

NBER WORKING PAPER SERIES

THE VALUE OF MEDICAID:
INTERPRETING RESULTS FROM THE OREGON HEALTH INSURANCE EXPERIMENT

Amy Finkelstein
Nathaniel Hendren
Erzo F.P. Luttmer

Working Paper 21308
<http://www.nber.org/papers/w21308>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2015

We are grateful to Lizi Chen for outstanding research assistance and to Isaiah Andrews, Liran Einav, Matthew Gentzkow, Jonathan Gruber, Conrad Miller, Jesse Shapiro, Matthew Notowidigdo, Ivan Werning, and seminar participants at Brown, Chicago Booth, Harvard Medical School, Michigan State, and the University of Houston for helpful comments. We gratefully acknowledge financial support from the National Institute of Aging under grants RC2AGO36631 and R01AG0345151 (Finkelstein) and the NBER Health and Aging Fellowship, under the National Institute of Aging Grant Number T32-AG000186 (Hendren). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2015 by Amy Finkelstein, Nathaniel Hendren, and Erzo F.P. Luttmer. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Value of Medicaid: Interpreting Results from the Oregon Health Insurance Experiment
Amy Finkelstein, Nathaniel Hendren, and Erzo F.P. Luttmer
NBER Working Paper No. 21308
June 2015
JEL No. H51,I13

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. Our baseline estimates of Medicaid's welfare benefit to recipients per dollar of government spending range from about \$0.2 to \$0.4, depending on the framework, with at least two-fifths – and as much as four-fifths – of the value of Medicaid coming from a transfer component, as opposed to its ability to move resources across states of the world. In addition, we estimate that Medicaid generates a substantial transfer, of about \$0.6 per dollar of government spending, to the providers of implicit insurance for the low-income uninsured. The economic incidence of these transfers is critical for assessing the social value of providing Medicaid to low-income adults relative to alternative redistributive policies.

Amy Finkelstein
Department of Economics, E17-228
MIT
77 Massachusetts Avenue
Cambridge, MA 02139
and NBER
afink@mit.edu

Erzo F.P. Luttmer
6106 Rockefeller Center, Room 305
Department of Economics
Dartmouth College
Hanover, NH 03755
and NBER
Erzo.FP.Luttmer@Dartmouth.Edu

Nathaniel Hendren
Harvard University
Department of Economics
Littauer Center Room 235
Cambridge, MA 02138
and NBER
nhendren@gmail.com

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

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

June 2015

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. Our baseline estimates of Medicaid’s welfare benefit to recipients per dollar of government spending range from about \$0.2 to \$0.4, depending on the framework, with at least two-fifths – and as much as four-fifths – of the value of Medicaid coming from a transfer component, as opposed to its ability to move resources across states of the world. In addition, we estimate that Medicaid generates a substantial transfer, of about \$0.6 per dollar of government spending, to the providers of implicit insurance for the low-income uninsured. The economic incidence of these transfers is critical for assessing the social value of providing Medicaid to low-income adults relative to alternative redistributive policies.

1 Introduction

Medicaid is the largest means-tested program in the United States. In 2011, public expenditures on Medicaid were over \$425 billion, compared to \$80 billion for food stamps (SNAP), \$50 billion for the Earned Income Tax Credit (EITC), \$50 billion for Supplemental Security Income (SSI), and \$33 billion for cash welfare (TANF).¹ Expenditures on Medicaid will increase even further with the 2014 Medicaid expansions under the Affordable Care Act.²

What are the welfare benefits of this large in-kind program? How do the welfare benefits from Medicaid compare to its costs? How do the welfare benefits from Medicaid per dollar of government spending compare to the welfare benefits from other, cash-based transfer programs?

*MIT, Harvard, and Dartmouth. We are grateful to Lizi Chen for outstanding research assistance and to Isaiah Andrews, Liran Einav, Matthew Gentzkow, Jonathan Gruber, Conrad Miller, Jesse Shapiro, Matthew Notowidigdo, Ivan Werning, and seminar participants at Brown, Chicago Booth, Harvard Medical School, Michigan State, and the University of Houston for helpful comments. We gratefully acknowledge financial support from the National Institute of Aging under grants RC2AGO36631 and R01AG0345151 (Finkelstein) and the NBER Health and Aging Fellowship, under the National Institute of Aging Grant Number T32-AG000186 (Hendren).

¹See Congressional Budget Office (2013)[45], Centers for Medicare and Medicaid Services (2012)[43], and Department of Health and Human Services (2012)[47].

²Congressional Budget Office (2014)[46].

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, standard practice in both academia and public policy is to either ignore the value of Medicaid – for example, in the calculation of the poverty line, or in analysis of income inequality (Gottschalk and Smeeding (1997)[28]) – or to make fairly ad hoc assumptions. For example, the Congressional Budget Office (2012)[44] values Medicaid at the average government expenditure per recipient. In practice, of course, an in-kind benefit like Medicaid may be valued at less, or at more, than its cost (see e.g. Currie and Gahvari (2008)[17]).

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.³ 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.⁴

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 precludes 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)[21] provide a review). Instead, one encounters the classic problem of valuing goods when prices are not observed (Samuelson (1954)[41]).

In this paper, we develop two main analytical frameworks for empirical welfare analysis of Medicaid coverage 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

³For more detail on these results, as well as on the experiment and affected population, see Finkelstein et al. (2012)[25], Baicker et al. (2013)[6], Taubman et al. (2014)[42], and Baicker et al. (2014)[4].

⁴The 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>).

specification of a normative utility function and estimates of the causal effect of Medicaid 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. This creates a potential methodological bias, as one can keep positing additional potential utility arguments until one is satisfied with the welfare estimates.

Our second approach, which we refer to as the “optimization” approach, is in the spirit of the “sufficient statistics” approach described by Chetty (2009)[13], 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 can significantly reduce the implementation requirements. In particular, it suffices to specify the marginal utility function over any single argument (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 covariance 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.

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. Absent a consumption survey in the Oregon context, we proxy for consumption as the difference between income and out-of-pocket medical expenditures, subject to a consumption floor; we also implement

an alternative version of the consumption-based optimization approach which measures consumption directly for a low-income sample in the Consumer Expenditure Survey. Our baseline health measure is self-reported health; we also report estimates based on alternative health measures, such as mortality and depression. In addition, we estimate the impact of Medicaid on government spending and on transfers to providers of partial, implicit insurance to the “uninsured”; these provide estimates of Medicaid’s gross and net program costs, which we compare to our estimates of the welfare benefits to Medicaid recipients.

All of our estimates indicate a welfare benefit from Medicaid to recipients that is below the government’s costs of providing Medicaid. Specifically, we estimate a welfare benefit to recipients per dollar of government spending of about \$0.4 from the complete-information approach and from the consumption-based optimization approach using a consumption proxy, and about \$0.2 from the other two optimization approaches. The differences in welfare estimates across the approaches primarily reflects different estimates of the “pure-insurance” value of Medicaid (i.e., its ability to move resources across states of the world). In all the approaches, at least two-fifths – and as much as four-fifths – of the value of Medicaid to recipients comes from a pure transfer component.

These findings indicate that if (counterfactually) Medicaid recipients had to pay the government’s average cost of Medicaid, they would rather be uninsured. Both moral hazard and crowd out of implicit insurance may reduce the value to recipients of insurance below its costs. In our setting, we find substantial Medicaid crowd out of implicit insurance. We estimate that the low-income uninsured pay only a small fraction (about 20 cents on the dollar) of their own medical expenses; external parties pay the remainder. As a result, we estimate that a substantial portion of the government’s Medicaid spending – about 60 cents on the dollar – represents a transfer to the providers of this implicit insurance, rather than a direct benefit for Medicaid recipients. If we instead compare Medicaid recipients’ value of Medicaid to its net costs (i.e., net of the transfers to the providers of implicit insurance), we find it is above 1 for the complete-information approach; it is 0.9 for the consumption-based optimization approach using the consumption proxy and 0.5 for the other two approaches. A ratio below 1 suggests that the moral hazard costs of Medicaid exceed the insurance value to Medicaid recipients, while a ratio above 1 suggests the converse.

Finally, we evaluate Medicaid as a redistributive tool, rather than as a potential instrument to correct a market failure, as in the preceding discussion. To do so, we compare Medicaid to other forms of redistribution – all of which also entail some resource cost – and also consider the incidence of the external transfers. We consider the hypothetical policy choice of eliminating Medicaid for low-income adults or making a budgetarily equivalent reduction in the Earned Income Tax Credit (EITC). We find that society’s preference between these two depends critically on the incidence of the \$0.6 per dollar of government spending that Medicaid generates in transfers to “external parties.” For example, assuming that EITC recipients and Medicaid recipients have the same social welfare weights, we find that society would prefer to cut Medicaid coverage than to make a budgetarily equivalent cut in the EITC if the incidence of these transfers is on the upper regions of the income distribution; however, if we assume that the incidence of these “external” transfers is on Medicaid

recipients themselves, then society would be roughly indifferent between the two (or slightly prefer cutting the EITC). Such indeterminacy highlights the importance of future work examining the incidence of Medicaid’s external transfers.

How seriously should our empirical welfare estimates be taken? Naturally, all of our quantitative results are sensitive to the framework used and to our specific implementation assumptions. We therefore explore the sensitivity of our baseline estimates to a variety of alternative “reasonable” assumptions, which produces estimates of the welfare benefits to recipients per dollar of government spending ranging from about \$0.15 to \$0.85. We discuss which modeling assumptions, features of the data, and parameter calibrations are quantitatively most important for the results.

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 to recipients than the implicit default of treating the value of Medicaid at zero or simply at cost, which occurs in so much existing work. Although we focus on the specific context of the value of Medicaid in the Oregon Health Insurance Experiment, 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. More generally, the basic challenges and tradeoffs we describe may also be of use for welfare analysis of other social insurance programs in settings where individuals do not reveal their willingness to pay through ex-ante choices.

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 these welfare estimates. The last section concludes.

2 Frameworks for Welfare Analysis

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

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

We assume health is produced according to:

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

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. 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.⁵

2.1 Complete-information approach

The complete-information approach to empirical welfare analysis assumes we observe the arguments of the utility function both with insurance and without insurance. It is then straightforward to define the welfare impact for Medicaid recipients $\gamma(1)$, as the implicit solution to:

$$E[u(c(0; \theta), h(0; \theta))] = E[u(c(1; \theta) - \gamma(1), h(1; \theta))], \quad (3)$$

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.⁶ Specifically, $\gamma(1) < G$ implies that if given a choice between losing Medicaid and having to give up G in consumption, the insured would choose to give up 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 .⁷

Estimation of equation (3) requires that we specify the normative utility function over all its arguments. We assume that the utility function takes the following form:

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

⁵We 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) [20].

⁶Note 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 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.

⁷As we discuss in Section 5.2 below, we would require additional information or assumptions to answer these alternative questions.

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.

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 $q = 1$:

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

We use equation (4) to solve for $\gamma(1)$. This requires observing the distributions of consumption and mean health outcomes that occur if the individual were on Medicaid ($c(1; \theta)$ and $h(1; \theta)$) and if he were not ($c(0; \theta)$ and $h(0; \theta)$). 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) [2]).⁸

2.2 Optimization approaches

We can reduce the implementation requirements of the complete-information approach through additional assumptions. Specifically, 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.

⁸Our 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) [23], 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.

We denote out-of-pocket spending on medical care by:

$$x(q, m) \equiv p(q)m. \quad (5)$$

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:

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

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

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)[3]) and where the complex nature of the decision problem may generate individually suboptimal decisions (Baicker, Mullainathan, and Schwartzstein (forthcoming)[5]).

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 (5) is now:

$$x(q, m) = qp(1)m + (1 - q)p(0)m. \quad (6)$$

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

$$-\frac{\partial x(q, m(q; \theta))}{\partial q} = (p(0) - p(1))m(q; \theta). \quad (7)$$

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 $q = 0$:

$$E[u(c(0; \theta), h(0; \theta))] = E[u(c(q; \theta) - \gamma(q), h(q; \theta))]. \quad (8)$$

2.2.1 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 (8):

$$\frac{d\gamma}{dq} = E \left[\frac{u_c}{E[u_c]} ((p(0) - p(1))m(q; \theta)) \right], \quad (9)$$

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 (9), 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.⁹

We can decompose the marginal value of Medicaid to recipients in equation (9) into a transfer term and a pure-insurance term. Empirical implementation will be based on estimating each term separately. The decomposition is:

$$\frac{d\gamma(q)}{dq} = \underbrace{(p(0) - p(1))E[m(q; \theta)]}_{\text{Transfer Term}} + \underbrace{\text{Cov} \left[\frac{u_c}{E[u_c]}, (p(0) - p(1))m(q; \theta) \right]}_{\substack{\text{Pure-Insurance Term} \\ \text{(consumption valuation)}}}. \quad (10)$$

The transfer term measures the recipients’ valuation of the expected transfer of resources from

⁹Intuitively, 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.

the rest of the economy to them. The “pure-insurance” term measures the benefit of reallocating resources (i.e., relaxing the recipient’s budget constraint) across different states of the world, θ .¹⁰ The movement of these resources is valued using the marginal utility of consumption in each state.

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

$$\begin{aligned} \gamma(1) &= \int_0^1 \frac{d\gamma(q)}{dq} dq \\ &= \underbrace{(p(0) - p(1)) \int_0^1 E[m(q; \theta)] dq}_{\text{Transfer Term}} + \underbrace{\int_0^1 \text{Cov} \left(\frac{u_c}{E[u_c]}, (p(0) - p(1))m(q; \theta) \right) dq}_{\substack{\text{Pure-Insurance Term} \\ \text{(consumption valuation)}}} \end{aligned} \quad (11)$$

which follows from the fact that $\gamma(0) = 0$, by definition. Figure 1 illustrates this conceptual integral computation, with the solid line representing $\frac{d\gamma}{dq}(q)$, and $\gamma(1)$ being the area under that curve.

Implementation

Pure-insurance term. Evaluation of the pure-insurance term in equation (10) 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 $v(\cdot)$ is the subutility function for health, which can be left unspecified.

With this assumption, the pure-insurance term in equation (10) can be re-written as:

$$\text{Cov} \left(\frac{c(q; \theta)^{-\sigma}}{E[c(q; \theta)^{-\sigma]}}, (p(0) - p(1))m(q; \theta) \right). \quad (12)$$

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 $\gamma(1)$

¹⁰This 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)[40] equation (3)).

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 (11) 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 the dashed line in Figure 1.

Transfer term. Evaluation of the transfer term in equation (10) does not require any assumptions about the utility function. However, integration in equation (11) to obtain an estimate of the transfer term requires that we know the path of $m(q; \theta)$ for interior values of q , which will not be directly observed. We will 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., $E[m(0; \theta)] \leq E[m(q; \theta)] \leq E[m(1; \theta)]$)¹¹, we obtain lower and upper bounds for the transfer value of Medicaid as the out-of-pocket price change of medical care due to Medicaid, $p(0) - p(1)$, times medical spending at, respectively, the uninsured and insured levels:

$$[p(0) - p(1)]E[m(0; \theta)] \leq (p(0) - p(1)) \int_0^1 \frac{E[m(q; \theta)]}{dq} dq \leq [p(0) - p(1)]E[m(1; \theta)]. \quad (13)$$

2.2.2 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 that individual choices satisfy a first-order condition:

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

$$u_c(c, h)p(q) = u_h(c, h) \frac{d\tilde{h}(m; \theta)}{dm} \quad \forall m, q, \theta. \quad (14)$$

The left-hand side equation (14) is the marginal cost of medical spending in terms of forgone consumption. The right-hand side of equation (14) is the marginal benefit of additional medical

¹¹A downward-sloping demand function for m would be sufficient for this assumption to hold.

