

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

Amy Finkelstein

Massachusetts Institute of Technology

Nathaniel Hendren

Harvard University

Erzo F. P. Luttmer

Dartmouth College

We develop frameworks for welfare analysis of Medicaid and apply them to the Oregon Health Insurance Experiment. Across different approaches, we estimate low-income uninsured adults' willingness to pay for Medicaid between \$0.5 and \$1.2 per dollar of the resource cost of providing Medicaid; estimates of the expected transfer Medicaid provides to recipients are relatively stable across approaches, but estimates of its additional value from risk protection are more variable. We also estimate that the resource cost of providing Medicaid to an additional recipient is only 40 percent of Medicaid's total cost; 60 percent of Medicaid spending is a transfer to providers of uncompensated care for the low-income uninsured.

We are grateful to Lizi Chen for outstanding research assistance and to Isaiah Andrews, David Cutler, Liran Einav, Matthew Gentzkow, Jonathan Gruber, Conrad Miller, Jesse Shapiro, Matthew Notowidigdo, Ivan Werning, three anonymous referees, Michael Greenstone (the editor), and seminar participants at Brown, Chicago Booth, Cornell, Harvard Medical School, Michigan State, Pompeu Fabra, Simon Fraser University, Stanford, the University of California Los Angeles, the University of California San Diego, the University of Houston, and the University of Minnesota for helpful comments. We gratefully acknowledge financial support from the National Institute of Aging under grants RC2AGO36631 and R01AG0345151 (Finkelstein) and the National Bureau of Economic Research Health

Electronically published November 5, 2019

[*Journal of Political Economy*, 2019, vol. 127, no. 6]

© 2019 by The University of Chicago. All rights reserved. 0022-3808/2019/12706-0006\$10.00

I. Introduction

Medicaid is the largest means-tested program in the United States. In 2015, public expenditures on Medicaid were over \$550 billion, compared to about \$70 billion for food stamps, \$70 billion for the Earned Income Tax Credit (EITC), \$55 billion for Supplemental Security Income, and \$30 billion for cash welfare.¹ How much would recipients be willing to pay for Medicaid, and how does this compare to Medicaid's costs? And how much of Medicaid's costs reflects a monetary transfer to nonrecipients who bear some of the costs of covering the low-income uninsured?

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, and risk exposure. But there has been little formal attempt to translate such estimates into statements about welfare. Absent other guidance, academic or public policy analyses often either ignore the value of Medicaid—for example, in the calculation of the poverty line or measurement of income inequality (Gottschalk and Smeeding 1997)—or make ad hoc assumptions. For example, the Congressional Budget Office (2012) values Medicaid at the average government expenditure per recipient. In practice, an in-kind benefit like Medicaid may be valued above or below its costs (see, e.g., Currie and Gahvari 2008).

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 from the first two years were (1) Medicaid increased health care use across the board—including outpatient care, preventive care, prescription drugs, hospital admissions, and emergency room visits; (2) Medicaid improved self-reported health and reduced depression but had no statistically significant impact on mortality or physical health measures; (3) Medicaid reduced the risk of large out-of-pocket medical expenditures; and (4) Medicaid had no economically or statistically significant impact on employment and earnings or on private health insurance coverage (Finkelstein et al. 2012, 2016; Baicker et al. 2013, 2014; Taubman et al. 2014). 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.²

and Aging Fellowship, under National Institute of Aging grant T32-AG000186 (Hendren). Data are provided as supplementary material online.

¹ See US Department of Health and Human Services (2015, 2016), US Department of Agriculture (2016), US Internal Revenue Service (2016), and US Social Security Administration (2016).

² The results of the Oregon Health Insurance Experiment received extensive media coverage, but the media drew a wide variety of conclusions, as the following headlines

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

We develop frameworks for empirically estimating the value of Medicaid to recipients in terms of the amount of current, nonmedical consumption they would need to give up to be indifferent between receiving and not receiving Medicaid; we refer to this as the recipient's "willingness to pay" for Medicaid. We focus on this normative measure because it is well defined even if individuals are not optimizing when making health care decisions. This allows us to incorporate various frictions—such as information frictions or behavioral biases—that could alter the individual's value of Medicaid relative to what a compensating variation approach would imply. Our approach, however, speaks directly to only the recipient's willingness to pay for Medicaid. An estimate of society's willingness to pay for Medicaid needs to take account of the social value of any redistribution that occurs through Medicaid, and, as is well known, such redistribution generally involves net resource costs that exceed the recipient's willingness to pay (Okun 1975).

We develop two main analytical frameworks for estimating recipient willingness to pay for Medicaid. Our first approach, which we refer to as the "complete-information" approach, requires a complete specification of a normative utility function and estimates of the causal effect of Medicaid on the distribution of all arguments of the utility function. The 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.

illustrate: "Medicaid Makes 'Big Difference' in Lives, Study Finds" (National Public Radio, July 11, 2011; www.npr.org/2011/07/07/137658189/medicaid-makes-big-difference-in-lives-study-finds) vs. "Spending on Medicaid Doesn't Actually Help the Poor" (*Washington Post*, May 2, 2013; 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, May 4, 2013; <http://www.ocpp.org/2013/05/04/blog20130504-oregon-lesson-nation-medicaid-works/>) vs. "Oregon Medicaid Study Shows Michigan Medicaid Expansion Not Worth the Cost" (Mackinac Center for Public Policy, May 3, 2013; <http://www.mackinac.org/18605>).

Our second approach, which we refer to as the “optimization” approach, is in the spirit of the “sufficient-statistics” approach described by Chetty (2009) and is the mirror image of the complete-information approach in terms of its strengths and weaknesses. We reduce the implementation requirements by parameterizing the way in which Medicaid affects the individual’s budget set and by assuming that individuals have the ability and information to make privately optimal choices with respect to that budget set. With these assumptions, it suffices to specify the marginal-utility function over any single argument. This is because the optimizing individual’s first-order condition allows us to value marginal impacts of Medicaid on any other potential arguments of the utility function through the marginal utility of that single argument. To make inferences about nonmarginal 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 substitutes for the structural assumptions about the utility function in the complete-information approach.

We implement these approaches for the Medicaid coverage provided by the Oregon Health Insurance Experiment. We use data from study participants to directly measure out-of-pocket medical spending, health care utilization, and health. The lottery’s random selection allows for causal estimates of the impact of Medicaid on the various outcomes. Our baseline health measure is a mapping of self-assessed health into quality-adjusted life years (QALYs) based on existing estimates of QALYs associated with different levels of self-assessed health; we also report estimates based on alternative health measures—such as self-reported physical and mental health, or a depression screen—combined with existing estimates of their associated QALYs. Absent a consumption survey in the Oregon context, we proxy for consumption by the difference between average consumption for a low-income uninsured population and out-of-pocket medical expenditures reported by study participants, subject to a consumption floor; we also report results that instead use consumption data for a low-income sample in the Consumer Expenditure Survey (CEX).

Our results reveal that Medicaid is best conceived of as consisting of two separate parts: a monetary transfer to external parties who would otherwise subsidize the medical care for the low-income uninsured and a subsidized insurance product for recipients. The experimental treatment effects of Medicaid on out-of-pocket spending and total medical spending imply that 60 percent of Medicaid’s gross expenditures—which we estimate to be \$3,600 per recipient—are a transfer to these external parties, leaving the net cost of Medicaid at about \$1,450 per recipient. Recipient willingness to pay for Medicaid could exceed this net cost because of the pure-insurance value it provides (reallocation toward states of the world

with high marginal utility) or could be less than its net cost because of recipients' moral-hazard response (induced medical spending valued below cost). Our different approaches reach different conclusions: willingness to pay for Medicaid by recipients per dollar of net cost ranges between \$0.5 and \$1.2; all approaches suggest that recipient willingness to pay for Medicaid is substantially below its gross cost (the value of Medicaid assumed by the Congressional Budget Office [2012]). For the approaches that provide point estimates of the sources of Medicaid's value to recipients, we estimate that between half and four-fifths of Medicaid's value to recipients comes from the increase in expected resources it provides rather than from its (budget-neutral) insurance value. Naturally, our estimates are specific to this particular Medicaid program for low-income adults and to the people for whom the lottery affected Medicaid coverage. Yet 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.

Our analysis complements other efforts to elicit a value of Medicaid to recipients through quasi-experimental variation in premiums (e.g., Dague 2014) or the extent to which individuals distort their labor earnings in order to become eligible for Medicaid (Keane and Moffitt 1998; Gallen 2015).³ These alternative approaches require their own, different sets of assumptions. Consistent with our results, these approaches also tend to indicate that Medicaid recipients place a low value on the program, relative to the government's gross cost of providing Medicaid. However, they do not generally estimate the monetary transfers to external parties or compare recipient value to net (of these monetary transfers) costs. Our finding that a large part of Medicaid spending represents a transfer to external parties complements related empirical work documenting the presence of implicit insurance for the uninsured (Mahoney 2015) and the role of formal insurance coverage in reducing the provision of uncompensated care by hospitals (Garthwaite, Gross, and Notowidigdo 2018) and unpaid medical bills by patients (Dobkin et al. 2018). Given the size of these external monetary transfers relative to Medicaid's value to recipients, our findings suggest that identifying the ultimate economic incidence of uncompensated care and assessing the relative efficiency of formal public insurance, versus an informal insurance system of uncompensated care, are important areas for further work.

The rest of the paper proceeds as follows. Section II develops the two theoretical frameworks for welfare analysis. Section III describes how we implement these frameworks for welfare analysis of the impact of the

³ Finkelstein, Hendren, and Shepard (2019) use variation in premiums for health insurance in Massachusetts to study the value of subsidized health insurance for low-income adults above the Medicaid eligibility threshold.

Medicaid expansion that occurred via lottery in Oregon. Section IV presents the results and discusses their interpretations and the trade-offs in our context across the alternative approaches. The last section presents our conclusions.

II. Frameworks for Welfare Analysis

A. A Simple Model of Individual Utility

Individuals derive utility from the consumption of nonmedical goods and services, c , and from health, h , according to

$$u = u(c, h). \quad (1)$$

We assume that health is produced according to

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

where m denotes the consumption of medical care and θ is an underlying state variable for the individual that includes medical conditions, other factors affecting health, and the productivity of medical spending. This framework is similar to that of Cardon and Hendel (2001). 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 refer to m as “medical spending” and to c as “consumption.”

