MIT Open Access Articles

What Does (Formal) Health Insurance Do, and for Whom?

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

As Published: 10.1146/ANNUREV-ECONOMICS-080217-053608

Publisher: Annual Reviews

Persistent URL: <https://hdl.handle.net/1721.1/135798>

Version: Original manuscript: author's manuscript prior to formal peer review

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike

NBER WORKING PAPER SERIES

WHAT DOES (FORMAL) HEALTH INSURANCE DO, AND FOR WHOM?

Amy Finkelstein
Neale Mahoney
Matthew J. Notowidigdo

Working Paper 23718
<http://www.nber.org/papers/w23718>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2017

We are grateful to Sarah Abraham, Craig Garthwaite, Andrew Goodman-Bacon, Colin Gray, Nathan Hendren, Erzo Luttmer, Mark Shepard, Johannes Spinnewijn, Jonathan Skinner, and Carolyn Stein for helpful comments. We gratefully acknowledge financial support from the National Institute of Aging under grants RC2AGO36631 and R01AG0345151 (Finkelstein). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed a financial relationship of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w23718.ack>

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2017 by Amy Finkelstein, Neale Mahoney, and Matthew J. Notowidigdo. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

What Does (Formal) Health Insurance Do, and For Whom?
Amy Finkelstein, Neale Mahoney, and Matthew J. Notowidigdo
NBER Working Paper No. 23718
August 2017
JEL No. H22,H42,H51,I11,I13

ABSTRACT

Health insurance confers benefits to the previously uninsured, including improvements in health, reductions in out-of-pocket spending, and reduced medical debt. But because the nominally uninsured pay only a small share of their medical expenses, health insurance also provides substantial transfers to non-recipient parties who would otherwise bear the costs of providing uncompensated care to the uninsured. The prevalence of uncompensated care helps explain the limited take-up of heavily-subsidized public health insurance and the evidence that many recipients value formal health insurance at substantially less than the cost to insurers of providing that coverage. The distributional implications of public subsidies for health insurance depend critically on the ultimate economic incidence of the transfers they deliver to providers of uncompensated care.

Amy Finkelstein
Department of Economics, E52-442
MIT
77 Massachusetts Avenue
Cambridge, MA 02139
and NBER
afink@mit.edu

Matthew J. Notowidigdo
Northwestern University
Department of Economics
2001 Sheridan Road
Evanston, IL 60208-2600
and NBER
noto@northwestern.edu

Neale Mahoney
Booth School of Business
University of Chicago
5807 South Woodlawn Avenue
Chicago, IL 60637
and NBER
Neale.Mahoney@chicagobooth.edu

What Does (Formal) Health Insurance Do, and For Whom?

Amy Finkelstein, Neale Mahoney, and Matthew J. Notowidigdo*

August 2017

Abstract

Health insurance confers benefits to the previously uninsured, including improvements in health, reductions in out-of-pocket spending, and reduced medical debt. But because the nominally uninsured pay only a small share of their medical expenses, health insurance also provides substantial transfers to non-recipient parties who would otherwise bear the costs of providing uncompensated care to the uninsured. The prevalence of uncompensated care helps explain the limited take-up of heavily-subsidized public health insurance and the evidence that many recipients value formal health insurance at substantially less than the cost to insurers of providing that coverage. The distributional implications of public subsidies for health insurance depend critically on the ultimate economic incidence of the transfers they deliver to providers of uncompensated care.

1 Introduction

The U.S. government heavily subsidizes health insurance. In 2013, the tax exclusion for employer-provided health insurance was the single largest federal tax expenditure (\$250 billion), and Medicare (\$550 billion) was the second-largest line item in the federal budget, behind only Social Security (CBO 2013; National Center for Health Statistics 2015). The 2010 Affordable Care Act (ACA) greatly expanded publicly-subsidized health insurance for low-income adults, through both the heavily-subsidized health insurance exchanges and state Medicaid expansions. Spending on Medicaid (at \$550 billion in 2015) dwarfs the size of the next largest means-tested programs – food stamps (SNAP) and the Earned Income Tax Credit (at \$70 billion each).¹ In this essay we ask: what do publicly-subsidized health insurance coverage expansions do, and for whom?

Economics textbooks provide a clear answer: health insurance helps recipients by allowing risk-averse individuals to smooth their marginal utility of consumption in the face of unanticipated, out-of-pocket medical expenses. Politicians echo this goal in their public statements. In enacting

*MIT, Chicago Booth, and Northwestern. We are grateful to Sarah Abraham, Craig Garthwaite, Andrew Goodman-Bacon, Colin Gray, Nathan Hendren, Erzo Luttmer, Mark Shepard, Johannes Spinnewijn, Jonathan Skinner, and Carolyn Stein for helpful comments. We gratefully acknowledge financial support from the National Institute of Aging under grants RC2AGO36631 and R01AG0345151 (Finkelstein).

¹See U.S. Department of Health and Human Services (2015), U.S. Department of Agriculture (2016), and U.S. Internal Revenue Service (2015).

Medicare into law in 1965, President Johnson declared, “No longer will older Americans be denied the healing miracle of modern medicine. No longer will illness crush and destroy the savings that they have so carefully put away over a lifetime.” At the signing ceremony for the Affordable Care Act in 2010, President Obama similarly declared, “We have now just enshrined, as soon as I sign this bill, the core principle that everybody should have some basic security when it comes to their health care.”²

But the history and politics of formal health insurance in the United States also suggest that the *recipients* of insurance are not the only important *beneficiaries* of this insurance. Since its inception, health insurance has conferred benefits to healthcare providers as well as patients. The first wide-scale formal health insurance plans in the United States, the Blue Cross hospital insurance plans, were created during the Great Depression to provide financial help not only to patients but also to the hospitals that served them. As one hospital executive from the time recalled, “I could remember the difficulties we had then, trying to keep our doors open . . . People brought chickens in and meat to pay their bills. They would paint or do work around the hospital of some kind...” (Cunningham and Cunningham, 1997). In more recent times, hospitals have been an important lobbying force for Medicaid expansions under the Affordable Care Act, and against their subsequent repeal, arguing that increases in the number of uninsured patients would be financially devastating (e.g., Ollove 2013, Goldstein 2016).