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 equation (14) to replace the marginal utility of consumption, u_c in equation (9) with a term depending on the marginal utility of health, u_h , yielding:

$$\frac{d\gamma}{dq} = E \left[\left(\frac{u_h}{E[u_c]} \frac{d\tilde{h}(m; \theta)}{dm} \frac{1}{p(q)} \right) ((p(0) - (p(1))m(q; \theta)) \right]. \quad (15)$$

We refer to equation (15) 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 expected 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 downward 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.¹² Assumption (6) is thus 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 (15) can be decomposed into a transfer term and a pure-insurance term:

$$\frac{d\gamma(q)}{dq} = \underbrace{(p(0) - p(1))E[m(q; \theta)]}_{\text{Transfer Term}} + \underbrace{Cov \left(\frac{u_h}{E[u_c]} \frac{d\tilde{h}(m; \theta)}{dm} \frac{1}{p(q)}, (p(0) - (p(1))m(q; \theta)) \right)}_{\substack{\text{Pure-Insurance Term} \\ \text{(health valuation)}}}. \quad (16)$$

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:

¹²If the individual is at a corner solution with respect to medical spending, then the first term between parentheses in equation (15) is less than the value that the individual puts on money in that state of the world (which, after all, is why the individual chooses to have zero medical spending in that state of the world and, instead, spends all resources on the consumption of other goods).

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 $\tilde{v}(\cdot)$ 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 (16) can be written as:

$$Cov\left(\frac{d\tilde{h}(m; \theta)}{dm} \frac{\phi}{p(q)}, (p(0) - (p(1))m(q; \theta))\right). \quad (17)$$

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)).

Comment: Endless Arguments

The option of using either a health-based optimization approach (equation 16) or a consumption-based optimization approach (equation 10) 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)[13]). 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)[15]). 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., Assumption 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.

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

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. Appendix A.2 provides additional details on our estimating equations.¹⁵

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. All of the objects we calculate are estimated for the compliers – i.e., all those who are covered by Medicaid if and only if they win the lottery (see, e.g., Angrist and Pischke (2009) [2]). Thus in our application “the insured” ($q = 1$) are treatment compliers and “the uninsured” ($q = 0$) are control compliers.

The data used here from the Oregon Health Insurance Experiment were previously analyzed by Finkelstein et al. (2012)[25] and are publicly available at www.nber.org/oregon. Data on Medicaid coverage (q) are taken from state administrative records. All of the other data elements 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.

¹³Eligibility 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.

¹⁴Throughout, 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) [25]).

¹⁵Finkelstein et al. (2012)[25] provide more details on the lottery, and supporting evidence on the assumptions required to use the lottery as an instrument for Medicaid coverage. 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)[25], Baicker et al. (2013)[6], (Baicker et al. (2014)[4]), and Taubman et al. (2014)[42]).

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 = 0.14). Panel B presents summary statistics on our key outcome measures, which we now discuss.

Health h

In our baseline specification, we measure health (h) as a binary variable, with $h = 1$ when the individual reports his health to be excellent, very good, or good (as opposed to fair or poor). About 61 percent of treatment compliers ($q = 1$) and 47 percent of control compliers ($q = 0$) report their health to be excellent, very good, or good. We explore sensitivity to other health measures below.

Medical spending m

For medical spending, we follow the approach used by Finkelstein et al. (2012)[25]. We estimate total annual medical spending for each individual based on their self-reports of utilization of prescription drugs, outpatient visits, ER visits, and inpatient hospital visits, weighting each type of use by its average cost (expenditures) among the low-income publicly insured adults in the Medical Expenditure Survey (MEPS). 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 ($q = 0$), as self-reported out-of-pocket medical expenditures in the last 6 months, multiplied by two. Average annual out-of-pocket medical expenditures for control compliers is $E[(x(0, m(0, \theta)))] = \489 .¹⁶

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

¹⁶The unadjusted mean of out-of-pocket spending for control compliers is \$481 per year. To be consistent with our treatment of out-of-pocket spending when we use it to estimate consumption (discussed below), we impose the same two adjustments here: a fitted distribution and a cap on out-of-pocket consumption that binds for less than 5 percent of control compliers. Both adjustments are described in more detail in Section 4.

¹⁷This 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.

Out-of-pocket prices

For the optimization approaches, we need to define the out-of-pocket price of medical care with Medicaid, $p(1)$, and without Medicaid, $p(0)$. We estimate $p(0)$ 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 $p(0) = 0.18$.

In other words, we estimate that 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

Consistent with our baseline assumption that $x(1,m) = 0$, we assume $p(1) = 0$; those with Medicaid pay nothing out of pocket for medical spending. We examine sensitivity to this assumption below.

Consumption c

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. As a result, we take two approaches to measuring consumption.

Approach #1: Proxy consumption using out-of-pocket expenditure We approximate c as the difference between income and out-of-pocket spending:

$$c = \bar{y} - x. \tag{18}$$

We use the definition of out-of-pocket spending above. We use average annual per capita income \bar{y} for our measure of income; this approach assumes away any direct impact of Medicaid on income, as well as heterogeneity across individuals in income.¹⁹

We construct average per capita income \bar{y} from individuals’ reported 12-month household income bin. We convert this to average annual household income by using income-bin midpoints and a top-coded value of \$50,000. We divide this household income measure by the total number of family members (adults and children) in the household to construct per capita annual income, and then

¹⁸The 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)[16], Figure 1), suggesting that on average the uninsured in the U.S. pay only 20% of their total medical expenses. 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.

¹⁹Prior analysis of the Oregon Health Insurance Experiment showed no evidence of a direct impact of Medicaid on income (Finkelstein et al. (2012) [25], Baicker et al. (2014)[4]). Heterogeneity of income is limited by the fact that the program requires income below the federal poverty line; while there is undoubtedly still some cross-sectional heterogeneity in income, as a practical matter we suspect that our measurement of it has a high noise-to-signal ratio. For example, about 13 percent of respondents report zero income.

take the average over our sample. Using this approach, we estimate average per capita income (\bar{y}) is \$3,808 per year for the compliers in our sample. This is quite low. Recall, however, that compliers must be below the federal poverty line. In 2009, the federal poverty line reflected per capita income of between \$5,513 and \$10,830 for family sizes of 4 individuals to 1 individual.²⁰

Because there is unavoidable measurement error in this approach to measuring 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)[9], Hoynes and Luttmer (2011)[33]) by imposing an annual consumption floor. Our baseline analysis imposes a consumption floor of \$1,000 per capita per year, which we implement by capping out-of-pocket expenditure at $\bar{y} - c_{floor}$. In the sensitivity analysis below, we consider how the results are affected by assuming higher values of average income \bar{y} , and thus of average consumption, than reported in these data. We also explore sensitivity to the assumed value of the consumption floor.

Approach #2: Measure consumption using national data from the CEX A concern with our consumption proxy 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 consumption floor. If individuals can borrow, draw down assets, or have other ways of smoothing consumption, this approach overstates the consumption smoothing benefits of Medicaid.

As an alternative approach, we use data on non-medical consumption (c) and out-of-pocket-spending (x) for low-income adults in the Consumer Expenditure Survey (CEX) to estimate the relationship between c and x for adults below the federal poverty line without formal insurance ($q = 0$). This allows us to estimate the pure-insurance term in equation (12) at $q = 0$. Because the CEX requests information on the health insurance status of only the household head, we restrict the sample to single adults with no children in the household so that we can accurately measure insurance status. Appendix A.4 provides more detail on the data and our sample definition. We estimate average consumption for the uninsured of \$13,541 (std. dev. of \$7,451).²¹

4 Welfare Results

Implementation of the complete-information approach requires an assumption for both the coefficient of relative risk aversion σ and the value of health ϕ . The consumption-based optimization approach requires only the former; the health-based optimization approach requires only the latter. For our baseline analyses, we assume $\sigma = 3$.

²⁰ The average complier had a family size of 3, for which the 2009 per-capita federal poverty line was about \$6,100. See <http://aspe.hhs.gov/poverty/09poverty.shtml>.

²¹ We impose the same baseline consumption floor as in the consumption proxy approach, but it never binds in the CEX data. Average consumption in the CEX is substantially higher than in the consumption proxy approach (\$13,300 compared to \$3,300). These differences likely reflect at least in part the well-known problem in the Consumer Expenditure Survey (CEX) that income tends to be under-reported, and thus average annual expenditures exceed income for some demographic groups (see see: <http://www.bls.gov/cex/faq.htm#q21>). We explore the implications of higher average consumption for our welfare estimates in Section (5.3) below.

We arrive at a baseline assumption for ϕ as follows. We (crudely) convert the estimated effect of Medicaid on self-reported health (see Table 1) into an estimated effect on 10-year life expectancy based on the 0.5 year cross-sectional difference in 10-year life expectancy between low-income adults with different self-reported health in the National Health Interview Survey.²² We further assume a baseline value of a statistical life year (VSLY) in our population of \$25,000.²³ Combining these assumptions yields a value for ϕ , the value of reporting excellent, very good, or good health (rather than fair or poor) of \$12,250.

4.1 Complete-information approach

We solve equation (4) for $\gamma(1)$. This requires us to estimate health outcomes and the distribution of consumption for control compliers ($q = 0$) and for treatment compliers ($q = 1$). We estimate that the probability that self-reported health is excellent, very good or good (as opposed to fair or poor) is 47 percent for control compliers and 61 percent for treatment compliers (see Table 1).

To estimate the impact of Medicaid on the distribution of c , we first estimate the impact of Medicaid on the distribution of out-of-pocket spending x , and then map out-of-pocket spending to consumption based on equation (18). To estimate the impact of Medicaid on the distribution of out-of-pocket spending, 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. We impose a consumption floor by capping out-of-pocket spending drawn from the log-normal distribution at $\bar{y} - c_{floor}$. Our baseline consumption floor of \$1,000 binds for less than 5 percent of control compliers. Appendix A.2 describes these techniques in more detail, and also reports that the parametric model fits the data quite well.

Figure 2 shows the resultant distributions of consumption for the the control compliers ($q = 0$)

²²We compare subsequent 10-year survival for low-income adults (aged 19-64) who report their health as excellent, very good or good to those who report fair or poor, using the linked mortality public-use data files of the National Health Interview Surveys from 1986-1996. We control in our analysis for year, gender, race, age and education level, but recognize that we estimate an association, not a causal effect. We find that low-income adults who report being in good, very good, or excellent health have a 0.49 year higher 10-year survival than low-income adults who report their health as fair or poor. Because we examine survival over only a 10-year window (due to data constraints), our estimate of ϕ will be biased down if high self-reported health results in survival gains beyond this 10-year window. On the other hand, we assume that survival gains from high self-reported health are additive over time, and our estimate of ϕ will be biased upwards if in fact these survival gains are subadditive.

²³Welfare analysis requires a measure of the VSLY that equals the (hypothetical) willingness to pay by the recipients out of their *own* income. We arrive at this figure by taking Cutler's (2004) [18] choice of \$100,000 for the VSLY, and adjusting for the fact that this VSLY is based on estimates for the general US population, while median household income for compliers in our population is \$9,000, roughly one-quarter of median US household income. In other words, we assume that the VSLY scales linearly with income. This produces a lower VSLY for lower income individuals, which is appropriate if we are considering an *individual's* willingness to pay. But it is distinct from the question of how *society's* willingness to pay for a statistical life year scales with income. The social welfare weights may well be higher for the poor. We consider Medicaid as redistribution explicitly in Section 5.2.2 below where we consider how the value of Medicaid to recipients per dollar of government spending compares to the feasible set of transfers to low-income individuals through the tax system. There, we assume that society places the same relative weights on life vs. consumption as the individual but allow for redistributive goals across individuals, which allows the social value of life of low-income individuals to be higher than the individuals' own private willingness to pay.

and treatment compliers ($q = 1$). Average consumption for control compliers ($q = 0$) is \$3,320 with a standard deviation of \$768. For treatment compliers ($q = 1$), consumption is simply average income, \$3,808, since by assumption $x(1, m) = 0$. The difference between the two lines in the figure shows the increase in consumption due to Medicaid for the compliers. On average, Medicaid increased consumption by \$489.

Results and decomposition.

The first column of Table 2 shows the resultant estimate: $\gamma(1) = \$1,576$. We estimate that a Medicaid recipient would be indifferent between giving up Medicaid and giving up \$1,576 in consumption. We define the welfare value of Medicaid to recipients operating through consumption, γ_C as:

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

and estimate $\gamma_C = 1,111$. We then infer the welfare value of Medicaid operating through health as $\gamma_M = \gamma(1) - \gamma_C = \465 .²⁴

To explore what aspects of Medicaid drive our estimate of $\gamma(1)$, we can decompose $\gamma(1)$ into transfer and “pure-insurance” terms (see Appendix A.3 for implementation details). In the complete-information approach, the transfer term, which we denote by T , is defined by the average change in consumption plus the effect on health if everyone has received the average increase in medical spending (i.e., $T = E[c(1; \theta) - c(0; \theta)] + \tilde{\phi} E \left[\frac{d\tilde{h}}{dm} \right] E[m(1, \theta) - m(0, \theta)]$). We estimate a welfare gain from the transfer of \$693, of which \$489 is the transfer coming from consumption (and \$205 is due to health).²⁵ Given estimates of the total value of Medicaid from consumption and health (γ_C and γ_M), and the transfer components of each, we can naturally also back out the value of the pure-insurance component in total (\$883) or operating via consumption (\$622) and via health (\$260).²⁶

These decompositions suggest that over two-thirds of the welfare gain from Medicaid operates through the impact on consumption, whereas less than a third can be attributed to the impact on health. They also suggest that a little less than half of the welfare benefit to Medicaid recipients comes from the transfer component, with slightly more than half coming from Medicaid’s ability to move resources across states of the world.

²⁴Because 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 \$1,195 and a welfare gain due to health of \$382.

²⁵Of course, this requires an estimate of $\frac{d\tilde{h}}{dm}$ which we have not yet discussed; we describe our estimation approach in sub-section 4.2, below.

²⁶The pure-insurance component operating through consumption smoothing is broadly similar to the approach taken by Feldstein and Gruber (1995)[22] to estimate the consumption-smoothing value of catastrophic health insurance, and Finkelstein and McKnight (2008)[24] to estimate the consumption-smoothing value of the introduction of Medicare.

4.2 Optimization approaches

Transfer component Recall that the transfer component for the optimization approaches is defined as $T = (p(0) - p(1)) \int_0^1 E[m(q; \theta)] dq$ (see equation (11)). The change in the out-of-pocket price for medical care due to insurance ($p(0) - p(1)$) is 0.18. Using the linear approximation Assumption 5, the transfer term in equation (11) is \$568. Without the linear approximation, we can derive lower and upper bounds for the transfer term. These bounds are \$488 and \$647, respectively (see equation (13)).

Consumption-based optimization approach with consumption proxy We estimate the pure-insurance value at $q = 0$ using equation (12). This requires an estimate of the joint distribution of c and $(p(0) - p(1))m(0; \theta)$ for control compliers. In our implementation, we use the same estimate of the distribution of out-of-pocket spending and of consumption for control compliers as described above in the complete-information approach. At $q = 1$, our assumption that $p(1) = 0$ together with our definition of consumption in equation (18) immediately implies that the marginal utility of consumption is constant, and hence the pure-insurance value of Medicaid is 0 on the margin; we relax this assumption in the sensitivity analysis below. Following the linear approximation in Assumption 5, the total pure-insurance component is therefore one-half of what we estimate at $q = 0$, or \$691.

Combining the estimate of the pure-insurance component with the transfer component yields an estimate of $\gamma(1) = \$1,260$. These results are shown in Table 2, column II. Similar to the finding of the complete-information approach, just under half of the welfare value comes from the transfer component, with slightly more than half coming from Medicaid’s ability to move resources across states of the world.