We assume that every Medicaid recipient faces the same distribution of θ . Conceptually, our welfare analysis can be thought of as conducted from behind the veil of ignorance (conditional on being a low-income adult). Empirically, we use the distribution of outcomes across individuals to measure the distribution of 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 on both Medicaid status, q , and the underlying state of the world, θ ; this dependence is denoted by $c(q; \theta)$, $m(q; \theta)$, and $h(q; \theta) \equiv \tilde{h}(m(q; \theta); \theta)$, respectively.⁴

We define $\gamma(1)$ as the 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:

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

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

where the expectations are taken with respect to the possible states of the world, θ . With some abuse of terminology, we refer to $\gamma(1)$ as the recipient's willingness to pay for Medicaid, even though it is measured in terms of forgone consumption rather than forgone income.

Importantly, $\gamma(1)$ is measured from the perspective of the individual recipient. A social welfare perspective would also account for the fact that Medicaid benefits a low-income group. Saez and Stantcheva (2016) show that, in general, this can be accomplished by scaling the individual valuation by a social marginal welfare weight, or the social marginal utility of income. For example, suppose that the social marginal utility of income of Medicaid beneficiaries is 20 times as high as the social marginal utility of income of the average person in the population. Then, society is willing to pay \$20 to deliver \$1 to a Medicaid beneficiary. To move from our estimates of the recipient's willingness to pay for Medicaid to society's willingness to pay, one would therefore scale our estimates of $\gamma(1)$ by 20; we return to this point in Section IV.C.

B. Complete-Information Approach

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

ASSUMPTION 1 (Full utility specification for the complete-information approach). The utility function takes the form

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

where σ denotes the coefficient of relative risk aversion and $\tilde{\phi}$ denotes the marginal utility of health. Scaling $\tilde{\phi}$ by the expected marginal utility of consumption yields the expected marginal rate of substitution (MRS) of health for consumption, $\phi (= \tilde{\phi}/E[c^{-\sigma}])$.

Utility has two additive components: a standard CRRA (constant relative risk aversion) function in consumption c with a coefficient of relative risk aversion of σ and a linear term in h . The assumption that utility is linear in health is consistent with our measure of health (QALYs, introduced below), which by construction is linear in utility. The assumption that consumption and health are additive is commonly made in the health literature, but it restricts the marginal utility of consumption to be independent of health. This assumption simplifies the implementation of our estimates (though our framework could in principle be applied with non-additive functions).

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

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

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

We can decompose $\gamma(1)$ into two economically distinct components: the increases in average resources for the individual and the (budget-neutral) reallocation of resources across states of the world. We refer to these as, respectively, the “transfer component” (T) and the “pure-insurance component” (I). The transfer component is given by the solution to the equation

$$\begin{aligned} \frac{E[c(0; \theta)]^{1-\sigma}}{1-\sigma} + \tilde{\phi}E[\tilde{h}(E[m(0; \theta)]; \theta)] = \\ \frac{(E[c(1; \theta)] - T)^{1-\sigma}}{1-\sigma} + \tilde{\phi}E[\tilde{h}(E[m(1; \theta)]; \theta)]. \end{aligned} \quad (5)$$

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

$$\frac{E[c(0; \theta)]^{1-\sigma}}{1-\sigma} - \frac{(E[c(1; \theta)] - T)^{1-\sigma}}{1-\sigma} = \tilde{\phi}E\left[\frac{d\tilde{h}}{dm}\right]E[m(1; \theta) - m(0; \theta)]. \quad (6)$$

Medicaid spending that increases consumption (c) increases T dollar for dollar; however, increases in medical spending (m) due to Medicaid may increase T by more or less than a dollar, depending on the health returns to medical spending as described by the health production function, $\tilde{h}(m; \theta)$.⁵ Relatedly, evaluating this equation requires an estimate of $E[d\tilde{h}/dm]$, the slope of the health production function between $m(1; \theta)$ and $m(0; \theta)$, averaged over all states of the world.

The pure-insurance term (I) is given by

$$I = \gamma(1) - T. \quad (7)$$

⁵ By the standard logic of moral hazard, if consumers optimally choose m , they would value the increase in health arising from the Medicaid-induced medical spending at less than its cost, since they chose not to purchase that medical spending at an unsubsidized price. Note, however, that we have not (yet) imposed consumer optimization.

The pure-insurance value will be positive if Medicaid moves resources toward states of the world with a higher marginal utility of consumption and a higher health return to medical spending.

C. Optimization Approaches

In the optimization approaches, we reduce the implementation requirements of the complete-information approach through two additional economic assumptions. We assume that Medicaid affects individuals only 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 $p(q)$.

This assumption rules out other ways in which Medicaid might affect c or h , such as through an effect of Medicaid on a provider's willingness to treat a patient. For implementation purposes, we assume that $p(q)$ is constant in m , although, in principle, we could allow for a nonlinear price schedule. We denote out-of-pocket spending on medical care by

$$x(q, m) \equiv p(q)m. \quad (8)$$

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

ASSUMPTION 3 (Individual optimization). Individuals choose m and c optimally, subject to their budget constraint. Individuals solve

$$\max_{c, m} u(c, \tilde{h}(m; \theta)),$$

subject to

$$c = y(\theta) - x(q, m) \quad \forall m, q, \theta,$$

where $y(\theta)$ denotes (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) and where the complex nature of the decision problem may generate individually suboptimal decisions (Baicker, Mullainathan, and Schwartzstein 2015).

Under these two assumptions, we consider the thought experiment of a "marginal" expansion in Medicaid. 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 (8) is now

$$x(q, m) = qp(1)m + (1 - q)p(0)m. \quad (9)$$

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

$$-\frac{\partial x(q, m(q; \theta))}{\partial q} = (p(0) - p(1))m(q; \theta). \quad (10)$$

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 $-(\partial x/\partial q)$ is a program parameter that holds behavior (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 $q = 0$:

$$E[u(c(0; \theta), h(0; \theta))] = E[u(c(q; \theta) - \gamma(q), h(q; \theta))]. \quad (11)$$

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 impact of insurance on recipient willingness to pay ($d\gamma/dq$) follows from applying the envelope theorem to equation (11):

$$\frac{d\gamma}{dq} = E\left[\frac{u_c}{E[u_c]}((p(0) - p(1))m(q; \theta))\right], \quad (12)$$

where u_c denotes the partial derivative of utility with respect to consumption. Appendix A.1 (apps. A and B are available online) provides the derivation. Note that 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. 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 utility.

The representation in equation (12), 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, $u_c/E[u_c]$ measures the value of money in each 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 each 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, this approach allows individuals to be at a corner with respect to their choice of medical spending.

We can decompose the marginal value of Medicaid to recipients in equation (12) into a transfer term (T) and a pure-insurance term (I):

$$\frac{d\gamma(q)}{dq} = \underbrace{(p(0) - p(1))E[m(q; \theta)]}_{\text{Transfer term}} + \underbrace{\text{Cov}\left[\frac{u_c}{E[u_c]}, (p(0) - p(1))m(q; \theta)\right]}_{\substack{\text{Pure-insurance term} \\ \text{(consumption valuation)}}}. \quad (13)$$

Although implemented differently, the transfer and pure-insurance terms are conceptually the same as in the complete-information approach above. The transfer term measures the recipients' valuation of the expected transfer of resources from the rest of the economy to them. Optimization implies that this value cannot exceed the cost of the transfer and will be below this cost because of the moral-hazard response to insurance. The "pure-insurance" term measures the benefit of a budget-neutral reallocation of resources across different states of the world, θ . The movement of these resources is valued by using the marginal utility of consumption in each state. The pure-insurance term will be positive for risk-averse individuals as long as Medicaid reallocates resources to states of the world with higher marginal utilities of consumption.

We integrate $d\gamma/dq$ from $q = 0$ to $q = 1$ to arrive at a nonmarginal estimate of the recipient's total willingness to pay for Medicaid, $\gamma(1)$, noting that $\gamma(0) = 0$:

$$\begin{aligned} \gamma(1) &= \int_0^1 \frac{d\gamma(q)}{dq} dq = \underbrace{(p(0) - p(1)) \int_0^1 E[m(q; \theta)] dq}_{\text{Transfer term}} \\ &\quad + \underbrace{\int_0^1 \text{Cov}\left(\frac{u_c}{E[u_c]}, (p(0) - p(1))m(q; \theta)\right) dq}_{\substack{\text{Pure-insurance term} \\ \text{(consumption valuation)}}}. \end{aligned} \quad (14)$$

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

Implementation.—Evaluation of the transfer term in equation (13) does not require any assumptions about the utility function. However, integration in equation (14) to obtain an estimate of the transfer term requires that we know the path of $m(q; \theta)$ for interior values of q , which are not directly observed. For our baseline implementation, we make the following statistical assumption:

ASSUMPTION 4 (Linear approximation). The integral expression for $\gamma(1)$ in equation (14) is well approximated by

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

While we use this assumption for our baseline results, we can bound the transfer term without it. 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., $E[m(0; \theta)] \leq E[m(q; \theta)] \leq E[m(1; \theta)]$), we obtain the following lower and upper bounds:

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

While the transfer term does not require specification of a utility function, evaluation of the pure-insurance term in equation (13) requires specifying the marginal utility of consumption. To do so, we assume that the utility function takes the following form:

ASSUMPTION 5 (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.

The pure-insurance term in equation (13) can then be rewritten as

$$\text{Cov} \left(\frac{c(q; \theta)^{-\sigma}}{E[c(q; \theta)^{-\sigma}]}, (p(0) - p(1))m(q; \theta) \right). \quad (16)$$

We can use the above equations to calculate the marginal value of the first and last units of insurance ($d\gamma(0)/dq$ and $d\gamma(1)/dq$, respectively). Assumption 4 then allows us to use estimates of $d\gamma(0)/dq$ and $d\gamma(1)/dq$ to form estimates of $\gamma(1)$.

