A Online Appendices

A.1 Derivations

To derive equation (12) from equation (11), consider the level of utility the individual obtains if she optimally chooses $m(q;\theta)$ subject to a budget constraint in which she must pay $\gamma(q)$ to obtain insurance:

$$V(q) = max_{m(q;\theta)}\, E\left[u\left(y(\theta) - x(q, m(q;\theta)) - \gamma(q), h(m(q;\theta);\theta)\right)\right], \tag{26}$$

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

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

Given that $V(q) = E\left[u\left(c(0;\theta), h(0;\theta)\right)\right]$ for all q by equation (11), it follows that $dV/dq = 0$. Using $dV/dq = 0$ to solve the equation above for $d\gamma/dq$ yields:

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

Using equation (10), $-\frac{\partial x}{\partial q} = (p(0) - p(1))m(q;\theta)$, we obtain:

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

which was what we wanted to show. Note that this derivation does not require medical spending to be strictly positive; in other words, individuals can be at a corner solution. The derivation also allows for cases where there are "lumpy" medical expenditures so that an individual is not indifferent between an additional \$1 of out-of-pocket medical spending and \$1 less consumption.[21] Intuitively, the individual values the mechanical relaxation of the budget constraint from Medicaid according to the marginal utility of consumption, regardless of the extent to which she has ability to substitute an increase in another good (e.g., medical care) for the increase in consumption.

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

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

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

Derivation in an alternative setup: insurance lotteries

In order to satisfy the maximization in equation (11), the relevant arguments of the marginal utility function, u_c, need to be the choices that individuals in state θ would actually make if they face price $p(q)$, have income $y(\theta)$, and pay $\gamma(q)$. Observed choices do not satisfy this maximization if individuals, in fact, do not pay $\gamma(q)$. Intuitively, for $q>0$, there would be income effects that cause people to change their allocation of c and m. In practice, in our baseline implementation this is not a problem, since we assume $x(1,m)=0$ – and therefore we know the pure-insurance term (as defined in equation (13)) must be zero on the margin for the fully insured – and linearly interpolate between our estimates of $\frac{d\gamma}{dq}$ at $q=0$ and $q=1$.

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

$$E\left[u\left(c\left(q;\theta\right)-\epsilon\gamma\left(q\right),h\left(q;\theta\right)\right)\right]=\left(1-\epsilon\right)E\left[u\left(c\left(q;\theta\right),h\left(q;\theta\right)\right)\right]+\epsilon E\left[u\left(c\left(0;\theta\right),h\left(0;\theta\right)\right)\right]. \quad (27)$$

We derive $\frac{d\gamma}{dq}$ in the insurance lotteries setup by considering the first-order condition for the choice of m in the special case when choices are continuously differentiable in q. This approach is detailed in the subsection below and shows that $\frac{dc}{dq}u_c+\frac{dh}{dq}u_h=\left(-\frac{\partial x}{\partial q}\right)u_c$. Taking the derivative of (27) with respect to q, and using $\frac{dc}{dq}u_c+\frac{dh}{dq}u_h=\left(-\frac{\partial x}{\partial q}\right)u_c$, we obtain:

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

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

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

Now, noting that $\left(-\frac{\partial x}{\partial q}\right) = (p(1) - p(0))\, m(q;\theta)$, we obtain equation (12).

Alternative derivation: using the first-order condition

Given the central role of equation (12) in the optimization approaches, we also derive equation (12) by exploiting the first-order condition. This derivation requires the first-order condition (equation (17)) to hold with equality and is therefore less general than our main derivation, which is based on the envelope theorem. However, the derivation based on the first-order condition very nicely shows the intuition behind the optimization approaches, and we therefore present it here.

To derive equation (12) from equation (11), it is useful to first derive two intermediate expressions. First, we differentiate the budget constraint $c(q;\theta) = y(\theta) - x(q, m(q;\theta))$ with respect to q:

$$\frac{dc}{dq} = -\frac{\partial x}{\partial q} - \frac{\partial x}{\partial m}\frac{dm}{dq} = -\frac{\partial x}{\partial q} - p(q)\frac{dm}{dq} \quad \forall q, \theta. \tag{28}$$

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

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

$$\frac{dh}{dq} = \frac{d\tilde{h}}{dm}\frac{dm}{dq} \quad \forall q, \theta. \tag{29}$$

We then totally differentiate equation (11) with respect to q, which yields the marginal impact of insurance on recipients' willingness to pay, $\frac{d\gamma}{dq}$, as the implicit solution to:

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

Rearranging, we obtain:

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

where the second line follows from substituting $\frac{dc}{dq}$ and $\frac{dh}{dq}$ (equations (28) and (29)).[22] Using $-\frac{\partial x}{\partial q} = (p(0) - p(1))m(q;\theta)$, we obtain:

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

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

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

This section provides some additional information on how we analyze the data from the Oregon Medicaid lottery. Much more detail on the data and the lottery can be found in Finkelstein et al. (2012)[11].

A.2.1 Estimation of impacts

As described in Section 3.1, we estimate the impact of Medicaid on outcomes by IV, where an indicator for being selected by the lottery is the instrument. When analyzing the mean impact of Medicaid on an individual outcome y_i (such as medical spending m_i, out-of-pocket spending x_i, or health h_i), we estimate equations of the following form:

$$y_i = \alpha_0 + \alpha_1 Medicaid_i + \epsilon_i, \tag{30}$$

where $Medicaid$ is an indicator variable for whether the individual is covered by Medicaid at any point in the study period. We estimate equation (30) by two-stage least squares, using the following first-stage equation:

$$Medicaid_i = \beta_0 + \beta_1 Lottery_i + \nu_i, \tag{31}$$

in which the excluded instrument is the variable "*Lottery*," which is an indicator variable for whether the individual was selected by the lottery. Winning the lottery increased the probability of being on Medicaid at any time during the subsequent year by about 30 percentage points. This "first-stage" effect of lottery selection on Medicaid coverage was below one because many lottery winners either did not apply for Medicaid or were deemed ineligible.

Previous work has used the lottery as an instrument for Medicaid to examine the impact of Medicaid on health care utilization, financial well-being, labor market outcomes, health, and private insurance coverage (Finkelstein et al. (2012)[11], Baicker et al. (2013)[5], Baicker et al. (2014)[4], and Taubman et al. (2014)[24]). Finkelstein et al. (2012)[11] provides supporting evidence on the

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

assumptions required to use the lottery as an instrument for Medicaid coverage.

One particular feature of the lottery design affects our implementation. Although the lottery selected individuals, any member of these individuals' households could apply for Medicaid. As a result, if more people from a household were on the waiting list, the household had more "lottery tickets" and a higher chance of being selected. The lottery was thus random conditional on the number of people in the household who were on the waiting list, which we refer to as the number of "lottery tickets." In practice, about 60 percent of the individuals on the list were in households with one ticket, and virtually all the remainder had two tickets. (We drop the less than 0.5 percent who had three tickets; no one had more). In households with two tickets, the variable "*Lottery*" is one if any household member was selected by the lottery. In all of our analysis, therefore, we perform the estimation separately for one-ticket and two-ticket households. Because there is no natural or interesting distinction between these two sets of households, all estimates presented in the paper consist of the weighted average of the estimates for these two groups.

