individuals can borrow, draw down assets, or have other ways of smoothing consumption, this approach overstates the consumption smoothing benefits of Medicaid. We therefore derive an alternative approach using national data on out-of-pocket spending (x) and non-medical consumption (c) from low-income individuals from the CEX. For the consumption-based optimization approach, these CEX data allow us to directly estimate the the pure-insurance term at $q = 0$ in equation (14), i.e., the covariance between the marginal utility of non-medical consumption c and out-of-pocket spending x among the uninsured.[28]

Appendix A.5.1 provides more detail on the data, sample definition, and summary statistics of the CEX data. To be broadly consistent with the Oregon sample, we limit the analysis to adults aged 19-64 who are below 100% of the federal poverty line. Because the CEX only requests information on the health insurance status of the household head, we restrict the sample to single adults with no children in the household, so that we can identify the individuals who are insured ($q = 1$) and uninsured ($q = 0$). Appendix Table 1 reports the summary statistics for the sample and compares it to the sample of compliers in the Oregon data. The CEX sample is slightly older, naturally has smaller family size (because of our limitation to singles), and tends to have somewhat lower out-of-pocket spending ($395 versus $569).

The ability in the CEX to directly observe consumption - and its covariance with out-of-pocket spending - makes it a valuable complement to the consumption proxy approach. Nevertheless, there are two important drawbacks to using the CEX in this manner. First, it cannot be used for the complete-information approach, because that requires a causal estimate of the impact of Medicaid on consumption which is not feasible to obtain in the CEX data.[29] Second, the data come from a national sample of low-income individuals, not the Oregon study data.

Estimation strategy. In principle, it is straightforward to directly estimate the correlation between the marginal utility of consumption and out-of-pocket medical spending for uninsured individuals in the CEX data. We wish to estimate equation (14). For $q = 0$, this reduces to

$$Cov\left(\frac{c\left(0;\theta\right)^{-\sigma}}{E\left[c\left(0;\theta\right)^{-\sigma}\right]}, x\left(0, m\left(0;\theta\right)\right) \right),$$

where c and x are observed non-medical consumption and out-of-pocket medical spending for the uninsured in the CEX. We impose the same consumption floor as in our prior approach.

In practice, we encounter an important practical challenge: the raw data show a negative covariance among the uninsured between the marginal utility of consumption and out-of-pocket spending (i.e., higher non-medical consumption is correlated with higher out-of-pocket medical

---

[28]Note that because we assume $p(1) = 0$ (the insured face a zero price for medical care), the second term of the covariance expression in equation (14) simplifies to $p(0)m(0)$ at $q = 0$. Hence, at $q = 0$, the second term is equal to out-of-pocket spending, x.

[29]For the pure-insurance term of the consumption-based optimization approach, we need to evaluate the covariance term of equation (14) only for $q = 0$ because we know that the covariance term is zero for $q = 1$, given our baseline assumption that the insured face no consumption risk from medical expenditures. Hence, we do not need a causal estimate of the impact of Medicaid on consumption.

spending). Moreover, this is not an idiosyncratic feature of the CEX. We verified that this same "wrong signed" negative correlation between the marginal utility of consumption and out-of-pocket medical spending exists in data from the Panel Study of Income Dynamics (PSID). This could be an accurate measure of the empirical covariance if it were driven, say, by unobserved income so that those with higher consumption had higher medical spending. In this case, the negative covariance would reflect the fact that a reduction in the marginal price of health expenditure is bringing resources to states of the world with a lower marginal utility of consumption, and the value of Medicaid would actually be below its transfer component. However, the negative covariance remains even after controlling for income and assets. We find it more likely that the covariance term is biased from measurement error that induces a negative correlation between $c(0;\theta)^{-\sigma}$ and $x(0;\theta)$. We therefore implement a measurement-error correction.

Measurement error correction approach. We develop a measurement error correction that allows for potentially nonclassical measurement error in out-of-pocket medical spending. We do so by exploiting a key implication of our model: the covariance between out-of-pocket medical spending and the marginal utility of consumption should be zero for the insured ($q = 1$) because they have no out-of-pocket medical spending. Under the assumption that measurement error in out-of-pocket medical spending is the same for the insured and uninsured, we use the estimated covariance term for the insured to infer the impact of measurement error on the covariance term for the uninsured.

We wish to infer the covariance between the marginal utility of consumption (normalized by its average), $\frac{c(0;\theta)^{-\sigma}}{E[c(0;\theta)^{-\sigma}]}$, and true out-of-pocket medical spending, $x(0;\theta)$, for uninsured individuals ($q = 0$). Here, our primary concern is mis-measurement of out-of-pocket spending, $x(0;\theta)$. In particular, we assume the observed out-of-pocket spending is given by

$$\hat{x}(q;\theta) = x(q;\theta) + \epsilon(q;\theta),$$

where $\epsilon(q;\theta)$ is a measurement-error shock to an individual of type θ that is drawn from a distribution with unknown functional form that, importantly, may be correlated with the marginal utility of consumption.

We identify the covariance term even under this fairly general measurement-error structure by making three assumptions. First, we assume (non-medical) consumption is measured without error. Second, we assume that the covariance of the marginal utility of consumption and the measurement error is the same for the insured and uninsured.[30] Third, we assume that true out-of-pocket medical spending is zero for the insured, so that $\hat{x}(1;\theta) = \epsilon(1;\theta)$. In other words, we allow for measurement error in $\hat{x}$ that is additive in x, arbitrarily correlated with c, and common

---

[30]Formally, we assume:

$$Cov\left(\frac{c(0;\theta)^{-\sigma}}{E\left[c(0;\theta)^{-\sigma}\right]}, \epsilon(0;\theta)\right) = Cov\left(\frac{c(1;\theta)^{-\sigma}}{E\left[c(1;\theta)^{-\sigma}\right]}, \epsilon(1;\theta)\right).$$

for the insured and uninsured. These assumptions would be satisfied, for example, if ϵ reflected consumption of uncovered healthcare for both the insured and uninsured (e.g., over-the-counter pain killers, or transportation costs associated with medical care) and these are consumed in equal amounts by both groups.

This approach yields an intuitive estimation strategy: we use the estimated covariance term for the insured as an estimate of the contribution of measurement error to the covariance term of the uninsured. In particular, our assumptions imply that the observed covariance between $\frac{c(0;\theta)^{-\sigma}}{E[c(0;\theta)^{-\sigma}]}$ and $\hat{x}(0;\theta)$ is the sum of the true covariance (which should be zero) and the measurement-error component:

$$Cov\left(\frac{c(0;\theta)^{-\sigma}}{E\left[c(0;\theta)^{-\sigma}\right]}, \hat{x}(0;\theta)\right) = Cov\left(\frac{c(0;\theta)^{-\sigma}}{E\left[c(0;\theta)^{-\sigma}\right]}, x(0;\theta)\right) + Cov\left(\frac{c(0;\theta)^{-\sigma}}{E\left[c(0;\theta)^{-\sigma}\right]}, \epsilon(0;\theta)\right).$$

We identify the measurement-error component of the covariance using the covariance term for those who are insured:

$$Cov\left(\frac{c(0;\theta)^{-\sigma}}{E\left[c(0;\theta)^{-\sigma}\right]}, \epsilon(0;\theta)\right) = Cov\left(\frac{c(1;\theta)^{-\sigma}}{E\left[c(1;\theta)^{-\sigma}\right]}, \hat{x}(1;\theta)\right).$$

Hence, the true covariance term for the uninsured is given by:

$$Cov\left(\frac{c(0;\theta)^{-\sigma}}{E\left[c(0;\theta)^{-\sigma}\right]}, x(0;\theta)\right) = Cov\left(\frac{c(0;\theta)^{-\sigma}}{E\left[c(0;\theta)^{-\sigma}\right]}, \hat{x}(0;\theta)\right) - Cov\left(\frac{c(1;\theta)^{-\sigma}}{E\left[c(1;\theta)^{-\sigma}\right]}, \hat{x}(1;\theta)\right).$$

(26)

Intuitively, we estimate the true covariance term as the difference in the covariance terms for the uninsured and insured, where the latter term removes the measurement error bias from the results.

Alternative measurement error correction. An alternative approach to correct for measurement error in $\hat{x}$ would be to use an instrumental variable. Unfortunately, no natural instrument presents itself in the CEX. However, in Appendix A.5.2, we use PSID data on whether the individual reports having gone to the hospital as an instrument for out-of-pocket spending. The drawback of this approach is that it may not recover the covariance of interest. By using hospitalization as an instrument, we obtain the correlation of consumption and out-of-pocket medical spending that is induced through hospitalization, which may not be representative of the overall correlation. For example, the consumption response to shocks, θ, that lead to hospitalization may be different from the response to shocks that don't result in hospitalization. We therefore view the approach using PSID data as a complement to the CEX analysis. In practice, we obtain similar results using the IV approach in the PSID to what is reported here using the CEX data.

4 Welfare Results

Panel A of Table 2 summarizes the welfare results from our baseline analysis.