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. 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 an assumption stronger than assumption 3, which states that individuals optimize; we now require that individual choices of m and c satisfy a first-order condition:

ASSUMPTION 6 (First-order condition holds). The individual's choices of m and c are at an interior optimum and hence satisfy the first-order condition

$$u_c(c, h)p(q) = u_h(c, h) \frac{d\tilde{h}(m; \theta)}{dm} \quad \forall m, q, \theta. \quad (17)$$

The left-hand side of equation (17) is the marginal cost of medical spending operating through forgone consumption. The right-hand side of equation (17) is the marginal benefit of additional medical spending, which equals the marginal utility of health, $u_h(c, h)$, multiplied by the increase in health provided by additional medical spending, $d\tilde{h}/dm$.

We use equation (17) to replace the marginal utility of consumption, u_c in equation (12) with a term depending on the marginal utility of health, u_h , yielding

$$\frac{d\gamma}{dq} = E \left[\left(\frac{u_h}{E[u_c]} \frac{d\tilde{h}(m; \theta)}{dm} \frac{1}{p(q)} \right) ((p(0) - p(1))m(q; \theta)) \right]. \quad (18)$$

We refer to equation (18) as the *health-based optimization approach*. The term $(u_h/E[u_c])(d\tilde{h}(m; \theta)/dm)(1/p(q))$ measures the value of money in each state of the world relative to its average value, and the term $(p(0) - p(1))m(q; \theta)$ measures how much Medicaid relaxes the individual's budget constraint in the current state of the world. From the health-based optimization approach's perspective, the program delivers greater value if it moves more resources to states of the world with a greater return to out-of-pocket spending.⁶

⁶ Unlike the consumption-based optimization approach, the health-based optimization approach will be biased 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. In this case, the first term between parentheses in eq. (18) is less than the true value that the individual puts on money in that state of the world (i.e., $(u_h/E[u_c])(d\tilde{h}(m; \theta)/dm)(1/p(q)) < u_c/E[u_c]$), generating upward bias in the covariance term in eq. (19) below because $(p(0) - p(1))m$ is below its mean at the corner solution $m = 0$.

The marginal value of Medicaid to recipients in equation (18) can be decomposed into a transfer term and a pure-insurance term:

$$\begin{aligned} \frac{d\gamma(q)}{dq} = & \underbrace{(p(0) - p(1))E[m(q; \theta)]}_{\text{Transfer term}} \\ & + \underbrace{\text{Cov}\left(\frac{u_h}{E[u_c]} \frac{d\tilde{h}(m; \theta)}{dm} \frac{1}{p(q)}, (p(0) - p(1))m(q; \theta)\right)}_{\substack{\text{Pure-insurance term} \\ \text{(health valuation)}}}. \end{aligned} \quad (19)$$

Implementation.—Since the transfer term does not depend on the utility function, it is the same as the transfer term in the consumption-based optimization approach. However, evaluation of the pure-insurance term requires a partial specification of the utility function:

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.

Assumption 7 allows us to write the pure-insurance term in the health-based optimization approach in equation (19) as

$$\text{Cov}\left(\frac{d\tilde{h}(m; \theta)}{dm} \frac{\phi}{p(q)}, (p(0) - p(1))m(q; \theta)\right). \quad (20)$$

The term $\phi \equiv \tilde{\phi}/E[\tilde{v}'(c)]$ is the marginal rate of substitution of health for consumption, as in the complete-information approach. Implementation of equation (20) requires that we estimate variation across states of the world in the marginal health return to medical spending, $d\tilde{h}/dm$. As with the consumption-based optimization approach, we estimate $\gamma(1)$ as the average of $d\gamma(0)/dq$ and $d\gamma(1)/dq$, using the linear approximation in assumption 4.

3. Comment: Endless Arguments

A key distinction between the complete-information and optimization approaches is that the optimization approaches allow one to consider marginal utility with respect to one argument of the utility function. In contrast, the complete-information approach requires “adding up” the impact of Medicaid on all arguments of the utility function. In the above model, we assumed that the only arguments were consumption and health. If we were to allow other potentially utility-relevant factors that Medicaid could affect (such as leisure, future consumption, children’s outcomes, or even

hassle costs when dealing with medical providers as an uninsured patient), we would also need to estimate the impact of Medicaid 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.

D. Gross and Net Costs

We benchmark our estimates of $\gamma(1)$ against Medicaid costs. We consider only medical expenditures when estimating program costs. This abstracts from potential administrative costs and from any labor supply responses to Medicaid, both of which could impose fiscal externalities on the government.⁷ Under these assumptions, the average cost to the government per recipient, which we denote by G , is given by

$$G = E[m(1; \theta) - x(1, m(1; \theta))]. \quad (21)$$

This gross cost per recipient, G , may be higher than the net resource cost to society; some component of Medicaid spending may replace costs previously borne by external parties (nonrecipients).

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

$$C = E[m(1; \theta) - m(0; \theta)] + E[x(0, m(0; \theta)) - x(1, m(1; \theta))]. \quad (22)$$

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

We decompose gross costs to the government, G , into net costs, C , and monetary transfers to external parties who provide implicit insurance to the uninsured, N :

$$G = C + N.$$

The monetary transfers to external parties (N) are given by

$$N = E[m(0; \theta)] - E[x(0, m(0; \theta))]. \quad (23)$$

E. Summary: Required Empirical Objects

Table 1 summarizes the empirical objects we need for each approach, highlighting the key trade-offs across approaches in terms of objects that

⁷ In the context of the Oregon Health Insurance Experiment, there is no evidence that Medicaid affected labor market activities (Baicker et al. 2014).

TABLE 1
OVERVIEW OF EMPIRICAL OBJECTS AND METHODOLOGY

USED FOR WHICH ESTIMATES IN WHICH FRAMEWORK(S)				
NOTATION (1)	MEANING (2)	Optimization		
		Complete Information (3)	Consumption Based (4)	Health Based (5)
A. Estimated Objects				
1.	$E[x(q, m(q, \theta))] \text{ for } q = 0, 1$	C, G, N	C, G, N	C, G, N
2.	$E[m(q; \theta)] \text{ for } q = 0, 1$	C, G, N	$C, G, N, \gamma(1)$	$C, G, N, \gamma(1)$
3.	$E[h(q; \theta)] \text{ for } q = 0, 1$	$\gamma(1)$	$\dots$	$\dots$
4.	$d\bar{h}(m; \theta)/dm$	$\dots$	$\dots$	$\gamma(1)$
5.	$c(q; \theta) \text{ for } q = 0, 1$	$\gamma(1)$	$\gamma(1)$	$\dots$
6.	$p(q) \text{ for } q = 0, 1$	$\dots$	$\gamma(1)$	$\gamma(1)$
7.	$x(q, m(q; \theta)) \text{ for } q = 0, 1$	$\dots$	$\gamma(1)$	$\gamma(1)$
B. Parameters of the Utility Function				
1.	Coefficient of relative risk aversion	$\gamma(1)$	$\gamma(1)$	$\dots$
2.	MRS of health for consumption	$\gamma(1)$	$\dots$	$\gamma(1)$

NOTE.— $\gamma(1)$: recipient willingness to pay (out of current consumption) for Medicaid; G : gross cost to the government of providing Medicaid; N : monetary transfers to nonrecipient external parties; $C = G - N$: net cost of Medicaid. All values are per Medicaid recipient. An empty cell in cols. 3–5 means that the required object listed in the row is not needed for the approach listed in the column.

need to be estimated and parameters that need to be calibrated. Estimating costs and transfers to external parties (G , C , and N) requires the same two objects for all three approaches: mean out-of-pocket spending and mean medical spending, both with and without Medicaid. The complete-information approach further requires estimates of mean health outcomes with and without Medicaid and the distribution of consumption with and without Medicaid. It also requires two calibrated parameters of the utility function: one to value health outcomes (ϕ) and one to value consumption outcomes (σ). By contrast, the optimization approaches require either estimates of the distribution of the health returns to medical spending and one calibrated parameter (ϕ) or information on the distribution of consumption, with and without Medicaid, and one calibrated parameter (σ). In addition, the optimization-based approaches both require the out-of-pocket price of Medicaid and the distribution of out-of-pocket spending, with and without Medicaid.

III. Application: The Oregon Health Insurance Experiment

We apply these approaches to the Medicaid expansion that occurred in Oregon in 2008 via a lottery. The Medicaid expansion covered low-income (below 100 percent of the federal poverty line), uninsured adults (aged 19–64) who were not already categorically eligible for Medicaid. The expansion provided comprehensive medical benefits with no patient cost sharing and no or low monthly premiums. We focus on the effects of Medicaid coverage after approximately 1 year.

We use this setting to implement the complete-information approach and two different variants of the consumption-based optimization approach. In the working-paper version (Finkelstein, Hendren, and Luttmer 2015), we also implemented the health-based optimization approach and found that we lacked the statistical power to credibly estimate heterogeneity in the return to medical spending, $d\bar{h}/dm$, and hence the pure-insurance component (I ; see eq. [20]).

A. Empirical Approach

In early 2008, the state opened a waiting list for the Medicaid expansion and then randomly selected 30,000 of the 75,000 people on the waiting list to be able to apply for Medicaid. Following the approach of previous work on the Oregon experiment, we use random assignment by the lottery as an instrument for Medicaid. As a result, all of our estimates of the impact of Medicaid are local average treatment effects of Medicaid for the compliers—that is, those who are covered by Medicaid if and only if they win the lottery. Thus, in our application, “the insured” ($q = 1$)

are treatment compliers and “the uninsured” ($q = 0$) are control compliers. More details can be found in appendix A.2.