Much of our analysis is based on estimates of characteristics of treatment and/or control compliers – i.e., those who are covered by Medicaid if and only if they win the lottery (see, e.g., Angrist and Pischke (2009)[3]). Our estimation of these characteristics is standard (see, e.g., Abadie (2002)[1]; Abadie (2003)[2]; and Angrist and Pischke (2009)[3]). For example, uninsured individuals who won the lottery provide estimates of characteristics of never-takers. Since uninsured individuals who lost the lottery include both control compliers and never-takers, with estimates of the never-taker sample and the share of individuals who are compliers, we can back out the characteristics of control compliers. Likewise, insured individuals who lost the lottery provide estimates of characteristics of always-takers. Since insured lottery winners include both treatment compliers and always-takers, we can in like manner identify the characteristics of treatment compliers. Differences between treatment and control compliers reflect the impact of Medicaid (i.e., α_1) in the IV estimation of equation (30).

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

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

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

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

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

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

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

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

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

where $LOGN\left(x|\mu, \nu\right)$ is the CDF of a log-normal distribution with mean and variance parameters, μ and ν, evaluated at $x > 0$. For $x = 0$, the CDF is given solely by ψ^g, so that this parameter captures the fraction of group g with zero out-of-pocket spending. Under standard IV assumptions, the 12 parameters are identified from the joint distribution of out-of-pocket spending, insurance status, and lottery status. (In practice, we estimate F^g separately for households with 1 and 2 lottery tickets, and therefore estimate 24 parameters).[23] We estimate all parameters jointly using

[23]We impose the consumption floor by capping the out-of-pocket spending distribution, as described in Section (3.4.1).

maximum likelihood using the approach laid out in subsection A.2.1. To assess the goodness of fit, Figure A1 plots the estimated and actual CDF separately based on lottery status (won or lost), insurance status, and number of tickets. As can be seen from these figures, the parametric model fits quite well.

A.2.3 Results and comparison to previous results

Our sample, variable definitions, and estimation approach are slightly different from those in Finkelstein et al. (2012)[11]. Appendix Table 3 walks through the differences in the approaches and shows that these differences are fairly inconsequential for the estimates reported in the two papers. Column I replicates the results from Finkelstein et al. (2012)[11]. In column II, we limit the data to the subsample used in our own analysis, which consists of about 15,500 individuals out of the approximately 24,000 individuals from Finkelstein et al. (2012)[11]. Our subsample excludes those who have missing values for any of the variables we use in the analysis. The primary reason for the loss of sample size is missing information on prescription drug utilization (a component of medical spending m). Missing data on self-reported health, household income, number of family members, out-of-pocket spending, and other health care use also contribute slightly to the reduction of sample size. We also exclude the few people who had three people in the household signed up for the lottery, as described above.

Column III reports the results on our subsample using our estimating equations above. These estimating equations differ from those used by Finkelstein et al. (2012)[11] in several ways. First, we stratify on the number of tickets and report weighted averages of the results rather than include indicator variables for the number of tickets, as in Finkelstein et al. (2012)[11]; we thus allow the effects of insurance to potentially differ by number of tickets. Second, we do not control for which of the 8 different survey waves the data come from as in Finkelstein et al. (2012)[11].[24] And finally, we do not up-weight the subsample of individuals in the intensive-follow-up survey arm. As shown in column III, these deviations do not meaningfully affect the results.

Finally, Column IV reports the results using our subsample and our estimating equation, adjusting the "raw" out-of-pocket data as described in Section 4. Specifically, we estimate the distribution of out-of-pocket spending by fitting the parametric distribution described above and shown in Figure A1; we set out-of-pocket spending to zero for the insured; and we impose a ceiling on out-of-pocket spending for the uninsured. Naturally, these adjustments only affect the estimated effect of Medicaid on out-of-pocket spending. The combination of these changes increases the estimated impact of Medicaid on out-of-pocket spending from -\$350 to -\$569, primarily as a result of setting out-of-pocket spending to zero for the insured. In particular, simply replacing $x = 0$ for the insured in the raw data in column III increases the estimated impact of Medicaid on out-of-pocket spending from -\$350 to -\$581. Imposing the parametric model and consumption floor (of \$1,977)

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

moves this estimate from -\$581 to -\$569.

A.3 Decomposition of $\gamma(1)$ in the complete-information approach

To provide insights into the drivers of the estimate of $\gamma(1)$, we decompose $\gamma(1)$ into γ_C and γ_h as described in Section (4.1.2). We can further decompose the components associated with consumption effects (γ_C) and effects on health (γ_h) into a transfer and a pure-insurance component. We estimate the consumption transfer term ($\gamma_{C,Transfer}$) as the mean increase in consumption due to the program so that

$$\gamma_{C,Transfer} = E[c(1;\theta) - c(0;\theta)]. \tag{32}$$

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

$$\gamma_{C,Ins} = \gamma_C - \gamma_{C,Transfer}. \tag{33}$$

By substituting the health production function (equation (2)) into the definition of γ (equation (4)), we can in principle decompose the component due to effects on health (γ_h) into a transfer component ($\gamma_{h,Transfer}$) and an insurance component ($\gamma_{h,Ins} = \gamma_h - \gamma_{h,Transfer}$). The transfer component in health ($\gamma_{h,Transfer}$) is given by:

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

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

$$\begin{aligned} E\left[\frac{c(0;\theta)^{1-\sigma} - \left(c(1;\theta) - \gamma_C - \gamma_{h,Transfer}\right)^{1-\sigma}}{1-\sigma}\right] &= \tilde{\phi}E\left[\tilde{h}\left(E[m(1;\theta)];\theta\right) - \tilde{h}\left(E[m(0;\theta)];\theta\right)\right] \\ &= \tilde{\phi}E\left[\frac{d\tilde{h}}{dm}\right]E\left[m(1;\theta) - m(0;\theta)\right]. \end{aligned} \tag{34}$$

Evaluating this equation requires an estimate of $E\left[\frac{d\tilde{h}}{dm}\right]$, the slope of the health production function between $m(1;\theta)$ and $m(0;\theta)$, averaged over all states of the world. As noted at the top of Section 3, we lack the statistical power to credibly estimate $\frac{d\tilde{h}}{dm}$ conditional on the state of the world. We therefore do not implement a decomposition of (γ_h) in to a transfer term and a pure-insurance term.

A.4 Health measures and their mapping to QALYs

Methodological approaches to mapping health measures to QALYs

There is an extensive literature on the measurement of QALYs. A good overview of this literature and the principal issues involved is given by Whitehead and Ali (2010)[27] and Brazier et al. (2010)[7]. QALYs are defined such that a life-year in perfect health is counted as one QALY and being dead in a given year counts as a QALY of zero. The challenge lies in assigning a QALY to a year lived in less-than-perfect health. Both principal methods used for this rely on self-reported preferences over hypotheticals. The "Standard-Gamble" method elicits a probability υ such that a respondent reports being indifferent between living in a particular health state and facing a gamble consisting of living in perfect health with probability υ and being dead with probability $1-\upsilon$. One year lived in this particular health state is assigned a QALY of υ. The "Time-Trade-Off" method elicits the value of υ for which the respondent is indifferent between living for some number of years (call this Z years) in a particular health state and living for νZ years in perfect health. One year lived in this particular health state is assigned a QALY of υ. In short, QALYs are designed to aggregate life-years taking the quality of those life-years into account.