4.1 Complete-information approach

We solve equation (4) for $\gamma(1)$. This requires us to estimate mean health outcomes and the distribution of consumption for control compliers ($q = 0$) and for treatment compliers ($q = 1$). Table 1 showed our estimates of the average health of control compliers and treatment compliers. Figure 2 showed the estimated distribution of consumption at $q = 0$ and $q = 1$. The complete-information approach requires that we assume a value of a statistical life year (ϕ) and a coefficient of relative risk aversion (σ). As discussed above, our baseline analysis assumes $\phi = \$5,000$. We also assume in our baseline analysis that $\sigma = 3$.[31]

The first column of Table 2 shows the resultant estimate: $\gamma(1) = \$1,675$. In other words, we estimate that a Medicaid recipient would be indifferent between giving up Medicaid and giving up $1,675 in consumption. We decompose the welfare value of Medicaid to recipients $\gamma(1)$ into transfer term of $699 (see equation (5)) and "pure-insurance" term of $976 (see equation (6)). This suggests that roughly 40 percent of the value of Medicaid comes from its transfer component, and about 60 percent comes from Medicaid's ability to move resources across states of the world.

Because the complete-information approach involves summing up over all the impacts of Medicaid on each argument of the utility function, it facilitates a natural decomposition into the component of the welfare value operating through health and the component operating through consumption. We define the welfare value of Medicaid to recipients operating through consumption, γ_C as:

$$E\left[\frac{c\left(0;\theta\right)^{1-\sigma}}{1-\sigma}\right] = E\left[\frac{(c\left(1;\theta\right) - \gamma_C)^{1-\sigma}}{1-\sigma}\right],$$

and estimate $\gamma_C = \$1,381$. We then infer the welfare value of Medicaid to recipients operating through health as $\gamma_M = \gamma(1) - \gamma_C = \$294$. Appendix A.3 provides implementation details.[32] This suggests that most of the value of Medicaid (80%) is coming through its impact on consumption as opposed to health.[33]

---

[31]A long line of simulation literature uses this value (see, e.g., Hubbard, Skinner, and Zeldes (1995)[46], Mitchell, Poterba, Warshawsky, and Brown (1999)[54], Scholz, Seshadri, and Khitatrakun (2006) [63], Brown and Finkelstein(2008)[14], and Einav et al. (2010)[29]. Naturally, though, there are a range of plausible estimates; a substantial consumption literature, summarized in Laibson, Repetto, and Tobacman (1998)[49], has found risk aversion levels closer to 1, while other papers report higher levels of relative risk aversion (Barsky, Kimball, Juster, and Shapiro (1997)[12], Palumbo (1999)[58]).

[32]Because of the curvature of the utility function, the order of operations naturally matters. If we instead directly estimate the welfare gain due to health and then infer the welfare gain due to consumption based on the difference between $\gamma(1)$ and the welfare gain due to health, we estimate a welfare gain due to consumption of $1059 and a welfare gain due to health of $615.

[33]We can also back out the fraction of the pure-insurance component, I, that operates through health versus consumption. We find that the majority operates through consumption ($812) versus health ($164). The pure-insurance component operating through consumption smoothing is broadly similar to the approach taken by Feldstein and Gruber (1995)[30] to estimate the consumption-smoothing value of catastrophic health insurance, and Finkelstein

4.2 Optimization approaches

For the optimization approaches, we estimate the transfer component and pure-insurance component separately, and then combine them for the overall welfare estimate. Estimation of the transfer component of Medicaid under the optimization approach is relatively straightforward, and is the same for the consumption- and health-based optimization approaches. Estimation of the "pure-insurance" component, however, is more complicated, and varies with the two approaches. The consumption-based optimization approaches require an assumption about the coefficient of risk aversion for the pure-insurance term; we use the same assumption as in the complete-information approach of $\sigma = 3$. The health optimization approaches require an assumption about the value of health (ϕ) for the pure-insurance term; we use the same baseline assumption as in the complete-information approach of $\phi = \$5,000$.

4.2.1 Transfer component

Without any assumptions about the utility function, the optimization approach allows us to estimate the value of the transfer component of Medicaid to recipients using only the estimates of m and p (see equation (13)). The change in the out-of-pocket price for medical care due to insurance $(p(0) - p(1))$ is 0.21. Using linear approximation (Assumption 5), the transfer term is \$661. We report this estimate in columns II through IV. Without the linear approximation, we can derive lower and upper bounds for the transfer term of \$566 and \$749, respectively (see equation (15)).

4.2.2 Pure-insurance term

Consumption-based optimization approach with consumption proxy. We estimate the pure-insurance value at $q = 0$ using equation (14). This requires an estimate of the joint distribution of consumption and out-of-pocket spending for control compliers (see footnote 23). The distribution of c for $q = 0$ was shown in Figure 2. The joint distribution follows immediately, given that the consumption proxy approach defines consumption as mean consumption adjusted for out-of-pocket spending (equation 25). At $q = 1$, our assumption that $p(1) = 0$ together with our definition of consumption implies that the marginal utility of consumption is constant, and hence the pure-insurance value of Medicaid is 0 on the margin. Following the linear approximation in Assumption 5, the total pure-insurance component is therefore one-half of what we estimate at $q = 0$, or \$760. We report this estimate in column II.

Consumption-based optimization approach with CEX consumption measure. We estimate the pure-insurance value at $q = 0$ using equation (26); this uses the difference in the observed covariance term for the uninsured and the observed covariance term for the insured to estimate the measurement-error corrected covariance term for the uninsured. Table 3 shows these components. Column I shows results for our baseline measure of non-health consumption, which is a

and McKnight (2008)[33] to estimate the consumption-smoothing value of the introduction of Medicare.

broad-based measure. It consists of total expenditure excluding individual expenditures for health care providers, prescription drugs, and medical devices.[34] In columns II and III, we show results based on alternative definitions of non-health consumption. Specifically, in column II, we exclude durables within each expenditure category because an expenditure on a durable good leads to a consumption flow over a longer period of time than that in which the expenditure occurred. In column III, we create a consumption measure that is limited to expenditures in categories that are relatively easy to adjust in the short run: food, entertainment, apparel, tobacco, alcohol, personal care, and reading.

Across all the consumption definitions, the covariances between the marginal utility of consumption and out-of-pocket spending for the uninsured are negative.[35] However, the covariance is more negative for the insured. Applying the measurement error correction approach from equation (26) yields a covariance between the marginal utility of consumption and out-of-pocket spending at $q = 0$ of \$265 for our baseline, broad-based measure of consumption (column I); the estimate based on the consumption measure excluding durables is similar (column II) while the estimate based on expenditures in relatively easily adjustable categories is substantially lower (column III). As before, the assumption that Medicaid provides full insurance implies that the pure-insurance value of Medicaid is 0 at the margin at $q = 1$ and, using the linear approximation to obtain an average covariance value over $q = 1$ to $q = 0$ yields a pure-insurance value of \$133. We report this estimate in Table 2, column III.

As noted in Section 3.4.2, we also implement the consumption covariance term using data from the PSID, and an alternative measurement error correction approach based on instrumenting for out-of-pocket medical spending with hospital admissions. Appendix A.5.2 provides more details on the data, implementation and results. Briefly: we replicate the CEX result of a negative OLS relationship between the marginal utility of consumption and out-of-pocket spending among the uninsured in the PSID. However, we find a relationship of the expected sign when instrumenting with hospitalization: hospitalization leads to both higher medical spending and lower consumption. The estimated IV relationship implies a covariance term (a pure-insurance value) that equals \$248 (s.e. \$138), which is higher but still of the same order of magnitude as the baseline estimate from the CEX of \$133.

Health-based optimization approach. We estimate the pure-insurance value at $q = 0$ using equation (19). To do so, we combine our estimates of the slope of the health production function conditional on θ^K, $E\left[\frac{\partial \tilde{h}}{\partial m} | \theta^K\right]$ (see Appendix Table 2), with estimates of the distribution of out-of-pocket spending among control compliers $x(0, m(0; \theta))$ conditional on θ^K.[36] As with the

---

[34]Contributions to private and public pension programs are part of the standard CEX expenditure measure but we exclude them from our measure because these correspond to savings rather than to current consumption.

[35]Although these results do not include any controls, this negative covariance persists even after controlling for a rich set of covariates including both time-invariant demographics and time-varying factors like income and wealth, as well as including consumer-unit fixed effects (results not shown).

[36]Specifically, we use $Cov_\theta[X, Y] = E_{\theta^K}\left[E_{\theta|\theta^K}[X, Y]\right] - E_\theta[X] E_\theta[Y]$, where X is $\frac{\partial \tilde{h}}{\partial m}$ and Y is $x(0, m(0, \theta))$. In other words, we calculate the first expectation of this covariance term first conditionally on θ^K, and then we take the

consumption-based optimization approaches, the assumption that Medicaid implies full insurance and the use of the linear approximation implies an estimated pure-insurance component that is half of what we estimate at $q = 0$, or \$30. We report this estimate in Table 2, column IV. The fact that the pure-insurance term is positive implies that the insurance program tends to increase medical spending more in states of the world with higher marginal health returns to medical spending.

4.2.3 Overall welfare

The first row of Table 2 shows overall welfare across the three different implementations of the optimization approach. This is the sum of the (common) \$661 transfer component and a pure-insurance value that ranges from \$30 in the health-based optimization approach to \$760 in the consumption-based optimization approach using the consumption proxy. As a result, our estimate of $\gamma(1)$ ranges from \$690 to \$1421. The transfer component represents a large share of this total value. In the health-based optimization approach virtually all the value of Medicaid to recipients comes from the transfer component, and in the consumption-based optimization approach using the CEX more than 80% of the value comes from the transfer component. In the consumption-based optimization approach using the consumption proxy, just under half of the welfare value comes from the transfer component, which is similar to what we found in the complete-information approach.