The data from the Oregon Health Insurance Experiment are publicly available at www.nber.org/oregon. Data on Medicaid coverage (q) are taken from state administrative records. In our baseline analysis, all other data elements from the Oregon Health Insurance Experiment come from mail surveys sent about 1 year after the lottery to individuals who signed up for the lottery. Table 2 presents descriptive statistics from this mail survey. Panel A presents demographic information. The population is 60 percent female and 83 percent white; about one-third are between the ages of 50 and 64. The demographic characteristics are balanced between treatment and control compliers (p -value = .12). Panel B presents summary

TABLE 2
SUMMARY STATISTICS

	Full Sample (1)	Treatment Compliers ($q = 1$) (2)	Control Compliers ($q = 0$) (3)	Impact of Medicaid (4)
A. Oregon Data Demographics				
Share female	.60	.57	.60	
Share age 50–64	.34	.36	.35	
Share age 19–49	.66	.64	.65	
Share white	.83	.84	.84	
Share black	.03	.03	.03	
Share Spanish/Hispanic/Latino	.11	.07	.08	
Mean family size, n	2.97	2.88	2.91	
B. Oregon Data Outcomes				
12-month medical spending, m :				
Mean medical spending, $E[m]$ (\$)	2,991	3,600	2,721	879
Fraction with positive medical spending, $E[m > 0]$	.74	.79	.72	.07
12-month out-of-pocket spending, x :				
Mean out-of-pocket spending, $E[x]$ (\$)	470	0	569	–569
Fraction with positive out-of-pocket spending, $E[x > 0]$	.38	0	.56	–.56
Health expressed in QALYs, $E[h]$	.77	.78	.74	.05
Share in poor health (QALY = .401)	.11	.10	.17	–.07
Share in fair health (QALY = .707)	.30	.29	.36	–.07
Share in good health (QALY = .841)	.36	.38	.28	.10
Share in very good health (QALY = .931)	.17	.18	.15	.03
Share in excellent health (QALY = .983)	.05	.05	.04	.02

NOTE.—This table reports data from a mail survey of participants in the Oregon Health Insurance Experiment ($N = 15,498$). Columns 2 and 3 report the implied means for treatment and control compliers, respectively, in the Oregon Health Insurance Experiment, and col. 4 reports the estimated impact of Medicaid (i.e., the difference between cols. 3 and 2). Since it cannot be directly observed whether any particular observation is a complier, the results in cols. 2 and 3 are estimated via the instrumental-variables techniques described in more detail in app. A.2.1, as are the results in col. 4. The Oregon health insurance lottery is used as an instrument for Medicaid coverage.

statistics on key outcome measures in the Oregon data; we now discuss their construction.

*B. Medical Spending, Out-of-Pocket Spending,
and Out-of-Pocket Prices*

1. Medical Spending m

Survey responses provide measures of utilization of prescription drugs, outpatient visits, emergency room visits, and inpatient hospital visits. To turn these into spending estimates, Finkelstein et al. (2012) annualized the utilization measures and summed them up, weighting each type by its average cost (expenditures) among low-income, publicly insured, non-elderly adults in the Medical Expenditure Survey (MEPS).⁸

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

2. Out-of-Pocket Spending x

We measure annual out-of-pocket spending for the uninsured ($q = 0$) based on self-reported out-of-pocket medical expenditures in the past 6 months, multiplied by 2.⁹ Average annual out-of-pocket medical expenditures for control compliers is $E[(x(0, m(0, \theta)))] = \569 .

Our baseline analysis assumes that the insured have zero out-of-pocket spending (i.e., $x(1, m(1; \theta)) = 0$). We make this assumption because Medicaid in Oregon has zero out-of-pocket cost sharing, no or minimal premiums, and comprehensive benefits.¹⁰ However, the insured do report positive spending, and we explore sensitivity to using these reports for $x(1, m(1; \theta))$; naturally, this reduces our estimate of the value of Medicaid to recipients.

⁸ The MEPS data on expenditures reflect actual payments (i.e., transacted prices) rather than contract or list prices (MEPS 2013, C-107)).

⁹ To be consistent with our treatment of out-of-pocket spending when we use it to estimate consumption (discussed below in Sec. III.D), we impose two adjustments. First, we fit a lognormal distribution on the out-of-pocket spending distribution. Second, we impose a per capita consumption floor by capping out-of-pocket spending so that per capita consumption never falls below the floor.

¹⁰ This assumes that the uninsured report their out-of-pocket spending without error but that the insured (some of whom report positive out-of-pocket spending in the data) do not. This is consistent with a model of reporting bias in which individuals are responding to the survey with their typical out-of-pocket spending, not the precise spending they have incurred since enrolling in Medicaid. In this instance, there would be little bias in the reported spending for those who are not enrolled in Medicaid (since nothing changed), but the spending for those recently enrolled via the lottery would be dramatically overstated because of recall bias.

3. Out-of-Pocket Prices p

The optimization approaches require that we define the out-of-pocket price of medical care with Medicaid, $p(1)$, and without Medicaid, $p(0)$. Our baseline analysis assumes that $p(1) = 0$; that is, that those with Medicaid pay nothing out of pocket toward medical spending. We measure $p(0)$ as the ratio of mean out-of-pocket spending to mean total medical spending for control compliers ($q = 0$), that is, $E[x(0, m(0; \theta))]/E[m(0; \theta)]$. We estimate $p(0) = 0.21$, which implies 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 is consistent with estimates from other contexts.¹¹

C. Health (h) Inputs

Both the complete-information approach and the health-based optimization approach require that we measure health and that we calibrate individuals' marginal rate of substitution of health for consumption. Our baseline measure of health is the widely used five-point self-assessed health question that asks "In general, would you say your health is . . . ?" and gives the following response options: excellent, very good, good, fair, poor. We conduct sensitivity analysis to using other measures of health.

A key challenge is how to express changes in a given measure of health in units of consumption. For nonmortality health measures, the standard approach involves two steps: first map these health measures into a cardinal utility scale, expressed in terms of QALYs, and then scale it by an estimate of the value of a statistical life year (VSLY) to express it in consumption units. The resulting marginal rate of substitution would be valid for a general population. To find the marginal rate of substitution for low-income individuals, we add a third step: we adjust the estimate for the general population to account for higher marginal utility of consumption in a low-income population. Intuitively, low-income populations have a lower willingness to pay out of consumption because they have lower consumption. We discuss these three steps in turn.

1. Mapping Self-Assessed Health to QALY Units

We map our baseline self-assessed health measure into QALYs, using the mapping that van Doorslaer and Jones (2003) estimated. Their mapping

¹¹ The Kaiser Commission on Medicaid and the Uninsured estimates that the average uninsured person in the United States pays only about 20 percent of their total medical expenses out of pocket (fig. 1 of Coughlin et al. [2014]). Hadley et al. (2008) estimate that the uninsured pay only 35 percent of their medical costs expenses out of pocket. In the 2009–11 MEPS, we estimate that uninsured adults aged 19–64 below 100 percent of the federal poverty line pay about 33 percent of their medical expenses out of pocket.

employs the widely used “Health Utilities Index Mark 3” scale, which applies the “standard gamble approach” to a random sample of 500 adults from the city of Hamilton, Ontario, Canada. Specifically, respondents make choices over hypothetical outcomes in order to find the probability v such that the respondent is indifferent between living in a particular health state and facing a gamble consisting of living in perfect health with probability v and being dead with probability $1 - v$. One year lived in this particular health state is assigned a QALY of v . Appendix A.4 provides more detail.

Panel B of table 2 shows results for our baseline self-assessed health measure, reported in QALY units. Treatment compliers are less likely to respond than control compliers that they are in poor or fair health and likelier to describe their health as good, very good, or excellent. Weighting the effect of Medicaid on each health state by the associated QALY of that health state, our estimates indicate that Medicaid increases health by 0.05 QALYs.

Although QALYs have been frequently used in the economics literature (e.g., Cutler, Huckman, and Kolstad 2010; Chandra, Jena, and Skinner 2011; García et al. 2017; and Lakdawalla, Malani, and Reif 2017), they have the unattractive feature of relying on stated preference. An additional limitation in our setting is that the mapping is estimated on a sample that differs from the population of the Oregon Health Insurance Experiment. To the extent that preferences over a probability of living in perfect health and living for sure in less-than-perfect health are reasonably stable across populations, the mapping will offer a reasonable measure of a QALY.

2. Choosing a VSLY for a General Population

Estimation of the value of a statistical life is also challenging. A large literature, reviewed by Viscusi (1993) and Cropper, Hammitt, and Robinson (2011), uses various approaches to do so. Some, but not all, of these approaches rely on stated preferences. We take as a “consensus” estimate from this literature Cutler’s (2004) choice of \$100,000 for the value of a VSLY for the general US population. In other words, we assume that the marginal rate of substitution of health (as measured by QALYs) for consumption to be \$100,000 in the general US population.

3. Adjusting the Marginal Rate of Substitution for a Low-Income Population

Our utility function assumes that the marginal utility of a QALY does not depend on the level of consumption. However, the marginal utility of consumption is higher in a low-income population because of their low levels

of consumption, and as a result, the marginal rate of substitution of health for consumption is lower in a low-income population. With CRRA utility over consumption (see assumption 1), our baseline assumption of a coefficient of relative risk aversion $\sigma = 3$ (see below) and per capita consumption for our population that is about 40 percent of the general population's (based on our estimates from the CEX), the MRS of health for consumption in our population is approximately 5 percent ($\approx 0.4^3$) of that in the general population. We therefore use a baseline ϕ value of \$5,000 for our population but report sensitivity to alternative values.

We emphasize that a ϕ of \$5,000 reflects the low-income individual's willingness to substitute on the margin their own consumption for QALYs. This does not imply that society's willingness to pay (i.e., reducing other people's consumption) for an additional QALY is only \$5,000 for low-income populations. We return to this distinction in Section IV.C.

D. Consumption (c) Inputs

Both the complete-information approach and the consumption-based optimization approach require that we measure consumption. Specifically, the complete-information approach requires that we estimate the impact of Medicaid on the distribution of consumption, while the consumption-based optimization approach requires that we estimate the joint distribution of consumption and out-of-pocket spending for the uninsured to measure the "pure-insurance" term.¹² For these approaches, we also need to calibrate a curvature of the utility function. In our baseline analysis, we calibrate the coefficient of relative risk aversion at $\sigma = 3$. Because the Oregon study does not contain consumption data, we take two different approaches to measuring consumption, which we now describe.

1. Consumption Proxy Approach

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

$$c = \bar{c} - (x - \bar{x})/n, \quad (24)$$

¹² Equation (16) suggests that we need to estimate the joint distribution of $c(0; \theta)$ and $(p(0) - p(1))m(0; \theta)$ at $q = 0$. Since $p(1) = 0$ by assumption, this reduces to the joint distribution of consumption c and out-of-pocket spending $x(0, m(0; \theta)) = p(0)m(0; \theta)$. We need to estimate this joint distribution only for the uninsured (so for $q = 0$), because our assumption that Medicaid provides full insurance (i.e., $p(1) = 0$) implies that the marginal value of additional insurance for the fully insured (so for $q = 1$) is zero.

where n denotes family size and $\bar{x}$ denotes average per capita out-of-pocket medical spending among control compliers; average family size among compliers is about 2.9 (see table 2). Our approach accounts for within-family resource sharing by assuming that consumption is shared equally within the family; that is, the impact of a given amount of out-of-pocket medical spending on nonmedical consumption is shared equally within families.¹³ This seems a reasonable assumption, given the joint nature of many components of consumption; however, in the sensitivity analysis, we consider that the out-of-pocket spending shock is borne entirely by the individual with the spending.