Consumption-based optimization approach with CEX consumption measure In principle, the pure-insurance term at $q = 0$ can be directly estimated in the CEX data from the covariance between $c^{-\sigma}$ and x for the uninsured. In practice, however, the raw data show a negative covariance 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 spending among the uninsured). This could be an accurate measure of the 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, in Appendix A.4 (and Appendix Table A.3), we illustrate that this covariance is less negative for the uninsured than the insured, which suggests that the covariance term may be biased from measurement error that induces a negative correlation between $c^{-\sigma}$ and x . Appendix A.4 provides a measurement-error model that shows how one can use the implication that this covariance should be zero if insurance is full at $q = 1$ to back out the unbiased covariance term at $q = 0$. In practice, this means using a simple dif-

ference in the observed covariance term for the uninsured and insured as a measure of the unbiased covariance for the uninsured ($q = 0$), which we estimate to be \$252. As before, the pure-insurance value is 0 on the margin at $q = 1$ because we model Medicaid as providing full insurance.²⁷

Column III presents the results. Relative to the results in the first two columns, the pure-insurance value, as measured by the covariance term, is an order of magnitude lower, at \$126. Using the same estimate of the transfer term as described above, this implies a total welfare impact to recipients of \$694.²⁸ Under this approach, about four-fifths of the value of Medicaid to recipients comes from the transfer component, and only about one-fifth comes from its ability to move resources across states of the world.

The CEX consumption measure thus yields a substantially lower estimate of the “pure-insurance” component of the value of Medicaid (column III) than the consumption proxy (column II).²⁹ This may reflect the fact that individuals can adjust their assets or have access to additional forms of informal insurance that prevent a \$1 increase in out-of-pocket spending from translating into a \$1 decrease in consumption, as is assumed in the consumption proxy approach. It is also possible that we underestimate the covariance between marginal utility of consumption and out-of-pocket medical spending using the CEX data despite our attempts to fit the measurement-error model, or that our sample of low-income uninsured adults in the CEX data is not representative of control compliers in the Oregon Experiment sample.³⁰

Health-based optimization approach As in the consumption-based optimization approach, we assume that out-of-pocket spending is zero at $q = 1$, which implies that the pure-insurance component is zero on the margin at $q = 1$ because at that point individuals are fully insured. To calculate the pure-insurance component of the health-based optimization approach at $q = 0$, we require an estimate of the marginal health return to medical spending, $\frac{d\tilde{h}}{dm}$ (see equation (17)). Estimating the health production function $\frac{d\tilde{h}}{dm}$ is notoriously challenging (see, e.g., Almond et al. (2010)[1] 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 .

We estimate the health returns to medical spending using the lottery as an instrument for med-

²⁷The fact that we need the distribution of consumption only at $q = 0$ enables us to implement the consumption-based optimization approach using the CEX data even though we do not have exogenous variation in insurance status in the CEX data. However, we cannot implement the complete-information approach with CEX data because this approach requires identifying variation in q to estimate what consumption would be at $q = 1$.

²⁸We use the transfer term estimated previously because estimating the transfer term in the CEX requires a causal estimate of the impact of Medicaid on m (or x), which we have in the Oregon data but not in the CEX data. By contrast, the covariance term in the consumption-based optimization approach, as we have implemented it, does not require a causal estimate of Medicaid on consumption because in our setting the covariance term is 0 at $q = 1$.

²⁹The value of Medicaid using the CEX consumption measure would be lower still if we had used the raw data (with negative covariances) as opposed to the measurement-error model.

³⁰The difference in estimates does not appear to reflect the fact that our CEX consumption measure (unlike our consumption proxy measure) is limited to singles; the consumption proxy approach yields somewhat higher estimates when limited to singles (results not reported). Interestingly, the higher average consumption obtained from the CEX consumption measure rather than the consumption proxy (see footnote (21)) tends to increase – not decrease – the estimate of the value of Medicaid to recipients. We discuss this in more detail in Section (5.3) below.

ical 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 .³¹ We use measures of baseline medical and financial status for θ^K . Appendix A.5 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 , $E\left[\frac{\partial \bar{h}}{\partial m} | \theta^K\right]$. By combining these estimates with estimates of the distribution of out-of-pocket spending among control compliers $x(0, m(0; \theta))$ conditional on θ^K , we can estimate the pure-insurance term in equation (17) for $q = 0$.³²

The results are shown in column IV of Table 2. Averaging the pure-insurance term estimated for $q = 0$ with a pure-insurance term of zero at $q = 1$, we estimate that the pure-insurance component is \$112. 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. Combining this pure-insurance value with the transfer term calculated above, we estimate a total welfare impact for Medicaid recipients from the health-based optimization approach of \$680. Note that the standard error for this estimate is considerably higher than in any other approach, reflecting the considerable uncertainty in our estimate of the health production function (see Appendix Table A.2). As with the consumption-based optimization approach based on the CEX consumption measure, about four-fifths of the value of Medicaid to recipients comes from the transfer component.

4.3 Comparison Across Approaches

The health-based optimization approach and the consumption-based optimization approach based on CEX data consistently deliver lower point estimates of $\gamma(1)$ than the consumption-based optimization approach based on the consumption proxy or the complete-information approach, both of which tend to produce similar estimates. Which approach one wants to rely on depends on how confident one is with the various assumptions required by each approach. We briefly discuss a few considerations here.

We are least confident in the results from the health-based optimization approach; the difficulties in estimating heterogeneity in health returns to out-of-pocket spending create considerable uncertainty around the results. The estimates contain greater statistical uncertainty than the other approaches, as reflected in the wider confidence intervals shown in Table 2. In addition, our implementation of the health-based optimization approach may produce downward biased estimates of

³¹This 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.

³²Specifically, we use $Cov_{\theta}[X, Y] = E_{\theta^K} [E_{\theta | \theta^K} [X, Y]] - E_{\theta} [X] E_{\theta} [Y]$, where X is $\frac{\partial \bar{h}}{\partial m}$ 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 expectation of these conditional expectations by weighting them by the fraction of control compliers of each type θ^K .

$\gamma(1)$ for individuals at a corner solution with respect to medical spending (see footnote 12) which is the case for about 20 percent of treatment compliers and 30 percent of control compliers. Another source of downward bias in the health-based optimization approach is that our estimation of the health production function conditional on θ^K misses any within- θ^K insurance (see footnote 31).

On the other hand, our consumption-based estimates 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 income minus 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.³³ Our CEX consumption measure (used in our other variant of the consumption-based optimization approach) shows a negative covariance between marginal utility of consumption and out-of-pocket medical expenditures among the uninsured in the raw (unadjusted) data, suggesting that Medicaid transfers resources to states of the world with lower marginal utilities of consumption; the estimates of $\gamma(1)$ that we report use a measurement-error correction to produce a positive “pure-insurance term” in this approach.

Our optimization approaches may also 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 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 [36]). A downward bias in our estimate of $p(0)$ reduces the estimate of $\gamma(1)$ (see equation 9) 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[37], 1999b[38])).³⁴ By contrast, the complete-information

³³Note that we 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.

³⁴Consider 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 decline thereafter. The linear approximation in Figure 1 would not capture the relatively large values of $\frac{d\gamma}{dq}$ that occur for intermediate values of q and would therefore underestimate the welfare effect of Medicaid on the recipient.

approach would accurately capture the value stemming from the liquidity Medicaid provides.³⁵ 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.

5 Interpretation

5.1 Medicaid costs

Approach

It is useful to benchmark our welfare estimates, $\gamma(1)$, against Medicaid costs. Since the government pays the medical bills for Medicaid recipients and Medicaid recipients have no out-of-pocket spending (i.e., $x(1, m) = 0$), the average cost to the government per recipient, which we denote by G , is given by:³⁶

$$G = E[m(1; \theta)]. \quad (19)$$

The (gross) cost to the government is likely higher than the net cost to society. Medicaid's net cost per recipient, which we denote by C , consists of the increase in medical spending induced by insurance, $m(1; \theta) - m(0; \theta)$, and the decrease in out-of-pocket spending due to insurance, $x(1, m(1; \theta)) - x(0, m(0; \theta))$. Hence,

$$C = E[m(1; \theta) - m(0; \theta)] - E[x(1, m(1; \theta)) - x(0, m(0; \theta))]. \quad (20)$$

Using $x(1, m) = 0$, we can decompose gross costs to the government, G , into net costs, C , and a net transfer to external parties:

$$G = C + N.$$

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

$$N = E[m(0; \theta)] - E[x(0, m(0; \theta))]. \quad (21)$$

In other words, N denotes 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

³⁵This 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)[31]) and approximate non-marginal welfare statements using statistical interpolations.

³⁶In addition, labor supply responses to Medicaid may impose fiscal externalities on the government via their impact on tax revenue. However, in the context of the Oregon Health Insurance Experiment, there is no evidence that Medicaid affected labor market activities (Baicker et al. (2014)[4]).

of uninsured individuals.

The identity of these providers of implicit insurance is a difficult and open question. The provision of uncompensated care by hospitals is a natural starting point; Garthwaite, Gross, and Ntowidigdo (2015)[27] recently estimated that increases in Medicaid coverage lead to large reductions in such uncompensated care. Family members or others who previously provided uncompensated informal insurance may also receive net transfers from Medicaid.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 net transfers, other parties bearing some of the incidence may include the privately insured, the government budget (for example, through reduced payments to providers for the provision of uncompensated care), or even the recipients themselves (for example, if reductions in unpaid medical debt improve their credit scores).³⁸ In principle, one could trace through these incidences by estimating the impact of Medicaid on the welfare of other parties, which may be above or below the dollar value of N . Doing so requires additional causal estimates beyond what is readily available in our setting. In practice, we assume the value of these transfers is equal to their costs, N , and explore different assumptions regarding their economic incidence.

Results

The top panel of Table 3 summarizes our estimates of Medicaid costs. We use the lottery as an instrument for Medicaid to obtain these estimates; Appendix A.2 provides more details. We estimate government costs G as total medical spending by treatment compliers ($q = 1$) of \$3,596 (see Table 1); this is broadly consistent with external estimates of the annual per-recipient spending in the subsequently-lotteried Medicaid program in 2001-2004 of about \$3,000 (Wallace et al. (2008)[48]). We estimate that Medicaid reduced out-of-pocket spending by \$489 and increased medical spending by \$885. This yields a net cost C of \$1,374 per recipient per year.

We estimate substantial transfers to external parties, N . We estimate that control compliers ($q = 0$) had total medical expenses of \$2,711 and out-of-pocket expenses of \$489 (see Table 1). Hence, equation (21) indicates that Medicaid delivers a transfer to external parties (N) of \$2,222. This is substantial relative to our estimates of Medicaid transfers to recipients and even relative to the total value (insurance plus transfer) of Medicaid to recipients (see Table 2). They will therefore be important to keep in mind in our benchmarking of the welfare estimates in the next subsection.

³⁷Another potential group to consider is the privately insured. While our framework takes crowd out of private insurance for Medicaid recipients into account in our estimate of $\gamma(1)$, such crowd out may also affect (with unknown sign) adverse selection in the private health insurance market and, hence, outcomes for non-recipients in that market. In the complete-information approach, such general equilibrium effects on non-Medicaid recipients can in principle be captured by estimating the impact of Medicaid on all arguments of the utility function of non-recipients as well. Likewise, in the optimization approach, these general equilibrium effects can be captured by estimating the impact of Medicaid on the budget constraint of non-recipients; Chetty and Saez (2010)[14] provide one modeling structure for this impact. In practice, in the context of the Oregon Health Insurance Experiment, there is no evidence of crowd out of private insurance by Medicaid (Finkelstein et al. (2012)[25]).

³⁸In 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)[25]).

5.2 Benchmarks

5.2.1 Value to recipients relative to costs

We first examine the potential efficiency-enhancing role of public insurance provision, by asking whether uninsured individuals would be willing to pay the cost to the government of providing Medicaid or, equivalently, whether Medicaid recipients would rather pay the government’s cost than give up Medicaid. The answer will be yes if $\gamma(1)$ exceeds G .³⁹ As noted in the Introduction, the Congressional Budget Office currently uses G to value Medicaid for recipients (Congressional Budget Office (2012)[44]). 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 . It need not however; even in the presence of adverse selection, moral hazard costs and crowd-out of implicit insurance may reduce $\gamma(1)$ below G . Ultimately these are empirical questions.

Across our baseline approaches, we consistently estimate that $\gamma(1)$ is less than G ; Medicaid recipients would rather give up Medicaid than pay the government’s costs of providing Medicaid. Specifically, we estimate a ratio of $\gamma(1)/G$ of about \$0.4 using the complete-information approach or the consumption-based optimization approach based on the consumption proxy, and about \$0.2 using the other consumption-based optimization approach or the health-based optimization approach (see Table 3, Panel B, row 1). We can produce a relatively robust lower bound on the value of Medicaid per dollar of government spending by considering only its transfer value. Without any assumptions about the utility function, we estimated a transfer value of Medicaid by the optimization approach of \$568, with a lower bound on the transfer value of \$488, or about \$0.15 per dollar of government spending. The key uncertainty is how much value one assigns to the pure-insurance component of Medicaid; as can be seen in Table 2, our different approaches yield a range of results here.

Since a large portion of G (about \$0.6 on the dollar) goes to pay for coverage the uninsured were effectively receiving prior to Medicaid, another natural benchmark is whether the uninsured would be willing to pay the *net* cost of Medicaid $C = G - N$. Row 2 reports our estimates of $\gamma(1)/C$. This is substantially higher than $\gamma(1)/G$, underscoring the importance of the large amount of implicit insurance for the uninsured (i.e., $p(0) = 0.18$) in reducing the welfare gains to recipients

³⁹As discussed in Section (2.1), 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 $E[u(c(0; \theta) + \tilde{\gamma}(1), h(0; \theta))] = E[u(c(1; \theta), h(1; \theta))]$. We find (in results not reported) that $\tilde{\gamma}(1)$ is considerably higher than $\gamma(1)$, namely \$4,156 compared to \$1,576. 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 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).

from Medicaid. However, our estimate of $\gamma(1)/C$ is above 1 only for the complete-information approach; it is 0.9 for the consumption-based optimization approach using the consumption proxy and 0.5 for the other two approaches. A ratio below 1 suggests that the moral hazard costs of Medicaid exceed the insurance value to individuals who have to pay to receive Medicaid; a ratio above 1 suggests the converse.

When the uninsured are not willing to pay the net cost of Medicaid to obtain it, a natural question is whether they would have been willing to purchase a lower level of Medicaid coverage at the associated costs. Intermediate forms of Medicaid are only well-defined in the optimization-based approaches (in which we have defined the program structure of Medicaid; see Assumption 2). We therefore also estimate the ratio of the marginal value of the first dollar of Medicaid to its marginal cost (i.e., $\frac{d(\gamma)/dq}{d(C)/dq}$, evaluated at $q = 0$) for the optimization-based approaches. If the uninsured were fully uninsured at $q = 0$ (i.e., $p(0) = 1$) this term would have to be larger than 1, since the first unit of insurance has a first-order effect on welfare for risk-averse individuals, but entails only second-order moral hazard costs. Given that the uninsured already have substantial informal insurance ($p(0) = 0.18$), it is an empirical question whether the uninsured have a willingness to pay for a first unit of coverage that exceeds its cost. The point estimates in row 3 indicate that this is the case only in the consumption-based optimization approach using the consumption-proxy. The point estimates of the other two approaches suggest that the willingness to pay for the first unit is only about one-third of its net cost.

5.2.2 Medicaid as redistribution

Comparing recipients' value of Medicaid to its costs addresses the (counterfactual) question of whether the low-income uninsured would be willing to purchase Medicaid if it were available on the private market. Such a "willingness-to-pay" framework however, ignores the fact that the recipients are low income (so perhaps it is not surprising that their willingness to pay is low), and that Medicaid is providing redistribution as well as insurance. Relatedly, redistribution is never "free"; because of the distortionary effects of redistributive policies, the amount of feasible redistribution through the tax system is also less than a dollar per dollar of government spending (Okun 1975 [39]).