This essay looks more rigorously at the claims behind these anecdotes. We show that the textbooks and the policymakers are certainly onto something: a compelling body of evidence shows that health insurance substantially reduces the risk to the recipient having large out-of-pocket expenditures; it also conveys other benefits to recipients, including improved health and reduced medical debt. But there is also growing evidence that a substantial share of the benefits from expanding formal coverage accrue to the healthcare providers and other parties who would otherwise bear the costs of providing uncompensated care to the uninsured. This is because the uninsured pay only a small fraction of their healthcare costs, on the order of one-fifth to one-third of their medical expenditures.

We then discuss some positive and normative implications. Substantial implicit insurance for the uninsured can explain low take-up of even heavily subsidized insurance on state and federal health insurance exchanges. It can also explain evidence showing that the low-income uninsured would be worse off if they had to finance the cost of their formal insurance through lump-sum taxes. Finally, it suggests that public subsidies for low-income individuals’ health insurance involve large transfers not only to the newly covered low-income individuals themselves, but also to those who previously bore the costs of their implicit insurance. The ultimate economic incidence of these transfers to non-recipients is an important, challenging, and as yet largely unanswered, question. But the possibility that these non-recipient beneficiaries may come from a very different part of the income distribution than the low income recipients raises important questions about the

²Johnson statement available here: <http://www.lbjlibrary.net/collections/selected-speeches/1965/07-30-1965.html>. Obama statement available here: <https://obamawhitehouse.archives.gov/photos-and-video/video/president-obama-signs-health-reform-law#transcript>

distributional implications of publicly-subsidized health insurance.

The rest of the essay proceeds as follows. In Section 2, we lay out a (fairly standard) stylized framework for analyzing health insurance, and define a number of empirical objects that will be the focus of the subsequent discussion. In Section 3, we discuss the descriptive evidence of the impact of health insurance on recipients and non-recipients. In Section 4, we discuss estimates of recipient willingness to pay for formal insurance, as well as positive and normative implications of these estimates. The final section offers some concluding thoughts and directions for further work.

2 A Framework for Analyzing Health Insurance

2.1 Set-up

We consider a two-period model. In the first period, a risk-averse, expected-utility-maximizing individual facing a stochastic distribution of health outcomes makes a decision of whether to purchase formal health insurance f or remain uninsured, which we denote by the insurance “contract” u . In the second period, the individual observes her realized health outcome and makes an optimal decision of how much medical care to consume and how much of the medical costs to pay out of pocket; the rest is left as unpaid debt. This is a standard framework (e.g., Cardon and Hendel 2001, Einav et al. 2013) with the added feature that, as in Dobkin et al. (2016), we let individuals choose not to pay some portion of their medical bills. We do this because unpaid medical bills are an empirically important phenomenon that is affected by insurance coverage. For simplicity, we write down the problem from the perspective of a representative individual. We briefly discuss individual heterogeneity at the end of this section.

We start by discussing the first period choice of insurance coverage and then turn to the second period utilization decision.

Insurance choice In the first period, the individual knows her income Y and the distribution of health shocks that she faces F_λ .³ The insurance contract j is characterized by an enrollee premium π_j and an out-of-pocket price p_j that the individual faces for medical spending. We let $p_j \geq 0$, with $p_j = 0$ denoting full coverage of medical spending and $p_j = 1$ denoting no coverage of medical spending under the “negotiated discounts” that characterize formal insurance coverage. We denote the amount the individual *owes* out of pocket by $p_j m_j$ and the amount she *pays* out of pocket by $x_j = p_j m_j - d_j$, where d_j is the amount of unpaid medical debt she chooses not to pay today (although she may end up repaying some of this debt in the future as debt collectors try to recover the unpaid medical debt).

We consider the choice between two contracts: (p_f, π_f) and (p_u, π_u) . It is natural to normalize $\pi_u = 0$, although we can think of a health insurance mandate as setting $\pi_u \geq 0$. Importantly, we

³It is standard to assume that the arrival of health shocks F_λ are unaffected by insurance. However, this ignores the possibility that insurance coverage might discourage the individual from undertaking non-medical investments in health (Ehrlich and Becker 1972).

allow for $p_u < 1$, to reflect the potential presence of “implicit insurance” for the uninsured.⁴ We also allow for $p_f > 0$, to reflect potential consumer cost-sharing in the formal insurance contract.⁵

Utilization choice Let $v(h, c, d)$ indicate von-Neuman Morgenstern utility, which we assume is increasing in health h , increasing in non-medical consumption c , and decreasing in medical debt d . In the second period, the individual receives health shock λ and chooses how much medical care $m \geq 0$ to consume and how much of the medical expenses to pay out of pocket ($p_j m$) rather than leave as unpaid debt d_j with $0 \leq d_j \leq p_j m$.

The individual chooses m , c , and d to maximize her second-period utility, subject to the health care production function $h(m - \lambda)$ and her budget constraint that consumption c must be weakly less than income Y net of out of pocket payments x_j and premium π_j . This yields the second period optimization problem:

$$\max_{m, c, d} v(h, c, d) = v(h(m - \lambda), c, d) \quad s.t. \quad c_j \leq Y - x_j(m, d) - \pi_j \quad (1)$$

A key object in the second period decision is the health production function $h()$ that governs how m and λ translate into h . We assume that $h()$ is increasing in $(m - \lambda)$ but with decreasing returns. As a result, individuals may choose not to fully restore their health and/or the health production function may be such that some health shocks cannot be fully undone by medical spending.

We assume that individuals are risk averse, so that $v(h, c, d)$ is concave in c . It is the concavity of the utility function that creates the economic rationale of insurance: risk-averse individuals would prefer to pay the expected cost of their medical care expenses than pay the realized cost m_j . The shape of the utility function over h and d , as well as how any interaction among the three arguments is not standardized in the literature.