Baseline: Mapping self-assessed health to QALYs

We use the estimates by Van Doorslaer and Jones (2003)[25] to map the responses to our baseline self-assessed health measure into QALYs. Van Doorslaer and Jones use about 15,000 observations from the 1994-1995 wave of the Canadian "National Population Health Survey" (NPHS). The NPHS contains the same self-assessed health question as in our Oregon data. In addition, the NPHS contains the standard set of more detailed health questions that are used to form the McMaster University "Health Utilities Index Mark 3" (HUI3). The HUI3 is a cardinal health scale that runs between zero and one and was constructed using the "Standard-Gamble" method based on responses of about 500 adults randomly selected from the general population of the City of Hamilton, Canada. This means that a year lived in a health state where HUI3=ν translates into ν QALYs. Details on the construction of the HUI3 are described in Furlong et al. (1998)[12]. Van Doorslaer and Jones (2003)[25] estimate how responses to the self-assessed health question correspond to the HUI3. We use their preferred estimates, which are based on interval regression methods to assign QALYs to responses to our self-assessed health variable. These estimates (their Table 4) show that a year lived in "poor health" corresponds to 0.4010 QALYs, "fair health" to 0.7070 QALYs, "good health" to 0.8410 QALYs, "very good health" to 0.9311 QALYs, and "excellent health" to 0.9833 QALYs.

Mapping responses to the Patient Health Questionnaire (PHQ) to QALYs

The Patient Health Questionnaire (PHQ) measures mental health. It consists of a number of questions that have the following structure: "Over the last two weeks, how often have you been bothered by X? Would you say it was..." Each question substitutes a different symptom for X, namely: "having little interest or pleasure in doing things," "feeling down, depressed, or hopeless," “trouble falling or staying asleep, or sleeping too much,” “feeling tired or having little energy,” “poor appetite or overeating,” “feeling bad about yourself, or that you’re a failure, or have let yourself or your family down,” “trouble concentrating on things, such as reading the newspaper or watching TV,” and “moving or speaking so slowly that other people could have noticed? Or the opposite - being so fidgety or restless that you have been moving around a lot more than usual?”. The 2-item PHQ-2 asked on the mail-in survey contains the first two symptoms listed above. The PHQ-8 asked during the in-person survey contains all 8 symptoms listed above.

Possible responses to each question are “not at all,” “several days,” “more than half the days,” “nearly every day,” “don’t know,” and “prefer not to answer.” The PHQ scoring method is that, on each item, “not at all” gets a score of 0, “several days” gets a score of 1, “more than half the days” gets a score of 2, and “nearly every day” gets a score of 3. The total score is simply the sum of the item scores.

We convert PHQ scores to QALYs by relying on estimates from Pyne et al. (2009)[20]. Pyne et al. (2009)[20] convert the PHQ-8 responses to QALYs using the “Standard Gamble” methodology. Specifically, they ask 95 randomly sampled adults from the Central Arkansas area to evaluate three vignettes of people with different degrees of depression. Each of these vignettes is described in terms of the frequency of the symptoms that constitute the items of the PHQ. Hence, each of these vignettes can readily be scored on the PHQ scale. Vignette A has a PHQ-8 score of 5, vignette B has a score of 10, and vignette C has a score of 24. Table 1 of Pyne et al. (2009)[20] show that living a year in the health condition described by vignette A corresponds to 0.78 QALYs. Vignette B corresponds to 0.70 QALYs and vignette C to 0.54 QALYs. These three vignettes therefore define QALYs for three scores on the PHQ-8 scale. By definition, a fully healthy person (with a PHQ-8 of zero) corresponds to a QALY of 1. We linearly interpolate the PHQ-8 to QALY correspondence for intermediate values of the PHQ-8 score. To map the PHQ-2 to QALYs, we first multiply the PHQ-2 score by 4 to make it comparable to the PHQ-8 score, and then apply the same PHQ-8 to QALY correspondence.

Mapping responses to the Short Form questions to QALYs

The in-person survey contains the questions from the 8-item version of Short Form health questionnaire (SF-8). The SF-8 questionnaire contains questions that ask the respondent about any experiences of physical and mental health issues in the past four weeks.[25] We translate responses to the SF-8 into a summary physical health score and a summary mental health score using the methodology developed by the designers of the SF-8 questionnaire (Ware et al. (2001)[26]). To map the two SF-8 summary measures in to a single measure on a QALY scale, we use the conversion formula estimated by Sullivan and Ghushchyan (2006)[23], as described on page 407 of their article. They estimate their conversion formula using 37,000 observations from the 2000 and 2002 waves of the Medical Expenditure Panel Survey (MEPS). These waves of the MEPS contain both

[25] The exact wording of the questions is described in http://www.nber.org/oregon/documents/survey/in-person/inperson-survey-interview-text.pdf

a SF health questionnaire and a cardinal health measure, the so-called EQ-5D index score, which is mapped to QALYs using the "Time-Trade-Off" method.[26]

A.5 Measuring consumption and estimating the pure-insurance term using national data

A.5.1 Consumption measure from the CEX

This subsection details the data used to estimate the covariance term in the consumption optimization approach using the CEX. The CEX consists of a series of short panels. Each "consumer unit" (CU) is interviewed every 3 months over 5 calendar quarters. In the initial interview, information is collected on demographic and family characteristics and on the consumer unit's inventory of major durable goods. Expenditure information is collected in the second through the fifth interviews using uniform questionnaires. Income and employment information is collected in the second and fifth interviews.

Our sample includes all CUs in 1996-2010 who have valid expenditure data in all 4 quarters (i.e., positive total expenditure and non-negative medical expenditure) and non-missing income data. To be broadly consistent with the Oregon sample, we further limit the analysis to families that are headed by an adult aged 19-64 and are below 100% of the federal poverty line. We measure insurance status q at the start of the survey[27], regardless of whether or not the individual obtains insurance later in the year (results are quite similar if we use concurrent insurance status). Because the CEX requests information on the health insurance status only of the household head, we restrict the sample to single adults with no children in the household, so that we can identify the individuals who are insured ($q = 1$) and uninsured ($q = 0$). We convert all dollar amounts to 2009-dollars, and impose an annual consumption floor (although in practice the baseline consumption floor of \$1,977 never binds).

Appendix Table 1 reports the summary statistics for the sample and compares it to the sample of compliers in the Oregon data. The CEX sample is slightly older, naturally has smaller family size (because of our limitation to singles), and tends to have somewhat lower out-of-pocket spending (\$395 versus \$569).

A.5.2 Measurement error correction approach

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

[26] Specifically, Sullivan and Ghushchyan (2006)[23] using the mapping from EQ-5D to QALY from Shaw et al. (2005)[22], who derive this mapping using the "Time-Trade-Off" method on a representative sample of about 4000 adults in the U.S.

[27] Insurance status is measured as any insurance, not just Medicaid.

particular, we assume the observed out-of-pocket spending is given by

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

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

We identify the covariance term even under this fairly general measurement-error structure by making three assumptions. First, we assume (non-medical) consumption is measured without error. Second, we assume that the covariance of the marginal utility of consumption and the measurement error is the same for the insured and uninsured.[28] Third, we assume that true out-of-pocket medical spending is zero for the insured, so that $\hat{x}(1;\theta) = \epsilon(1;\theta)$. In other words, we allow for measurement error in $\hat{x}$ that is additive in x, arbitrarily correlated with c, and common for the insured and uninsured. These assumptions would be satisfied, for example, if ϵ reflected consumption of uncovered healthcare for both the insured and uninsured (e.g., over-the-counter pain killers, or transportation costs associated with medical care) and these are consumed in equal amounts by both groups.

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

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

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

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

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

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

Intuitively, we estimate the true covariance term as the difference in the covariance terms for the

[28] Formally, we assume:

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

uninsured and insured, where the latter term removes the measurement error bias from the results.

A.5.3 Alternative measurement error correction.