5 Interpretation

5.1 Benchmarks

We benchmark these welfare estimates against the costs of Medicaid and against the monetary transfers to external parties. Government costs, G, are simply total medical spending for treatment compliers ($q = 1$) given that treatment compliers have no out-of-pocket spending (see equation (20)). Table 1 indicates that G is \$3,600 per recipient year. This is broadly consistent with external estimates of annual per-recipient spending in the Medicaid program in Oregon (Wallace et al. (2008)[75]). The net cost of Medicaid, C, is \$1,448, which according to equation (21) equals the impact of Medicaid on total medical spending (\$879) plus the reduction in out-of-pocket spending (\$569).

The monetary transfer from Medicaid to external parties, N, is the difference between G and C, or \$2,152 (see equation (22)). Thus, about 60 cents of every dollar of government spending on Medicaid is a transfer to external parties ($N/G \approx 0.6$).

With these estimates in hand, we discuss a number of potentially illuminating comparisons. These are summarized in panel B of Table 2.

Recipients vs. external parties. One striking finding is the magnitude of the monetary transfer to external parties relative to our estimates of the value of Medicaid to recipients. Depending on

---

expectation of these conditional expectations by weighting them by the fraction of control compliers of each type θ^K.

the approach used, the results in Table 2 indicate that the value of Medicaid to recipients, $\gamma(1)$, is roughly between one-third and three-quarters of the monetary transfer to external parties, N.

The identity of these "external parties" is an important and open question. The provision of uncompensated care by hospitals is a natural starting point. Recent evidence indicates that hospital visits by the uninsured are associated with very large unpaid bills (Dobkin et al. (2016)[26]) and that increases in Medicaid coverage lead to large reductions in uncompensated care by hospitals (Garthwaite, Gross, and Notowidigdo (2015)[37]).

The ultimate economic incidence of the transfers to external parties is even more complicated; while some of the incidence may fall on the direct recipients of the monetary transfers, other parties may bear some of the incidence, including the privately insured, the recipients themselves (for example, if reductions in unpaid medical debt improve their credit scores)[37], and the public sector budget; indeed, Hadley et al. (2008)[41] estimate that 75 percent of "uncompensated care" for the uninsured is paid for by the government. In addition, if one were to trace through these incidences and estimate the impact of Medicaid on the welfare of other parties, this might be above or below the dollar value N that we have used here. Given the magnitude of these external transfers, we consider a better understanding of their ultimate economic incidence an important area for future work. Among other things, it has important implications for the redistributive nature of Medicaid.

Value to recipients relative to gross cost. As noted in the Introduction, the Congressional Budget Office currently uses G for the value of Medicaid for recipients (Congressional Budget Office (2012)[67]). However, a priori, $\gamma(1)$ may be less than or greater than G. If rational individuals have access to a well-functioning insurance market and choose not to purchase insurance, $\gamma(1)$ will be less than G. However, if market failures such as adverse selection (e.g., knowledge of θ when choosing insurance) result in private insurance not being available at actuarially fair prices, $\gamma(1)$ could exceed G, although it might not if moral hazard costs and crowd-out of implicit insurance (i.e., N) sufficiently reduce $\gamma(1)$. Ultimately these are empirical questions.

Across our baseline approaches, we consistently estimate that $\gamma(1)$ is less than G. Depending on the approach used, we estimate a ratio of $\gamma(1)/G$ between \$0.2 and \$0.5. This implies that Medicaid recipients would rather give up Medicaid than pay the government's costs of providing Medicaid; likewise, an uninsured person would choose the status quo over giving up G in consumption to obtain Medicaid. However, $\gamma(1) < G$ does not answer the question of whether an uninsured person would prefer receiving Medicaid to receiving G in additional consumption or, equivalently, whether an insured person would be willing to give up Medicaid in exchange for a consumption increase of G.[38]

---

[37]In our context, the evidence from the Oregon Health Insurance experiment indicates that only about half the sample had revolving credit, and Medicaid receipt did not affect credit scores (Finkelstein et al. (2012)[34]).

[38]As discussed in Section 2.2, an estimate of $\gamma(1)$ below G does *not* imply that the uninsured would prefer an increase in consumption of G to receiving Medicaid. To investigate the latter question, we could replace the baseline definition of $\gamma(1)$ in equation (3) by $\tilde{\gamma}(1)$, defined as the implicit solution to $E\left[u\left(c\left(0;\theta\right) + \tilde{\gamma}(1), h\left(0;\theta\right)\right)\right] = E\left[u\left(c\left(1;\theta\right), h\left(1;\theta\right)\right)\right]$. We find (in results not reported) that $\tilde{\gamma}(1)$ is considerably higher than $\gamma(1)$, namely \$4,217 compared to \$1,675 for the complete-information approach. Note however that, while $\tilde{\gamma}(1)$ is in the spirit of an equivalent variation of gaining Medicaid, it is an overestimate of the true equivalent variation of gaining Medicaid

Value to recipients relative to net costs. Comparing $\gamma(1)$ to G also does not answer the question of whether the insured would choose to give up Medicaid if they had to cover the *net* (rather than gross) costs of Medicaid. The gross costs of Medicaid (G) greatly exceed its *net* costs ($C = G - N$) since, as noted, about \$0.6 on the dollar of G goes to pay for coverage the uninsured were effectively receiving from external parties prior to Medicaid (i.e., N).

We therefore examine whether the value to recipients exceeds the net costs of Medicaid (i.e., $\gamma(1)/C$)).[39] We think of this as a useful thought exercise even though it is not clear that this question always has a corresponding practical implementation option. For example, it is not obvious how to deliver Medicaid without the transfer to external parties. However, to the extent that the government is itself a major recipient of the transfers to "external parties" (Hadley et al. (2008)[41]) our net cost estimate C may approximate the "true" cost of Medicaid to the public sector.

The value of Medicaid to recipients may be higher than its net cost due to its insurance value, or it may be lower because the moral hazard effects of insurance increase net costs in excess of the value to recipients. Thus $\gamma(1)$ above C implies that the insurance value of Medicaid to recipients, I, exceeds the moral hazard costs of Medicaid, $G - N - T$, while $\gamma(1)$ below C implies the converse.[40] Estimates of the insurance value of Medicaid vary considerably across approaches - from \$30 to \$976. The bottom rows of Table 2 report $\gamma(1)/C$ and the moral hazard costs of Medicaid ($G - N - T$) implied by each approach. The results indicate that, depending on the approach, the value of Medicaid relative to net costs (i.e., $\gamma(1)/C$) varies from about 0.5 to 1.2.

Of course, even when the private valuation of Medicaid to recipients is lower than net costs, it may still be the case that the social valuation of Medicaid exceeds net costs because Medicaid provides benefits to a low-income group. Indeed, a relevant question is whether Medicaid is less or more costly as a method of redistribution as compared to other redistributive programs. For example, Hendren (2016)[44] estimates that recipients of the EITC - a different low-income population - value the EITC at \$0.88 per dollar of net cost.

Value to recipients if the uninsured had no implicit insurance. Another way to illustrate the importance of the monetary transfers to the external parties is to estimate the value of Medicaid to recipients in the hypothetical scenario in which the uninsured have to pay the full cost of their medical care, i.e., $p(0) = 1$. Creating this counterfactual scenario requires extrapolating out of

---

because we have not allowed the uninsured individuals to re-optimize their choice of m versus c after the hypothetical receipt of $\tilde{\gamma}(1)$ in additional consumption. Such a problem does not arise under our baseline measure in which $\gamma(1)$ is subtracted from individuals with Medicaid, since medical care is free for those with Medicaid (and assuming that the income elasticity of the demand for m is zero when individuals face a zero price of m).

[39] Alternatively in a more general setup, one could try to account for the incidence of N and the heterogeneous costs to the government of taxing back N from external parties when providing Medicaid; we pursued such an exercise in the working paper version of this paper (see Finkelstein et al., (2015)[31]).

[40] By definition, $\gamma(1) = T + I$ and $C = G - N$. Therefore a comparison of I to $G - N - T$ is equivalent to a comparison of $\gamma(1)$ to C. The transfer value T is not simply an estimate of a monetary transfer (as with N) but rather of the value that recipients place on the transfer of resources. This may be less than its monetary cost because - as discussed in Section 2 above - part of the transfer component arises from the increase in medical spending induced by Medicaid; this "moral hazard" (or induced medical spending from subsidizing its price) is presumably valued at less than cost.

sample from the observed demands for medical care at $p = 0$ (for treatment compliers) and at $p = 0.21$ (for control compliers) to the demand for medical care at $p = 1$; we do this by assuming that the demand for medical care is log-linear in p.[41] As this is far out of sample from the prices we observe, the results of this exercise should be taken with the appropriate grains of salt.

The results suggest that the value of Medicaid to beneficiaries would likely be much larger if the uninsured faced an environment where they paid all of their medical costs. We estimate a value of Medicaid under the complete-information approach of \$2,749 if there were no implicit insurance for the uninsured, compared to our baseline estimate of \$1,675 (Table 2). Our estimates under the other approaches are higher as well: \$3,875 for the consumption-based optimization approach (compared to \$1,421 for our baseline estimate in Table 2) and \$2,233 for our health-based optimization approach (compared to \$690 in our baseline).[42]

5.2 Tradeoffs across alternative approaches

Which approach - or combination of approaches - to estimating $\gamma(1)$ one prefers depends on how confident one is with the various assumptions required by each approach. We highlight a few considerations and opinions here that are specific to our empirical application.