We are agnostic about the underlying source of the disutility from medical debt d . Lower consumption in future (un-modeled) periods in which the debt is (partially or fully) repaid is one natural source. Medical debt may also raise borrowing costs and hence the costs of smoothing consumption inter-temporally. Finally, medical debt may exert direct negative effects on individual utility through its effects on psychic stress, guilt, or “peace of mind”, including perhaps the anticipated psychic costs of eventual default.

The trade-off the individual with $p_j > 0$ faces is that increased m produces higher h but lower non-medical consumption c and/or higher debt d . The insurance contract $j \in \{f, u\}$ determines the parameters of this trade-off by setting the price of medical care to the individual p_j .⁶

⁴Our framework also allows $p_u > 1$ to capture a scenario where the uninsured pay more out-of-pocket than the combined insurance and out-of-pocket payments by the insured. This would occur if the uninsured end up paying more (net of implicit insurance) than the “negotiated discounts” insurers obtain through bargaining.

⁵Characterizing insurance contracts with a single price is a simplification. In reality, the insured face non-linear contracts, with deductibles and out-of-pocket maximums, along with cost-sharing exemptions for preventative services. The uninsured face different degrees of implicit insurance based on whether the care they seek is emergency or routine, the charity care preferences of the health care provider, and so on.

⁶It is standard – but not entirely inconsequential – to model formal insurance as affecting only the price the individual faces for medical care p_j . In practice, of course, formal insurance may further affect the recipient by,

We denote the individual's second-period utility-maximizing choices of medical spending and debt as m_j^* and d_j^* respectively. The individual's resultant non-medical consumption under contract j is therefore given by the budget constraint: $c_j^* = Y - x_j(m_j^*, d_j^*) - \pi_j$. We denote the (maximized) second-period utility by v_j^* .

Expected utility from contract j is given by $V_j^* \equiv \int v_j^* dF_\lambda(\lambda)$. In the first period, the individual chooses the insurance contract that maximizes her expected utility – i.e., $\max_{j \in \{f, u\}} V_j^*$.

2.2 Empirical Objects of Interest

Impact of health insurance on recipients In Section 3.1, we will briefly describe (some of) the extensive reduced form evidence on the impact of formal insurance on each of the three arguments of the utility function: h , c , and d . We will also briefly discuss the evidence of the impact of insurance on medical spending m . We will use Δ to denote the differences between formal health insurance f and the uninsured contract u , and capital letters to denote expectations; thus, we will examine $\Delta H \equiv E_\lambda[h_f - h_u]$, $\Delta C \equiv E_\lambda[c_f - c_u]$, $\Delta D \equiv E_\lambda[d_f - d_u]$, and $\Delta M \equiv E_\lambda[m_f - m_u]$.

There may be additional (un-modeled) arguments of the utility function that are affected by formal health insurance, such as impacts on family members, or “peace of mind”; these could have first order effects on recipient welfare. Formal health insurance may also affect the budget constraint in equation (1). The expansion of formal health insurance is usually associated with a lowering of its price, perhaps from infinity (or some very high price) to a lower amount (such as a price of zero for a free public insurance program). This in turn can affect other un-modeled decisions such as human capital investments or labor market participation. As long as individuals are making optimal choices, however, responses to changes in the budget constraint will have only second-order impacts on recipient welfare.

Impact of health insurance on non-recipients In Section 3.2, we will describe reduced form evidence of the impact of formal health insurance coverage on non-recipients. This has received comparatively less empirical attention than the venerable and extensive literature on the impact of health insurance on recipients. But it is an active and growing area of work with important implications.

We begin by documenting the substantial amount of implicit insurance payments received by the uninsured. We denote the implicit insurance payment for an uninsured individual by I_u :

$$I_u \equiv (1 - p_u)m_u + (1 - \delta)d_u, \quad (2)$$

where δ is the expected recovery rate on the (initial) choice of the uninsured's medical debt d_u .⁷

for example, affecting access to particular networks of medical providers. These could be modeled at the cost of introducing additional notation and complexity. It would also presumably introduce additional empirical objects of interest. For example, Garthwaite et al. (2017) present evidence that the ACA Medicaid expansion increased access to a broader range of hospital facilities for the newly insured.

⁷To see the derivation of equation (2), note that I_u can be written as $I_u = m_u - x_u - \delta d_u$, where $x_u = p_u m_u - d_u$ reflects the (initial) choice of out-of-pocket spending and medical debt and δd_u reflects the subsequent recovery of

Implicit insurance payments for the uninsured are thus the difference between their medical expenditures (m_u) and the amount they pay out of pocket, either up front or by eventually paying down their debt $p_u m_u - (1 - \delta)d_u$. Note the two components of implicit insurance: ex-ante charity care $(1 - p_u)m_u$ and ex-post bad debt $(1 - \delta)d_u$.

We assume that the insured receive no ex-ante charity care but may also have ex-post bad debt $(1 - \delta)d_f$.⁸ Implicit insurance payments for an insured individual is therefore $I_f = (1 - \delta)d_f$. Formal insurance therefore affects (and likely reduces) the amount of payments from informal insurers by

$$\Delta I \equiv E_\lambda[I_u - I_f] = E_\lambda[(1 - p_u)m_u] - (1 - \delta)\Delta D. \quad (3)$$

There are two components of the term ΔI . The first term ($E_\lambda[(1 - p_u)m_u]$) is the “monetary” value of the expected charity care provided to the uninsured, such as free care provided by non-profit hospitals or public health clinics. Note that the “cost” of the charity care may be lower than the monetary value $E_\lambda[(1 - p_u)m_u]$. To the extent that the price of medical care reflects a markup, part of $E_\lambda[(1 - p_u)m_u]$ is a reduction in profits that should not be interpreted as charity care “costs.” The second term $-(1 - \delta)\Delta D$ is the expected decline in bad debt due to formal insurance coverage.