We therefore also consider Medicaid as a redistributive tool – which by necessity entails some resource cost – and consider the value of Medicaid to recipients per dollar of government expenditure relative to the value achieved by other redistributive policies. As an illustrative example, we consider society's preference between the hypothetical policy choice of eliminating Medicaid coverage for low-income adults or a budgetarily equivalent reduction in the EITC. We focus on the EITC because it is the primary form of tax subsidies to low-income individuals in the United States and because – not unrelatedly – there is a large empirical literature that has estimated its distortionary effect on labor supply. Drawing on this existing literature, Hendren (2014) [32] estimates that the EITC generates roughly \$0.9 of welfare to EITC recipients for every dollar of government spending; this ratio – which is the EITC equivalent of our estimates of $\gamma(1)/G$ for Medicaid – is less than 1 due

to the labor supply distortions created by the EITC.⁴⁰

To compare this estimate for the EITC to the value of Medicaid per dollar of government expenditure, we must account for the transfers Medicaid generates to external parties of \$0.6 per dollar of G . This in turn requires knowing the incidence of where these transfers fall in the income distribution. Absent such knowledge, we present stylized calculations under two alternative, extreme scenarios.

First, we assume the recipients of Medicaid themselves bear the ultimate incidence of the transfers to external parties. In this approach, $\gamma(1) + N$ is the total welfare benefit accruing to the recipients.⁴¹ As shown in row 4, this produces a ratio of the value of Medicaid per dollar of government spending of \$0.8 to \$1 depending on the approach. Second, we assume that the incidence of the external-party transfers, N , is on the group from whom it is most costly to redistribute from, which is generally at the upper regions of the income distribution. Hendren (2014) [32] estimates that \$1 falling near the top of the income distribution can be turned into \$0.50 to an EITC beneficiary through modifications in the income tax schedule.⁴² In this case, only 50% of the 0.6 transfer to external parties can be brought back to the bottom of the income distribution. As shown in row 5, this produces a ratio of the value of Medicaid per dollar of government spending of \$0.5 to \$0.7 depending on the approach. If recipients of Medicaid and of the EITC are (largely) the same individuals, these results suggest that if the incidence of Medicaid's external transfers is on the top of the income distribution, society would prefer cutting Medicaid coverage to a budgetarily equivalent reduction in the EITC (since the former is valued at \$0.5 to \$0.7 on the dollar, compared to \$0.9 for the EITC); if however the incidence of Medicaid's external transfers is on Medicaid recipients themselves, society would be roughly indifferent between the two policies.

Of course, Medicaid and the EITC recipients are not exactly the same population. Although both programs are targeted at the bottom of the income distribution, EITC recipients reach higher in the income distribution than the Medicaid program we are studying here, which requires an income below the federal poverty line. In addition the average Medicaid beneficiary likely differs from the average EITC beneficiary in other dimensions such as health status or ability to work. As a result, the relative social welfare weights on EITC and Medicaid recipients may differ; for example, one may put higher social welfare weight on those who, conditional on income, have higher health

⁴⁰Note that in the Medicaid context, G should include any impacts of labor supply distortions from the provision of Medicaid on the government budget. As noted above, Baicker et al. (2014)[4] find no statistically significant impact on labor supply.

⁴¹This would also be the relevant welfare metric even when transfers fall to other regions of the income distribution if those transfers induce behavioral responses that have offsetting fiscal externalities (Kaplow 2008 [34]); for example, people may increase labor supply because the effective tax rate on high incomes is decreased when the transfers fall on the top of the income distribution.

⁴²In other words, 50 cents are "lost" because the behavioral responses to the additional taxes required for the transfer reduce taxable income. Of course, the amount of the "loss" naturally depends on the assumed value of the taxable income elasticity. The \$0.50 estimate was calculated assuming a taxable income elasticity of value of around 0.3. Alternatively, rather than thinking about hypothetical transfers across the income distribution, if one assumes that the tax system is set in response to a social welfare function, the finding that \$1 in the hands of the top of the income distribution can be turned into \$0.50 for the bottom of the income distribution implies a relative social welfare weight on the top of the income distribution of 0.5.

expenditures.

An alternative way to interpret these illustrative calculations, therefore, is to use them to solve for the relative social welfare weights on the two populations that would leave the social planner indifferent between the two policies. Under the assumption that the incidence of N is on the top of the income distribution, spending \$1 on Medicaid and taxing the \$.6 transfer of N back to EITC recipients at a cost of 0.5 per dollar generates a welfare impact to recipients of $\frac{\gamma(1)}{G}\eta^M + 0.3\eta^{EITC}$, where η^M is the social welfare weight on the Medicaid beneficiaries and η^{EITC} is the social welfare weight on the EITC beneficiaries. Conversely, \$1 of spending on EITC generates $0.9\eta^{EITC}$ units of welfare. For a given estimate of $\frac{\gamma(1)}{G}$, we can therefore solve for the minimum ratio of social welfare weights $\frac{\eta^M}{\eta^{EITC}}$ that implies $\frac{\gamma(1)}{G}\eta^M + 0.3\eta^{EITC} = 0.9\eta^{EITC}$. This ratio has a simple interpretation in terms of social preferences: society is indifferent between giving \$1 to a Medicaid recipient versus $\frac{\eta^M}{\eta^{EITC}}$ dollars to an EITC recipient.

Given our estimated range of \$0.2 to \$0.4 for $\frac{\gamma(1)}{G}$, this implies that society would prefer cutting Medicaid to a budgetarily equivalent cut in the EITC if it values \$1 in the hands of a Medicaid recipient less highly than \$1.5 to \$3 in the hands of an EITC recipient. By contrast, if the incidence of N is on the Medicaid recipients themselves, society would prefer cutting Medicaid coverage to a budgetarily equivalent reduction in the EITC if it values \$1 in the hands of Medicaid recipients less highly than \$0.9 to \$1.10 in the hands of an EITC recipient.⁴³ The large range of potential welfare weights generated by different assumptions about the incidence of the external transfers highlights the importance of future work examining the incidence of these transfers.

5.3 Sensitivity Analysis

Table 4 explores the sensitivity of our results within each framework (shown in different rows) to a number of different assumptions (shown in different columns). For the sake of brevity, we focus on two main results: the value of Medicaid to recipients per dollar of government spending ($\gamma(1)/G$), and the value of Medicaid to recipients per dollar of net costs ($\gamma(1)/C$). Column I shows our baseline results.

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 \$500 or \$2,000, compared to our baseline assumption of \$1,000). Higher risk aversion or a lower consumption floor increases the value of the program using the complete-information approach or the consumption-based optimization approaches.⁴⁴

⁴³To see this, note that 0.9 to 1.1 is the range of values for the ratio of welfare benefits to Medicaid recipients per dollar of government spending (\$0.8 to \$1) to the welfare benefit to EITC recipients per dollar of government spending (\$0.9).

⁴⁴The health-based optimization approach is mechanically unaffected by our assumption regarding 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 (i.e., we cap the distribution of out-of-pocket spending at $\bar{y} - c_{floor}$) and because we use the same estimates of out-of-pocket spend-

Alternative health measures In column VI, we assume a VSLY of \$100,000 instead of our baseline assumption of \$25,000; A VSLY of \$100,000 implicitly assumes no adjustment to the value in Cutler (2004)[18] for that fact that low-income individuals have a higher marginal utility of consumption, and therefore a lower willingness to pay, than an average person. In column VII, we assume no health benefits; this seems a reasonable alternative to our baseline approach of valuing the improvements in self-reported health, since we are unable to reject the null hypothesis of no impact of Medicaid on mortality (Finkelstein et al. (2012)[25]) or on specific measures of physical health (Baicker et al. (2013)[6]).

Column VIII uses a depression measure as our health measure, instead of our baseline self-reported health measure. We estimate (see Appendix Table A.1) that Medicaid reduces the probability of screening positive for depression by about 6 percentage points (about 20 percent). The World Health Organization estimates that a moderate depression episode corresponds to a loss of 0.4 life years.⁴⁵ Combining this with our baseline VSLY of \$25K, we assign a value of \$10K to a reduction in depression. For the complete-information approach, we can use only the depression measure of health (as reported in Table 4) or, if we believe the self-reported improvements in health capture health gains that are not entirely due to depression, we could instead include both subjective health and depression outcomes in the utility function as two separate arguments.⁴⁶ In contrast to the complete-information approach, which requires evaluating all health outcomes affected by Medicaid, the health-based optimization approach can be implemented using any single health measure; this also highlights the versatility of the optimization-based approaches, as they can be implemented using any argument of the utility function, even within health or consumption.

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 non-trivial out-of-pocket spending (Finkelstein et al. (2012)[25]).⁴⁷ In Column IX, 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.⁴⁸ We now have to estimate the “pure-insurance” term in equation (10) 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

ing 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.

⁴⁵See Annex Table D, http://www.who.int/healthinfo/statistics/GlobalDALYmethods_2000_2011.pdf

⁴⁶Our estimate of $\gamma(1)/G$ increases to 0.47 (not reported) if we do so, compared to 0.36 using only depression, or our baseline estimate of 0.44 using only self-reported health.

⁴⁷This 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)[29]).

⁴⁸In constructing $-\frac{\partial x}{\partial q} = x(0, m(q; \theta)) - x(1, m(q; \theta))$, 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.6.

income effects.⁴⁹

Allowing for $x(1, m) > 0$ necessarily reduces our estimates of $\gamma(1)$ but it also reduces our estimates of G and of C , so that the net effect on $\gamma(1)/G$ or $\gamma(1)/C$ is a priori ambiguous. On practice, Column IX shows that it lowers our estimates of these ratios of the value of Medicaid to its (gross or net) cost.

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.7 provides implementation details. In Column X, we assume instead that the demand for medical care is linear in price.⁵⁰ In column XI, we calculate an upper bound for $\gamma(1)$ over possible interpolation assumptions by searching for the (non-parametric) 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 $q = 1$.⁵¹

Alternative assumption about average consumption Measuring consumption is both notoriously difficult and important for welfare estimates. Our baseline approach used average per capita self-reported income of compliers for \bar{y} in computing consumption in equation (18). Yet formal income may substantially understate consumption in a low-income population. Indeed, our baseline estimate of average c for control compliers of \$3,320 is considerably below the CEX's \$13,541 estimate of non-medical per-capita consumption for uninsured adults below the federal poverty line.

We therefore considered how our welfare estimates would change if we assume an average level of consumption corresponding to the CEX estimates. Specifically, we assume that \bar{y} is \$14,030, which is the sum of mean non-medical consumption in the CEX (\$13,541) plus out-of-pocket spending of control compliers in the Oregon sample (\$489). We set the consumption floor at \$3,704 which corresponds to about the 1st percentile of reported consumption for our uninsured sample in the CEX (and represents the same proportional increase as average income).

Column XII reports the results of this large increase in the assumed level of average consumption, which are substantial. For example, the estimate of $\gamma(1)/G$ almost doubles (to about 0.85 per dollar of government spending) under either the complete-information approach or the consumption-based optimization approach using the consumption proxy.⁵² This increase partially reflects the fact

⁴⁹Specifically, 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 non-zero 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 $x(1, m) = 0$.

⁵⁰Because 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.

⁵¹We do not report an upperbound 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.

⁵²The results from the other two approaches are barely affected. In principle, the estimates from the health-based

that in the presence of a consumption floor, increases in average consumption increase exposure to downside risk and hence increase the value of insurance. In addition, with higher average consumption, the marginal utility of consumption is lower and therefore the “willingness to pay” for Medicaid is higher.⁵³

Summary Looking across the columns gives a sense of the scope and drivers of our estimates. In the complete-information approach, we can almost double our estimate of the value of Medicaid (relative to either gross or net costs) by using an average income level based on average consumption reported in the CEX rather than based on self-reported income (column XII); we can reduce the estimates by up to one-third by allowing for out-of-pocket spending among the insured ($x(1, m) > 1$), or with alternative parametric assumptions regarding risk aversion, the consumption floor, or the VSLY. These same alternative specifications also produce the extremes of the estimates under the consumption-based optimization approach using the consumption proxy. (The other two approaches are relatively stable across alternative assumptions, although – as noted above – the health-based optimization approach has quite large standard errors on it).

Under all the assumptions we have examined, all of the frameworks consistently estimate $\gamma(1) < G$; in other words, we estimate that the low-income uninsured would prefer not to have Medicaid if they had to pay the government’s costs of the program. The different frameworks and assumptions produce a range of estimates of $\gamma(1)/G$ of about 0.15 to 0.85. We also consistently estimate that the low-income uninsured would prefer not to have Medicaid if they had to pay the net costs of the program (i.e., $\gamma(1) < C$) under the consumption based optimization approach using the CEX consumption measure and the health-based optimization approach; the complete-information approach and the consumption-based optimization approach using the consumption proxy tend to deliver an estimate of $\gamma(1)/C$ around 1, although under particular assumptions this ratio can get substantially above or below 1.

6 Conclusion

Welfare estimation of non-market goods is important, but also challenging. As a result of these challenges, the welfare benefits from Medicaid are often ignored in academic and public policy discourse, or based on ad-hoc approaches. In this paper, we have developed, implemented, and

optimization approach and the consumption-based optimization approach based on the CEX consumption measure should not be affected by this alternative assumption regarding \bar{y} . In practice, the estimates from these other two approaches are affected because we implement the consumption floor (used in the complete-information and consumption-based optimization approach with the consumption proxy) by capping the distribution of out-of-pocket spending at $\bar{y} - c_{floor}$ and, for consistency, we use the same distribution of out-of-pocket spending for all approaches (even ones that do not depend on consumption). In addition, the change in the assumed consumption floor can have a direct effect on the consumption-based optimization approach using the CEX consumption measure.

⁵³Interestingly, at the higher level of assumed average consumption in Column XII, it makes virtually no difference whether we calculate welfare as giving up consumption in the insured state (as in our baseline definition of $\gamma(1)$) or as increasing consumption in the uninsured state (as in the $\tilde{\gamma}(1)$ metric described in footnote (39)). The estimates are $\gamma(1) = \$2,969$ and $\tilde{\gamma}(1) = \$3,012$.

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. We also emphasize that the interpretation depends critically on the metric used to evaluate Medicaid.

If we ask whether Medicaid recipients would be willing to cover the government’s costs of Medicaid, the answer is a fairly robust “no”. Our baseline estimates indicate that the welfare benefit to recipients per dollar of government spending is roughly between \$0.2 and \$0.4, depending on the framework used. Across a variety of alternative implementing assumptions, we estimate welfare benefits to recipients per dollar of government spending ranging from about \$0.15 to \$0.85. This stands in contrast to the current approach used by the Congressional Budget Office to value Medicaid at government cost. It is, however, not inconsistent with the few other attempts we know of to formally estimate a value for Medicaid; these are based on using choices to reveal ex-ante willingness to pay, and tend to find estimates (albeit for different populations) in the range of 0.3 to 0.5.⁵⁴

A key factor behind our finding of a low recipient value of Medicaid per dollar of government spending is our estimate that the uninsured 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. Substantial implicit insurance for the uninsured reduces the potential pure insurance value of Medicaid for recipients; it also reduces the transfers to them from Medicaid. 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; this transfer to external parties is larger than our baseline estimates of the value of Medicaid to recipient. It is consistent with recent findings of Garthwaite et al. (2015)[27] who find that Medicaid significantly reduces the provision of uncompensated care by hospitals.

If instead we ask whether Medicaid represents a preferable form of redistribution to low-income individuals compared to the EITC, the answer is “it depends.” While the result is naturally affected by our framework and modeling assumptions, a critical component is the ultimate economic incidence of the transfers from Medicaid to the previous providers of implicit insurance for the low-income uninsured. Depending on our assumptions, we find that the social planner would prefer eliminating Medicaid coverage over a budgetarily equivalent reduction in the EITC if he values \$1 in the hands of Medicaid recipients less than \$0.9 to \$3 in the hands of an EITC recipient. The relative social welfare weights on Medicaid compared to EITC recipients could be higher if, for example, there is a higher social welfare weight on those who, conditional on income, have higher health expenditures. Alternatively, society may have a non-utilitarian preference for individuals

⁵⁴Researchers have attempted to elicit willingness to pay for Medicaid in other contexts through quasi-experimental variation in premiums (Dague (2014)[19]), and the extent to which individuals distort their labor earnings in order to become eligible for Medicaid (Gallen (2014)[26], Keane and Moffitt (1998)[35]). These approaches – which face their own challenges – are not available in the context of the Oregon Health Insurance Experiment, since it randomized the ability to apply for Medicaid but not the premium at which one could do so or the eligibility requirements.

being able to access to health care; indeed, this may explain the large amount of implicit insurance in the absence of Medicaid.