This consumption proxy approach makes several simplifying assumptions. First, it assumes that the only channel by which Medicaid affects consumption is by reducing out-of-pocket spending; it rules out Medicaid affecting consumption by changing income, which seems empirically reasonable in our context (Finkelstein et al. 2012; Baicker et al. 2014). Second, it assumes that per capita consumption would be the same for all individuals in the Oregon study if they had the same out-of-pocket spending. This is an assumption made for convenience and unlikely to be literally true. However, it approximates reality to the extent to which heterogeneity in nonmedical consumption is limited within our low-income population. Finally, it does not allow for the possibility of any intertemporal consumption smoothing through borrowing or saving. Such opportunities are likely limited in our low-income study population but presumably are not zero; by not allowing for this possibility, we likely bias upward our estimate of $\gamma(1)$.

Implementation.—We use the Oregon survey data to measure x (as described above) and also family size n . We estimate $\bar{c}$ as mean per capita nonmedical consumption in a population that has characteristics similar to those of participants in the Oregon study, namely, families that live below the federal poverty line, have an uninsured household head, and are in the CEX. To estimate the impact of Medicaid on the distribution of out-of-pocket spending x , we make the parametric assumption that out-of-pocket spending is a mixture of a mass point at zero and a lognormal spending distribution and then estimate the distribution of out-of-pocket spending x for control compliers, using standard, parametric quantile instrumental-variables techniques; see appendix A.2 for more detail.

Because there is unavoidable measurement error in estimating consumption and because marginal utility is sensitive to low values of consumption,

¹³ This same logic implies that the benefits from Medicaid are also shared among family members. This is captured in the optimization approach by eq. (14); this equation values any dollar flowing to the family by the marginal utility of consumption of the individual irrespective of whether the dollar is used to benefit the individual or other family members. However, for the complete-information approach, it requires that we replace $\gamma(1)$ by $\gamma(1)/n$ when estimating eq. (3).

we rule out implausibly low values of c by imposing an annual consumption floor. Our baseline analysis imposes a consumption floor at the 1st percentile of nonmedical consumption for low-income uninsured individuals in the CEX (i.e., \$1,977). We impose the consumption floor by capping the out-of-pocket spending drawn from the fitted lognormal distribution at $\bar{x} + n(\bar{c} - c_{\text{floor}})$, where $\bar{x}$ is average per capita out-of-pocket medical spending, as in equation (24). Our baseline consumption floor binds for fewer than 0.3 percent of control compliers. In the sensitivity analysis, we explore sensitivity to the assumed value of the consumption floor. Finally, we map the fitted, capped out-of-pocket spending distribution to consumption, using equation (24).

Figure 1 shows the resultant distributions of consumption for control compliers ($q = 0$) and treatment compliers ($q = 1$). Average nonmedical consumption for control compliers is \$9,214, with a standard deviation of \$1,089. For treatment compliers, consumption is simply the average nonmedical consumption for the insured (\$9,505), since by assumption $x(1, m) = 0$.¹⁴ The difference between the two lines in the figure shows the increase in consumption due to Medicaid for the compliers.

2. CEX Approach

The consumption proxy approach 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 or draw down assets or have other ways of smoothing consumption, this approach overstates the consumption-smoothing benefits of Medicaid. We therefore also employ an alternative approach that uses national data on out-of-pocket spending x and nonmedical consumption c for low-income individuals from the CEX. For the consumption-based optimization approach, the CEX data allow us to directly estimate the pure-insurance term at $q = 0$ in equation (16), that is, the covariance between the marginal utility of nonmedical consumption c and out-of-pocket spending x among the uninsured. Appendix A.5.1 provides more detail on the data, sample definition, and summary statistics in the CEX data and compares that sample to the sample of compliers in the Oregon data.

The key advantage of the CEX approach over the consumption proxy approach is the ability to directly observe consumption and its covariance with out-of-pocket spending. But it has two important drawbacks. First, it cannot be used for the complete-information approach because this

¹⁴ Average nonmedical consumption for the low-income uninsured (i.e., $\bar{c}$) is \$9,214 in the CEX. To account for the fact that nonmedical consumption for the uninsured is presumably lowered because of out-of-pocket medical costs, which are 0 for the insured, we assume that average nonmedical consumption for the insured is $\bar{c} + \bar{x}/n$ (see eq. [24]), where $\bar{x}$ denotes average out-of-pocket spending for the uninsured.




FIG. 1.—Consumption distribution for treatment and control compliers. Color version available as an online enhancement.

approach requires a causal estimate of the impact of Medicaid on consumption, which cannot be estimated in the CEX data.¹⁵ Second, the data come from a national sample of low-income individuals, not the Oregon study data.

In principle, it is straightforward to directly estimate the correlation between the marginal utility of consumption and out-of-pocket medical spending for uninsured individuals in the CEX data. We wish to estimate equation (16). For $q = 0$, this reduces to

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

where c and x are, respectively, observed nonmedical consumption and out-of-pocket medical spending for the uninsured in the CEX. We impose the same consumption floor as in the consumption proxy approach.

¹⁵ For the pure-insurance term of the consumption-based optimization approach, we need to evaluate the covariance term of eq. (16) only for $q = 0$, because we know that the covariance term is zero for $q = 1$, given our baseline assumption that the insured face no consumption risk from medical expenditures. Hence, we do not need a causal estimate of the impact of Medicaid on consumption.

In practice, we face an additional challenge that the raw data show a negative covariance between the marginal utility of consumption and out-of-pocket spending among the uninsured. This is not an idiosyncratic feature of the CEX; we also estimate a negative covariance in the Panel Study of Income Dynamics. The negative covariance remains even after income and assets are controlled for, and we suspect that the covariance term is biased from measurement error that induces a negative correlation between $c(0; \theta)^{-\sigma}$ and $x(0; \theta)$.

We therefore implement a measurement-error correction that allows for potentially nonclassical measurement error in out-of-pocket medical spending. We do so by exploiting a key implication of our model: the covariance between out-of-pocket medical spending and the marginal utility of consumption should be zero for the insured ($q = 1$) because they have no out-of-pocket medical spending. Under the assumption that measurement error in out-of-pocket medical spending is the same for the insured and the uninsured, we use the estimated covariance term for the insured to infer the impact of measurement error on the covariance term for the uninsured. Appendix A.5.2 provides more detail on our approach.

IV. Results

A. Baseline Results

1. Utility-Free Estimates: Medicaid Costs and Transfers

Without any assumptions about the utility function, the experimental estimates deliver several key objects. The gross cost of Medicaid (G) equals total medical spending for treatment compliers ($q = 1$), since treatment compliers have no out-of-pocket spending (see eq. [21]). Table 2 indicates that G is \$3,600 per recipient year. This is broadly consistent with external estimates of annual per-recipient spending in the Medicaid program in Oregon (Wallace et al. 2008).

The net cost of Medicaid (C) equals the average increase in medical spending due to Medicaid plus the average decrease in out-of-pocket spending due to Medicaid (see eq. [22]). Table 2 shows that the impact of Medicaid on medical spending is \$879 and that the impact on out-of-pocket spending is $-\$569$. Hence, $C = \$1,448$. The monetary transfer from Medicaid to external parties, N , is the difference between G and C (see eq. [23]), or \$2,152. Thus, about 60 cents of every dollar of government spending on Medicaid is a transfer to external parties ($N/G \approx 0.6$).

Finally, the optimization approach allows us to estimate the value of the transfer component of Medicaid to recipients using only the estimates of the impact of Medicaid on m and p (see eq. [14]). The change in the out-of-pocket price for medical care due to insurance ($p(0) - p(1)$) is

0.21. Using linear approximation (assumption 4) and the estimates of $E[m(0, \theta)]$ and $E[m(1, \theta)]$ of \$2,721 and \$3,600, respectively (see table 2), we calculate a transfer term of \$661. Without the linear approximation, we can derive lower and upper bounds for the transfer term of \$569 and \$752, respectively (see eq. [15]).

2. Complete-Information Approach

As shown in equation (4), the complete-information approach requires us to estimate mean health outcomes and the distribution of consumption for control compliers ($q = 0$) and treatment compliers ($q = 1$). Table 2 shows the estimates for mean health outcomes, while figure 1 shows the estimated distribution of consumption at $q = 0$ and $q = 1$. The complete-information approach further requires that we calibrate the marginal rate of substitution of health for consumption (ϕ) and a coefficient of relative risk aversion (σ). As discussed above, our baseline specification assumes $\phi = \$5,000$ and $\sigma = 3$.

This implementation of the complete-information approach yields an estimate of $\gamma(1) = \$1,675$. In other words, a Medicaid recipient would be indifferent between giving up Medicaid and giving up \$1,675 in consumption. The complete-information approach lends itself to decomposing $\gamma(1)$ into the component operating through health (γ_h) and the component operating through consumption (γ_c). We define γ_c by

$$E\left[\frac{c(0; \theta)^{1-\sigma}}{1-\sigma}\right] = E\left[\frac{(c(1; \theta) - \gamma_c)^{1-\sigma}}{1-\sigma}\right], \quad (25)$$

and estimate $\gamma_c = \$1,381$. We then infer the value of Medicaid to recipients operating through health as $\gamma_h = \gamma(1) - \gamma_c = \294 .¹⁶ In other words, about 80 percent of the recipient willingness to pay for Medicaid comes through its impact on consumption, as opposed to health.

Decomposition of $\gamma(1)$ into a transfer term and a pure-insurance term requires estimates of heterogeneity in the return to medical spending, $d\tilde{h}/dm$ (see eqq. [6], [7]). As mentioned above, we do not have the statistical power to estimate heterogeneity in $d\tilde{h}/dm$. However, because the estimates of $d\tilde{h}/dm$ are needed only for the decomposition of the health component γ_h , we can still find the transfer term of the consumption component γ_c .¹⁷ By setting the right-hand side of equation (6) to zero, we

¹⁶ Because of the curvature of the utility function, the order of operations naturally matters. If we instead directly estimate γ_h and infer γ_c from $\gamma(1) - \gamma_h$, we estimate $\gamma_c = \$1,059$ and $\gamma_h = \$615$.