Recipient willingness to pay In Section 4.1 we discuss estimates of recipient willingness to pay for formal insurance, γ . We define γ as the implicit solution to:

$$\int v(h(m_u^* - \lambda), Y - x_u^*, d^*)dF_\lambda = \int v(h(m_f^* - \lambda), Y - x_f^* - \gamma, d^*)dF_\lambda. \quad (4)$$

In other words, γ denotes the maximum amount an individual could pay that would leave her expected utility under formal insurance (the right hand side of equation (4)) equal to her expected utility when uninsured (the left hand side of equation (4)). By definition, for the marginal enrollee, γ is equal to the insurance premium π_f . For infra-marginal enrollees, $\gamma > \pi_f$, reflecting the consumer surplus from formal insurance.

Estimates of γ can be used for positive analysis of take-up of formal health insurance under alternative public subsidy schemes. If one is willing to interpret the estimates of γ normatively, they can also be used for welfare analysis of alternative subsidy schemes.

It will be useful to benchmark γ against the (gross) cost of formal insurance G . We define G to be the expected medical claim payments by the formal insurer:

$$G \equiv (1 - p_f)E_\lambda[m_f]. \quad (5)$$

This definition of insurance costs abstracts from any administrative costs of providing insurance. Actuarially fair insurance has $\pi_f = G$, while subsidized insurance has $\pi_f < G$.

It is useful to rewrite the expression for G to emphasize the mechanical and behavioral compo-

part of the unpaid medical debt.

⁸This assumes that the recovery rate δ is the same for debt incurred by formally insured and uninsured patients, which could of course be easily relaxed at the cost of more notational complexity.

nents:

$$G = (1 - p_f)E_\lambda[m_u] + (1 - p_f)E_\lambda[m_f - m_u].$$

The mechanical component $(1 - p_f)E_\lambda[m_u]$ represents the costs to the insurer of paying for the medical spending the individual consumed when uninsured. The behavioral component $(1 - p_f)E_\lambda[m_f - m_u]$ represents the costs to the insurer from the increase in medical spending due to insurance; if insurance is full ($p_f = 0$) then this behavioral component is simply ΔM . This behavioral component is commonly referred to as “moral hazard” (Pauly 1968).⁹ When the government finances the insurance, this behavioral response represents the “fiscal externality” of the insurance on the government (Hendren 2016b). Any re-distributional impacts of the publicly financed insurance can generate additional fiscal externalities beyond medical spending; we abstract from this for now but return to it in Section 4.2 below.

Finally, it will also be useful to benchmark γ against the *net* cost of formal insurance C :

$$C = G - \Delta I. \tag{6}$$

which measures the gross cost of formal insurance net of any transfer to non-recipients ΔI .

The discussion has abstracted from individual heterogeneity. In practice, individuals are likely to be heterogeneous on a number of relevant dimensions. These include heterogeneity in the utility function (v), the distribution of potential health risks F_λ , the health care production function $h(m - \lambda)$, and income Y . Such heterogeneity is important to bear in mind in our discussion of the descriptive evidence of the impact of health insurance in Section 3. These estimates all reflect the impacts of insurance for people whose coverage is affected by the various policies studied; impacts on other populations could differ. It will also create heterogeneity in willingness to pay for insurance (γ) and potentially affect the normative implications, both of which we discuss in Section 4.

3 Descriptive Evidence of The Impact of Health Insurance

3.1 The Impact of Health Insurance On Recipients

A large body of empirical work has produced evidence on the effects of health insurance on the elements in our framework: medical spending (ΔM), health (ΔH), consumption (ΔC), and medical debt (ΔD). We make no attempt to comprehensively review this voluminous – and ever-growing – literature. Instead, we briefly summarize our reading of what we know (and do not know) from this literature.

⁹Note that in the canonical set-up, it is assumed that the choice of m is socially optimal absent formal health insurance. In practice, there are a number of reasons to think that, absent insurance, the choice of m may not be socially optimal. First, if health care is not priced at social cost, individual consumption choices under $p_u = 1$ may not be optimal. For example, the patent system creates mark ups on the price of prescription drugs above the (static) social marginal cost, so that uninsured individuals presumably consume sub-optimally low levels of prescription drugs (Lakdawalla and Sood, 2009). Second, individual healthcare choices may not be privately optimal; liquidity constraints and “behavioral biases” may inhibit private optimization (Baicker et al. 2015).

Much of the work we describe uses quasi-experimental variation induced by public health insurance expansions. These include the introduction of Medicare coverage for the elderly both in calendar time (1965) and over the life-cycle (at age 65), the staggered introduction of Medicaid coverage by states in the 1960s, and the staggered expansions by states in Medicaid coverage to low-income women and children in the 1980s and 1990s.¹⁰ More recently, researchers have examined the expansion in health insurance in Massachusetts for prime age adults and children that occurred under the 2006 mandate, and the 2014 Medicaid expansions that occurred in some states under the Affordable Care Act.¹¹

We also discuss, where relevant, results from three randomized evaluations of health insurance coverage in the United States, all of which focus on non-elderly populations, primarily adults. The 1970s RAND Health Insurance Experiment randomly assigned different cost sharing provisions (p_j in our above framework) across about 2,000 non-elderly families for three to five years (Newhouse et al., 1993; Aron-Dine et al., 2013). The 2008 Oregon Health Insurance Experiment randomly assigned Medicaid for about two years to about 10,000 uninsured adults below the federal poverty line (Finkelstein et al., 2012; Baicker et al., 2013; Baicker et al., 2014; Taubman et al., 2014, Finkelstein et al., 2016).¹² The 2007-2009 Accelerated Benefits Demonstration Project randomly assigned public health insurance to about 1,000 uninsured adults on Social Security Disability Insurance during their two-year waiting period for Medicare (Michalopoulos et al., 2011).

Medical spending (ΔM)

An enormous literature has examined the moral hazard effects of health insurance (i.e., ΔM). There is compelling evidence rejecting the null hypothesis that medical utilization is not affected by insurance (i.e., that demand for medical care is completely inelastic). All three randomized evaluations show that insurance coverage increases medical utilization, as do a large number of quasi-experimental studies (e.g., Card et al. 2008, Anderson et al. 2012).