Our findings thus highlight the importance of further work on who bears both the immediate and ultimate economic incidence of the large Medicaid transfers to external parties. A related question is whether there are efficiency gains – such as through reduced transaction costs – in having formal public insurance provision rather than implicit insurance, which presumably entails various efforts to recoup some costs from the uninsured and/or other potential payers. Our analysis has ignored any such potential efficiency gains from formal insurance.

Our empirical findings are naturally specific to our setting. Fortunately, the approaches we have developed can be applied to studying similar questions in other settings. For example, the nature of the Oregon Health Insurance Experiment only permits empirical examination of the effects of Medicaid over the first one to two years. An interesting direction for future work would be to apply the approaches developed here to estimating the value of Medicaid over longer time horizons, drawing perhaps on recent empirical estimates of the long-run impacts of Medicaid (see, e.g., Wherry et al. (2015)[49] and Brown et al. (2015)[8]).

Relatedly, it would be interesting to apply the approaches we have developed to studying the value of Medicaid for other Medicaid 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)[30] and Buchmueller et al. (2015)[10] 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)[11, 12], Barcellos and Jacobson (forthcoming)[7]) 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] 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.

- [2] Joshua D. Angrist and Jörn-Steffen Pischke. Mostly harmless econometrics: An empiricist's companion. *Princeton University Press*, 2009.
- [3] Kenneth J. Arrow. Uncertainty and the welfare economics of medical care. *American Economic Review*, 53(5):941–973, 1963.
- [4] 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.
- [5] Katherine Baicker, Sendhil Mullainathan, and Joshua Schwartzstein. Behavioral hazard in health insurance. *Quarterly Journal of Economics*, forthcoming.
- [6] 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.
- [7] Silvia Helena Barcellos and Mireille Jacobson. The effects of Medicare on medical expenditure risk and financial strain. *American Economic Journal: Economic Policy*, forthcoming.
- [8] David W. Brown, Amanda E. Kowalski, and Ithai Z. Lurie. Medicaid as an investment in children: What is the long-term impact on tax receipts? NBER Working Paper 20835, 2015.
- [9] 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.
- [10] 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.
- [11] 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.
- [12] David Card, Carlos Dobkin, and Nicole Maestas. Does Medicare save lives? *Quarterly Journal of Economics*, 124(2):597–636, 2009.
- [13] Raj Chetty. Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods. *Annual Review of Economics*, 1(1):451–488, 2009.
- [14] Raj Chetty and Emmanuel Saez. Optimal taxation and social insurance with endogenous private insurance. *American Economic Journal: Economic Policy*, 2(2):85–114, 2010.
- [15] John H. Cochrane. Asset pricing. *Princeton University Press*, 2005.

- [16] 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.
- [17] Janet Currie and Firouz Gahvari. Transfers in cash and in-kind: Theory meets the data. *Journal of Economic Literature*, 46(2):333–383, 2008.
- [18] David M. Cutler. Your money or your life: Strong medicine for America’s health care system. *New York: Oxford University Press*, 2004.
- [19] Laura Dague. The effect of Medicaid premiums on enrollment: A regression discontinuity approach. *Journal of Health Economics*, 37:1–12, 2014.
- [20] Isaac Ehrlich and Gary S. Becker. Market insurance, self-insurance, and self-protection. *Journal of Political Economy*, 80(4):623–648, 1972.
- [21] Liran Einav, Amy Finkelstein, and Jonathan Levin. Beyond testing: Empirical models of insurance markets. *Annual Review of Economics*, 2(1):311–336, 2010.
- [22] 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.
- [23] 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.
- [24] 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.
- [25] 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.
- [26] Trevor S. Gallen. Measuring the value of Medicaid using TennCare. *Working Paper*, 2014.
- [27] Craig Garthwaite, Tal Gross, and Matthew J. Notowidigdo. Hospitals as insurers of last resort. *Working Paper*, 2015.
- [28] Peter Gottschalk and Timothy M. Smeeding. Cross-national comparisons of earnings and income inequality. *Journal of Economic Literature*, 35(2):633–687, 1997.
- [29] 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.

- [30] Jonathan Gruber. Medicaid. In Robert A. Moffitt, editor, *Means-tested transfer programs in the United States*, pages 15–78. University of Chicago Press, 2003.
- [31] Arnold Harberger. The measurement of waste. *American Economic Review*, 54(3):58–76, 1964.
- [32] Nathaniel Hendren. The inequality deflator: Interpersonal comparisons without a social welfare function. *Working Paper*, 2014.
- [33] 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.
- [34] Louis Kaplow. *The Theory of Taxation and Public Economics*. Princeton University Press, 2008.
- [35] Michael Keane and Robert Moffitt. A structural model of multiple welfare program participation and labor supply. *International Economic Review*, 39(3):553–589, 1998.
- [36] Neale Mahoney. Bankruptcy as implicit health insurance. *American Economic Review*, 105(2):710–46, 2015.
- [37] John A. Nyman. The economics of moral hazard revisited. *Journal of Health Economics*, 18(6):811–824, 1999.
- [38] John A. Nyman. The value of health insurance: The access motive. *Journal of Health Economics*, 18(2):141–152, 1999.
- [39] Arthur Okun. *Equality and Efficiency*. Brookings Institution Press, 1975.
- [40] Thomas Piketty and Emmanuel Saez. Optimal labor income taxation. *Handbook of Public Economics*, 5:391–474, 2013.
- [41] Paul A. Samuelson. The pure theory of public expenditure. *Review of Economics and Statistics*, 36(4):387–389, 1954.
- [42] 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.
- [43] U.S. Centers for Medicare and Medicaid Services. 2012 Actuarial report on the financial outlook for Medicaid. 2012.
- [44] U.S. Congressional Budget Office. The distribution of household income and federal taxes, 2008 and 2009. 2012.
- [45] U.S. Congressional Budget Office. Growth in means-tested programs and tax credits for low-income households. 2013.

- [46] U.S. Congressional Budget Office. Insurance coverage provisions of the Affordable Care Act - CBO's February 2014 baseline. 2014.
- [47] U.S. Department of Health and Human Services. Fiscal year 2011 TANF financial data. 2012.
- [48] 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.
- [49] Laura R. Wherry, Sarah Miller, Robert Kaestner, and Bruce D. Meyer. Childhood medicaid coverage and later life health care utilization. Working Paper 20929, 2015.

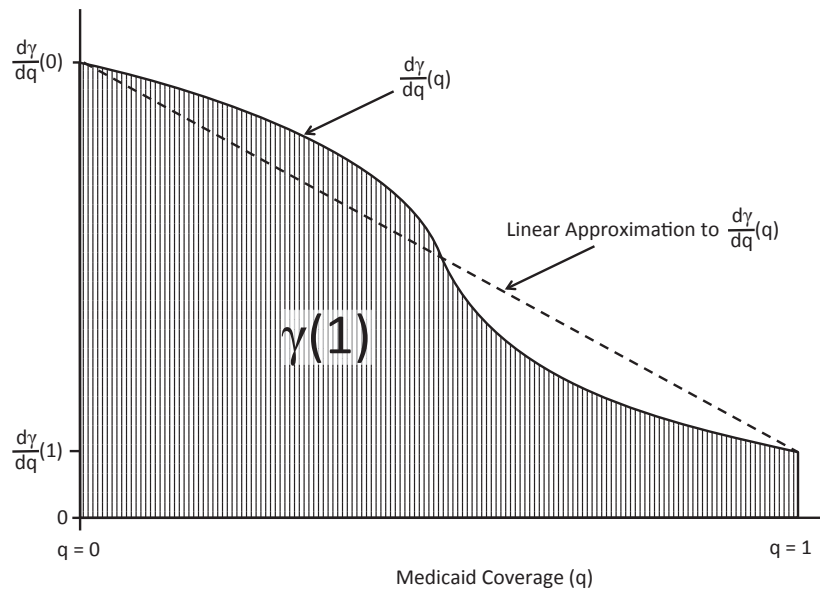


Figure 1: From $\frac{d\gamma}{dq}(q)$ to $\gamma(1)$

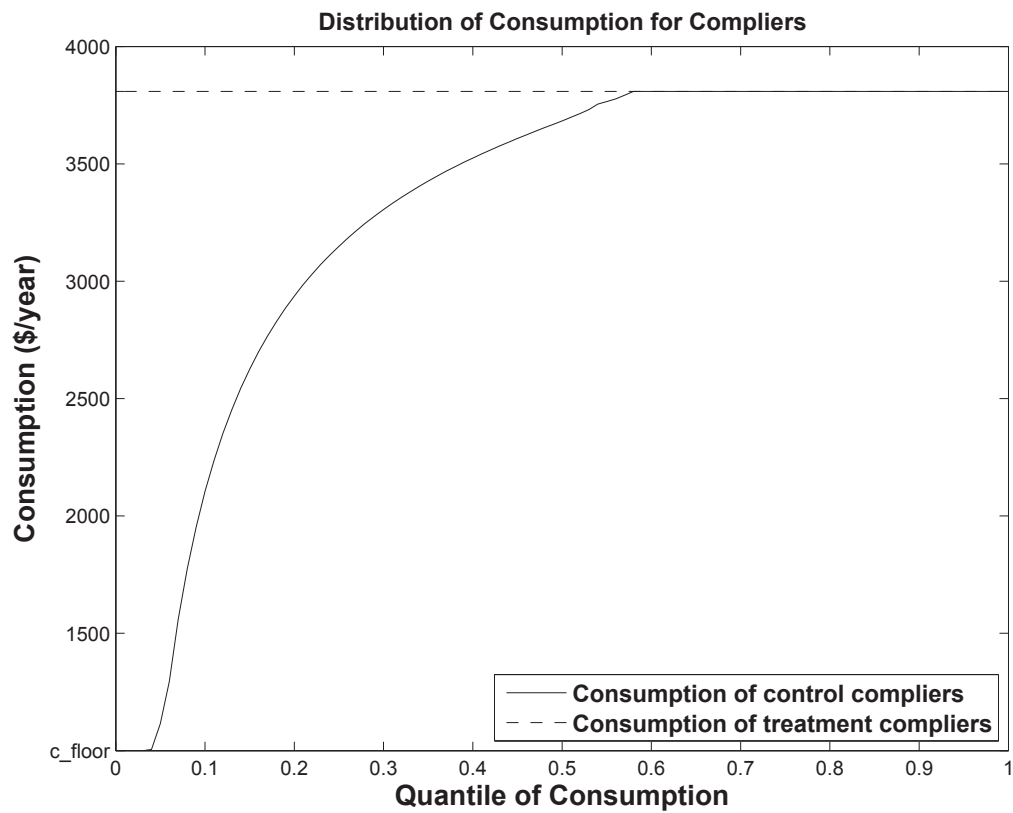


Figure 2: Consumption Distribution for Treatment and Control Compliers

Table 1: Summary Statistics

	I	II	III	IV
	Full Sample	Treatment Compliers ($q=1$)	Control Compliers ($q=0$)	Impact of Medicaid
Panel A: Oregon Data Demographics				
Female	0.60	0.57	0.60	
Age 50-64	0.34	0.36	0.35	
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	0.08	0.08	
Panel B: Oregon Data Outcomes				
Medical spending (m)				
Mean medical spending, $E[m]$	2990	3596	2711	885
Fraction with positive medical spending, $E[m > 0]$	0.74	0.79	0.72	0.07
Out-of-pocket spending (x)				
Mean out-of-pocket spending, $E[x]$	351	0	489	-489
Fraction with positive out-of-pocket spending, $E[x > 0]$	0.38	0	0.56	-0.56
Fraction in good, very good or excellent health, $E[h]$	0.59	0.61	0.47	0.14

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 - all based on using the lottery as an instrument for Medicaid coverage (See Appendix A.1 for details on how these estimates are constructed). See the text for detailed definitions of variable construction in Panel B. Spending and consumption are expressed in dollars per year per person.

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, $\gamma(1)$	1576	1260	694	680
(standard error)	(91)	(260)	(391)	(1063)
Operating through consumption	1111			
Operating through health	465			
B. Decomposition of Welfare Effect on Recipients				
Transfer component, T	693	568	568	568
Operating through consumption	489			
Operating through health	205			
Pure-Insurance (non-transfer) component	883	691	126	112
Operating through consumption	622			
Operating through health	260			

Notes: All estimates are expressed in dollars per year per Medicaid recipient. Panel A presents estimates of the welfare impact on recipients, $\gamma(1)$, for each of the approaches. Panel B decomposes this number into components. The sample in column III is limited to single individuals. Standard errors are calculated based on bootstrapping with 500 repetitions.

Table 3: Benchmarking the Welfare Estimates

	I	II	III	IV
	Complete-Information Approach	Consumption-Based Optimization Approach (Consumption Proxy)	Consumption-Based Optimization Approach (CEX Consumption Measure)	Health-Based Optimization Approach
Panel A: External Transfers and Costs				
Net Resource Costs , C	1374 (269)	1374 (269)	1374 (269)	1374 (269)
Government Costs , G	3596	3596	3596	3596
Transfers to External Parties, N	2222	2222	2222	2222
Panel B: Benchmarks				
$\gamma(1) / G$	0.44 (0.03)	0.35 (0.07)	0.19 (0.11)	0.19 (0.83)
$\gamma(1) / C$	1.15 (0.16)	0.92 (0.10)	0.51 (0.15)	0.49 (1.38)
$(dy/dq)/(dc/dq)$, at $q=0$	-	1.91 (1.33)	0.35 (1.81)	0.31 (1.84)
$(\gamma(1)+N) / G$	1.06 (0.09)	0.97 (0.07)	0.81 (0.12)	0.81 (1.04)
$(\gamma(1)+0.5N) / G$	0.75 (0.06)	0.66 (0.07)	0.50 (0.12)	0.50 (0.60)

Notes: All estimates in Panel A are expressed in dollars per year per Medicaid recipient. Standard errors are calculated based on 1000 bootstrap repetitions.

Table 4: Sensitivity of Welfare Estimates to Assumptions

	I	II	III	IV	V	VI	VII	VIII	IX	X	XI	XII
	Baseline	Coefficient of Relative Risk Aversion	5	Consumption Floor	Consumption Floor	Value of a Statistical Life Year (VSLY)	0	Alternative health measure (depression)	Allow for out-of-pocket spending at q=1	Alternative Interpolations	Upperbound	Use CEX for assumption of mean income
Panel A: $\gamma(I) / G$												
Complete-Information Approach	0.44	0.36	0.49	0.52	0.37	0.61	0.31	0.36	0.37	-	-	0.83
Consumption-Based Optimization Approach (Consumption Proxy)	0.35	0.24	0.42	0.52	0.23	0.35	0.35	0.35	0.21	0.43	0.69	0.93
Consumption-Based Optimization Approach (CEX Consumption Measure)	0.19	0.18	0.22	0.20	0.17	0.19	0.19	0.19	0.14	-	-	0.22
Health-Based Optimization Approach	0.19	0.19	0.19	0.19	0.18	0.28	0.16	0.18	0.13	0.22	0.31	0.21
Panel B: $\gamma(I) / C$												
Complete-Information Approach	1.15	0.94	1.28	1.34	1.02	1.59	0.81	0.95	1.12	-	-	2.04
Consumption-Based Optimization Approach (Consumption Proxy)	0.92	0.62	1.11	1.34	0.64	0.92	0.92	0.92	0.62	1.12	1.82	2.30
Consumption-Based Optimization Approach (CEX Consumption Measure)	0.51	0.48	0.58	0.51	0.48	0.51	0.51	0.51	0.43	-	-	0.55
Health-Based Optimization Approach	0.49	0.49	0.49	0.48	0.50	0.74	0.41	0.46	0.39	0.57	0.81	0.52

Notes: This table presents the sensitivity of the baseline estimates in Table 3 to alternative assumptions and specifications. Column I presents our baseline results, which assume a coefficient of relative risk aversion of 3, a consumption floor of \$500, a VSLY of \$25,000, no out-of-pocket spending for the insured, and mean income per capita per year of \$3,808. Each subsequent column presents a one-off deviation from these baseline assumptions. Columns II-III adjust the coefficient of relative risk aversion from 1 to 5, as compared to our baseline parameter of 3. Columns IV-V adjust the consumption floor between 500 and 2000, as opposed to our baseline parameter of 1000. Columns VI-VII set the value of a statistical life year at \$100,000 or \$0, as compared to our baseline value of \$25,000. The estimate for the health-based optimization approach in column VII is based on the limit when the VSLY approaches zero. Column VIII uses the reduction in depression symptoms to value health improvements. Column IX reports results from a specification that drops the assumption that individuals have zero out-of-pocket spending under Medicaid but instead treats self-reported out-of-pocket medical spending as the estimate of out-of-pocket health spending under Medicaid. Columns X and XI report results for alternative interpolations between q=0 and q=1 in the optimization approaches. Column XII reports results from the specification in which mean income is set at \$14,030 - the sum of mean out-of-pocket medical spending of Oregon control compliers (\$489) and mean reported non-health consumption among uninsured individuals in baseline CEX sample (\$13,541) - as compared to our baseline value of \$3808; we set the consumption floor at \$3,704, the same fraction of mean income as in our baseline specification.