We are the least confident in the results from the health-based optimization approach, due to the conceptual and empirical challenges in estimating heterogeneity in health returns to out-of-pocket spending. In addition, our implementation of the health-based optimization approach may be biased upward for individuals at a corner solution with respect to medical care (see footnote 11) and biased downward if our estimation of the health production function conditional on θ^K misses any within-θ^K insurance (see footnote 22).[43]

We highlight a few areas of possible concern that apply to the other approaches. First, our consumption-based estimates based on the consumption proxy measure may be biased upward. The consumption proxy measure (used in the complete-information approach and in one variant of the consumption-based optimization approach) models consumption as average consumption adjusted for out-of-pocket expenses and therefore ignores the possibility of the uninsured smoothing consumption through other means such as savings, borrowing, or transfers from friends or family.[44]

Second, the optimization approaches may be biased downward because we assume a constant out-of-pocket price of medical care for the uninsured. If, however, the uninsured face a range of

---

[41]Once we have an estimate of the (counterfactual) distribution of m at $p = 0$, this straightforwardly implies counterfactual distributions of x and of c (in our consumption-proxy based approach). For the complete-information approach we also need a counterfactual estimate for the mean of h, which we get by simple linear extrapolation.

[42]We can also consider a lower bound for the value of Medicaid with no implicit insurance for the uninsured, under the assumption that for demand for m is sufficiently elastic that no medical care is consumed at $p = 1$. Even at the lower bound, the estimate for $\gamma(1)$ for the optimization-based approaches is \$1800, higher than the corresponding baseline estimates in Table 2. The lower bound for the complete-information approach is \$1151, which is not binding because it is lower than the corresponding estimate in Table 2.

[43]In practice, bias from corner solutions is likely small as it shows up only in the pure-insurance term and therefore at most biases the estimate of $\gamma(1)$ up by our estimate of the pure-insurance term of \$30.

[44]Note that our estimates do reflect any direct payment of bills by other parties (so that the out-of-pocket spending, x, is reduced); what we are ruling out is the ability of the individual to smooth the impact of x with help from other parties.

out-of-pocket prices across different treatments and are more likely to undergo treatments with a low price, then our estimate of the impact of Medicaid on the out-of-pocket spending schedule will be biased down because it is based on the selected sample of treatments undergone.

A related issue for our optimization approaches – which could create bias of either sign – is that our estimate of $p(0)$ is based on the average price for the uninsured, while the relevant price for welfare analysis is the marginal price of medical care for the uninsured. The marginal price may be higher than the average price – if the uninsured tend to avoid treatments for which they would have to pay a higher out-of-pocket price – or it might be lower than the average price – if above a certain level of expenditures, the uninsured effectively face no out-of-pocket costs (Mahoney (2015)[52]). A downward bias in our estimate of $p(0)$ reduces the estimate of $\gamma(1)$ (see equation 11) and, incidentally, creates an upward bias in the effect on external parties, N. An upward bias in $p(0)$ has the opposite effect.

Finally the linear interpolation between $\frac{d\gamma(0)}{dq}$ and $\frac{d\gamma(1)}{dq}$ that we use to implement the optimization approaches (see Assumption 5) may downward bias our estimates of $\gamma(1)$ since it does not allow for the possibility that some of the benefit of health insurance may operate via an "access motive" in which additional income (or liquidity) allows for discontinuous or lumpy changes in health care consumption (Nyman (1999a,b)[55, 56]).[45] By contrast, the complete-information approach would accurately capture the value stemming from the liquidity Medicaid provides.[46] Of course, since the complete-information approach requires specifying all arguments of the utility function while the optimization approaches do not, omission of any utility-relevant outcomes that are affected by Medicaid may cause the estimate of $\gamma(1)$ from the complete-information approach to be biased either up or down.

6 Sensitivity Analysis

We explore sensitivity to a number of key assumptions. Such "specification" uncertainty is not reflected in the standard errors in Table 2; the sensitivity analysis here provides an informal way of gauging sensitivity to these assumptions.

6.1 Alternative assumptions unrelated to health

Table 4 explores the sensitivity of our results within each framework (shown in different rows) to a number of different non-health assumptions (shown in different columns). For the sake of brevity,

---

[45]Consider an extreme example in which there is a single expensive medical procedure that individuals may undergo in the event of a health shock. Individuals are sufficiently liquidity constrained that they will undertake this procedure only if $q \geq 0.4$. As a result $\frac{d\gamma}{dq}$ would be zero until $q = 0.4$, jump up at $q = 0.4$ and possibly decline thereafter. The linear approximation in Figure 1 would not capture the relatively large values of $\frac{d\gamma}{dq}$ that could occur for intermediate values of q and would in that case underestimate the welfare effect of Medicaid on the recipient.

[46]This is an important difference relative to the complete-information approach which, because it specifies a full utility function, can deliver non-marginal welfare estimates directly. In contrast, the optimization approaches follow the spirit of Harberger's classic triangle (Harberger (1964)[42]) and approximate non-marginal welfare statements using statistical interpolations.

we focus the discussion on two main results: the value of Medicaid to recipients relative to the transfer to external parties ($\gamma(1)/N$), and the value of Medicaid to recipients per dollar of net costs ($\gamma(1)/C$). Appendix Table 3 reports results from the alternative specifications for all the main elements shown in Table 2 so that other comparisons are easily made. Column I shows our baseline results from Table 2.

Risk aversion and consumption floor. Columns II through V explore alternative choices for risk aversion (coefficients of relative risk aversion of 1 and 5, compared to our baseline of 3) and the consumption floor (of \$1,000 or \$5,000, compared to our baseline assumption of \$1,977). A lower consumption floor increases the value of Medicaid to recipients $\gamma(1)$ using the complete-information approach or the consumption-based optimization approaches, and does not affect our estimates of C or N.[47] In the full-information approach and in the consumption-based optimization approaches, higher risk aversion raises our estimate of $\gamma(1)$ and lower risk aversion lowers it (as expected). The estimates of the health-based optimization approach do not depend on risk aversion by construction.

Alternative measure of out-of-pocket spending for those on Medicaid $(x(1,m))$**.** In our baseline analysis, we assume that, consistent with Medicaid rules, the insured have no out-of-pocket spending ($x(1,m) = 0$). In practice, however, the insured in our data report nontrivial out-of-pocket spending (Finkelstein et al. (2012)[34]).[48] In column VI, we therefore present estimates from an alternative approach in which we re-estimate all of our fitted consumption and out-of-pocket spending distributions based on self-reported out-of-pocket spending for treatment compliers as well as control compliers.[49] We now have to estimate the "pure-insurance" term in equation (12) at $q = 1$, since we no longer assume full insurance at $q = 1$ as in the baseline analysis; our estimate of this term is not exact due to a technical complication relating to re-optimization in response to income effects.[50]

Allowing for $x(1,m) > 0$ necessarily reduces our estimates of $\gamma(1)$ but it also reduces our estimates of C (and hence N), so that the net effect on $\gamma(1)/N$ or $\gamma(1)/C$ is a priori ambiguous.

---

[47]The health-based optimization approach is mechanically unaffected by our assumption regarding the the coefficient of relative risk aversion because it does not use this parameter. In principle, it should likewise not be affected by our assumption regarding the consumption floor. In practice, because – as discussed in Section 4 – we implement the consumption floor by adjusting out-of-pocket spending such that the consumption floor holds and because we use the same estimates of of out-of-pocket spending for all approaches, the health-based optimization approach is indirectly affected by our assumption regarding the consumption floor since it affects the estimates of out-of-pocket spending.

[48]This does not appear to be an artifact of our data or setting; in the Medical Expenditure Panel Survey, Medicaid recipients also self-report substantial out-of-pocket spending (Gross and Notowidigdo (2011)[39]).

[49]In constructing $-\frac{\partial x}{\partial q} = (p(0)-p(1))m(q,\theta) = x(0, m(q;\theta)) - x(1, m(q;\theta))$, we assume quantile stability so that we can approximate $-\frac{\partial x}{\partial q}$ using the difference in out-of-pocket spending quantiles for the given distribution of medical spending, $m(q,\theta)$, at insurance level q. Further details on the construction of $-\frac{\partial x}{\partial q}$ when at least some Medicaid recipients have strictly positive out-of-pocket expenditures ($x(1,m) > 0$) can be found in Appendix A.7.

[50]Specifically, under the conceptual thought experiment in which individuals "pay" $\gamma(1)$ units of consumption, they will re-optimize over m and c if m has a nonzero income elasticity. In Appendix A.1, we show that failure to take this income effect into account corresponds to omitting a term from the definition of $\frac{d\gamma(q)}{dq}$ that captures the individual's willingness to pay to re-optimize; this additional term is zero by construction at $q = 0$, and is also zero at $q = 1$ under our baseline assumption that $x(1,m) = 0$.

In practice, column VI shows that it substantially lowers our estimates of the value of Medicaid relative to either transfers to external parties or to net costs.

Alternative assumption about within-family smoothing. Our baseline consumption proxy approach assumed that out-of-pocket medical spending reduced consumption of each family member by the same amount. Substantial within-family risk smoothing seems likely, given how much of consumption is joint (e.g., housing). But the extreme of full smoothing within the family (i.e., the effect on an individual's consumption is the same regardless of whether the individual or a family member incurred the out-of-pocket medical spending) may not be warranted. In column VII, therefore, we examine the sensitivity of our results to the alternative extreme assumption: that the out-of-pocket spending affects consumption only for the individual who incurred the expenses. This substantially raises our estimates of the value of Medicaid relative to their transfers to external parties or net costs under the two approaches that use the consumption proxy: the complete-information approach and the consumption-based optimization approach using the consumption proxy.