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

Sample and variables

For our baseline specification, we consider the sample of all household heads between the ages of 25-64 with non-missing reports for hospitalization and consumption data drawn from the biennial waves of the PSID from 2003-2013. To better align with the low-income population in the Oregon sample, we restrict the sample to households with per capita household income below $20,000. This yields 6,600 observations from 3,715 unique household heads, as reported in Appendix Table 6.

We define consumption expenditure as the sum of all expenditures available in the PSID excluding health expenditures (food, rent/home expenses, car expenses) normalized where appropriate to arrive at an annual expenditure measure.[29] We divide this consumption by the number of household members to arrive at per capita consumption, c.[30] We define out-of-pocket medical expenditure as the sum of hospital and nursing costs, doctor and dental costs, and prescription drug costs for the past two years divided by two. Again, we divide this expenditure number by the number of household members to arrive at per capita out-of-pocket medical expenditures. For our baseline specification, we winsorize per capita consumption, c, and per capita out-of-pocket medical expenditure, x, at the 1st and 99th percentiles. All dollar variables are deflated to 2009 using the CPI-U-RS. Finally, we let Z denote an indicator for any hospitalization of the household head in the past 12 months.

Appendix Table 6 reports the summary statistics for the sample. Average per capita consumption expenditure is $5,351, which is substantially lower than per capita consumption in the Oregon

[29] For example, the survey asks about transportation costs, such as parking fees, for the last month; in contrast, it generally asks food costs in the past year. We scale transportation costs by 12 to arrive at an annual measure of expenditure.

[30] As noted in previous literature (Li et al. (2010)[17]), the PSID contains roughly 70% of the consumption expenditure that is captured by the CEX. We have also conducted results using solely food expenditure to estimate the relationship between consumption and medical spending, and obtained similar results.

sample. This large difference in consumption expenditure likely reflects the fact that the PSID does not capture a comprehensive measure of consumption.

Setup

We implement our measurement error correction as follows. We wish to estimate $cov\left(\frac{u'(c)}{E[u'(c)]}, x\right)$, where $u'(c) = c^{-\sigma}$. We construct a log-linearization of the utility function:

$$\frac{c^{-\sigma}}{E\left[c^{-\sigma}\right]} - 1 \approx log\left(c^{-\sigma}\right) = \alpha + \beta log\left(x\right) + \epsilon, \tag{36}$$

so that β is a log-log regression coefficient of $c^{-\sigma}$ on x, $\beta = \frac{cov\left(log\left(c^{-\sigma}\right), log(x)\right)}{var(log(x))}$. With this approximation,

$$\begin{aligned} cov\left(\frac{c^{-\sigma}}{E\left[c^{-\sigma}\right]}, x\right) &\approx cov\left(log\left(c^{-\sigma}\right), x\right) \\ &= cov\left(log\left(c^{-\sigma}\right), \frac{x}{E\left[x\right]}\right) E\left[x\right] \\ &\approx cov\left(log\left(c^{-\sigma}\right), log\left(x\right)\right) E\left[x\right] \\ &\approx \beta var\left(log\left(x\right)\right) E\left[x\right] \\ &\approx \beta var\left(\frac{x}{E\left[x\right]}\right) E\left[x\right], \end{aligned}$$

so that

$$cov\left(\frac{c^{-\sigma}}{E\left[c^{-\sigma}\right]}, x\right) \approx \beta \frac{var\left(x\right)}{E\left[x\right]}. \tag{37}$$

Equation (37) provides a method to recover the covariance term using (a) data on the distribution of x and (b) data on the relationship between x and c. To most closely align with the Oregon sample, we take $var\left(x\right)$ and $E\left[x\right]$ from control compliers in the Oregon sample. These are $E\left[x\right] = \$569$ and $std\left(x\right) = \$543$. But, we use data from the PSID to calculate β. To do so, we regress $log\left(c^{-\sigma}\right)$ on $log\left(x\right)$, including controls for an age cubic, quadratic controls for household size, and year dummies. We use our baseline value of $\sigma = 3$. Because there is no reason to expect the effect of out-of-pocket spending on consumption to depend on insurance status, the estimation sample for β need not be restricted to the uninsured, and we do not impose this restriction. As insurance status is only measured in the PSID for household heads, limiting the sample to the uninsured would only be possible by examining uninsured household heads, which would dramatically reduce our sample size.

A.5.4 Consumption Covariance Estimates

Appendix Table 2 shows the estimates of the consumption covariance in the CEX data corrected for measurement error. Column I shows results for our baseline measure of non-health consumption, which is a broad-based measure. It consists of total expenditure excluding individual expenditures for health care providers, prescription drugs, and medical devices.[31] In columns II and III, we show results based on alternative definitions of non-health consumption. Specifically, in column II, we exclude durables within each expenditure category because an expenditure on a durable good leads to a consumption flow over a longer period of time than that in which the expenditure occurred. In column III, we create a consumption measure that is limited to expenditures in categories that are relatively easy to adjust in the short run: food, entertainment, apparel, tobacco, alcohol, personal care, and reading.

Across all the consumption definitions, the covariances between the marginal utility of consumption and out-of-pocket spending for the uninsured are negative.[32] However, the covariance is more negative for the insured. Applying the measurement error correction approach from equation (35) yields a covariance between the marginal utility of consumption and out-of-pocket spending at $q = 0$ of \$265 for our baseline, broad-based measure of consumption (column I); the estimate based on the consumption measure excluding durables is similar (column II) while the estimate based on expenditures in relatively easily adjustable categories is substantially lower (column III). As before, the assumption that Medicaid provides full insurance implies that the pure-insurance value of Medicaid is 0 at the margin at $q = 1$ and, using the linear approximation to obtain an average covariance value over $q = 1$ to $q = 0$ yields a pure-insurance value of \$133.

As noted above, we also implement the consumption covariance term using data from the PSID, and an alternative measurement error correction approach based on instrumenting for out-of-pocket medical spending with hospital admissions. In practice, we obtain similar results using the IV approach in the PSID to what is reported here using the CEX data. Appendix Table 7 reports the results from the PSID approach and illustrates the calculation of the consumption covariance term. Column I illustrates the OLS relationship between $log\,(x)$ and $log\,(c^{-\sigma})$.[33] As in the CEX, we find a negative relationship, with $\beta \approx -0.2$. Taken literally, it would imply a consumption covariance of -\$52. To deal with potential measurement error in the elicitations, we instrument $log\,(x)$ with an indicator for hospitalization of the household head in the past 12 months. Column II presents the results from this IV strategy. We estimate that being hospitalized is associated with an 18% increase in the marginal utility of consumption (this is a 6% drop in consumption, as $\sigma = 3$) and a 19% increase in out-of-pocket medical spending. This suggests $\beta \approx 1$. Combining this estimate with $E\,[x]$ and $var\,(x)$ from the Oregon data yields a value of the consumption covariance term at $q = 0$ of \$495. Dividing by 2 for the linear interpolation between $q = 0$ and $q = 1$ results in an estimated consumption covariance term of \$248 (s.e.: \$138), which is larger but statistically indistinguishable from our baseline estimate using the CEX approach of \$133.

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

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

[33] Standard errors are clustered by household head.