A Online Appendices

A.1 Derivations

To derive equation (9) from equation (8), consider the level of utility the individual obtains if she optimally chooses $m(q; \theta)$ subject to a budget constraint in which she must pay $\gamma(q)$ to obtain insurance:

$$V(q) = \max_{m(q; \theta)} E[u(y(\theta) - x(q, m(q; \theta)) - \gamma(q), h(m(q; \theta); \theta))], \quad (22)$$

where the expectation is taken with regard to θ . The envelope theorem implies:

$$\frac{dV}{dq} = E\left[\left(-\frac{\partial x}{\partial q}\right)u_c\right] - \frac{d\gamma}{dq}E[u_c].$$

Given that $V(q) = E[u(c(0; \theta), h(0; \theta))]$ by equation (8), it follows that $dV/dq = 0$. Using $dV/dq = 0$ to solve the equation above for $d\gamma/dq$ yields:

$$\frac{d\gamma}{dq} = \frac{E\left[\left(-\frac{\partial x}{\partial q}\right)u_c\right]}{E[u_c]}.$$

Using equation (7), $-\frac{\partial x}{\partial q} = (p(0) - p(1))m(q; \theta)$, we obtain:

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

which was what we wanted to show. Note that this derivation does not require medical spending to be strictly positive and allows for cases where there are “lumpy” medical expenditures so that an individual is not indifferent between an additional \$1 of out-of-pocket medical spending and \$1 less consumption.⁵⁵ Intuitively, the individual values the mechanical increase in consumption from Medicaid according to the marginal utility of consumption, regardless of the extent to which she has ability to substitute an increase in another good (e.g., medical care) for the increase in consumption.

This derivation is based on individuals facing a budget constraint in which they pay $\gamma(q)$ for insurance. If in reality, they do not have to pay $\gamma(q)$ and the demand for m has a non-zero income elasticity, the observed choices of m will differ from those used in this derivation and, as a result, this derivation is not exact. In practice, however, this issue does not affect our estimates because

⁵⁵ Although the optimization requires individuals to equalize the marginal cost and marginal benefit of additional medical spending, we did not require concavity in the health production function, and we allow for insurance to affect medical spending in a discontinuous or lumpy fashion. Non-concavities in the health production function and non-convexities in the out-of-pocket spending schedule could lead to discontinuities in the marginal utilities (e.g., the marginal utility of consumption may jump up at the point of deciding to increase medical spending by a discontinuous amount in order to undergo an expensive medical procedure), but the equation for $\gamma(q)$ in integral form will remain continuous because, when the individual is at the margin of undertaking the jump, the individual will be indifferent to undertaking the jump or not.

our estimates are based on a linear interpolation of $d\gamma/dq$ between $q = 0$ and $q = 1$. At $q = 0$, γ is zero and therefore the observed choices of m are the same whether or not individuals have to pay γ . At $q = 1$, individuals are fully insured and the marginal insurance value is zero in any case. The only estimate for which this issue does arise is the specification check discussed in column IX of Table 4, in which we do not assume that individuals are fully insured at $q = 1$ (i.e., where they have strictly positive out-of-pocket spending at $q = 1$). The next subsection contains an alternative setup, in which individuals face insurance lotteries. In this alternative setup, the above formula for $d\gamma/dq$ holds exactly.

Derivation in an alternative setup: insurance lotteries

In order to satisfy the maximization in equation (8), the relevant arguments of the marginal utility function, u_c , need to be the choices that individuals in state θ would actually make if they face price $p(q)$, have income $y(\theta)$, and pay $\gamma(q)$. Observed choices do not satisfy this maximization if individuals, in fact, do not pay $\gamma(q)$. Intuitively, for $q > 0$, there would be income effects that cause people to change their allocation of c and m . In practice, in our baseline implementation this is not a problem, since we assume $x(1, m) = 0$ – and therefore we know the pure-insurance term (as defined in equation (10)) must be zero on the margin for the fully insured – and linearly interpolate between our estimates of $\frac{d\gamma}{dq}$ at $q = 0$ and $q = 1$.

More generally though, we can derive the optimization implementation for a thought experiment in which we consider the willingness to pay to avoid an ϵ -chance of losing Medicaid (and returning to $q = 0$). In this alternative setup, we define $\gamma(q)$ as $1/\epsilon$ times the willingness to pay to avoid an ϵ -chance of losing q units of insurance, with $\epsilon \rightarrow 0$. Formally: $\gamma(q)$ solves for $\epsilon \rightarrow 0$:

$$E[u(c(q; \theta) - \epsilon\gamma(q), h(q; \theta))] = (1 - \epsilon) E[u(c(q; \theta), h(q; \theta))] + \epsilon E[u(c(0; \theta), h(0; \theta))]. \quad (23)$$

We derive $\frac{d\gamma}{dq}$ in the insurance lotteries setup by considering the first-order condition for the choice of m in the special case when choices are continuously differentiable in q . This approach is detailed in the subsection below and shows that $\frac{dc}{dq}u_c + \frac{dh}{dq}u_h = \left(-\frac{\partial x}{\partial q}\right)u_c$. Taking the derivative of (23) with respect to q , and using $\frac{dc}{dq}u_c + \frac{dh}{dq}u_h = \left(-\frac{\partial x}{\partial q}\right)u_c$, we obtain:

$$\begin{aligned} E[u_c(c(q; \theta) - \epsilon\gamma(q), h(q; \theta))] \left(-\epsilon \frac{d\gamma}{dq}\right) + E\left[u_c(c(q; \theta) - \epsilon\gamma(q), h(q; \theta)) \left(-\frac{\partial x}{\partial q}\right)\right] \\ = (1 - \epsilon) E\left[u_c(c(q; \theta) - \epsilon\gamma(q), h(q; \theta)) \left(-\frac{\partial x}{\partial q}\right)\right]. \end{aligned}$$

Rearranging and taking the limit $\epsilon \rightarrow 0$ yields:

$$\frac{d\gamma}{dq} = \lim_{\epsilon \rightarrow 0} \frac{E\left[u_c(c(q; \theta) - \epsilon\gamma(q), h(q; \theta)) \left(-\frac{\partial x}{\partial q}\right)\right]}{E[u_c(c(q; \theta) - \epsilon\gamma(q), h(q; \theta))]} = \frac{E\left[u_c(c(q; \theta), h(q; \theta)) \left(-\frac{\partial x}{\partial q}\right)\right]}{E[u_c(c(q; \theta), h(q; \theta))]}.$$

Now, noting that $\left(-\frac{\partial x}{\partial q}\right) = (p(1) - p(0))m(q; \theta)$, we obtain precisely equation (9).

Alternative derivation: using the first-order condition

Given the central role of equation (9) in the optimization approaches, we also derive equation (9) by exploiting the first-order condition. This derivation requires the first-order condition (equation (14)) to hold and is therefore less general than our main derivation, which is based on the envelope theorem. However, the derivation based on the first-order condition very nicely shows the intuition behind the optimization approaches, and we therefore present it here.

To derive equation (9) from equation (8), it is useful to first derive two intermediate expressions. First, we differentiate the budget constraint $c(q; \theta) = y(\theta) - x(q, m(q; \theta))$ with respect to q :

$$\frac{dc}{dq} = -\frac{\partial x}{\partial q} - \frac{\partial x}{\partial m} \frac{dm}{dq} = -\frac{\partial x}{\partial q} - p(q) \frac{dm}{dq} \quad \forall q, \theta. \quad (24)$$

The total change in consumption from a marginal change in Medicaid benefits, $\frac{dc}{dq}$, equals the impact on the budget constraint, $-\frac{\partial x}{\partial q}$, plus the impact through the behavioral response in the choice of m , $-\frac{\partial x}{\partial m} \frac{dm}{dq} = -p(q) \frac{dm}{dq}$.

Second, we use the health production function (equation (2)) to express the marginal impact of Medicaid on health, $\frac{dh}{dq}$, as:

$$\frac{dh}{dq} = \frac{d\tilde{h}}{dm} \frac{dm}{dq} \quad \forall q, \theta. \quad (25)$$

We then totally differentiate equation (8) with respect to q , which yields the marginal welfare impact of insurance on recipients, $\frac{d\gamma}{dq}$, as the implicit solution to:

$$0 = E \left[\left(\frac{dc}{dq} - \frac{d\gamma}{dq} \right) u_c + \frac{dh}{dq} u_h \right].$$

Rearranging, we obtain:

$$\begin{aligned} \frac{d\gamma}{dq} &= \frac{1}{E[u_c]} E \left[\frac{dc}{dq} u_c + \frac{dh}{dq} u_h \right] \\ &= \frac{1}{E[u_c]} E \left[\left(-\frac{\partial x}{\partial q} - p(q) \frac{dm}{dq} \right) u_c + \left(\frac{d\tilde{h}}{dm} \frac{dm}{dq} \right) u_h \right] \\ &= \frac{1}{E[u_c]} E \left[\left(-\frac{\partial x}{\partial q} \right) u_c + \underbrace{\left[-p(q) u_c + \frac{d\tilde{h}}{dm} u_h \right]}_{=0 \text{ by the FOC}} \frac{dm}{dq} \right] \\ \frac{d\gamma}{dq} &= E \left[\left(\frac{u_c}{E[u_c]} \right) \left(-\frac{\partial x}{\partial q} \right) \right]. \end{aligned}$$

where the second line follows from substituting $\frac{dc}{dq}$ and $\frac{dh}{dq}$ (equations (24) and (25)).⁵⁶ Using $-\frac{\partial x}{\partial q} = (p(0) - p(1))m(q; \theta)$, we obtain:

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

which is identical to the expression derived using the envelope theorem.

A.2 Instrumental variable analysis of the Oregon Health Insurance Experiment data

Our application uses the Oregon Medicaid lottery and data previously analyzed by Finkelstein et al. (2012)[25] and publicly available at www.nber.org/oregon. This section provides some additional information on how we analyze the data. Much more detail on the data and the lottery can be found in Finkelstein et al. (2012)[25].

A.2.1 Estimation of impacts

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:

$$y_i = \alpha_0 + \alpha_1 \text{Medicaid}_i + \epsilon_i, \quad (26)$$

where *Medicaid* is an indicator variable for whether the individual is covered by Medicaid at any point in the study period. We estimate equation (26) by two-stage least squares, using the following first-stage equation:

$$\text{Medicaid}_i = \beta_0 + \beta_1 \text{Lottery}_i + \nu_i, \quad (27)$$

in which the excluded instrument is the variable “*Lottery*” which is an indicator variable for whether the individual was selected by the lottery.

One particular feature of the lottery design affects our implementation. The lottery selected individuals, but if an individual was selected, any household member could apply for Medicaid. As a result, if more people from a household were on the waiting list, the household had more “lottery tickets” and a higher chance of being selected. The lottery was thus random conditional on the number of people in the household who were on the waiting list, which we refer to as the number of “lottery tickets.” In practice, about 60 percent of the individuals on the list were in households with one ticket, and virtually all the remainder had two tickets. (We drop the less than 0.5 percent who had three tickets; no one had more). In households with two tickets, the variable “*Lottery*”

⁵⁶Note that the first-order condition requires that the arguments of u_c and u_h be the choices that the individual makes facing $p(q)$ and paying $\gamma(q)$; in general, one would also subtract $\gamma(q)$ from their income and allow individuals to re-optimize; but as discussed above, we abstract from these income effect issues and instead motivate γ with a insurance lottery interpretation.

is one if any household member was selected by the lottery. In all of our analysis, therefore, we perform the estimation separately for one-ticket and two-ticket households. Because there is no natural or interesting distinction between these two sets of households, all estimates presented in the paper consist of the weighted average of the estimates for these two groups.

Much of our analysis is based on estimates of characteristics of treatment and/or control compliers – i.e., those who are covered by Medicaid if and only if they win the lottery (see, e.g., Angrist and Pischke (2009) [2]). Our estimation of these characteristics is standard. For example, uninsured individuals who won the lottery provide estimates of characteristics of never-takers. Since uninsured individuals who lost the lottery include both control compliers and never-takers, with estimates of the never-taker sample and the share of individuals who are compliers, we can back out the characteristics of control compliers. Likewise, insured individuals who lost the lottery provide estimates of characteristics of always takers. Since insured lottery winners include both treatment compliers and always-takers, we can in like manner identify the characteristics of treatment compliers. Differences between treatment and control compliers reflect the impact of Medicaid (i.e., α_1) in the IV estimation of equation (26).

To make this more concrete, let $f_g(x)$ denote the probability density function (pdf) x for group $g \in \{TC, CC, AT, NT\}$ where TC are the treatment compliers, CC are the control compliers, AT are the always-takers, and NT are the never-takers. We observe, $f_{NT}(x)$, the distribution of x for the never-takers, as the distribution of x for those who choose not to take up in the treatment group. The population fraction of never-takers, π_{NT} , is given by the fraction of the treatment group that did not take up the program. Similarly, $f_{AT}(x)$, the distribution of x for the always-takers, is given by the observed distribution of x for those who choose to take up in the control group, and the population fraction of always-takers, π_{AT} , is given by the fraction of the control group that took up the program.

The population fraction of compliers is given by: $\pi_C = 1 - \pi_{NT} - \pi_{AT}$. However, the distribution of x for compliers requires more work to calculate and differs for compliers in the control group and those in the treatment group. In the control group, those choosing not to take up are a mixture of never-takers and control compliers (those who would take up if offered). Using the observed distribution of x for never-takers (see above), we can back out $f_{CC}(x)$, the distribution of x for the compliers in the control group, by noting that the distribution of x for the those who don't take up the program in the control group is given by: $\frac{\pi_C}{\pi_C + \pi_{NT}} f_{CC}(x) + \frac{\pi_{NT}}{\pi_C + \pi_{NT}} f_{NT}(x)$. Similarly, those who take up the program in the treatment group are a mixture of always-takers and treatment compliers. Using the observed distribution of x for always-takers (see above), we can back out $f_{TC}(x)$, the distribution of x for the compliers in the treatment group, by noting that the distribution of x for those who take up the program in the treatment group is given by: $\frac{\pi_C}{\pi_C + \pi_{AT}} f_{TC}(x) + \frac{\pi_{AT}}{\pi_C + \pi_{AT}} f_{AT}(x)$. So, for example, one can solve for the treatment complier mean, μ_{TC} , using the equation $\frac{\pi_C}{\pi_C + \pi_{AT}} \mu_{TC} + \frac{\pi_{AT}}{\pi_C + \pi_{AT}} \mu_{AT} = \mu_{TT}$, where μ_{TT} is the observed mean of x of those in the treatment group who take up the program and μ_{AT} is the observed mean of those who take up the program in the control group. This yields:

$$\mu_{TC} = \frac{(\pi_C + \pi_{AT})\mu_{TT} - \pi_{AT}\mu_{AT}}{\pi_C}.$$

Similarly, the formula for control complier means is given by:

$$\mu_{CC} = \frac{(\pi_C + \pi_{NT})\mu_{CN} - \pi_{NT}\mu_{NT}}{\pi_C},$$

where μ_{CN} denotes the observed mean of x among those in the control group who do not take up the program and μ_{NT} denotes the observed mean of x among those in the treatment group who do not take up the program. These formulas were used to compute the the complier means presented in the text.