Alternative interpolations in the optimization approaches. In the baseline optimization approaches, we assumed $d\gamma/dq$ was linear in q to interpolate between $q = 0$ and $q = 1$ (see Assumption 5). Here, we explore the sensitivity of our results to alternative interpolations; Appendix A.8 provides implementation details. In column VIII, we assume instead that the demand for medical care is linear in price.[51] For the consumption-optimization approach, this yields a ratio of $\gamma(1)/N$ of 0.60, as opposed to the baseline estimate of 0.66, and a ratio of $\gamma(1)/C$ of 0.89 as opposed to the baseline estimate of 0.98. For the health optimization approach, we estimate these ratios to be 0.34 and 0.50 respectively, similar to the baseline specification. In column IX, we calculate an upper bound for $\gamma(1)$ over possible interpolation assumptions by searching for the (nonparametric) functional form for the demand for medical care that maximizes $\gamma(1)$, with the restriction that demand at values of $q \in (0, 1)$ must lie somewhere between demand at $q = 0$ and at $q = 1$.[52] For the consumption-optimization approach, the upper bound for $\gamma(1)/N$ is 1.42. This suggests that even under fairly extreme assumptions about the shape of demand and the utility function, the impact on external parties is still a substantial component of Medicaid. The estimate of 2.12 for the ratio $\gamma(1)/C$ suggests that it is possible for the value of Medicaid to significantly exceed its net resource cost under alternative assumptions about the shape of demand and the utility function.

---

[51]Because the transfer term is linear in m and because q is linear in p, the transfer term is unaffected; only the pure-insurance term is affected by this alternative assumption.

[52]We do not report an upper bound interpolation for the consumption-based optimization approach using the CEX consumption measure because these alternative interpolations require knowledge of m which is not observed in the CEX data.

6.2 Alternative health values and measures

Table 5 explores the robustness of the estimates of $\gamma(1)/N$ and $\gamma(1)/C$ to different assumptions on the valuation of health and different health measures; Appendix Table 4 shows the robustness of all estimates from Table 2. Column I replicates our baseline estimates. Naturally, assumptions regarding health valuation matter only for the complete-information approach and the health-based optimization approach; one attraction of the consumption-based optimization approach is that it does not require us to estimate and value health improvements.

The results suggest that the estimates are reasonably stable across specifications. We focus our discussion on the complete-information results, given that the health-based optimization approach is, not surprisingly, less sensitive to health assumptions because most of the welfare benefit in this approach comes from the transfer component, which does not depend on the health assumption.

Valuing health improvements (ϕ). In column II, we assume the value of health benefits is 0. This is motivated by the fact that while many measures of self-reported health improved, we are unable to reject the null of no impact of Medicaid on mortality (Finkelstein et al. (2012)[34]) or on our specific measures of physical health (Baicker et al. (2013)[10]). Therefore, an alternative of "no health benefits" seems a not unreasonable bound. Since the health component of the value of Medicaid was fairly small relative to the consumption component in the complete-information approach, this has a relatively small effect on the estimates.

Our baseline implementation assumed a VSLY (ϕ) of $5,000$ for our low-income population. This came from scaling the "consensus" estimate of a VSLY of $100,000$ for the general population by the ratio of the marginal utility of consumption for our population and the general population. In column III, we instead assume that the VSLY scales linearly with consumption; we therefore use a VSLY of $40,000$ rather than $5,000$. In the complete-information approach $\gamma(1)/N$ rises from 0.8 to 1.4, while $\gamma(1)/C$ rises from 1.2 to 2.1. These are substantial quantitative changes, although they do not alter the main qualitative conclusions from the complete-information approach: monetary transfers to external parties are nontrivial relative to the value to recipients and the value of Medicaid to recipients exceeds its net costs. While assuming a linear scaling with consumption is ad hoc and conceptually inconsistent with our assumption of $\sigma = 3$, it is closer than our baseline assumption to the findings of Kniessner et al. (2010)[48], who estimate an elasticity of the VSLY with respect to income of about 1.4.

As an alternative to using an external estimate of the VSLY, we can also estimate the VSLY from the first-order condition for m in equation (16). In other words, given our estimates of the return to medical spending $(d\tilde{h}/dm)$ and the price p, we can estimate what u_h/u_c must be for the first-order condition to hold for the observed choices of m. Given that we express h in QALYs, the ratio u_h/u_c is the VSLY. We find that a VSLY of $5,364$ causes the first-order condition to hold on average for compliers. This estimate is biased upward due to corner solutions because those who choose $m = 0$ place a lower value on a statistical life year than is implied by our estimate of u_h/u_c. We found the closeness of this internally-derived estimate of the VSLY to our baseline assumption

from external estimates ($5000) to be broadly reassuring regarding our baseline assumption.

Alternative health measures. In the remaining columns, we return to our baseline VSLY and explore robustness to alternative health measures. In column IV, we examine the 2-item Patient Health Questionnaire (PHQ-2) measure, which is commonly used as a depression screen. The PHQ-2 asks respondents about the prevalence in the last two weeks of having been "bothered by little interest or pleasure in doing things" and of having been "bothered by feeling down, depressed, or hopeless." We rely on the estimates from Pyne et al. (2009)[60], which are based on the "Standard Gamble" approach, to convert the PHQ-2 responses to QALYs; Appendix A.4 provides details of this conversion. We estimate that Medicaid increases health by 0.027 QALYs based on the PHQ-2 health measure, as compared to 0.045 QALYs under our baseline self-reported health measure. The welfare estimates for the PHQ-2 health measure are correspondingly lower for the complete-information approach, which is not surprising given that the complete-information approach requires a comprehensive health measure wheres PHQ-2 only measures mental health. The welfare estimates for the health-based optimization approach are relatively stable, which is consistent with the fact that the optimization approach does not require a comprehensive health measure.

We also draw on a separate data source - based on a series of in-person interviews conducted in the Portland metro area about two years after the lottery - that has additional health measures not available in our baseline data: the 8-item Patient Health Questionnaire (PHQ-8) and the "Short Form" health survey (SF-8).[53] Columns V through VIII show the results. In column V, we replicate our baseline self-assessed health measure for the subsample of the in-person data that answered that question as well as all questions of the PHQ-8. Columns V and I show how results compare when we use the same self-assessed health measure in the mail-in survey and the in-person survey (which is limited to the Portland area). In the in-person data, we can measure health using the same PHQ-2 measure used in the mail survey (column VI) and the richer PHQ-8 measure (column VII). Welfare estimates using the PHQ-8 measure are quite similar to those using the PHQ-2 measure (compare columns VI and VII).

We also use data from the in-person interviews to see how our estimates change if we use the SF-8. The SF-8 is a general health survey that captures both physical and mental health. We convert the SF-8 to QALYs using the mapping from Sullivan and Ghushchyan (2006)[65]. Unlike the previous mappings to QALYs we used, which were all based on the "Standard Gamble" method, this last mapping uses the other principal method in the literature: the "Time-Trade-Off" method; Appendix A.4 provides more detail. Column VIII shows that results using the SF-8 measure are similar to results with the PHQ measures.

---

[53]The in-person interview data are available for fewer individuals than the mail survey, and are limited to individuals in the Portland metro area. However, in practice, empirical estimates from the two data sets are quite similar (see Finkelstein et al. (2012)[34] and Baicker et al. (2013)[10] respectively).

Summary. Looking across the columns of Tables 4 and 5 gives a sense of the scope and drivers of our estimates. In the complete-information approach, the biggest impact on the estimates comes from assuming $\sigma = 5$; this raises our estimate of $\gamma(1)/N$ from 0.8 to 1.9, and our estimate of $\gamma(1)/C$ from 1.2 to 2.8. The next biggest effect on the estimates comes from assuming a VSLY of \$40,000 instead of \$5,000. Still, the quantitatively large role of external transfers relative to the value of Medicaid to recipients remains; even at its smallest relative size, the size of external transfers remains over half the size of the value of Medicaid to recipients.

Under the consumption-based optimization approach using the consumption proxy, the biggest change comes from assuming that the shock is borne entirely by the individual. This more than doubles our estimate of $\gamma(1)/N$ - from 0.7 to 1.5 - but still results in a size of external transfers that is over two-thirds as large as the value of Medicaid to recipients; it also more than doubles our estimate of $\gamma(1)/C$ from 1.0 to 2.2. The consumption-based optimization approach using the CEX consumption measure is more stable. The biggest impact comes from assuming $\sigma = 5$; this raises $\gamma(1)/N$ from 0.7 to 0.8, and $\gamma(1)/C$ from 0.55 to 0.60.

7 Conclusion

Welfare estimation of non-market goods is important, but also challenging. As a result, the welfare benefits from Medicaid are often ignored in academic and public policy discourse, or based on ad-hoc approaches. In this paper, we developed, implemented, and compared the results from alternative formal frameworks for valuing a Medicaid expansion for low-income, uninsured adults that occurred by random assignment in Oregon.

Not surprisingly, the "bottom line" is open to interpretation. We have endeavored to describe how the results vary with the framework used as well as the specific implementing assumptions.

However, one key, robust result that emerges is that the monetary transfers from Medicaid to external parties are quantitatively important relative to the welfare benefits of Medicaid to recipients. In our baseline estimates, the value of Medicaid to recipients is consistently lower than - roughly one-third to three-quarters of - the size of the monetary transfers to external parties. The key driving factor behind this result is our related finding that the low-income uninsured in our sample pay only a small fraction of medical expenditures; we confirmed that this holds not only in our context but more generally in national survey data as well. As a result, we estimate that about \$0.6 of every dollar of government Medicaid spending does not accrue directly to recipients but instead replaces implicit partial insurance for the low-income uninsured. These findings highlight the importance of further work on who bears both the immediate and ultimate economic incidence of the large Medicaid transfers to external parties; they also raise an interesting question of the relative efficiency of formal provision of public insurance through Medicaid as opposed to the provision of implicit insurance by third parties.