A.6 Health production function, $E_{\theta|\theta^K}\left[\frac{\partial \tilde{h}}{\partial m}\right]$

To implement the health-based optimization approach and to exactly decompose the complete-information estimate into a transfer and a pure-insurance term, we would need to estimate the health returns to medical spending conditional on medical spending, m, and state of the world, θ. To get causal estimates, we use the Medicaid lottery as an instrument for medical spending. The challenge lies in capturing heterogeneity in the health returns to medical spending. In the working paper version (Finkelstein, Hendren, and Luttmer (2015)[10]), we estimated heterogeneity in $\frac{\partial \tilde{h}}{\partial m}$ by assuming that differences in the returns to medical spending can be captured by differences in state variables, θ^K, that were based on measures of financial and health states from an initial survey (fielded essentially concurrently with the lottery). Even though we allowed for only four states of the world, we lacked sufficient statistical power to estimate a sufficiently significant first stage separately by state. As a result, the IV estimates of $\frac{\partial \tilde{h}}{\partial m}$ by θ^K , and the estimates of $\gamma(1)$ that depend on them, are not sufficiently credible to report in the current paper.

B Sensitivity Analysis

In this Section, we explore sensitivity to a number of key assumptions. Such "specification" uncertainty is not reflected in the standard errors in Table 3; the sensitivity analysis here provides an informal way of gauging sensitivity to these assumptions.

B.1 Alternative assumptions unrelated to health

Appendix Table 4 explores the sensitivity of our results within each framework (shown in different rows) to a number of different non-health assumptions (shown in different columns). For the sake of brevity, we focus the discussion on two main estimates: recipients' willingness to pay out of consumption to obtain Medicaid (i.e., $\gamma(1)$), and the ratio of recipient willingness to pay to net costs (i.e., $\gamma(1)/C$). Column I shows our baseline results from Table 3.

Risk aversion and consumption floor. Our baseline analysis uses a coefficient of relative risk aversion of 3 and a consumption floor of \$1,977.[34] Columns II through V explore alternative choices for risk aversion (coefficients of relative risk aversion of 1 and 5) and the consumption floor (of \$1,000 or \$5,000). A lower consumption floor increases $\gamma(1)$ using the complete-information approach or the consumption-based optimization approach using the consumption proxy because it exposes recipients to greater downside risk. There is no effect on the estimate from the consumption-based optimization approach using the CEX consumption proxy because the consumption floor is not

[34] A long line of simulation literature uses a value of 3 for the coefficient of relative risk aversion (see, e.g., Hubbard, Skinner, and Zeldes (1995)[14], Mitchell, Poterba, Warshawsky, and Brown (1999)[18], Scholz, Seshadri, and Khitatrakun (2006) [21], Brown and Finkelstein(2008)[8], and Einav et al. (2010)[9]. Naturally, though, there are a range of plausible estimates; a substantial consumption literature, summarized in Laibson, Repetto, and Tobacman (1998)[16], has found risk aversion levels closer to 1, while other papers report higher levels of relative risk aversion (Barsky, Kimball, Juster, and Shapiro (1997)[6], Palumbo (1999)[19]).

binding for the baseline estimates for this approach. Lowering the consumption floor does not measurably affect our estimates of C or N, and the ratio of $\gamma(1)/C$ rises as a result.[35] In all approaches, higher risk aversion raises our estimates of $\gamma(1)$ and $\gamma(1)/C$, and lower risk aversion decreases it (as expected).

Alternative measure of out-of-pocket spending for those on Medicaid $(x(1, m))$**.** In our baseline analysis, we assume that, consistent with Medicaid rules, the insured have no out-of-pocket spending ($x(1, m) = 0$). In practice, however, the insured in our data report nontrivial out-of-pocket spending (Finkelstein et al. (2012)[11]).[36]

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

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

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

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

We evaluate this expression by taking the difference in the distributions of out-of-pocket expenditures of control compliers ($p(0)m(0;\theta)$) and treatment compliers ($p(1)m(1;\theta)$) and add to this the price faced by the uninsured times the difference in the distributions of medical spending of treatment compliers minus medical spending of control compliers ($p(0)\left(m(1;\theta) - m(0;\theta)\right)$). The price faced by the uninsured is calculated as the ratio of mean out-of-pocket spending to mean total spending for the control compliers. In the construction of differences in distributions, we assume quantile stability. In other words, we take the difference in distributions assuming an individual with a given θ that puts him at quantile r in the control complier distribution would have been at quantile r in the treatment complier distribution if he had been in the treatment group. In column VI, we present estimates from this alternative approach, in which we re-estimate all of our fitted consumption and out-of-pocket spending distributions based on self-reported out-of-pocket spending for treatment compliers as well as control compliers. In addition, we now have to estimate the "pure-insurance" term in equation (13) at $q = 1$, since we no longer assume full insurance at $q = 1$

[35] The results are opposite for an increase in the consumption floor, though now the consumption-based optimization approach using the CEX is affected (because the higher floor starts to bind for some CEX observations) and we now observe slight changes in C and N because the higher consumption floor is implemented by lowering the cap on out-of-pocket spending.

[36] This does not appear to be an artifact of our data or setting; in the Medical Expenditure Panel Survey, Medicaid recipients also self-report substantial out-of-pocket spending (Gross and Notowidigdo (2011)[13]).

as in the baseline analysis; our estimate of this term is not exact due to a technical complication relating to re-optimization in response to income effects.[37]

Allowing for $x(1, m) > 0$ necessarily reduces our estimates of $\gamma(1)$ but it also reduces our estimates of C (and hence N), so that the net effect on $\gamma(1)/C$ is a priori ambiguous. In practice, column VI shows that it substantially lowers our estimates of the willingness to pay for Medicaid relative to net costs.

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

Alternative interpolations in the optimization approach. In the baseline optimization approaches, we assumed $d\gamma/dq$ was linear in q to interpolate between $q = 0$ and $q = 1$ (see Assumption 4). Here, we explore the sensitivity of our results to two alternative interpolations. We cannot implement these alternatives for the approach that relies on the CEX consumption measure because they require we observe medical spending both for treatment and control compliers, which is only the case for our Oregon data.

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

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

given by: $\frac{E[m(1;\theta)]}{E[m(0;\theta)]}x(0;\theta)/p(0)$. Using the assumption that the demand for medical care is linear in q, we have:

$$m(q;\theta) = q\frac{E(m(1;\theta))}{E(m(0;\theta))}x(0;\theta)/p(0) + (1-q)x(0;\theta)/p(0). \tag{38}$$

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

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

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

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

We use equation (24) to infer the distribution of consumption from the distribution of out-of-pocket spending. From the distribution of consumption, we calculate the distribution of the marginal utility of consumption using $u_c = (c\,(q;\theta) - \gamma(q))^{-\sigma}$. We calculate the distribution of the marginal relaxation of the budget constraint, $-\partial x/\partial q = (p(1) - p(0))m(q;\theta)$, for each value of q by substituting in the expression for the demand of medical care (equation (38)) and noting that $p(1) = 0$. This yields:

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

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

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

and solve this differential equation using Picard's method (using 1000 iterations) to obtain $\gamma(1)$. Column VIII in Appendix Table 4 presents the results. We find that the interpolation based on a linear demand for m yields moderately lower estimates for $\gamma(1)$ (\$1283 instead of \$1421) and for $\gamma(1)/C$ (0.89 instead of 0.98).

The second alternative interpolation allows any functional form for the demand for medical care (rather than assuming it is linear in price) and finds the functional form that maximizes $\gamma(1)$. We allow for arbitrary (nonparametric) functional forms for the demand for medical care with the restriction that demand at values of $q \in (0,1)$ must lie somewhere between demand at $q = 0$ and at $q = 1$. Specifically, we define the distribution of medical care at insurance level q to be some linear combination of the distribution of medical care at $q = 0$ and at $q = 1$, where these distributions are given by equation (38). Formally, the distribution of medical care at insurance level q is given by $\hat{m}(\lambda(q);\theta) = \lambda m(0;\theta) + (1-\lambda)m(1;\theta)$ for some $\lambda(q) \in [0,1]$.