A more challenging problem has been how to model the medical spending response under alternative health insurance contracts not observed in the data. In our stylized framework, health insurance contract j determines a (constant) price of medical utilization, p_j . However, most real-world health insurance contracts involve a non-linear mapping from medical utilization during the year to the amount owed out of pocket. The researcher must therefore model how an individual's healthcare choice responds dynamically as health shocks arrive sequentially and the current

¹⁰For the introduction of Medicare see, e.g., Finkelstein (2007) and Finkelstein and McKnight (2008). For the impact of Medicare at age 65 see, e.g., Card et al. (2008), Card et al. (2009), and Barcellos and Jacobson (2015). For the introduction of Medicaid see, e.g., Goodman-Bacon (forthcoming). For expansions of Medicaid coverage see, e.g., Currie and Gruber (1996), LoSasso and Buchmueller (2004), Wherry and Meyer (2016), and Wherry et al. (forthcoming).

¹¹For impacts of the Massachusetts expansions see, e.g., Simon (2005), Kolstad and Kowalski (2012), and Mazumder and Miller (2016). For some early evidence of the effect of the 2014 Medicaid expansions, see, e.g., Hu et al. (2016) and Miller and Wherry (2016, 2017).

¹²For an overview of the results of the Oregon Health Insurance Experiment see https://www.povertyactionlab.org/sites/default/files/publications/Insuring_the_Uninsured.pdf. For more details see www.nber.org/oregon.

and future price of medical care to the consumer evolve. Einav and Finkelstein (2017) provide a discussion of some of these issues.

Health (ΔH)

With the exception of mortality, data on objective, physical health are scarce. As a result, the quasi-experimental literature has focused almost exclusively on the impact of health insurance on mortality and, relatedly, on high-mortality populations such as infants and the elderly. The randomized evaluations – which focused on non-elderly adults – examined non-mortality health measures, including self-reported health, depression, and physical health measures such as blood pressure and cholesterol collected specifically in the course of the experiment.

Once again, the empirical evidence compellingly rejects the null hypothesis that health insurance does not affect health. For example, Card et al. (2009) document a reduction in mortality for non-discretionary hospitalizations associated with the discrete increase in insurance coverage that occurs with Medicare eligibility at age 65.¹³ Results from the randomized evaluations indicate that health insurance improves self-reported health (Finkelstein et al. 2012, Michalopoulos et al 2011) and reduces depression (Finkelstein et al. 2012).

However, a comprehensive assessment of the health impact of health insurance remains elusive. The set of potential clinical, non-mortality health benefits is large, and only a subset of them have proven measurable. Where clinical health measures are available – which primarily has been in the context of randomized evaluations – results have not been encouraging. The RAND Health Insurance Experiment found little-to-no physical health benefits from reduced consumer cost-sharing, with the possible exception of reductions in hypertension (Newhouse et al. 1993). The Oregon Health Insurance Experiment found no evidence of improvements in measured physical health – including blood pressure, blood sugar, or cholesterol (Baicker et al. 2013). In some cases, such as blood sugar, the estimates were too imprecise to rule out meaningful effects; in other cases, such as blood pressure, the experiment rejected the beneficial effects found in the previous, quasi-experimental literature (Baicker and Finkelstein 2013).

Another challenge for estimating health impacts is that health is most naturally modeled as a stock; for example Grossman (1972) models the individual’s “health capital” as evolving over time in response to various medical and behavioral investments. As a result, the long-run impacts of health insurance may differ from the short-run impacts. Longer-run impacts are more difficult to estimate credibly. Perhaps the best hope is from randomized evaluations, but these experiments have been limited in duration: about two years in the case of the Oregon Experiment and the Accelerated Benefits Demonstration Project, and 3-5 years in the case of the RAND experiment.

Overall, the evidence to date suggests that health insurance is an element – but certainly not the only and perhaps not the most important element – in the health production function. This

¹³According to the authors, the reduction in mortality at age 65 in their sample likely reflects a combination of the discrete increase in coverage for the previously uninsured individuals as well as the change from private insurance to Medicare for those with private insurance before age 65.

is consistent with studies examining the main correlates of cross-sectional differences in population health, which have emphasized an important role for non-medical factors - and particularly health behaviors - in determining population health. Early work by Fuchs (1974) famously compared the health of the populations in two neighboring states (Utah and Nevada) with similar levels of income and medical care but enormous differences in health and health behaviors. More recent evidence on the cross-sectional correlates of mortality across different areas of the United States has similarly suggested a potentially important role for health behaviors (Chetty et al. 2016).

Consumption (ΔC)

Compared to the tradition of studying the impact of health insurance on health, studies of the impact of health insurance on consumption are relatively recent. The famous RAND Health Insurance Experiment did not consider impacts on consumption or economic security. To our knowledge, the quasi-experimental empirical work related to these outcomes dates back only to Finkelstein and McKnight (2008).

This literature has focused primarily on measuring the impact of health insurance on out-of-pocket medical expenditures (e.g. Finkelstein and McKnight, 2008; Engelhard and Gruber, 2011; Finkelstein et al. 2012; Baicker et al., 2013; Michalopoulos et al., 2011; Barcellos and Jacobson 2015). Both the quasi-experimental evidence and results from randomized evaluations provide compelling evidence against the null hypothesis that health insurance has no effect on out-of-pocket medical spending. Moreover, the results suggest reductions in right-tail out of pocket expenditures, of the kind that might portend large consumption declines that could be particularly costly to risk-averse individuals. For example, the Oregon Health Insurance Experiment found that Medicaid virtually eliminated the possibility of catastrophic out-of-pocket medical expenditures, defined as 30 percent or more of household income (Finkelstein et al., 2012; Baicker et al., 2013).

However, evidence of the impact of health insurance on out-of-pocket spending likely gives a very incomplete picture of the impact of health insurance on *consumption* for two reasons. First, consumption will fall less than out-of-pocket spending if households are able to borrow to smooth out expenditure shocks. Second, adverse health events reduce labor supply and future income (Dobkin et al., 2016) and insurance could have an impact on these effects.