¹⁷ Appendix A.3 provides implementation details of how we decompose γ_c into a transfer component and a pure-insurance component. The pure-insurance component, operating through consumption smoothing, is broadly similar to the approach taken by Feldstein and Gruber (1995) to estimate the consumption-smoothing value of catastrophic

obtain an estimate of the consumption component of the transfer term of \$569. Thus, the lower bound for the entire transfer term is \$569. By assuming that the entire health component ($\gamma_h = \$294$) is part of the transfer term, we obtain an upper bound for the transfer term of \$863 (\$569 + \$294). The resulting bounds on the pure-insurance component are \$812 and \$1,106. This suggests that roughly a third to a half of the value of Medicaid comes from its transfer component, with the remainder coming from Medicaid's ability to move resources across states of the world.

3. Consumption-Based Optimization Approach

We estimate the transfer component and the pure-insurance component separately and combine them for our estimate of $\gamma(1)$. Estimation of the transfer component is straightforward and, as described above, produced an estimate of \$661. Estimation of the "pure-insurance" component, however, is more complicated. We undertake two approaches; for both, we assume $\sigma = 3$.

Consumption-based optimization approach with consumption proxy.—We estimate the pure-insurance value at $q = 0$ by using equation (16) on the Oregon sample. This requires an estimate of the joint distribution of consumption and out-of-pocket spending for control compliers (see n. 12). The distribution of c for $q = 0$ is shown in figure 1, and the joint distribution of consumption and out-of-pocket spending follows from equation (24). At $q = 1$, the pure-insurance value of Medicaid is zero because the marginal utility of consumption is constant. Following the linear approximation in assumption 4, the total pure-insurance component is therefore one-half of what we estimate at $q = 0$, or \$760. Adding this to the previously estimated transfer component implies that $\gamma(1) = \$1,421$.

Consumption-based optimization approach with CEX consumption measure.—We also estimate the pure-insurance value at $q = 0$ by using low-income individuals in the CEX. As explained in detail in appendix A.5 and shown in appendix table 2, we use the difference between the observed covariance term for the uninsured and that for the insured to estimate the measurement-error-corrected covariance term for the uninsured. The resulting measurement-error-corrected covariance between the marginal utility of consumption and out-of-pocket spending at $q = 0$ is \$265 for our baseline measure of consumption. As above, the assumption that Medicaid provides full insurance implies that the pure-insurance value of Medicaid is zero at the margin at $q = 1$. The linear approximation over $q = 1$ and $q = 0$ yields a pure-insurance value of \$133. Adding this to our previously estimated transfer component implies that $\gamma(1) = \$793$.

health insurance and by Finkelstein and McKnight (2008) to estimate the consumption-smoothing value of the introduction of Medicare.

B. *Summary and Sensitivity*

The first row of table 3 summarizes our estimates of recipient willingness to pay for Medicaid $\gamma(1)$. The estimates range from \$1,675 (standard error = \$60) in the complete-information approach to \$1,421 (standard error = \$180) in the consumption-based optimization approach using a consumption proxy, to \$793 (standard error = \$417) in the consumption-based optimization approach using the CEX consumption measure. The next two rows summarize the decomposition of $\gamma(1)$ into a transfer and a pure-insurance component. The results suggest that the transfer component represents a large share of $\gamma(1)$. Under the optimization approach, the transfer component contributes between one-half and four-fifths of total willingness to pay; the bounds in the complete-information approach suggest that the transfer component accounts for at least a third and as much as half of $\gamma(1)$.

Panel B provides some benchmarks. The first row shows that the net cost of providing Medicaid to recipients is only 40 percent of government spending on Medicaid; the majority of government spending on Medicaid goes to external parties who, in the absence of Medicaid, would have given the recipients medical care without being fully paid. The second row compares recipient willingness to pay to net cost ($C = \$1,148$). It shows that whether or not recipient willingness to pay exceeds net costs depends on the approach—with $\gamma(1)/C$ ranging from 0.55 to 1.16. A finding of $\gamma(1)$ above C implies that the insurance value of Medicaid to

TABLE 3
WILLINGNESS TO PAY (WTP) FOR MEDICAID BY RECIPIENTS

	COMPLETE- INFORMATION APPROACH (1)	OPTIMIZATION APPROACHES	
		Consumption-Based (Consumption Proxy) (2)	Consumption-Based (CEX Consumption Measure) (3)
A. Recipient WTP for Medicaid			
$\gamma(1)$ (standard error)	1,675 (60)	1,421 (180)	793 (417)
Transfer component, T	569–863	661	661
Pure-insurance component, I	812–1,106	760	133
B. Benchmarks			
Net costs as fraction of gross cost, C/G	.40	.40	.40
Recipient WTP as fraction of net cost, $\gamma(1)/C$	1.16	.98	.55
Moral hazard cost, $G - T - N$	585–879	787	787

NOTE.—Estimates of WTP and moral-hazard costs are expressed in dollars per year per Medicaid recipient. Standard errors are bootstrapped with 500 repetitions.

recipients, I , exceeds the moral-hazard costs of Medicaid, $G - N - T$, while $\gamma(1)$ below C implies the converse.¹⁸ The final row of panel B shows that the moral hazard of Medicaid is substantial across all approaches.

Naturally, all of our quantitative results are sensitive to the framework used and to our specific implementation assumptions. We explored sensitivity to a variety of alternative assumptions, including the assumed level of risk aversion, the assumed consumption floor, the measurement of out-of-pocket spending for those on Medicaid, the assumed amount of within-family risk smoothing, and the method of interpolation in the optimization approach. We also explored sensitivity to alternative ways of valuing health improvements and alternative health measures. As table 1 indicates, these will affect certain estimates and not others.

Appendix B describes our sensitivity analyses and results. Across specifications, recipient willingness to pay is roughly of the same order of magnitude as net costs, and the transfer value to recipients is always substantial, but the estimates of the pure-insurance value are more sensitive. For the complete-information approach, the biggest impact on the estimates comes from assuming $\sigma = 5$, which raises our estimate of $\gamma(1)/C$ from 1.2 to 2.8. The next-biggest effect comes from replacing our baseline calibration of a marginal rate of substitution of health for consumption of \$5,000 with a value of \$40,000, which would result if willingness to pay for health scales linearly with consumption. Under the consumption-based optimization approach using the consumption proxy, the biggest change comes from assuming that the shock is borne entirely by the individual. This more than doubles our estimate of $\gamma(1)/C$, from 1.0 to 2.2. The consumption-based optimization approach using the CEX consumption measure is more stable. The biggest impact comes from assuming $\sigma = 5$; this raises $\gamma(1)/C$ from 0.55 to 0.60.

C. Discussion

1. External Parties

A striking finding is that a major beneficiary of Medicaid expansions are nonrecipients, who receive 60 percent of each dollar of government Medicaid spending. An open and important question concerns the identity of these “nonrecipients.” The provision of uncompensated care by hospitals is a natural starting point. Recent evidence indicates that hospital visits by the uninsured are associated with very large unpaid bills (Dobkin et al. 2018) and that increases in Medicaid coverage lead to large reductions in uncompensated care by hospitals (Garthwaite et al. 2018).

¹⁸ By definition, $\gamma(1) = T + I$ and $C = G - N$. Therefore, a comparison of I to $G - N - T$ is equivalent to a comparison of $\gamma(1)$ to C .

The ultimate economic incidence of the transfers to external parties is more complicated. While some of the incidence may fall on the direct recipients of the monetary transfers, other parties, including the privately insured, the recipients themselves (e.g., if reductions in unpaid medical debt provide benefits to recipients), and the public-sector budget, may bear some of the incidence.

2. Recipient Willingness to Pay If (Counterfactually) the Uninsured Had No Implicit Insurance

We assess how much higher recipient willingness to pay for Medicaid would be if (counterfactually) the low-income uninsured had no implicit insurance. To do so, we extrapolate (grossly) out of sample from the observed demand for medical care at $p = 0$ (for treatment compliers) and at $p = 0.21$ (for control compliers) to the demand for medical care at $p = 1$ (i.e., if the uninsured had to pay the full cost of their medical care). We do this by assuming that the demand for medical care is log-linear in p .¹⁹ This out-of-sample exercise suggests that, if the low-income uninsured had to pay all of their medical costs, recipient willingness to pay for Medicaid would increase to \$2,749 under the complete-information approach (compared to our baseline estimate of \$1,675) and to \$3,875 for the consumption-based optimization approach (compared to \$1,421 for our baseline estimate in table 3).

3. Recipient Willingness to Pay Relative to Medicaid Cost

As noted in the introduction, the Congressional Budget Office currently uses G for the value of Medicaid for recipients (Congressional Budget Office 2012). 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 . If market failures such as adverse selection (e.g., knowledge of θ when choosing insurance) result in private insurance not being available at actuarially fair prices, $\gamma(1)$ could exceed G , although it might not if moral-hazard costs and crowd-out of implicit insurance (i.e., N) sufficiently reduce $\gamma(1)$. Ultimately, these are empirical questions.

Across the different approaches, we consistently estimate that $\gamma(1)$ is less than G , with our estimates of $\gamma(1)/G$ ranging between \$0.2 and \$0.5. This implies that Medicaid recipients would rather give up Medicaid than

¹⁹ Once we have an estimate of the (counterfactual) distribution of m at $p = 1$, this straightforwardly implies counterfactual distributions of x and c (in our consumption-proxy-based approach). For the complete-information approach, we also need a counterfactual estimate for the mean of h , which we get by simple linear extrapolation.

pay the government's costs of providing Medicaid; likewise, an uninsured person would choose the status quo over giving up G in consumption to obtain Medicaid. This contrasts with the current approach used by the Congressional Budget Office to value Medicaid at government cost.

However, since the gross costs of Medicaid (G) greatly exceed its net costs ($C = G - N$), it is also instructive to compare $\gamma(1)$ to C . We think of this as a useful thought exercise even though it is not clear that there is a corresponding practical implementation option of delivering Medicaid without the transfer to external parties. Of course, if the government is itself a major recipient of the transfers to "external parties" (Hadley et al. 2008), our net cost estimate C may approximate the "true" cost of Medicaid to the public sector.