The distribution of out-of-pocket spending for each value of q and λ is given by $p(q)\hat{m}(\lambda(q);\theta) = (1-q)p(0)\hat{m}(\lambda(q);\theta)$. We use equation (24) to infer the distribution of consumption from the distribution of out-of-pocket spending; we denote the resulting consumption level by $\hat{c}(\lambda(q);\theta)$. From the distribution of consumption, we calculate the distribution of the marginal utility of consumption using $u_c = \left(\hat{c}\left(\lambda(q);\theta\right) - \gamma(q)\right)^{-\sigma}$. We calculate the distribution of the marginal relaxation of the budget constraint as $-\partial x/\partial q = p(0)\hat{m}(\lambda(q),\theta)$.

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

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

We solve this differential equation using Picard's method to find the upper bound for $\gamma(1)$. Column IX in Appendix Table 4 shows that the upper bound for $\gamma(1)$ is \$3065 and the upper bound for $\gamma\left(1\right)/C$ is 2.12. This latter estimate suggests that it is possible for the value of Medicaid to significantly exceed its net resource cost under alternative assumptions about the shape of demand and the utility function.

B.2 Alternative health values and measures

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

The marginal rate of substitution of health for consumption (ϕ). In column II, we assume the MRS of health for consumption is 0. This is motivated by the fact that while many measures of self-reported health improved, we are unable to reject the null of no impact of Medicaid on mortality (Finkelstein et al. (2012)[11]) or on our specific measures of physical health (Baicker et al. (2013)[5]). Therefore, an alternative of "no health benefits" seems a not unreasonable bound. Since the health component of the value of Medicaid was fairly small relative to the consumption component in the complete-information approach, this has a relatively small effect on the estimates.

Our baseline implementation assumed a MRS of health for consumption (ϕ) of $\$5,000$ for our low-income population. This came from scaling the "consensus" estimate of a VSLY of $\$100,000$ for the general population by the ratio of the marginal utility of consumption for our population and the general population. In column III, we instead assume that the MRS scales linearly with consumption; we therefore use a MRS of health for consumption of \$40,000 rather than \$5,000. As a result, the estimate for $\gamma(1)$ almost doubles while $\gamma(1)/C$ rises from 1.2 to 2.1. These are substantial changes that underscore the importance of the calibrated value of the MRS of health for consumption in a low-income population. While assuming that the MRS scales linearly with consumption is ad hoc and conceptually inconsistent with our assumption of $\sigma = 3$, it is closer than our baseline assumption to the findings of Kniessner et al. (2010)[15], who estimate an income elasticity of about 1.4 for this parameter.

As an alternative to using an external estimate of the MRS of health for consumption, we can also estimate the MRS of health for consumption from the first-order condition for m in equation (17). In other words, given our estimates of the return to medical spending ($d\tilde{h}/dm$) and the price p, we can estimate what u_h/u_c must be for the first-order condition to hold for the observed choices of m. Given that we express h in QALYs, the ratio u_h/u_c is the MRS of health for consumption. We find that a MRS of health for consumption of \$6,343 causes the first-order condition to hold on average for compliers. We note two caveats to this internal estimate. First, this estimate requires us to estimate $d\tilde{h}/dm$ for the entire sample by IV (using the lottery indicator as an instrument). While the first-stage of this IV regression is statistically significant, a t-statistic of 2.1 nevertheless indicates a weak first stage implying that the resulting IV estimate may not be reliable. Second, the internal estimate for the MRS is biased upward due to corner solutions because those who choose $m = 0$ place a lower value on a statistical life year than is implied by our estimate of u_h/u_c. Nevertheless, we found the finding that the internally-derived estimate of the MRS of health for consumption is of the same order of magnitude as our baseline assumption from external sources (\$5000) to be broadly reassuring.

Alternative health measures. In the remaining columns, we return to our baseline MRS of health for consumption (ϕ) of $\$5,000$ and explore robustness to alternative health measures. Our baseline measure was the five-point self-assessed health question. In prior work on the Oregon Health Insurance Experiment, Medicaid coverage was estimated to improve this measure of self-reported health, in addition to two other measures: the Patient Health Questionnaire (PHQ) depression screen, and the Short Form health questionnaire's (SF-8) measures of recent physical and mental health problems (Finkelstein et al. (2012)[11]; Baicker et al. (2013)[5]).

In column IV, we examine the 2-item Patient Health Questionnaire (PHQ-2) measure, which is commonly used as a depression screen. The PHQ-2 asks respondents about the prevalence in the last two weeks of having been "bothered by little interest or pleasure in doing things" and of having been "bothered by feeling down, depressed, or hopeless." We rely on the estimates from Pyne et al. (2009)[20], which are based on the "Standard Gamble" approach, to convert the PHQ-2 responses to QALYs; Appendix A.4 provided details of this conversion. We estimate that Medicaid increases health by 0.027 QALYs based on the PHQ-2 health measure, as compared to 0.045 QALYs under our baseline self-reported health measure. The estimates for the PHQ-2 health measure are correspondingly lower, which is not surprising given that the complete-information approach requires a comprehensive health measure whereas PHQ-2 only measures mental health.

We also draw on a separate data source - based on a series of in-person interviews conducted in the Portland metro area about two years after the lottery - that has additional health measures not available in our baseline data: the 8-item Patient Health Questionnaire (PHQ-8) and the "Short

Form" health survey (SF-8).[38] Columns V through VIII show the results. In column V, we replicate our baseline self-assessed health measure for the subsample of the in-person data that answered that question as well as all questions of the PHQ-8. Columns V and I show how results compare when we use the same self-assessed health measure in the mail-in survey and the in-person survey (which is limited to the Portland area). In the in-person data, we can measure health using the same PHQ-2 measure used in the mail survey (column VI) and the richer PHQ-8 measure (column VII). Estimates using the PHQ-8 measure are quite similar to those using the PHQ-2 measure (compare columns VI and VII).

We also use data from the in-person interviews to see how our estimates change if we use the SF-8. The SF-8 is a general health survey that captures both physical and mental health. We convert the SF-8 to QALYs using the mapping from Sullivan and Ghushchyan (2006)[23]. Unlike the previous mappings to QALYs we used, which were all based on the "Standard Gamble" method, this last mapping uses the other principal method in the literature: the "Time-Trade-Off" method; Appendix A.4 provided more detail on this method. Column VIII shows that results using the SF-8 measure are similar to results with the PHQ measures.

Appendix-only references

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

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

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

[4] Katherine Baicker, Amy Finkelstein, Jae Song, and Sarah Taubman. The impact of Medicaid on labor force activity and program participation: Evidence from the Oregon Health Insurance Experiment. *American Economic Review: Papers and Proceedings*, 104(5):322–328, 2014.

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

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

[38] The in-person interview data are available for fewer individuals than the mail survey, and are limited to individuals in the Portland metro area. However, in practice, empirical estimates from the two data sets are quite similar (see Finkelstein et al. (2012)[11] and Baicker et al. (2013)[5] respectively).