Thus, while the evidence of declines in out of pocket spending associated with health insurance coverage is evidence against the null hypothesis that health insurance has no direct impact on consumption, the magnitude of the consumption effects of health insurance remains uncertain. We know of no existing empirical work directly measuring the impact of health insurance on consumption in the United States. Obtaining high-quality consumption data at scale is a pervasive problem for empirical work on many topics, and the impact of health insurance is no exception. Better evidence on the impact of health insurance on consumption likely awaits progress on obtaining high-quality consumption data for the relevant populations.

Medical debt (ΔD)

Even more recently, there has been growing interest in the impact of health insurance on consumer medical debt. Medical debt is substantial: the Consumer Financial Protection Bureau (CFPB) found in 2014 that unpaid medical bills constitute more than half of *all* collections lines on consumers credit reports. Based on reports from the CFPB and the Federal Trade Commission (FTC), we estimate the total value of overdue medical debt to be roughly \$75 billion, spread across 43 million Americans (FTC 2012; CFPB 2014).¹⁴ This estimate of the “stock” of overdue medical debt seems broadly consistent with existing estimates of the “flow” of unpaid hospital debt, which is roughly \$20 billion a year according to Dranove et al. (2015).

The literature has typically measured medical debt through self-reports, credit reports of unpaid medical bills that were sold to collection agencies (“medical collections”), and bankruptcy (arguably the ultimate non-payment of bills). Medical debt is, not surprisingly, disproportionately concentrated among the uninsured. For example, Dobkin et al. (2016) estimate that in California, the average (non-pregnancy) hospital admission for a prime age adult results, four years later, in \$300 in medical collections for an individual with health insurance, but over \$6,000 in medical collections for an individual without health insurance. Importantly, the vast majority of medical collection balances remain unpaid (Avery et al. 2003).

Once again, the empirical literature rejects the null hypothesis that health insurance does not affect medical debt. The results from randomized evaluations indicate that health insurance reduces medical collections and self-reported medical debt (Finkelstein et al., 2012; Michalopoulos et al. 2011). The quasi-experimental estimates likewise point to declines in medical debt from health insurance expansions (Barcellos and Jacobson 2015, Hu et al., 2016, Mazumder and Miller 2016), as well as reductions in bankruptcy rates from health insurance expansions (Gross and Notowidigdo 2011, Mazumder and Miller 2016).¹⁵

Reductions in medical debt tend to be heralded – in the media and in academic work (including our own) – as benefits to the newly insured. Empirical evidence of these benefits, however, remains elusive. As noted, most medical debt remains unpaid; therefore, reductions in this debt are unlikely to substantially increase future consumption. Dobkin et al. (2016) find that hospital admissions for the uninsured generate substantial medical debt but have no meaningful impact on credit scores,

¹⁴Our estimate is based on two different calculations. The first uses the 2014 CFPB report on medical debt, which finds that 19.4% of 220 million credit reports had medical collections; these are unpaid medical bills sold to collection agencies and represent only a fraction of medical debt. For those credit reports with only medical collections, the average amount of overdue medical debt was \$1,766. For those with both medical and non-medical collections (unpaid bills of other kinds sent to collection agencies), the average amount of overdue debt was \$5,638. Assuming the same amount of overdue medical debt in this latter population (i.e., that \$1,766 of \$5,638 is medical), this gives us a total amount of overdue medical debt of \$75.9 billion (43 million $\times$ \$1,766). Alternatively, we can use information from the FTC report on (unconditional) average number of collections, which is 1.13 per credit report. The CFPB report notes that 52 percent of all collections are medical, and the average medical collection was \$579 (many individuals have multiple medical collections). This gives us a total value of overdue medical debt of \$74.8 billion (220 million $\times$ 1.13 $\times$.52 $\times$ \$579).

¹⁵The inability to reject the null of no impact of health insurance on bankruptcy in the Oregon experiment (Finkelstein et al. 2012) likely reflects its low power to detect an effect on this rare outcome, relative to the quasi-experimental literature.

and hence presumably little impact on borrowing costs. Less tangible benefits from reductions in medical debt – such as increased “peace of mind” or reduced “guilt and stress” – could be important, but so far have not been measured in quasi-experimental studies. Such “intangible” benefits do appear in a range of interviews with industry professionals (e.g. debtor and creditor attorneys, and judges) discussed in Mann and Porter (2010).

In addition to any benefits to recipients, reduced medical debt presumably benefits whomever was the ultimate holder of this debt. This group usually receives less attention, but it is the one to which we now turn.

3.2 The Impact of Health Insurance on Non-Recipients (ΔI)

The impact of (formal) health insurance on the amount of payments from informal insurance is given by the reduction in payments from informal insurers: $\Delta I \equiv E_\lambda[I_u - I_f]$ (see equation (3)). We focus our discussion on implicit insurance payments for the uninsured (I_u). We are unaware of any research systematically quantifying implicit insurance payments to the insured (I_f). To the extent that I_f is small, we can approximate the effect of insurance coverage on non-recipients by the implicit insurance payments for the uninsured: $\Delta I \approx I_u$. This seems a reasonable approximation, given the evidence in Dobkin et al. (2016) discussed above that hospital admissions have a relatively small impact on medical debt for those with health insurance.

The uninsured receive substantial implicit insurance. Mahoney (2015, Figure 1A, reproduced below), shows how payments by the privately insured and the uninsured scale with medical charges. Medical charges are the “list” price of medical care and proxy for the quantity of medical care received. Payments by the privately insured scale up proportionally with charges. By contrast, payments by the uninsured scale up at the same rate as the privately insured until about \$2,000 in charges, and then flatten out abruptly. This suggests that the uninsured have substantial implicit insurance against large health shocks.