A.2.2 Estimation of impact on out-of-pocket spending distribution

To estimate the distribution of out-of-pocket spending for the treatment and control compliers in our relatively small sample, we follow a parametric IV technique. Fortunately, reported out-of-pocket spending closely follows a log-normal distribution combined with a mass at zero spending. Therefore, we approximate the distribution of out-of-pocket spending by assuming that out-of-pocket spending is a mixture of a mass point at zero and a log-normal spending distribution for strictly positive values. We allow the parameters of this mixture distribution to differ across four groups: treatment compliers (TC), control compliers (CC), always-takers (AT), and never-takers (NT). Specifically, let F_x^g denote the CDF of out-of-pocket spending for group g :

$$F_x^g(x|\phi^g, \mu^g, \nu^g) = \phi^g + (1 - \phi^g) \text{LOGN}(x|\mu^g, \nu^g) \quad \text{for } g \in \{TC, CC, AT, NT\}$$

where $\text{LOGN}(x|\mu, \nu)$ is the CDF of a log-normal distribution with mean and variance parameters, μ and ν , evaluated at $x > 0$. For $x = 0$, the CDF is given solely by ϕ^g , so that this parameter captures the fraction of group g with zero out-of-pocket spending. Under standard IV assumptions, the 12 parameters are identified from the joint distribution of out-of-pocket spending, insurance status, and lottery status. (In practice, we estimate F^g separately for households with 1 and 2 lottery tickets, and therefore estimate 12 parameters).⁵⁷ We estimate all parameters jointly using maximum likelihood using the approach laid out in subsection A.2.1. Except for the specification in column IX of Table 4, we set the distribution of out-of-pocket expenditures to zero for treatment compliers because Medicaid does not require copayments and charges zero or negligible premiums.

To the assess the goodness of fit, Figure A1 plots the estimated and actual CDF separately based on lottery status (won or lost), insurance status, and number of tickets. As can be seen from these figures, the parametric model fits quite well.

⁵⁷To ensure consistency with the consumption floor and to ensure that the relationship $x + c = \bar{y}$ always holds for the uninsured in our data, we redefine out-of-pocket spending x as $\min(x^f, \bar{y} - c_{floor})$, where x^f is fitted out-of-pocket spending in the model and c_{floor} denotes the consumption floor. For the insured, out-of-pocket spending is always consistent with the consumption floor because we have imposed $x(1, m) = 0$.

A.2.3 Results and comparison to previous results

Our sample, variable definitions, and estimation approach are slightly different from those in Finkelstein et al. (2012)[25]. Table A.1 walks through the differences in the approaches and shows that these differences are fairly inconsequential for the estimates reported in the two papers. Column I replicates the results from Finkelstein et al. (2012)[25]. In column II, we limit the data to the subsample used in our own analysis, which consists of about 15,500 individuals out of the approximately 24,000 individuals from Finkelstein et al. (2012)[25]. Our subsample excludes those who have missing values for any the variables we use in the analysis. The primary reason for the loss of sample size is missing information on prescription drug utilization (a component of medical spending m). Missing data on self-reported health, household income, number of family members, out-of-pocket spending, and other health care use also contribute slightly to the reduction of sample size. We also exclude the few people who had three people in the household signed up for the lottery, as described above.

Column III reports the results on our subsample using our estimating equations above. These estimating equations differ from those used by Finkelstein et al. (2012)[25] in several ways. First, we stratify on the number of tickets and report weighted averages of the results rather than include indicator variables for the number of tickets, as in Finkelstein et al. (2012) [25]; we thus allow the effects of insurance to potentially differ by number of tickets. Second, we do not control for which of the 8 different survey waves the data come from as in Finkelstein et al. (2012)[25].⁵⁸ And finally, we do not up-weight the subsample of individuals in the intensive-follow-up survey arm. As shown in column III, these deviations do not meaningfully affect the results.

Finally, column IV reports the results using our subsample and our estimating equation, adjusting the “raw” out-of-pocket data as described in Section 4. Specifically, we estimate the distribution of out-of-pocket spending by fitting the distribution described above, set out-of-pocket spending to zero for the insured, and impose a ceiling on out-of-pocket spending for the uninsured. Naturally, these adjustments only affect the estimated effect of Medicaid on out-of-pocket spending.⁵⁹

A.3 Decomposition of welfare effects in the complete-information approach

To provide insights into the drivers of the estimates of $\gamma(1)$ in equation (4), we decompose $\gamma(1) = \gamma_C + \gamma_M$, where γ_C denotes the welfare component associated with the effects of the program on consumption and γ_M the component due to changes in health. We can then obtain an implicit expression for $\gamma_C + \gamma_M$ that is in terms of consumption and health:

⁵⁸Covariates are more difficult to handle in our estimates of the distributional impact of Medicaid on out-of-pocket spending (and, hence, consumption), so we stratify by ticket size in the analyses of effects on distributions. We do the same thing for our mean estimates for consistency.

⁵⁹Starting from the raw data in column III, imposing the fitted distribution has very little effect (the estimated impact of Medicaid on out-of-pocket spending changes from -\$346 to -\$349). Imposing that the insured have zero out-of-pocket spending raises this estimate to -\$564, and from there the imposition of the cap on out-of-pocket spending of $\bar{y} - c_{floor}$ reduces the estimate to \$489, given our baseline consumption floor of \$1,000.

$$E[u(c(0; \theta), h(0; \theta))] = E[u(c(1; \theta) - \gamma_C - \gamma_M, h(1; \theta))]. \quad (28)$$

Given the additive separability of the utility function, we can estimate γ_C just based on the consumption term in the utility function:

$$E\left[\frac{c(0; \theta)^{1-\sigma}}{1-\sigma}\right] = E\left[\frac{(c(1; \theta) - \gamma_C)^{1-\sigma}}{1-\sigma}\right], \quad (29)$$

and calculate γ_M from our estimates of $\gamma(1)$ and γ_C using the identity:

$$\gamma_M = \gamma(1) - \gamma_C. \quad (30)$$

Of course, due to the curvature of the utility function, the order of operations can matter.

We can further decompose the welfare components associated with consumption effects (γ_C) and effects on health (γ_M) into a transfer and a pure-insurance component. We estimate the consumption transfer term ($\gamma_{C,Transfer}$) as the mean increase in consumption due to the program so that

$$\gamma_{C,Transfer} = E[c(1; \theta) - c(0; \theta)]. \quad (31)$$

The pure-insurance component operating through consumption ($\gamma_{C,Ins}$) is then:

$$\gamma_{C,Ins} = \gamma_C - \gamma_{C,Transfer}. \quad (32)$$

By substituting the health production function (equation (2)) into the definition of γ (equation (3)), we can similarly decompose the welfare components due to effects on health (γ_M) into a transfer component ($\gamma_{M,Transfer}$) and an insurance component ($\gamma_{M,Ins}$). We estimate the transfer component in health ($\gamma_{M,Transfer}$) by:

$$E\left[\frac{c(0; \theta)^{1-\sigma}}{1-\sigma} + \tilde{\phi}\tilde{h}(E[m(0; \theta)]; \theta)\right] = E\left[\frac{(c(1; \theta) - \gamma_C - \gamma_{M,Transfer})^{1-\sigma}}{1-\sigma} + \tilde{\phi}\tilde{h}(E[m(1; \theta)]; \theta)\right]$$

so that $\gamma_{M,Transfer}$ is the additional welfare benefit for the health improvements that would come with an average increase in medical spending due to the program. Approximating this health improvement by $E\left[\frac{d\tilde{h}}{dm}\right] E[m(1; \theta) - m(0; \theta)]$, we implement this calculation of $\gamma_{M,Transfer}$ as the solution to:

$$\begin{aligned} E\left[\frac{c(0; \theta)^{1-\sigma} - (c(1; \theta) - \gamma_C - \gamma_{M,Transfer})^{1-\sigma}}{1-\sigma}\right] &= \tilde{\phi}E\left[\tilde{h}(E[m(1; \theta)]; \theta) - \tilde{h}(E[m(0; \theta)]; \theta)\right] \\ &= \tilde{\phi}E\left[\frac{d\tilde{h}}{dm}\right] E[m(1; \theta) - m(0; \theta)]. \end{aligned} \quad (33)$$

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 above, and then take its expectation here. Finally, the pure-insurance component operating through health ($\gamma_{M,Ins}$) is given by:

$$\gamma_{M,Ins} = \gamma_M - \gamma_{M,Transfer}. \quad (34)$$

A.4 Consumption measure from the CEX

Data and sample The CEX consists of a series of short panels. Each “consumer unit” (CU) is interviewed every 3 months over 5 calendar quarters. In the initial interview, information is collected on demographic and family characteristics and on the consumer unit’s inventory of major durable goods. Expenditure information is collected in the second through the fifth interviews using uniform questionnaires. Income and employment information is collected in the second and fifth interviews.

Our sample includes all CUs in 1996-2010 who have valid expenditure data in all 4 quarters (i.e., positive total expenditure and non-negative medical expenditure) and non-missing income data. To be broadly consistent with the Oregon sample, we further limit the analysis to adults aged 19-64 who are below 100% of the federal poverty line. We measure insurance status q at the start of the survey, regardless of whether or not the individual obtains insurance later in the year (results are quite similar if we use concurrent insurance status). 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 and uninsured. We convert all dollar amounts to 2009-dollars, and impose an annual consumption floor (although in practice the baseline consumption floor of \$1,000 never binds).

Measuring the consumption covariance We observe three variables: reported consumption, \hat{c} , reported out-of-pocket medical spending, \hat{x} , and an indicator for insurance status, $i \in \{0, 1\}$. We wish to estimate

$$Cov \left(\frac{c^{-\sigma}}{E[c^{-\sigma}]}, x | i = 0 \right),$$

where c and x are actual consumption and out-of-pocket medical spending for those without formal insurance, $i = 0$. Table A.3 presents the results for our baseline specification and for alternative definitions of consumption, c . For both the insured and uninsured, we compute $c^{-\sigma}$ using $\sigma = 3$ and then compute the covariance between $c^{-\sigma}$ and out-of-pocket medical spending, x , normalized by the mean value of $c^{-\sigma}$, $E[c^{-\sigma}]$.

As Table A.3 illustrates, these covariances for the uninsured are negative across all of our specifications. 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). However, as can be seen in Table A.3, the covariance is more negative for the insured. We infer from this that the basic problem is that self-reported consumption and health spending may not equal the actual consumption and health spending. We outline a model that allows us to correct for measurement error.

The core idea behind our particular measurement-error model is that individuals may misreport their out-of-pocket medical spending. If our model is correctly specified, the covariance between out-of-pocket medical spending and the marginal utility of consumption should be zero for the insured. 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.

More formally, we observe non-medical consumption, \hat{c} , and out-of-pocket medical spending, \hat{x} , and wish to infer the covariance between the marginal utility of consumption (normalized by its average), $\frac{c^{-\sigma}}{E[c^{-\sigma}]}$, and true out-of-pocket medical spending, x . Here, our primary concern is mis-measurement of out-of-pocket spending, x . Therefore, we opt to allow for an arbitrary functional form on the shape of the distribution of this measurement error. In particular, we assume

$$\hat{x} = x + \epsilon$$

where ϵ is a measurement-error shock 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 consumption is measured without error, $\hat{c} = c$, which implies that the marginal utility of consumption is also 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. Third, we assume true out-of-pocket medical spending is zero for the insured, so that $\hat{x} = \epsilon$ for the insured. These assumptions would be satisfied if ϵ reflected consumption of uncovered healthcare for both the insured and uninsured (e.g., over-the-counter pain killers) and these are consumed in equal amounts by both groups.

Under these assumptions, the observed covariance between $\frac{\hat{c}^{-\sigma}}{E[\hat{c}^{-\sigma}]}$ and \hat{x} for the insured provides an estimate of the bias induced by measurement error when estimating this covariance for the uninsured. The observed covariance between $\frac{\hat{c}^{-\sigma}}{E[\hat{c}^{-\sigma}]}$ and \hat{x} is the sum of the true covariance and the measurement-error component:

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

Under our three assumptions, we can identify the measurement-error component for the insured:

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

Hence, the true covariance term for the uninsured is given by the difference between the observed

covariance for the uninsured and the insured:

$$Cov\left(\frac{c^{-\sigma}}{E[c^{-\sigma}]}, x|i=0\right) = Cov\left(\frac{\hat{c}^{-\sigma}}{E[\hat{c}^{-\sigma}]}, \hat{x}|i=0\right) - Cov\left(\frac{\hat{c}^{-\sigma}}{E[\hat{c}^{-\sigma}]}, \hat{x}|i=1\right). \quad (35)$$

Of course, this is one particular model of measurement error, and the true measurement error could be of a different form. But, our approach has the advantage of allowing for an arbitrary shape to the unknown distribution of measurement error in out-of-pocket spending, ϵ , and leads to an intuitive estimation strategy of using the estimated covariance term for the insured (which should be zero) to provide information about the true covariance term.

Table A.3 shows the results. Taking the difference between the covariance estimates for the insured and uninsured, as illustrated in equation (35), yields a covariance value of \$252 in the baseline specification. Dividing by 2 to form the linear approximation to the average covariance value over $q = 1$ to $q = 0$, we have a pure-insurance value of \$126 for the consumption-based optimization approach using the CEX data, as illustrated in Table 2. This estimate is largely similar if one chooses alternative measures of consumption, such food, education, reading, entertainment, and personal care (Column II), and all non-health consumption, excluding alcohol and tobacco (Column III).

A.5 Health production function, $E_{\theta|\theta^K} \left[\frac{\partial \bar{h}}{\partial m} \right]$

To implement the health-based optimization approach, we must estimate the health returns to medical spending conditional on medical spending, m . To do so, we use the Medicaid lottery as an instrument for medical spending. To capture heterogeneity, we assume differences in m can be captured by differences in state variables, θ^K , that consist of measures of financial and health states from an initial survey (fielded essentially concurrently with the lottery). We construct a binary “financial constraint” variable that takes the value of 1 if the individual responded affirmatively to any of these questions: (i) whether or not the individual had to forgo medical treatment because of financial conditions (ii) whether or not the individual had to forgo prescription drugs because of financial conditions, and (iii) whether or not the individual was refused medical treatment due to inability to pay. Approximately 36% report having a financial constraint in this initial survey. We construct a binary health state variable that takes the value 1 if the individual was previously diagnosed with diabetes, asthma, high blood pressure, emphysema, congestive heart failure, or depression. Approximately 45% report having a major health diagnosis in this initial survey. The state variables θ^K consist of four dummy variables, each of which corresponds to one of the four values that the interaction of the financial state variable and the health state variable can take on.⁶⁰ For each value of the state variables, we estimate the expected return to medical spending using the lottery as an instrument for medical spending.

Table A.2 reports the estimated IV results of the effect of medical spending on the health

⁶⁰In principle, one could use more than four state variables; however, our estimates are already fairly imprecise with only four state variables and additional variables would further increase the already considerable noise in the estimates.

indicator. Consistent with the hypothesis that the value of insurance is higher to those who are more constrained, the IV estimates of the impact of medical spending on health are largest for those with financial constraints (columns III and IV). However, all of our estimates are very imprecise and none are statistically different from zero. Moreover, one should bear in mind that our measure of health is self-reported, and our measures of the state variables are quite coarse.