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

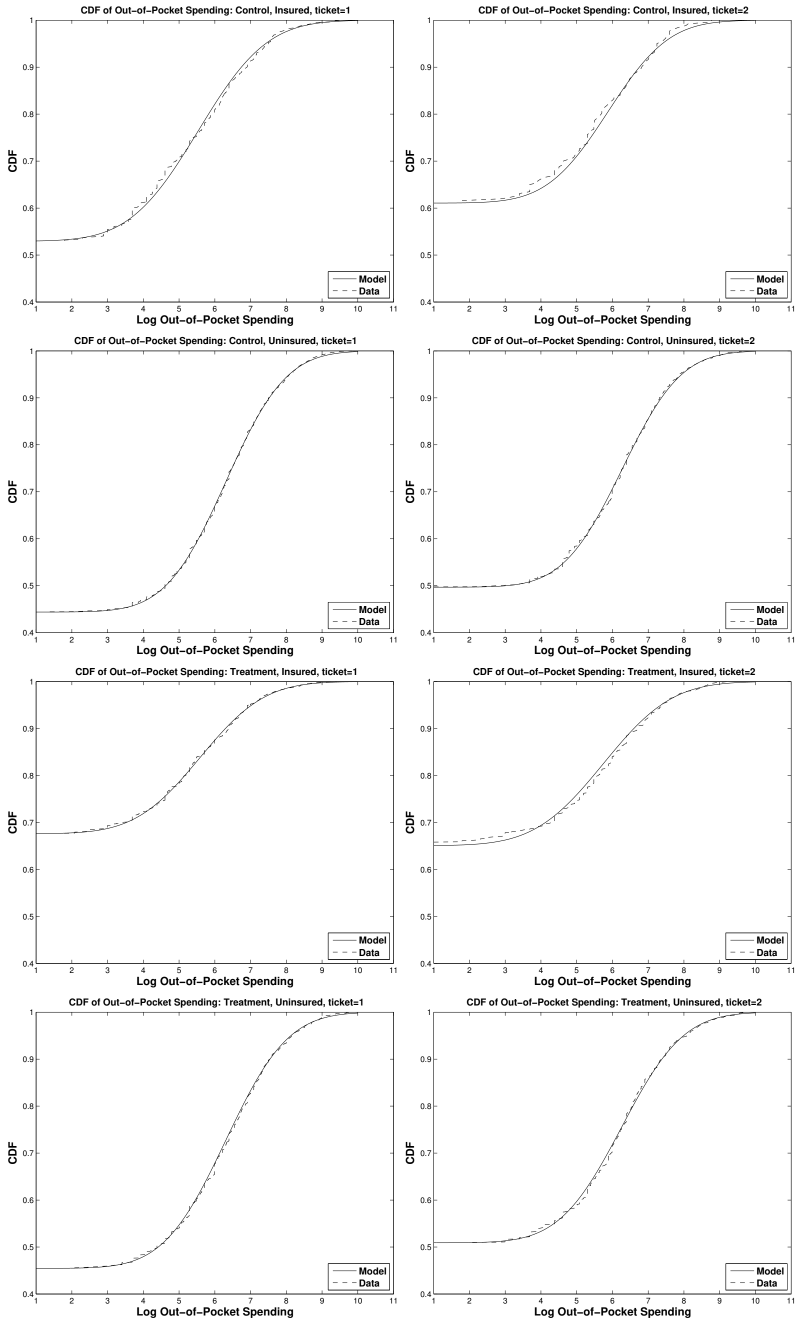

Appendix Table 1: Summary Statistics for the Oregon and CEX samples

	Oregon Sample (Control Compliers)	CEX Sample
Share female	0.60	0.54
Share age 50-64	0.35	0.46
Share age 19-49	0.65	0.54
Mean family size, n	2.91	1.00
Mean out-of-pocket spending, E[x]	569	395
Mean consumption, E[c]	9,214	13,542
Number of individuals	2,374	371

Notes: This table compares summary statistics among the uninsured that we analyze in two different data sets: the Oregon sample and the CEX. The statistics for the Oregon sample are for control compliers and are estimated using the IV techniques described in Appendix A.2.1. In the Oregon sample, we report, for mean non-medical consumption, the results from the consumption proxy approach (see equation (24)). The CEX sample is limited to single adults aged 19-64 without health insurance and living below the federal poverty line. Further details of the sample construction are in Appendix A.5.1. The sample size for the CEX in this table is smaller than the sample size for the CEX in Appendix Table 2 because the latter sample also includes individuals with health insurance. Mean consumption in the Oregon sample is based on CEX data; it is mean per capita non-medical consumption in families living below the federal poverty line and headed by an uninsured adult. Other details for the sample used to calculate mean consumption for the Oregon sample are identical to those described in Appendix A.5.1.

Appendix Table 2: Measurement of Consumption Covariance in CEX Consumption Approach

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

Notes: This table presents our baseline estimates for the pure-insurance term in the consumption-based optimization approach that uses the consumption measure from the Consumer Expenditure Survey. The sample consists of single adults aged 19-64 below 100% of the federal poverty line (N=1,065). The numbers reported in the table are the covariances of marginal utility of non-health consumption (using a coefficient of relative risk aversion of 3) and out-of-pocket medical spending, normalized by the mean value of the marginal utility of consumption for the relevant population; see equation (13). The consumption measure in the first column consists of all non-health expenditure in the CEX (excluding contributions to private and public pension programs), where we define health expenditure as individual expenditures for health care providers, prescription drugs, and medical devices. The consumption measure in column II is the same as that in column I but excludes expenditures on durables: vehicle purchases, major household appliances, house furnishings and equipment, and entertainment equipment (including TVs and radios). The consumption measure in column III consists of non-health expenditures in categories that can be relatively easily adjusted: food, entertainment, apparel, tobacco, alcohol, personal care, and reading.

Appendix Table 3: Comparison with Prior Estimates from Finkelstein et al. (QJE, 2012)

	I	II	III	IV
Sample	QJE Sample	Restricted Sample	Restricted Sample	Restricted Sample
Data	Raw Data	Raw Data	Raw Data	Adjusted Data
Estimation Method	QJE Paper	QJE Paper	This Paper	This Paper
First Stage: Lottery impact on Insurance				
Lottery indicator	0.290	0.290	0.302	0.302
(s.e.)	(0.007)	(0.008)	(0.007)	(0.007)
IV: Impact of Medicaid on…				
12-month medical spending ($), m	778	903	879	879
(s.e.)	(371)	(434)	(365)	(365)
12-month out-of-pocket spending ($), x	-244	-364	-350	-569
(s.e.)	(86)	(104)	(78)	(73)
Self-reported health binary indicator	0.133	0.103	0.141	0.141
(s.e.)	(0.026)	(0.032)	(0.028)	(0.028)
Number of observations	23,741	15,498	15,498	15,498

Notes: This table compares our baseline estimates of the impact of Medicaid with the baseline estimates of Finkelstein et al. (2012), which we refer to as "QJE." Self-reported health is a dummy variable that equals 1 if the individual reports being in good, very good, or excellent health. Column I replicates the QJE results. In column II, we use the same regressions as in column I but restrict the QJE sample to respondents living in households that have at most 2 lottery tickets and that have non-missing data on all the required variables (see Appendix A.2.1 for more details). In column III, we use the same sample as in column II but apply the regression approach of this paper (see Appendix A.2.1 for more details). In column IV, we use the estimation method and sample from this paper, applied to the "adjusted data" for out-of-pocket spending. "Adjusted data" refers to the out-of-pocket spending data after (i) estimating it by fitting a lognormal distribution with a mass point at zero for the distribution of out-of-pocket spending, (ii) adjusting the out-of-pocket spending of the insured to be 0, and (iii) imposing a ceiling on out-of-pocket spending for the uninsured such that consumption does not fall below the consumption floor (see Appendix A.2.2 for more details). All dollar amounts are per Medicaid recipient per year.