Other studies corroborate this finding of substantial implicit insurance for the uninsured. The Kaiser Commission on Medicaid and the Uninsured estimates that the average uninsured person paid \$500 out of pocket but incurred total medical expenses of \$2,443 (Coughlin et al. 2014, Figure 1), suggesting that on average the uninsured pay only 20 percent of their total medical expenses. Finkelstein, Hendren and Luttmer (2015) estimate in the Medical Expenditure Panel Survey that the average uninsured adult below 100 percent of the federal poverty line (who are primary recipients of Medicaid expansions under the Affordable Care Act) pays about 33 percent of their medical expenses out of pocket.¹⁶ Hadley et al. (2008) estimate that out-of-pocket payments by the uninsured are about 35 percent of what a provider would have received from a fully insured patient for the same medical care.

This implicit insurance arises because of a host of factors. Uninsured patients with limited ability-to-pay likely receive medical care both because of a sense of ethical obligation felt by the

¹⁶Importantly, both these estimates consider out of pocket payments relative to an estimate of what total medical expenditures based on transacted, rather than list, prices.

medical profession and because of legal requirements. The federal Emergency Medical Treatment and Labor Act (EMTALA) requires hospitals to provide *emergency* care on credit and prohibits them from delaying treatment to inquire about insurance status or means of payment. As a matter of practice, many hospitals report providing non-emergency medical care on credit as well. For instance, an Internal Revenue Service (IRS) survey (IRS 2007) of nonprofit hospitals found that 90 percent reported never denying any medical services to uninsured patients.

Medical care for uninsured patients could, in principle, be billed to the uninsured ex-post. In practice, however, much of it is provided by ex-ante charity care. Some hospitals, particularly those with religious affiliations, see the provision of charity care as central to their institutional mission. Moreover, non-profit hospitals – which constitute about 70 percent of all hospitals – are required to provide a “community benefit” in exchange for federal, state, and local tax exemptions; charity care, along with medical research and teaching, is one way that hospitals can fulfill this requirement (Nicholson et al. 2000; GAO 2008). In addition, a number of states have charity care pools that redistribute funding to hospitals based on the volume of uncompensated care (Dranove et al. 2015).

Finally, even when medical providers wish to seek payment for medical services, a number of factors limit their recovery rates, generating ex-post charity care -- i.e., bad debt. The uninsured are disproportionately low income and many have very few assets. Debt collectors that contract with medical providers have limited leverage, because most uninsured individuals can shield virtually all of the assets in personal bankruptcy (Mahoney 2015). For these reasons, medical providers typically recover only about 10–20 percent of bills submitted to uninsured individuals (LeCuyer and Singhal 2007).

We return in Section 4.3 below to consider what we know about the direct and ultimate economic incidence of this implicit insurance for the uninsured.

4 Willingness to Pay and its Implications

The previous section described the types of benefits that subsidized health insurance expansions confer on recipients and providers. We now turn to evidence on whether recipients would be willing to pay for these benefits. A number of studies suggest that recipients would *not* be willing to fund the relevant health insurance expansions, and that this can likely be explained by the prevalence of uncompensated care if they remain uninsured. We then discuss the positive and normative implications of this finding.

4.1 Recipient Willingness to Pay (γ)

Approaches to estimating γ

Conceptually, estimating willingness to pay for formal health insurance (γ) is straightforward: Equation (4) defines γ as an object directly revealed by the demand curve for formal insurance. If one were to randomly vary the premium at which formal insurance were offered (π_f), the share of

individuals who enrolled when the premium decreased from π_f to $\pi_f - d\pi_f$ have willingness to pay (γ) between π_f and $\pi_f - d\pi_f$.

Historically, however, such an approach has faced considerable practical challenges. Until recently, most of the uninsured were either not offered health insurance, or faced prices that were difficult to measure, while the publicly insured faced either no price or a uniform price for their coverage. In a creative effort to surmount this substantial obstacle, Krueger and Kuziemko (2013) conducted a survey experiment designed to elicit willingness to pay for hypothetical plan offerings among a large sample of the uninsured.

In another attempt, Finkelstein, Hendren, and Luttmer (2015) show how one can use descriptive evidence on the impacts of Medicaid on various outcomes to infer willingness-to-pay for a health insurance product, such as Medicaid for most populations, that is not traded in a marketplace. They specify a normative utility function v over c and h (ignoring changes in medical debt d) and define γ as the implicit solution to:

$$\int v(h(m_u - \lambda), c(m_u))dF_\lambda(\lambda) = \int v(h(m_f - \lambda), c(m_f - \gamma))dF_\lambda(\lambda). \quad (7)$$

Thus γ is defined in equation (7) as the amount of consumption that the individual would need to give up in the world with formal insurance that would leave him at the same level of expected utility as in the world without formal insurance.¹⁷ They then use empirical estimates of the distribution of c_j and h_j , plus a specific utility function, to solve equation (7) for γ .

Empirically, this approach requires estimates of the distribution of each argument of the utility function under formal insurance and under no insurance. Since one of these is counterfactual, we are in the familiar territory of estimating the distribution of “potential outcomes” under treatment and control (e.g., Angrist and Pischke, 2009). The authors use empirical estimates from the Oregon Health Insurance experiment to derive and compare the distributions of h_j and c_j under formal insurance (f) and uninsurance (u), and thus derive estimates of willingness-to-pay for Medicaid among a low-income adult population. Conceptually, this approach requires the researcher to specify (and observe) all elements of the utility function that may be affected by insurance.¹⁸

Several recent policy developments are opening new opportunities for researchers to directly estimate demand for health insurance among the previously uninsured low income population. The introduction of premiums for certain Medicaid populations is one such development; Dague (2014) examines how the length of enrollment spells in a Wisconsin Medicaid program change when monthly Medicaid premiums increase from \$0 to \$10 at specific income threshold. Mandate

¹⁷This approach will yield the same estimate of willingness to pay as direct estimation of a demand curve for formal insurance as long as (1) the utility function specified here is correct and complete, and (2) the two estimates are made under the same information set about risk – i.e., the same distribution of potential health risks F_λ (Hendren 2016a).