Another finding that emerges is that Medicaid recipients would rather lose Medicaid than forego consumption equal to the government's costs of providing Medicaid. Our baseline estimates indicate

that the welfare benefit to recipients per dollar of government spending is between \$0.2 and \$0.5, depending on the framework used. This result - which is directly related to our finding that a large share of Medicaid spending accrues to non-recipients - contrasts with the current approach used by the Congressional Budget Office to value Medicaid at government cost.

A more subtle question is whether recipients value Medicaid at more than its *net* (of external transfers) resource cost. This thought exercise addresses the efficiency gains from Medicaid: the value of Medicaid to recipients will exceed its net cost if the (budget-neutral) insurance value to recipients exceeds the moral hazard costs of Medicaid. Our findings suggest that the answer varies with the approach used; our baseline estimates of the value of Medicaid to recipients relative to its net costs range from 0.5 to 1.2. Most of our specifications give a number below 1, suggesting that recipients are not willing to pay the net costs of Medicaid. Of course, from a redistributive perspective, it is possible that Medicaid may be an efficient method of redistribution relative to the available alternatives that also generally involve some social cost.

Our empirical findings are naturally specific to our setting. Fortunately, the approaches we have developed can be applied to studying the value of Medicaid in other contexts. In particular, as noted in Finkelstein et al. (2012)[34], the impact of Medicaid may well differ when it is mandatory rather than voluntary, when it is expanded to cover a larger number of individuals, or when it is provided over a longer time horizon. It would also be interesting to apply the approaches we have developed to studying the value of Medicaid for other populations. The low-income adult population covered by Medicaid through the Oregon Health Insurance Experiment is of particular interest, given its similarity to those newly covered by the 2014 Medicaid expansions under the Affordable Care Act. However, the welfare effects of Medicaid are potentially very different for other Medicaid populations, such as children, the disabled, or the elderly, for whom there is also a large empirical literature on Medicaid's effects (see Gruber (2003)[40] and Buchmueller et al. (2015)[15] for reviews). Future work could also consider the value of other public health insurance programs; for example, there is a large empirical literature examining the impacts of Medicare on health care use, health, and out-of-pocket medical expenditures (e.g., Card et al. (2008, 2009)[16, 17], Barcellos and Jacobson (2015)[11]) to which our frameworks could be applied.

Our paper illustrates the possibilities – but also the challenges – in doing welfare analysis even with a rich set of causal program effects. Behavioral responses are not prices and do not reveal willingness to pay without additional assumptions. We provide a range of potential pathways to welfare estimates under various assumptions, and offer a range of estimates that analysts can consider, rather than the common defaults of zero valuation or valuation at government cost. We hope the flexibility offered by these approaches provides guidance to future research examining the welfare impact of the public provision of other non-market goods.

References

[1] Alberto Abadie. Bootstrap tests for distributional treatment effects in instrumental variable models. *Journal of the American Statistical Association*, 97(457):284–292, 2002.

[2] Alberto Abadie. Semiparametric instrumental variable estimation of treatment response models. *Journal of Econometrics*, 113:231–263, 2003.

[3] Joseph E. Aldy and W. Kip Viscusi. Adjusting the value of a statistical life for age and cohort effects. *Review of Economics and Statistics*, 90(3):573–581, 2008.

[4] Douglas Almond, Joseph J. Doyle Jr., Amanda E. Kowalski, and Heidi Williams. Estimating marginal returns to medical care: Evidence from at-risk newborns. *Quarterly Journal of Economics*, 125(2):591–634, 2010.

[5] Joshua D. Angrist and Jörn-Steffen Pischke. Mostly harmless econometrics: An empiricist's companion. *Princeton University Press*, 2009.

[6] Kenneth J. Arrow. Uncertainty and the welfare economics of medical care. *American Economic Review*, 53(5):941–973, 1963.

[7] Orley Ashenfelter and Michael Greenstone. Using mandated speed limits to measure the value of a statistical life. *Journal of Political Economy*, 112(1):226–267, 2004.

[8] Katherine Baicker, 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. *American Economic Review: Papers and Proceedings*, 104(5):322–328, 2014.

[9] Katherine Baicker, Sendhil Mullainathan, and Joshua Schwartzstein. Behavioral hazard in health insurance. *Quarterly Journal of Economics*, 130(5):1623–1667, 2015.

[10] Katherine Baicker, Sarah Taubman, Heidi Allen, Mira Bernstein, Jonathan Gruber, Joseph Newhouse, Eric Schneider, Bill Wright, Alan Zaslavsky, Amy Finkelstein, and the Oregon Health Study Group. The Oregon experiment – effects of Medicaid on clinical outcomes. *New England Journal of Medicine*, 368:1713–1722, 2013.

[11] Silvia Helena Barcellos and Mireille Jacobson. The effects of Medicare on medical expenditure risk and financial strain. *American Economic Journal: Economic Policy*, 7(4):41–70, 2015.

[12] Robert B. Barsky, F. Thomas Muster, Miles S. Kimball, and Matthew D. Shapiro. Preference parameters and behavioral heterogeneity: An experimental approach in the Health and Retirement Study. *Quarterly Journal of Economics*, 112(2):537–579, 1997.

[13] John E. Brazier, Yaling Yang, Aki Tsuchiya, and Donna Louise Rowen. A review of studies mapping (or cross walking) non-preference based measures of health to generic preference-based measures. *European Journal of Health Economics*, 11(2):215–225, 2010.

[14] Jeffrey R. Brown and Amy Finkelstein. The interaction of public and private insurance: Medicaid and the long-term care insurance market. *American Economic Review*, 98(3):1083–1102, 2008.

[15] Thomas Buchmueller, John Ham, and Lara Shore-Sheppard. Medicaid. In Robert A. Moffitt, editor, *Means-tested transfer programs in the United States, II*. University of Chicago Press, in progress.

[16] David Card, Carlos Dobkin, and Nicole Maestas. The impact of nearly universal insurance coverage on health care utilization: Evidence from Medicare. *American Economic Review*, 98(5):2242–2258, 2008.

[17] David Card, Carlos Dobkin, and Nicole Maestas. Does Medicare save lives? *Quarterly Journal of Economics*, 124(2):597–636, 2009.

[18] James H. Cardon and Igal Hendel. Asymmetric information in health insurance: evidence from the National Medical Expenditure Survey. *RAND Journal of Economics*, pages 408–427, 2001.

[19] Raj Chetty. Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods. *Annual Review of Economics*, 1(1):451–488, 2009.

[20] John H. Cochrane. Asset pricing. *Princeton University Press*, 2005.

[21] Teresa A. Coughlin, John Holahan, Kyle Caswell, and Megan McGrath. Uncompensated care for uninsured in 2013. Technical report, Kaiser Commission on Medicaid and the Uninsured, May 2014.

[22] Maureen Cropper, James K. Hammitt, and Lisa A. Robinson. Valuing mortality risk reductions: Progress and challenges. *Annual Review of Resource Economics*, 3:313–336, 2011.

[23] Janet Currie and Firouz Gahvari. Transfers in cash and in-kind: Theory meets the data. *Journal of Economic Literature*, 46(2):333–383, 2008.

[24] David M. Cutler. Your money or your life: Strong medicine for America's health care system. *New York: Oxford University Press*, 2004.

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

[26] Carlos Dobkin, Amy Finkelstein, Raymond Kluender, and Matthew J. Notowidigdo. The economic consequences of hospital admissions. NBER Working Paper 22288, 2016.

[27] Isaac Ehrlich and Gary S. Becker. Market insurance, self-insurance, and self-protection. *Journal of Political Economy*, 80(4):623–648, 1972.

[28] Liran Einav, Amy Finkelstein, and Jonathan Levin. Beyond testing: Empirical models of insurance markets. *Annual Review of Economics*, 2(1):311–336, 2010.

[29] Liran Einav, Amy Finkelstein, and Paul Schrimpf. Optimal mandates and the welfare cost of asymmetric information: Evidence from the U.K. annuity market. *Econometrica*, 78(3):1031–1092, May 2010.

[30] Martin Feldstein and Jonathan Gruber. A major risk approach to health insurance reform. In James M. Poterba, editor, *Tax Policy and the Economy, Volume 9*, pages 103–130. MIT Press, 1995.

[31] Amy Finkelstein, Nathaniel Hendren, and Erzo F.P. Luttmer. The value of Medicaid: Interpreting results from the Oregon Health Insurance Experiment. *NBER Working Paper No. 21308*, 2015.

[32] Amy Finkelstein, Erzo F.P. Luttmer, and Matthew J. Notowidigdo. What good is wealth without health: The effect of health on the marginal utility of consumption. *Journal of the European Economic Association*, 11(s1):221–258, 2013.

[33] Amy Finkelstein and Robin McKnight. What did Medicare do? The initial impact of Medicare on mortality and out of pocket medical spending. *Journal of Public Economics*, 92(7):1644–1668, 2008.

[34] Amy Finkelstein, Sarah Taubman, Bill Wright, Mira Bernstein, Jonathan Gruber, Joseph P. Newhouse, Heidi Allen, and Katherine Baicker. The Oregon Health Insurance Experiment: Evidence from the first year. *Quarterly Journal of Economics*, 127(3):1057–1106, 2012.