It is theoretically ambiguous whether $\gamma(1)$ will be higher or lower than C . Recipient willingness to pay for Medicaid may be higher than its net cost because of its insurance value, or it may be lower because of moral-hazard effects. The results indicate that, depending on the approach, recipient willingness to pay for Medicaid relative to its net costs (i.e., $\gamma(1)/C$) varies from about 0.5 to 1.2.

4. Recipient versus Societal Willingness to Pay

The fact that our population has low levels of consumption (because of low levels of income and/or liquidity constraints) implies that they have a high marginal utility of consumption, which contributes to a low willingness to give up consumption for other goods. As we emphasized at the outset, societal willingness to pay may be considerably higher, given the redistributive nature of Medicaid. We can derive a societal willingness to pay for Medicaid by multiplying $\gamma(1)$ by the relevant social welfare weight.²⁰

Consider, for example, a utilitarian social welfare function over individual utilities. Social willingness to pay is therefore recipient willingness to pay multiplied by the ratio of the marginal utility of consumption of the recipient to the marginal utility of consumption of the average person in the population. A rough calculation from the CEX suggests that the median consumption in the recipient population in the Oregon Health Insurance Experiment is about 40 percent of the median consumption level of the general population. Given our assumption of CRRA individual utility with a coefficient of relative risk aversion of $\sigma = 3$ (i.e., a marginal

²⁰ Societal willingness to pay may also be higher than individual willingness to pay for two other reasons not captured in our analysis. First, Medicaid may provide insurance value to those who are not currently eligible for Medicaid but would become eligible if they experienced a sufficiently large negative shock. Second, government provision of insurance may reduce inefficiencies stemming from the Samaritan's dilemma (Coate 1995). Both channels are beyond the scope of our paper.

utility of consumption of $1/c^3$), this would suggest a societal willingness to pay for Medicaid that is nearly 20 times recipient willingness to pay; even with log utility, societal willingness to pay would be 2.5 times recipient willingness to pay.

An alternative approach that does not require assuming a specific social welfare function would be to compare recipient willingness to pay per dollar of government expenditure to their willingness to pay for alternative redistributive instruments, such as a tax cut (Hendren 2016, 2017). This asks whether redistributing through Medicaid is a more or less costly way to transfer resources to a low-income population than other transfer programs. To make this comparison, one can compare the marginal value of public funds (MVPF) of Medicaid spending to the MVPF of other policies, such as the EITC. For comparison, Hendren (2016) estimates an MVPF of about \$0.9 for the EITC; the recipients would be willing to pay roughly \$0.9 for every dollar of government spending on the EITC.

Whether Medicaid's MVPF compares favorably to the EITC's depends critically on the ultimate economic incidence of the transfers to external parties ($N \approx 0.6G$). If the government bears the incidence of uncompensated-care payments, then the cost to the government of providing Medicaid would be C , so that the MVPF would be γ/C , which ranges from 0.5 to 1.2. If the low-income individuals bear the ultimate incidence of the transfers—as would be the case if reductions in uncompensated-care costs for the newly insured allowed medical providers to provide better care to the remaining low-income uninsured—then the relevant comparison of the 0.9 estimate for the EITC is to $(\gamma + 0.6G)/G$. This value ranges from 0.8 to 1.1. Finally, if the ultimate incidence of the transfers is on the high end of the income distribution—such as hospital owners or the privately insured—then Hendren (2017) shows that one can down-weight these gains by the marginal cost of moving \$1 from the top to the bottom of the income distribution through modifications to the tax schedule, which yields an estimated weight of 0.5. In this case, the relevant EITC comparison could be to $(\gamma + 0.3G)/G$, or from 0.5 to 0.8. Thus, if affluent populations are the ultimate beneficiaries of reductions in uncompensated care, it suggests that the MVPF, or “bang for the buck,” for Medicaid spending is lower than that for the EITC.

D. Trade-offs across Alternative Approaches

We highlight some of the trade-offs and limitations across the various approaches; we also highlight which limitations could be surmounted with better data and which are fundamental limitations of each approach.

First, as noted in Section III.C, we rely on stated preferences to translate our measured impacts of Medicaid on health into QALYs, which we then value by using an estimate of the MRS of health for consumption among a low-income population. Reasonable people may well have

serious concerns about either the reliability of QALYs or the assumptions we make to translate a VSLY estimate into an MRS for our low-income population. In other settings, direct estimates of the mortality impact of Medicaid would allow the researcher to avoid QALYs but would still require an assumption about the MRS for a low-income population. As a result, readers wishing to avoid estimates of the VSLY or its translation into an MRS may prefer the consumption-based optimization approach, which does not require such assumptions.

Second, with better consumption data for the Oregon study population, one could avoid using a consumption proxy in both the complete-information approach and the consumption-proxy variant of the consumption-based optimization approach. Willingness-to-pay estimates using our consumption proxy may be biased upward, since the proxy assumes that the uninsured have no means of smoothing consumption through savings or borrowing.

Third, our estimates from the optimization approaches may be biased (in either direction) because of imperfect measurement of prices. 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 the uninsured effectively face no out-of-pocket costs above a certain level of expenditures (Mahoney 2015). A downward bias in our estimate of $p(0)$ reduces the estimate of $\gamma(1)$ (see eq. [12]) and creates an upward bias in the effect on external parties, N . An upward bias in $p(0)$ has the opposite effect.

Fourth, the linear interpolation between $d\gamma(0)/dq$ and $d\gamma(1)/dq$ used in the optimization approaches may downwardly 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, 1999b). This limitation cannot be addressed with better data (short of observing a program that would give individuals partial Medicaid coverage).²¹ By contrast, the complete-information approach would accurately capture the value stemming from the liquidity Medicaid provides.

We also highlight some general trade-offs between the optimization and complete-information approaches. Since the complete-information

²¹ Consider an extreme example in which there is a single expensive medical procedure that individuals may undergo in the event of a health shock and in which individuals are sufficiently liquidity constrained that they will undertake this procedure only if $q \geq 0.4$. As a result, $d\gamma/dq$ would be zero until $q = 0.4$, jump up at $q = 0.4$, and decline thereafter. The linear approximation would not capture the relatively large values of $d\gamma/dq$ that would occur for intermediate values of q and would in this case underestimate $\gamma(1)$.

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 bias the complete-information estimate of $\gamma(1)$ (in either direction). On the other hand, the optimization approaches assume that individuals have the ability and information to maximize their utility. If this assumption is violated, we will mismeasure how much they value the relaxation of the budget constraint provided by Medicaid, which can bias our estimate of $\gamma(1)$ in either direction.

Fundamentally, even with ideal data, all of the approaches developed here require assumptions about the shape of the utility function. Other methods that substitute alternative assumptions can provide a useful complement to the approaches developed here. For example, Krueger and Kuziemko (2013) directly survey individuals about their stated willingness to pay for hypothetical health insurance plan offerings, while Finkelstein et al. (2019) estimate demand for private health insurance in a low-income population and interpret the demand curve through the lens of revealed preference.

A final distinction between the optimization approaches and the complete-information approach that could be relevant in other settings is that the complete-information approach is better suited to do welfare analysis when there are externalities or when the social welfare function does not solely take individuals' utilities as arguments. In such cases, we could put different social weights on different components of utility (e.g., the consumption component vs. the health component) to capture externalities or paternalistic social welfare preferences. This is obviously not possible with the optimization approaches, because those do not allow for a decomposition of utility into its components.

V. Conclusion

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

Our analysis uncovers that Medicaid is best conceived of as having two distinct parts: a subsidized health insurance product for low-income individuals and a transfer to external parties who would otherwise subsidize medical care for the low-income uninsured. We estimate that 60 cents of every dollar of government Medicaid spending is a transfer to these external parties. This suggests the importance of future work studying their immediate and ultimate economic incidence.

A priori, recipient willingness to pay for Medicaid could be higher or lower than the net (of transfer to external parties) cost of Medicaid. Our

results here are sensitive to the framework used, with estimates of recipient willingness to pay per dollar of net cost ranging from 0.5 to 1.2. Across approaches, recipient willingness to pay coming from Medicaid transfer of resources (since the insurance is heavily subsidized) is relatively stable around 0.5 dollars per dollar of net cost, whereas estimates of recipient willingness to pay for the pure (budget-neutral)-insurance component vary considerably.

Our empirical findings are naturally specific to our setting. In particular, as noted in Finkelstein et al. (2012), the impact of Medicaid may well differ when it is mandatory rather than voluntary, when it is expanded to cover a larger number of individuals, or when it is provided over a longer time horizon. The value of Medicaid may also differ for Medicaid populations other than the low-income adult population studied here, such as children, the disabled, or the elderly, for whom there is also a large empirical literature on Medicaid's effects (see Buchmueller, Ham, and Shore-Sheppard [2016] for a recent review). However, the approaches we have developed can be applied to studying the value of Medicaid in other contexts.

Our approach can also be adapted to study the welfare impact of other social insurance programs. For example, for disability insurance, an existing literature uses approaches analogous to our "complete-information approach" for welfare analysis (e.g., Autor et al. 2019). Our frameworks clarify the role of the modeling assumptions in these welfare analyses and provide potential pathways to use the optimization-based approaches to relax some of these assumptions. Likewise, our frameworks could be applied to Medicare, where there is a large empirical literature examining the impacts of Medicare on welfare-relevant outcomes such as health care use, health, and out-of-pocket medical expenditures (e.g., Card, Dobkin, and Maestas 2008, 2009; Barcellos and Jacobson 2015).

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. These approaches advance beyond common defaults of zero valuation or valuation at government cost. We hope that 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

- Angrist, Joshua D., and Jörn-Steffen Pischke. 2009. *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton, NJ: Princeton Univ. Press.
- Arrow, Kenneth J. 1963. "Uncertainty and the Welfare Economics of Medical Care." *A.E.R.* 53 (5): 941–73.