A.6 Construction of $-\partial x/\partial q$ when Medicaid recipients have positive out-of-pocket expenditures

When at least some Medicaid recipients have strictly positive out-of-pocket spending, the expression for the relaxation of the budget constraint at $q = 1$ becomes:

$$-\frac{\partial x}{\partial q}|_{q=1} = p(0)m(1; \theta) - p(1)m(1; \theta).$$

The second term, $p(1)m(1; \theta)$, is the distribution of out-of-pocket spending of the insured, which is given by the distribution of out-of-pocket spending by treatment compliers. The first term, $p(0)m(1; \theta)$, is the distribution of out-of-pocket spending that the uninsured would have had if they had incurred the medical spending of the insured. We rewrite the expression for the relaxation of the budget constraint at $q = 1$ as:

$$-\frac{\partial x}{\partial q}|_{q=1} = p(0)m(0; \theta) - p(1)m(1; \theta) + p(0)(m(1; \theta) - m(0; \theta)).$$

We evaluate this expression by taking the difference in the distributions of out-of-pocket expenditures of control compliers ($p(0)m(0; \theta)$) and treatment compliers ($p(1)m(1; \theta)$) and add to this the price faced by the uninsured times the difference in the distributions of medical spending of treatment compliers minus control compliers ($p(0)(m(1; \theta) - m(0; \theta))$). The price faced by the uninsured is calculated as the ratio of mean out-of-pocket spending to mean total spending for the control compliers. In the construction of differences in distributions, we assume quantile stability. In other words, we take the difference in distributions assuming an individual with a given θ that puts him at quantile r in the control complier distribution would have been at quantile r in the treatment complier distribution if he had been in the treatment group.

The expression for the relaxation of the budget constraint at $q = 0$ is derived analogously:

$$-\frac{\partial x}{\partial q}|_{q=0} = p(0)m(0; \theta) - p(1)m(0; \theta).$$

The first term, $p(0)m(0; \theta)$, is the distribution of out-of-pocket spending of the uninsured, which is given by the distribution of out-of-pocket spending by control compliers. The second term, $p(1)m(0; \theta)$, is the distribution of out-of-pocket spending that the insured would have had if they had incurred the medical spending of the uninsured. We rewrite the expression for the relaxation of the budget constraint at $q = 0$ as:

$$-\frac{\partial x}{\partial q}|_{q=0} = p(0)m(0; \theta) - p(1)m(1; \theta) + p(1)(m(1; \theta) - m(0; \theta)).$$

We evaluate this expression by taking the difference in the distributions of out-of-pocket expenditures of control compliers ($p(0)m(0; \theta)$) and treatment compliers ($p(1)m(0; \theta)$) and add to this the price faced by the insured times the difference in the distributions of medical spending of treatment compliers minus control compliers ($p(1)(m(1; \theta) - m(0; \theta))$). The price faced by the insured is calculated as the ratio of mean out-of-pocket spending to mean total spending for the treatment compliers.

A.7 Relaxation of the linear interpolation assumption for $d\gamma/dq$

Linear demand for medical care Given our definition of $p(q) \equiv qp(1) + (1 - q)p(0)$, the assumption that the demand for medical care, m , is linear in price implies that the demand is also linear in q . Because the empirical distribution of medical care is measured imprecisely, we infer the distribution of $m(0; \theta)$ by the distribution of out-of-pocket expenditure divided by the price that uninsured individuals pay for medical care, $x(0; \theta)/p(0)$, where $x(0; \theta)$ denotes the empirical distribution of out-of-pocket spending among the uninsured. We infer the distribution of medical care for the insured from the distribution of medical care for the uninsured by assuming that each point in the distribution scales up proportionally to the overall increase in medical care due Medicaid coverage, $E[m(1; \theta)]/E[m(0; \theta)]$. Thus, the distribution of medical care for the insured is given by: $\frac{E[m(1; \theta)]}{E[m(0; \theta)]}x(0; \theta)/p(0)$. Using the assumption that the demand for medical care is linear in q , we have:

$$m(q; \theta) = q \frac{E(m(1; \theta))}{E(m(0; \theta))} x(0; \theta)/p(0) + (1 - q)x(0; \theta)/p(0). \quad (36)$$

The distribution of out-of-pocket spending for each value of q is given by:

$$x(q; \theta) = p(q)m(q; \theta) = (1 - q)p(0)m(q; \theta),$$

where the latter equality follows from the fact that Medicaid recipients face a zero price of medical care, i.e., $p(1) = 0$. Substituting the expression for $m(q; \theta)$ into this equation yields the expression for out-of-pocket spending that we use in our implementation:

$$x(q; \theta) = (1 - q)x(0; \theta) \left(q \frac{E(m(1; \theta))}{E(m(0; \theta))} + (1 - q) \right).$$

We use equation (18) to infer the distribution of consumption from the distribution of out-of-pocket spending. From the distribution of consumption, we calculate the distribution of the marginal utility of consumption using $u_c = (c(q; \theta) - \gamma(q))^{-\sigma}$. We calculate the distribution of the marginal relaxation of the budget constraint, $-\partial x/\partial q = (p(1) - p(0))m(q; \theta)$, for each value of q by substituting in the expression for the demand of medical care (equation (36)) and noting that

$p(1) = 0$. This yields:

$$-\frac{\partial x}{\partial q} = x(0; \theta) \left(q \frac{E(m(1; \theta))}{E(m(0; \theta))} + (1 - q) \right)$$

We then use the distributions of consumption and the marginal relaxation of the budget constraint to calculate $d\gamma/dq$ at each value of q :

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

and solve this differential equation using Picard's method to obtain $\gamma(1)$. Column (X) in Table 4 presents the results.

Upper bound for $\gamma(1)$ for arbitrary functional form of the demand for medical care

Rather than assuming that demand for medical care is linear in price, we now allow any functional form for the demand for medical care and find the functional form that maximizes $\gamma(1)$. We allow for arbitrary (non-parametric) functional forms for the demand for medical care with the restriction that demand at values of $q \in (0, 1)$ must lie somewhere between demand at $q = 0$ and at $q = 1$. Specifically, we define the distribution of medical care at insurance level q to be some linear combination of the distribution of medical care at $q = 0$ and at $q = 1$, where these distributions are given by (36). Formally, the distribution of medical care at insurance level q is given by $\hat{m}(\lambda(q); \theta) = \lambda m(0; \theta) + (1 - \lambda)m(1; \theta)$ for some $\lambda(q) \in [0, 1]$.

The distribution of out-of-pocket spending for each value of q and λ is given by $p(q)\hat{m}(\lambda(q); \theta) = (1 - q)p(0)\hat{m}(\lambda(q); \theta)$. We use equation (18) to infer the distribution of consumption from the distribution of out-of-pocket spending; we denote the resulting consumption level by $\hat{c}(\lambda(q); \theta)$. From the distribution of consumption, we calculate the distribution of the marginal utility of consumption using $u_c = (\hat{c}(\lambda(q); \theta) - \gamma(q))^{-\sigma}$. We calculate the distribution of the marginal relaxation of the budget constraint as $-\partial x/\partial q = p(0)\hat{m}(\lambda(q), \theta)$.

We search for the value of $\lambda(q) \in [0, 1]$ that maximizes $d\gamma/dq$ at each value of q :

$$\frac{d\gamma}{dq}(q) = \max_{\lambda(q)} E \left[\frac{u_c}{E[u_c]} \left(-\frac{\partial x}{\partial q} \right) \right] = \max_{\lambda(q)} E \left[\frac{(\hat{c}(\lambda(q); \theta) - \gamma(q))^{-\sigma}}{E[(\hat{c}(\lambda(q); \theta) - \gamma(q))^{-\sigma}]} (p(0)\hat{m}(\lambda(q), \theta)) \right].$$

We solve this differential equation using Picard's method to find the upper bound for $\gamma(1)$. Column (XI) in Table 4 presents the results.

Figure A1: Fitted and actual CDFs of out-of-pocket spending

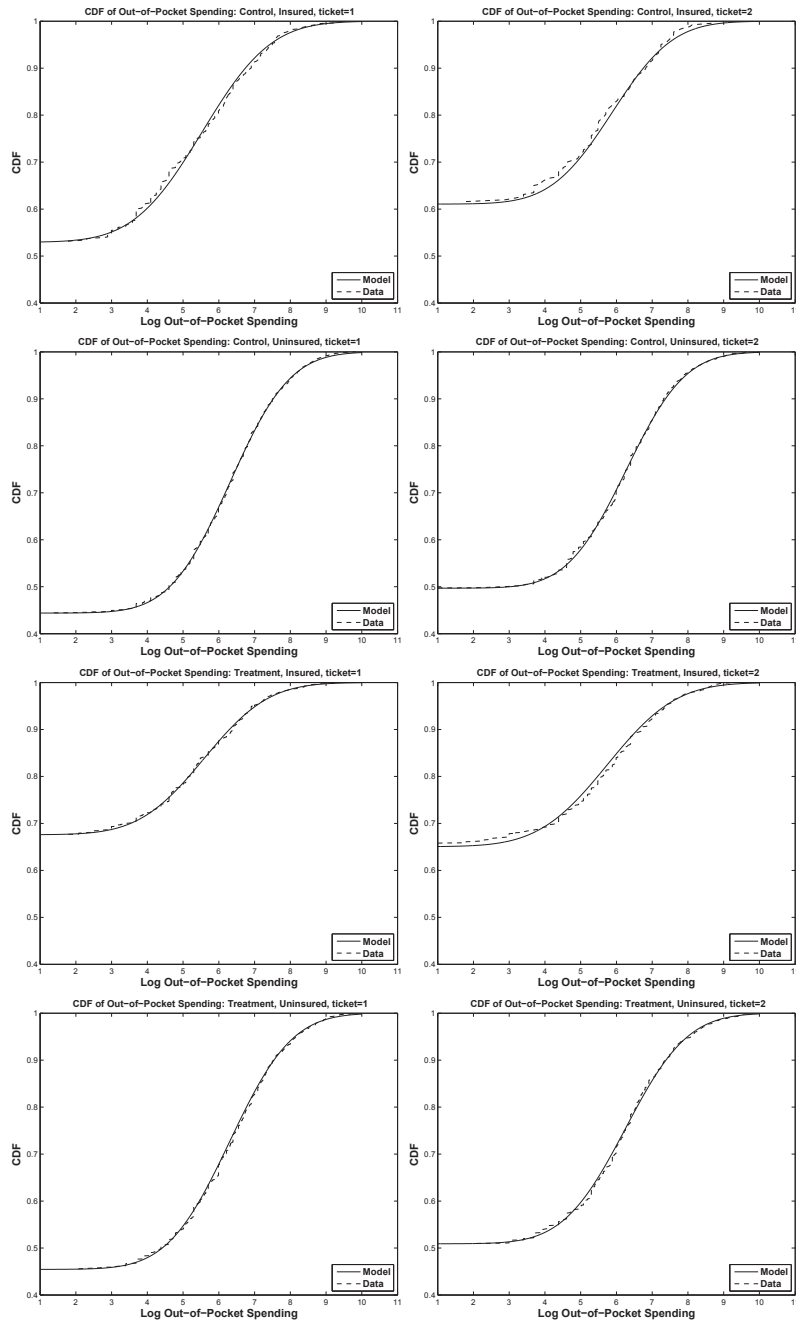


Table A.1: Comparison with Prior Estimates from Finkelstein et al. (QJE, 2012)

	I	II	III	IV
Sample	QJE sample	Restricted sample	Restricted sample	Restricted sample
Data	Raw data	Raw data	Raw data	Adjusted data
Estimation Method	QJE paper	QJE paper	This paper	This paper
<i>First Stage: Lottery impact on Insurance</i>				
Lottery Indicator (s.e.)	0.290 (0.007)	0.290 (0.008)	0.302 (0.007)	0.302 (0.007)
<i>IV: Impact of Medicaid on...</i>				
12-month medical spending (s.e.)	778 (371)	903 (434)	885 (366)	885 (366)
12-month out-of-pocket spending (s.e.)	-244 (86)	-364 (104)	-349 (79)	-489 (38)
self-reported health (s.e.)	0.133 (0.026)	0.103 (0.032)	0.142 (0.028)	0.142 (0.028)
did not screen positive for depression (s.e.)	0.078 (0.025)	0.060 (0.030)	0.059 (0.026)	0.059 (0.026)
N	23,741	15,498	15,498	15,498

Notes: This table compares our baseline estimates of the impact of Medicaid with the baseline estimates of Finkelstein et al. (2012), which we refer to as "QJE." Self-reported health is a dummy variable that equals 1 if the individual reports being in good, very good, or excellent health. Column I replicates the QJE results. In column II, we use the same regressions as in column I but restrict the QJE sample to respondents living in households that have at most 2 lottery tickets and that have non-missing data on all the required variables (see Appendix A.2 for more details). In column III, we use the same sample as in column II but apply the regression approach of this paper (again see Appendix A.2 for more details). In column IV, we use the estimation method and sample from this paper, applied to the "adjusted data" for out-of-pocket spending. "Adjusted data" refers to the out-of-pocket spending data after (i) estimating it by fitting a lognormal distribution with a mass point at zero for the distribution of out-of-pocket spending, (ii) adjusting the out-of-pocket spending of the insured to be 0, and (iii) imposing a ceiling on out-of-pocket spending of (mean(income per capita) - consumption floor) for the uninsured; see text for more details. Column IV represents the data and approach used in this paper. All amounts are in dollars per Medicaid recipient per year.

Table A.2: Health Production Function Estimates

	I	II	III	IV
Variable:	No major health diagnosis, no financial constraint	Major health diagnosis, no financial constraint	No major health diagnosis, financial constraint	Major health diagnosis, financial constraint
<i>First Stage: Lottery impact on Medical Spending</i>				
Lottery indicator	64	268	364	566
(s.e.)	(147)	(236)	(286)	(240)
<i>Reduced Form: Lottery Impact on Health</i>				
Lottery indicator	0.032	0.006	0.040	0.058
(s.e.)	(0.011)	(0.018)	(0.025)	(0.015)
<i>IV: Impact of Medical Spending on Health</i>				
\$1000 in annual medical spending	-0.097	0.031	0.129	0.124
(s.e.)	(0.134)	(0.120)	(0.141)	(0.088)
Sample size (N)	6993	2886	1457	4162
Mean out-of-pocket spending	352	417	329	541

Notes: Columns show results for four different subsamples, as defined in the text. As we do throughout the paper, all models are separately estimated for households with one ticket and households with two tickets, and we report the weighted average of these two estimates in our tables. Because of this procedure, the IV estimate reported in the table is not exactly equal to 1000 times the reduced-form estimate divided by the first-stage estimate (even the signs of the estimates don't need to match). Medical spending and out-of-pocket spending are measured in dollars per year per Medicaid recipient. Health is a dummy variable that equals 1 if the individual reports being in good, very good, or excellent health.

Table A.3: Measurement of Consumption Covariance in CEX Consumption Approach

	Baseline	II	III
Consumption covariance			
Insured	-318	-411	-345
Uninsured	-66	-71	-53
Difference (= Measurement-error corrected covariance)	252	340	293
Definition of non-health consumption	All non-health Consumption	Food, education, reading, entertainment, personal care	All non-health excluding alcohol and tobacco
Mean of non-health consumption (in annual \$ per capita)	13,310	11,130	12,559

Notes: This table presents 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 includes all single, childless consumption units in 1996-2010 who are complete income reporters and have valid expenditure data in all 4 quarters (i.e., strictly positive total expenditure and weakly positive out-of-pocket medical spending), and whose household head is between 19 and 65 years old in all 4 quarters, and with incomes not exceeding 100% of the Federal Poverty Line (N=1056). 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. To be consistent with the Oregon data, we impose a \$1000 per capita annual consumption floor and convert all dollar amounts to 2009 dollars.