[35] William Furlong, David Feeny, George W. Torrance, Charles H. Goldsmith, Sonja DePauw, Zenglong Zhu, Margaret Denton, and Michael Boyle. Multiplicative multi-attribute utility function for the Health Utilities Index mark 3 (HUI3) system: A technical report. McMaster University Center for Health Economics and Policy Analysis Working Paper 98-11, McMaster University, 1998.

[36] Trevor S. Gallen. Measuring the value of Medicaid using TennCare. *Working Paper*, 2014.

[37] Craig Garthwaite, Tal Gross, and Matthew J. Notowidigdo. Hospitals as insurers of last resort. *Working Paper*, 2015.

[38] Peter Gottschalk and Timothy M. Smeeding. Cross-national comparisons of earnings and income inequality. *Journal of Economic Literature*, 35(2):633–687, 1997.

[39] Tal Gross and Matthew J. Notowidigdo. Health insurance and the consumer bankruptcy decision: Evidence from expansions of Medicaid. *Journal of Public Economics*, 95(7-8):767–778, 2011.

[40] Jonathan Gruber. Medicaid. In Robert A. Moffitt, editor, *Means-tested transfer programs in the United States*, pages 15–78. University of Chicago Press, 2003.

[41] Jack Hadley, John Holahan, Teresa Coughlin, and Dawn Miller. Covering the uninsured in 2008: Current costs, sources of payment, and incremental costs. *Health Affairs*, 27(5):w399–w415, 2008.

[42] Arnold Harberger. The measurement of waste. *American Economic Review*, 54(3):58–76, 1964.

[43] Nathaniel Hendren. The inequality deflator: Interpersonal comparisons without a social welfare function. NBER Working Paper 20351, 2014.

[44] Nathaniel Hendren. The policy elasticity. *Tax Policy and the Economy*, 30, 2016.

[45] Hilary W. Hoynes and Erzo F.P. Luttmer. The insurance value of state tax-and-transfer programs. *Journal of Public Economics*, 95(11-12):1466–1484, 2011.

[46] R. Glenn Hubbard, Jonathan Skinner, and Stephen Zeldes. Precautionary saving and social insurance. *Journal of Political Economy*, 103(2):360–399, 1995.

[47] Michael Keane and Robert Moffitt. A structural model of multiple welfare program participation and labor supply. *International Economic Review*, 39(3):553–589, 1998.

[48] Thomas J. Kniesner, W. Kip Viscusi, and James P. Ziliak. Policy relevant heterogeneity in the value of statistical life: New evidence from panel data quantile regressions. *Journal of Risk and Uncertainty*, 40(1):15–31, 2010.

[49] David Laibson, Andrea Repetto, and Jeremy Tobacman. Self-control and saving for retirement. *Brookings Papers on Economic Activity*, 1:91–196, 1998.

[50] Geng Li, Robert F. Schoeni, Sheldon Danziger, and Kerwin Kofi Charles. New expenditure data in the PSID: Comparisons with the CE. *Monthly Labor Review*, 133(2):29–39, 2010.

[51] Henrik Lindhjem, Ståle Navrud, Nils Axel Braathen, and Vincent Biausque. Valuing mortality risk reductions from environmental, transport, and health policies: A global meta-analysis of stated preference studies. *Risk Analysis*, 31(9):1381–1407, 2011.

[52] Neale Mahoney. Bankruptcy as implicit health insurance. *American Economic Review*, 105(2):710–746, 2015.

[53] MEPS. MEPS HC-147 2011 full year consolidated data file. `https://meps.ahrq.gov/mepsweb/data_stats/download_data_files_detail.jsp?cboPufNumber=HC-147`, 2013.

[54] Olivia S. Mitchell, James M. Poterba, Mark J. Warshawsky, and Jeffrey R. Brown. New evidence on the money's worth of individual annuities. *American Economic Review*, 89(5):1299–1318, 1999.

[55] John A. Nyman. The economics of moral hazard revisited. *Journal of Health Economics*, 18(6):811–824, 1999.

[56] John A. Nyman. The value of health insurance: The access motive. *Journal of Health Economics*, 18(2):141–152, 1999.

[57] Arthur Okun. *Equality and Efficiency*. Brookings Institution Press, 1975.

[58] Michael G. Palumbo. Uncertain medical expenses and precautionary saving near the end of the life cycle. *Review of Economic Studies*, 66(2):395–421, 1999.

[59] Thomas Piketty and Emmanuel Saez. Optimal labor income taxation. *Handbook of Public Economics*, 5:391–474, 2013.

[60] Jeffrey M. Pyne, John C. Fortney, Shanti Tripathi, David Feeny, Peter Ubel, and John Brazier. How bad is depression? Preference score estimates from depressed patients and the general population. *Health Services Research*, 44(4):1406–1423, 2009.

[61] Emmanuel Saez and Stefanie Stantcheva. Generalized social welfare weights for optimal tax theory. *American Economic Review*, 106(1):24–45, 2016.

[62] Paul A. Samuelson. The pure theory of public expenditure. *Review of Economics and Statistics*, 36(4):387–389, 1954.

[63] John K. Scholz, Ananth Seshadri, and Surachai Khitatrakun. Are Americans saving 'optimally' for retirement? *Journal of Political Economy*, 114(4):607–643, 2006.

[64] James W. Shaw, Jeffrey A. Johnson, and Stephen Joel Coens. US valuation of the EQ-5D health states: Development and testing of the D1 valuation model. *Medical Care*, 43(3):203–220, 2005.

[65] Patrick W. Sullivan and Vahram Ghushchyan. Mapping the EQ-5D index from the SF-12: US general population preferences in a nationally representative sample. *Medical Decision Making*, 26(4):1401–409, 2006.

[66] Sarah Taubman, Heidi Allen, Bill Wright, Katherine Baicker, Amy Finkelstein, and the Oregon Health Study Group. Medicaid increases emergency-department use: Evidence from Oregon's health insurance experiment. *Science*, 343:263–268, 2014.

[67] U.S. Congressional Budget Office. The distribution of household income and federal taxes, 2008 and 2009. 2012.

[68] U.S. Department of Agriculture. The food assistance landscape: FY 2015 annual report. 2016.

[69] U.S. Department of Health and Human Services. 2015 actuarial report on the financial outlook for Medicaid. 2015.

[70] U.S. Department of Health and Human Services. FY 2015 Federal TANF & state MOE financial data. 2016.

[71] U.S. Internal Revenue Service. EITC calendar year report. 2015.

[72] U.S. Social Security Administration. Social Security Income program 2016 technical materials. February 2016.

[73] Eddy Van Doorslaer and Andrew M. Jones. Inequalities in self-reported health: Validation of a new approach to measurement. *Journal of Health Economics*, 22(1):61–87, 2003.

[74] W. Kip Viscusi. The value of risks to life and health. *Journal of Economic Literature*, 31(4):1912–1946, 1993.

[75] Neal T. Wallace, K. John McConnell, Charles A. Gallia, and Jeanene A. Smith. How effective are copayments in reducing expenditures for low-income adult Medicaid beneficiaries? Experience from the Oregon health plan. *Health Services Research*, 43(2):515–530, 2008.

[76] J.E. Ware, M. Kosinski, J.E. Dewey, and B. Gandek. How to score and interpret single-item health status measures: A manual for users of the SF-8 health survey. Technical report, QualityMetric Incorporated, Lincoln RI, 2001.

[77] Sarah J. Whitehead and Shehzad Ali. Health outcomes in economic evaluation: the QALY and utilities. *British Medical Bulletin*, 96(1):5–21, 2010.

Table 1: Summary Statistics

	I Full Sample	II Treatment Compliers ($q=1$)	III Control Compliers ($q=0$)	IV Impact of Medicaid
Panel A: Oregon Data Demographics				
Share female	0.60	0.57	0.60	
Share age 50-64	0.34	0.36	0.35	
Share age 19-49	0.66	0.64	0.65	
Share white	0.83	0.84	0.84	
Share black	0.03	0.03	0.03	
Share Spanish / Hispanic / Latino	0.11	0.07	0.08	
Mean family size, n	2.97	2.88	2.91	
Panel B: Oregon Data Outcomes				
12-month medical spending, m				
Mean medical spending ($), E[$m$]	2991	3600	2721	879
Fraction with positive medical spending, E[$m>0$]	0.74	0.79	0.72	0.07
12-month out-of-pocket spending, x				
Mean out-of-pocket spending ($), E[$x$]	470	0	569	-569
Fraction with positive out-of-pocket spending, E[$x>0$]	0.38	0	0.56	-0.56
Health expressed in QALYs, E[h]	0.77	0.78	0.74	0.05
Share in poor health (QALY=0.401)	0.11	0.10	0.17	-0.07
Share in fair health (QALY=0.707)	0.30	0.29	0.36	-0.07
Share in good health (QALY=0.841)	0.36	0.38	0.28	0.10
Share in very good health (QALY=0.931)	0.17	0.18	0.15	0.03
Share in excellent health (QALY=0.983)	0.05	0.05	0.04	0.02

Notes: This table reports data from a mail survey of participants in the Oregon Health Insurance Experiment (N=15,498). Columns II and III report the implied means for treatment and control compliers in the Oregon Health Insurance Experiment, and Column IV reports the estimated impact of Medicaid. Columns II, III, and IV use the Oregon health insurance lottery as an instrument for Medicaid coverage.