- Autor, David, Andreas Kostol, Magne Mogstad, and Bradley Setzler. 2019. "Disability Benefits, Consumption Insurance, and Household Labor Supply." *A.E.R.* 109 (7): 2613–54.
- Baicker, Katherine, Amy Finkelstein, Jae Song, and Sarah Taubman. 2014. "The Impact of Medicaid on Labor Force Activity and Program Participation: Evidence from the Oregon Health Insurance Experiment." *A.E.R.* 104 (5): 322–28.
- Baicker, Katherine, Sendhil Mullainathan, and Joshua Schwartzstein. 2015. "Behavioral Hazard in Health Insurance." *Q.J.E.* 130 (4): 1623–67.
- Baicker, Katherine, Sarah Taubman, Heidi Allen, et al. 2013. "The Oregon Experiment—Effects of Medicaid on Clinical Outcomes." *New England J. Medicine* 368:1713–22.
- Barcellos, Silvia Helena, and Mireille Jacobson. 2015. "The Effects of Medicare on Medical Expenditure Risk and Financial Strain." *American Econ. J.: Econ. Policy* 7 (4): 41–70.
- Buchmueller, Thomas, John C. Ham, and Lara D. Shore-Sheppard. 2016. "The Medicaid Program." In *Economics of Means-Tested Transfer Programs in the United States*, vol. 1, edited by Robert A. Moffitt, 21–136. Chicago: Univ. Chicago Press (for NBER).
- Card, David, Carlos Dobkin, and Nicole Maestas. 2008. "The Impact of Nearly Universal Insurance Coverage on Health Care Utilization: Evidence from Medicare." *A.E.R.* 98 (5): 2242–58.
- . 2009. "Does Medicare Save Lives?" *Q.J.E.* 124 (2): 597–636.
- Cardon, James H., and Igal Hendel. 2001. "Asymmetric Information in Health Insurance: Evidence from the National Medical Expenditure Survey." *RAND J. Econ.* 32 (3): 408–27.
- Chandra, Amitabh, Anupam B. Jena, and Jonathan S. Skinner. 2011. "The Pragmatist's Guide to Comparative Effectiveness Research." *J. Econ. Perspectives* 25 (2): 27–46.
- Chetty, Raj. 2009. "Sufficient Statistics for Welfare Analysis: A Bridge between Structural and Reduced-Form Methods." *Ann. Rev. Econ.* 1:451–88.
- Coate, Stephen. 1995. "Altruism, the Samaritan's Dilemma, and Government Transfer Policy." *A.E.R.* 85 (1): 46–57.
- Congressional Budget Office. 2012. "The Distribution of Household Income and Federal Taxes, 2008 and 2009." Publication no. 4441. <https://www.cbo.gov/publication/43373>.
- Coughlin, Teresa A., John Holahan, Kyle Caswell, and Megan McGrath. 2014. "Uncompensated Care for Uninsured in 2013: A Detailed Examination." Technical report, Kaiser Commission on Medicaid and the Uninsured, Washington, DC.
- Cropper, Maureen, James K. Hammitt, and Lisa A. Robinson. 2011. "Valuing Mortality Risk Reductions: Progress and Challenges." *Ann. Rev. Resource Econ.* 3:313–36.
- Currie, Janet, and Firouz Gahvari. 2008. "Transfers in Cash and In-Kind: Theory Meets the Data." *J. Econ. Literature* 46 (2): 333–83.
- Cutler, David M. 2004. *Your Money or Your Life: Strong Medicine for America's Health-care System*. New York: Oxford Univ. Press.
- Cutler, David M., Robert S. Huckman, and Jonathan T. Kolstad. 2010. "Input Constraints and the Efficiency of Entry: Lessons from Cardiac Surgery." *American Econ. J.: Econ. Policy* 2 (1): 51–76.
- Dague, Laura. 2014. "The Effect of Medicaid Premiums on Enrollment: A Regression Discontinuity Approach." *J. Health Econ.* 37:1–12.
- Dobkin, Carlos, Amy Finkelstein, Raymond Kluender, and Matthew J. Notowidigdo. 2018. "The Economic Consequences of Hospital Admissions." *A.E.R.* 108 (2): 308–52.

- Ehrlich, Isaac, and Gary S. Becker. 1972. "Market Insurance, Self-Insurance, and Self-Protection." *J.P.E.* 80 (4): 623–48.
- Einav, Liran, Amy Finkelstein, and Jonathan Levin. 2010. "Beyond Testing: Empirical Models of Insurance Markets." *Ann. Rev. Econ.* 2:311–36.
- Feldstein, Martin, and Jonathan Gruber. 1995. "A Major Risk Approach to Health Insurance Reform." In *Tax Policy and the Economy*, vol. 9, edited by James M. Poterba, 103–30. Cambridge, MA: MIT Press.
- Finkelstein, Amy, Nathaniel Hendren, and Erzo F. P. Luttmer. 2015. "The Value of Medicaid: Interpreting Results from the Oregon Health Insurance Experiment." Working Paper no. 21308 (June), NBER, Cambridge, MA.
- Finkelstein, Amy, Nathaniel Hendren, and Mark Shepard. 2019. "Subsidizing Health Insurance for Low-Income Adults: Evidence from Massachusetts." *A.E.R.* 109 (4): 1530–67.
- Finkelstein, Amy, and Robin McKnight. 2008. "What Did Medicare Do? The Initial Impact of Medicare on Mortality and Out of Pocket Medical Spending." *J. Public Econ.* 92 (7): 1644–68.
- Finkelstein, Amy, Sarah Taubman, Bill Wright, et al. 2012. "The Oregon Health Insurance Experiment: Evidence from the First Year." *Q.J.E.* 127 (3): 1057–1106.
- Finkelstein, Amy N., Sarah L. Taubman, Heidi L. Allen, Bill J. Wright, and Katherine Baicker. 2016. "Effect of Medicaid Coverage on ED Use—Further Evidence from Oregon's Experiment." *New England J. Medicine* 375:1505–7.
- Gallen, Trevor S. 2015. "Using Participant Behavior to Measure the Value of Social Programs: The Case of Medicaid." Working paper, Dept. Econ., Purdue Univ.
- García, Jorge Luis, James J. Heckman, Duncan Ermini Leaf, and María José Prados. 2017. "Quantifying the Life-Cycle Benefits of a Prototypical Early Childhood Program." Working Paper no. 23479 (June), NBER, Cambridge, MA.
- Garthwaite, Craig, Tal Gross, and Matthew J. Notowidigdo. 2018. "Hospitals as Insurers of Last Resort." *American Econ. J.: Appl. Econ.* 10 (1): 1–39.
- Gottschalk, Peter, and Timothy M. Smeeding. 1997. "Cross-National Comparisons of Earnings and Income Inequality." *J. Econ. Literature* 35 (2): 633–87.
- Hadley, Jack, John Holahan, Teresa Coughlin, and Dawn Miller. 2008. "Covering the Uninsured in 2008: Current Costs, Sources of Payment, and Incremental Costs." *Health Affairs* 27 (5): w399–w415.
- Hendren, Nathaniel. 2016. "The Policy Elasticity." *Tax Policy and Econ.* 30:51–89.
- . 2017. "Efficient Welfare Weights." Working Paper no. 20351 (July), NBER, Cambridge, MA.
- Keane, Michael, and Robert Moffitt. 1998. "A Structural Model of Multiple Welfare Program Participation and Labor Supply." *Internat. Econ. Rev.* 39 (3): 553–89.
- Krueger, Alan, and Ilyana Kuziemko. 2013. "The Demand for Health Insurance among Uninsured Americans: Results of a Survey Experiment and Implications for Policy." *J. Health Econ.* 32 (5): 780–93.
- Lakdawalla, Darius, Anup Malani, and Julian Reif. 2017. "The Insurance Value of Medical Innovation." *J. Public Econ.* 145:94–102.
- Mahoney, Neale. 2015. "Bankruptcy as Implicit Health Insurance." *A.E.R.* 105 (2): 710–46.
- MEPS (Medical Expenditure Panel Survey). 2013. "MEPS HC-147 2011 Full Year Consolidated Data File." https://meps.ahrq.gov/mepsweb/data_stats/download_data_files_detail.jsp?cboPufNumber=HC-147.
- Nyman, John A. 1999a. "The Economics of Moral Hazard Revisited." *J. Health Econ.* 18 (6): 811–24.
- . 1999b. "The Value of Health Insurance: The Access Motive." *J. Health Econ.* 18 (2): 141–52.

- Okun, Arthur. 1975. *Equality and Efficiency: The Big Tradeoff*. Washington, DC: Brookings Inst. Press.
- Saez, Emmanuel, and Stefanie Stantcheva. 2016. "Generalized Social Welfare Weights for Optimal Tax Theory." *A.E.R.* 106 (1): 24–45.
- Samuelson, Paul A. 1954. "The Pure Theory of Public Expenditure." *Rev. Econ. and Statis.* 36 (4): 387–89.
- Taubman, Sarah, Heidi Allen, Bill Wright, Katherine Baicker, and Amy Finkelstein. 2014. "Medicaid Increases Emergency-Department Use: Evidence from Oregon's Health Insurance Experiment." *Science* 343 (6168): 263–68.
- US Department of Agriculture. 2016. "The Food Assistance Landscape: FY 2015 Annual Report." Econ. Information Bull. EIB-150. <https://www.ers.usda.gov/publications/pub-details/?pubid=44062>.
- US Department of Health and Human Services. 2015. "2015 Actuarial Report on the Financial Outlook for Medicaid." <https://www.medicaid.gov/medicaid/finance/downloads/medicaid-actuarial-report-2015.pdf>.
- . 2016. "FY 2015 Federal TANF & State MOE Financial Data." https://www.acf.hhs.gov/sites/default/files/ofa/tanf_financial_data_fy_2015.pdf.
- US Internal Revenue Service. 2016. "Statistics of Income Bulletin Historical Table 1." <https://www.irs.gov/statistics/soi-tax-stats-historical-table-1>.
- US Social Security Administration. 2016. "Annual Report of the Supplemental Security Income Program." <https://www.ssa.gov/OACT/ssir/SSI16/ssi2016.pdf>.
- van Doorslaer, Eddy, and Andrew M. Jones. 2003. "Inequalities in Self-Reported Health: Validation of a New Approach to Measurement." *J. Health Econ.* 22 (1): 61–87.
- Viscusi, W. Kip. 1993. "The Value of Risks to Life and Health." *J. Econ. Literature* 31 (4): 1912–46.
- Wallace, Neal T., K. John McConnell, Charles A. Gallia, and Jeanene A. Smith. 2008. "How Effective Are Copayments in Reducing Expenditures for Low-Income Adult Medicaid Beneficiaries? Experience from the Oregon Health Plan." *Health Services Res.* 43 (2): 515–30.