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

	I	II	III	IV	V	VI	VII	VIII	IX
		Coefficient of Relative Risk Aversion		Consumption Floor		Allow for out-of-pocket spending	Shock borne entirely by	Alternative Interpolations	
	Baseline	1	5	\$1,000	\$5,000	at $q=1$	individual	Linear demand for m	Upper-bound
Panel A. Recipient WTP for Medicaid, $\gamma(1)$									
Complete information	1675	886	4053	2810	856	738	2342	-	-
Consumption-based optimization, consumption proxy	1421	727	1658	1719	954	918	3131	1283	3065
Consumption-based optimization, CEX consumption measure	793	752	873	793	786	580	791	-	-
Panel B: Recipient WTP Relative to Net Costs, $\gamma(1)/C$									
Complete information	1.16	0.61	2.80	1.94	0.59	0.60	1.62	-	-
Consumption-based optimization, consumption proxy	0.98	0.50	1.15	1.19	0.66	0.74	2.17	0.89	2.12
Consumption-based optimization, CEX consumption measure	0.55	0.52	0.60	0.55	0.55	0.47	0.55	-	-
Panel C: Components of Recipient WTP									
Consumption-based optimization, consumption proxy									
Transfer component, T	661	661	661	661	654	447	658	661	749
Pure-insurance component, I	760	66	998	1058	301	470	2473	622	2316
Consumption-based optimization, CEX consumption measure									
Transfer component, T	661	661	661	661	654	447	658	-	-
Pure-insurance component, I	133	92	213	133	133	133	133	-	-
Panel D: Other Metrics									
Net costs, C	1448	1448	1448	1448	1442	1235	1446	1448	1448
Gross costs, G	3600	3600	3600	3600	3600	3387	3600	3600	3600
Monetary transfer to external parties, N	2152	2152	2152	2152	2158	2152	2154	2152	2152

Notes: This table presents the sensitivity of our baseline results to alternative assumptions unrelated to health valuations. Alternative assumptions about health specifications are reported in Appendix Table 5. Because the WTP estimates from the complete-information approach cannot be exactly decomposed into a transfer component and a pure-insurance component (only relatively wide ranges for these components can be provided), the table does not report a decomposition of the complete-information approach. Column I reports the baseline specification. We report the estimates under alternative assumptions about risk aversion (columns II-III) and the consumption floor (columns IV-V). We also consider specifications that allow for out-of-pocket spending to be positive for the insured (column VI) and assume that the health expenditure shock is borne solely by the individual instead of being shared equally within families (column VII). Columns VIII and IX report alternative interpolation assumptions, including linear demand (column VIII) and the upper bound procedure described in the text (column IX).

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

	I	II	III	IV	V	VI	VII	VIII
	Mail-In Survey				In-Person Survey			
MRS of QALYs for consumption, φ:	\$5K	\$0	\$40K	\$5K	\$5K	\$5K	\$5K	\$5K
Health measure:	SAH	SAH	SAH	PHQ-2	SAH	PHQ-2	PHQ-8	SF-8
Panel A. Recipient WTP for Medicaid, $\gamma(1)$								
Complete information	1675	1068	3039	1240	1326	1223	1168	1139
Panel B: Recipient WTP Relative to Net Costs, $\gamma(1)/C$								
Complete information	1.16	0.74	2.10	0.81	1.53	1.41	1.35	1.53
Panel C: Other Metrics								
Net costs, C	1448	1448	1448	1539	868	868	868	742
Gross costs, G	3600	3600	3600	3610	3033	3033	3033	2969
Monetary transfer to external parties, N	2152	2152	2152	2070	2165	2165	2165	2227
Number of observations	15,498	15,498	15,498	15,201	11,546	11,546	11,546	11,625

Notes: Because the WTP estimates from the complete-information approach cannot be exactly decomposed into a transfer component and a pure-insurance component (only relatively wide ranges for these components can be provided), the table does not report this decomposition. Consequently, the table does not report an estimate of the moral hazard cost, which depends on this decomposition. Column I reports our baseline estimates, while other columns report alternatives as indicated in the column headings. The MRS of Quality-Adjusted Life Years (QALYs) for consumption (φ) is the amount of consumption a low-income person is willing to forgo on the margin for a QALY. The health measure is the measure on which the QALY score is based. SAH denotes the "Self-Assessed Health" measure; PHQ denotes the "Patient Health Questionnaire" score, which measures mental health. PHQ-2 denotes a score based on two items from the PHQ where as PHQ-8 denotes a score based on 8 items of the PHQ. SF-8 denotes the 8-item "Short Form" health measure, which measures both physical and mental health. Additional information on these health measures and their mapping to QALYs is provided in Appendix A.4. In our baseline specification, we estimate that Medicaid increases QALYs by 0.045 (see Table 2). For QALYs measured by different health measures or in different samples, the estimated impact of Medicaid on QALYs becomes 0.027 in column IV, 0.049 in column V, 0.032 in column VI, 0.027 in column VII, and 0.027 in column VIII.

Appendix Table 6: Summary Statistics for PSID Sample

	PSID Sample
Share female	0.37
Share age 50-64	0.22
Share age 19-49	0.78
Mean family size, n	3.35
Mean out-of-pocket spending, $E[x]$	388
Mean consumption, $E[c]$	5,351
Mean hospitalization rate, $E[Z]$	0.11
Number of unique individuals	3,715
Number of person-year observations	6,600

Notes: This table presents summary statistics for the PSID sample used to estimate the relationship between out-of-pocket medical spending and consumption in Appendix Table 7. The sample consists of household heads aged 25-64 with non-missing reports for hospitalization and consumption data. We restrict the sample to households with per capita household income below $20,000. Consumption is per capita, expressed in dollars per year, and based on the consumption categories collected in the PSID, which are not comprehensive. The hospitalization rate is the fraction of individuals that were hospitalized in the past 12 months.

Appendix Table 7: Consumption-Based Optimization Approach using the PSID

	I OLS	II IV
Regressions		
Reduced-form regression of log(c^{-3}) on hospitalization indicator	0.18	-
(s.e.)	(0.06)	-
First-stage regression of log(x) on hospitalization indicator	0.19	-
(s.e.)	(0.07)	-
β, coefficient of regression of log(c^{-3}) on log(x)	-0.20	0.96
(s.e.)	(0.01)	(0.53)
Intermediate Steps for Covariance Calculation		
E[x] in Oregon control compliers	569	569
Std(x) in Oregon control compliers	543	543
Covariance term at q=0: β* Std(x)^2/E[x]	-105	495
(s.e.)	(7)	(277)
Covariance term at q=1 (by definition)	0	0
Implied Covariance Term (Averaged over *q*=0 and *q*=1)	**-52**	**248**
(s.e.)	(3)	(138)
Number of unique individuals	3,715	3,715
Number of person-year observations	6,600	6,600

Notes: This table presents the calculation of the pure-insurance term for the consumption-based optimization approach using data from the PSID. Column I presents the calculation based on an OLS regression of log(c^{-3}) on log(x). Column II presents the results from an IV strategy that uses hospitalization of the household head as an instrument for log out-of-pocket medical spending. Both regressions include a cubic in age, a quadratic function in household size, and year dummies as controls. Standard errors for the covariance term reflect sampling uncertainty in β but treat E[x] and Std(x) as non-stochastic. Standard errors are clustered by the individual household head.