Table 2: Welfare Benefit Per Recipient

	I	II	III	IV
			Optimization Approaches	
	Complete-Information Approach	Consumption-Based (Consumption Proxy)	Consumption-Based (CEX Consumption Measure)	Health-Based
A. Welfare Effect on Recipients, $\gamma(1)$	1675	1421	793	690
(standard error)	(60)	(180)	(417)	(420)
Transfer component, T	699	661	661	661
Pure-insurance component, I	976	760	133	30
B. Benchmarks				
Welfare effects on recipients relative to:				
monetary transfer to external parties, $\gamma(1)/N$	0.78	0.66	0.37	0.32
gross costs, $\gamma(1)/G$	0.47	0.39	0.22	0.19
net costs, $\gamma(1)/C$	1.16	0.98	0.55	0.48
Moral hazard cost, $G\text{-}T\text{-}N$	749	787	787	787

Notes: Estimates of welfare effects and moral hazard costs are expressed in dollars per year per Medicaid recipient. Standard errors are bootstrapped with 500 repetitions.

Table 3: Measurement of Consumption Covariance in CEX Consumption Approach

	Baseline	II	III
Consumption covariance			
Insured	-305	-309	-179
Uninsured	-40	-41	-86
Difference (= measurement-error corrected covariance)	265	268	93
Definition of non-health consumption	All non-health expenditure	All non-health expenditure excluding durables	Relatively easily adjustable non-health expenditure categories
Mean of non-health consumption (in annual $ per capita)	13,310	11,789	5,174

Notes: This table presents our baseline estimates for the pure-insurance term in the consumption-based optimization approach that uses the consumption measure from the Consumer Expenditure Survey. The sample consists of single adults aged 19-64 below 100% of the federal poverty line (N=1,065). The numbers reported in the table are the covariances of marginal utility of non-health consumption (using a coefficient of relative risk aversion of 3) and out-of-pocket medical spending, normalized by the mean value of the marginal utility of consumption for the relevant population; see equation (14). The consumption measure in the first column consists of all non-health expenditure in the CEX (excluding contributions to private and public pension programs), where we define health expenditure as individual expenditures for health care providers, prescription drugs, and medical devices. The consumption measure in column II is the same as that in column I but excludes expenditures on durables: vehicle purchases, major household appliances, house furnishings and equipment, and entertainment equipment (including TVs and radios). The consumption measure in column III consists of non-health expenditures in categories that can be relatively easily adjusted: food, entertainment, apparel, tobacco, alcohol, personal care, and reading.

Table 4: Sensitivity of Welfare Estimates, Part I (Non-Health Assumptions)

	I	II	III	IV	V	VI	VII	VIII	IX
						Allow for out-of-pocket spending	Shock borne entirely by	Alternative Interpolations	
		Coefficient of Relative Risk Aversion		Consumption Floor				Linear demand	Upper-
	Baseline	1	5	$1,000	$5,000	at q=1	individual	for m	bound
Panel A: Welfare Relative To Transfers To External Parties, $\gamma(1)/N$									
Complete information	0.78	0.41	1.88	1.31	0.40	0.34	1.09	-	-
Consumption-based optimization, consumption proxy	0.66	0.34	0.77	0.80	0.44	0.43	1.45	0.60	1.42
Consumption-based optimization, CEX consumption measure	0.37	0.35	0.41	0.37	0.36	0.27	0.37	-	-
Health-based optimization	0.32	0.32	0.32	0.32	0.32	0.22	0.32	0.34	0.40
Panel B: Welfare Relative to Net Costs, $\gamma(1)/C$									
Complete information	1.16	0.61	2.80	1.94	0.59	0.60	1.62	-	-
Consumption-based optimization, consumption proxy	0.98	0.50	1.15	1.19	0.66	0.74	2.17	0.89	2.12
Consumption-based optimization, CEX consumption measure	0.55	0.52	0.60	0.55	0.55	0.47	0.55	-	-
Health-based optimization	0.48	0.48	0.48	0.48	0.48	0.38	0.47	0.50	0.59

Notes: This table presents the sensitivity of our baseline results to alternative assumptions unrelated to health valuations. Alternative assumptions about health specifications are reported in Table 5. Column I reports the baseline specification. We report the estimates under alternative assumptions about risk aversion (columns II-III) and the consumption floor (columns IV-V). We also consider specifications that allow for out-of-pocket spending to be positive for the insured (column VI) and assume that the health expenditure shock is borne solely by the individual instead of being shared equally within families (column VII). Columns VIII and IX report alternative interpolation assumptions, including linear demand (column VIII) and the upper bound procedure described in the text (column IX).

Table 5: Sensitivity of Welfare Estimates, Part II (Assumptions Related to Valuing Health)

	I	II	III	IV	V	VI	VII	VIII
		Mail-In Survey				In-Person Survey		
VSLY, φ:	$5K	$0	$40K	$5K	$5K	$5K	$5K	$5K
Health measure:	SAH	SAH	SAH	PHQ-2	SAH	PHQ-2	PHQ-8	SF-8
Panel A: Welfare Relative To Transfers To External Parties, $\gamma(1)/N$								
Complete information	0.78	0.50	1.41	0.60	0.61	0.56	0.54	0.51
Health-based optimization	0.32	0.31	0.42	0.33	0.25	0.29	0.22	0.21
Panel B: Welfare Relative to Net Costs, $\gamma(1)/C$								
Complete information	1.16	0.74	2.10	0.81	1.53	1.41	1.35	1.53
Health-based optimization	0.48	0.46	0.62	0.45	0.62	0.73	0.55	0.62
Number of observations	15,498	15,498	15,498	15,201	11,546	11,546	11,546	11,625

Notes: Column I reports our baseline estimates, while other columns report alternatives as indicated in the column headings. The VSLY (φ) is the value of a statistical life year as evaluated by a low-income person. The health measure is the measure on which the QALY score is based. SAH denotes the "Self-Assessed Health" measure; PHQ denotes the "Patient Health Questionnaire" score, which measures mental health. PHQ-2 denotes a score based on two items from the PHQ where as PHQ-8 denotes a score based on 8 items of the PHQ. SF-8 denotes the 8-item "Short Form" health measure, which measures both physical and mental health. Additional information on these health measures and their mapping to QALYs is provided in Appendix A.4. In our baseline, we estimate that Medicaid increases QALYs by 0.045 (see Table 1). For QALYs measured by different health measures or in different samples, the estimated impact of Medicaid on QALYs becomes 0.027 in column IV, 0.049 in column V, 0.032 in column VI, 0.027 in column VII, and 0.027 in column VIII.
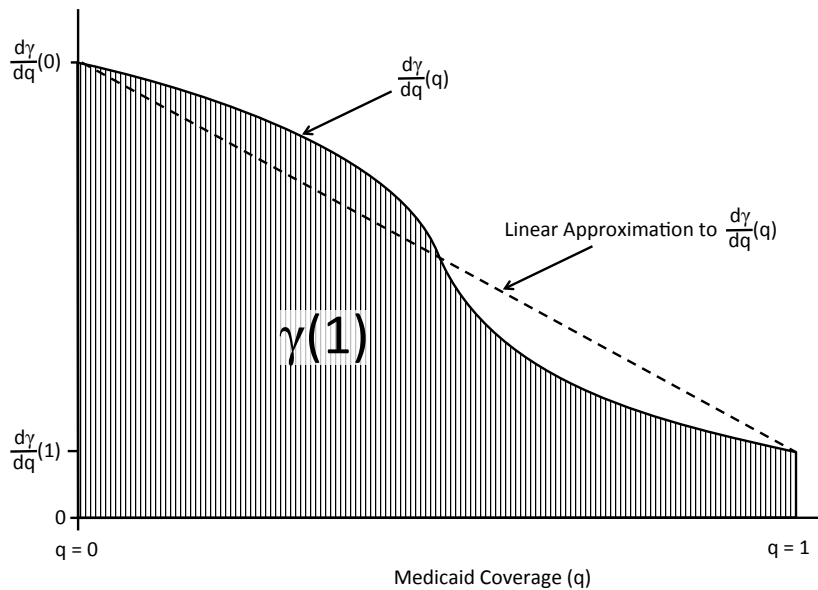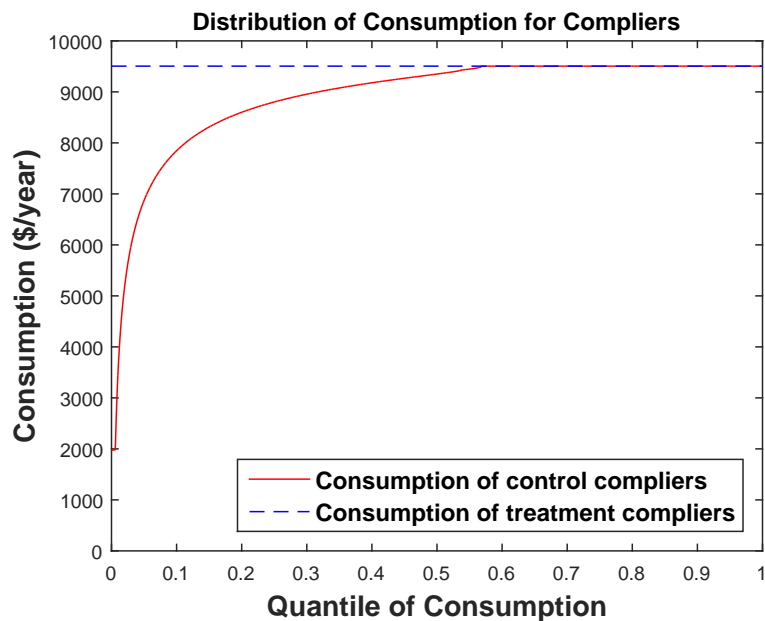
Figure 1: From $\frac{d\gamma}{dq}(q)$ to $\gamma(1)$

Figure 2: Consumption Distribution for Treatment and Control Compliers