¹⁸This is not only a high hurdle but a potentially limitless target; unsatisfied by the results of a particular analysis, one could simply posit additional (unmeasured) elements of the utility function that health insurance could conceivably impact – such as on marital stability, leisure, or children’s outcomes. To reduce these high information requirements, Finkelstein, Hendren and Luttmer (2015) also developed an alternative, “optimization-based” approach to estimating γ where additional economic assumptions allow them to significantly reduce the information requirements about the utility function.

penalties for the uninsured are another relevant development; Hackman, Kolstad and Kowalski (2015) use the 2006 introduction of a mandate penalty in Massachusetts to estimate demand for unsubsidized health insurance among the uninsured above 300 percent of the poverty line. Finally, the advent of publicly-subsidized health insurance exchanges for low income individuals provides another useful platform for estimating demand; Finkelstein, Hendren, and Shepard (2017) exploit discontinuities at income thresholds in subsidized enrollee premiums for low-income adults under Massachusetts' pioneer health insurance exchange to trace out the demand curve for formal insurance. Other recent work has analyzed newly-insured low-income individuals' price sensitivity in their choice among health plans (e.g., Chan and Gruber 2010, Tebaldi 2016).

Findings

The results from different approaches indicate that low-income adults' willingness to pay (γ) is far below the costs of formal insurance (G). For example, using the impacts of Medicaid from the Oregon Health Insurance Experiment combined with a specification of the utility function, Finkelstein, Hendren and Luttmer (2015) estimate $\frac{\gamma}{G}$ of 0.2 to 0.5 for Medicaid for adults below the poverty line. In other words, they estimate that adults below the poverty line would be willing to pay only 20 to 50 cents per dollar of Medicaid spending; put differently, they would rather give up Medicaid than pay the insurance costs of providing it. Using more direct estimates of demand for insurance among adults in Massachusetts between 135 and 300 percent of the poverty line, Finkelstein, Hendren and Shepard (2017) similarly find γ to be less than half of G for a Medicaid-like health insurance contract.

These findings suggest that even modest enrollee premiums will be a major deterrent to enrollment for most low income individuals. Finkelstein, Hendren and Shepard (2017) estimate that if premiums were 75 percent of insurers' average costs (G), at most half of potential enrollees would buy insurance; premiums subsidized to 90 percent of insurers' average costs would still leave at least 20 percent uninsured. Such findings can explain coverage outcomes in the ACA exchanges, where enrollee premiums were higher than in these calculations, and where early evidence suggests highly incomplete take-up (Tebaldi 2016, Avalere 2016).

Why is willingness to pay for insurance by low income adults so far below insurer average costs? One natural explanation is adverse selection, in which marginal enrollees are lower cost than average enrollees. Evidence indicates that adverse selection exists, even in the presence of a coverage mandate such as the one introduced in 2006 in Massachusetts or in 2014 by the ACA (e.g. Chandra, Gruber and McKnight 2011; Hackman, Kolstad and Kowalski 2015; Finkelstein, Hendren and Shepard 2017; Panhans 2017). However, the estimates in Finkelstein, Hendren and Shepard (2017) suggest that willingness to pay lies far below *own* costs, not just average costs of all those with higher willingness to pay. This suggests that adverse selection cannot be the primary explanation of low take-up.

Another natural explanation is that moral hazard raises the costs of insurance without a corresponding increase in its value to recipients. However, the results in Finkelstein, Hendren and

Shepard (2017) also suggest that standard estimates of moral hazard effects in health insurance are an order of magnitude too low to close the gap between willingness to pay and costs.

Yet another possible explanation for the finding is that behavioral biases or liquidity constraints reduce demand-based estimates of willingness to pay below “true value”. Finkelstein, Hendren and Shepard (2017) find no “smoking gun” evidence of such effects in their data, but that does not preclude their existence or importance in that context or in others. However, the alternative approach taken by Finkelstein, Hendren and Luttmer (2015) to estimating willingness to pay by wrapping a researcher-imposed utility function over estimates of utility-relevant arguments with and without insurance would not be affected by such factors; and this approach similarly finds willingness to pay that is far below insurer costs.

An explanation that *can* rationalize the finding of willingness to pay that is substantially below (gross) insurer costs is the substantial amount of implicit insurance payments available to the low income uninsured: both Finkelstein, Hendren and Luttmer (2015) and Finkelstein, Hendren and Luttmer (2017) estimate that willingness to pay is in the ballpark of *net* costs (C).

Implicit insurance for the uninsured reduces willingness to pay for two reasons. First, a large portion of the formal insurance contract is paying for infra-marginal coverage. Finkelstein, Hendren and Luttmer (2015) estimate that 60 cents of every dollar of adult Medicaid spending in the Oregon experiment represent a transfer to the previous providers of implicit insurance payments for the uninsured (i.e. $\frac{\Delta I}{G} = \$0.6$).¹⁹ Second, due to concavity of the utility function, the marginal value of additional insurance is lower than the marginal value of the first unit.

Consistent with an important role for transfers to non-recipients in explaining willingness to pay below insurer costs for low income adults, Hackman, Kolstad, and Kowalski (2015) estimate that willingness to pay is above insurer costs for individuals above 300 percent of the poverty line in Massachusetts. These higher income individuals may well have less access to uncompensated care; Mahoney (2015) finds that uninsured individuals with more wealth receive less implicit insurance.

This evidence of the importance of uncompensated care in explaining willingness to pay below costs is suggestive, but indirect. An important area for further work is to try to more directly estimate the impact of the availability of uncompensated care on willingness to pay for insurance. To date, the only such estimates we know of are Herring’s (2005) cross-sectional estimates of the relationship between the availability of charity care and the extent of private insurance coverage and Mahoney (2015)’s analysis of the relationship between the implicit insurance from personal bankruptcy and private insurance coverage.

¹⁹Their estimate follows directly from the finding that low income uninsured adults in the Oregon Experiment paid only about 20 cents on the dollar for their medical care (using transacted prices for those on Medicaid rather than list prices for the uninsured to estimate medical expenses). The remaining 80 cents per dollar of medical care was paid by “non-recipients”. (This translates into less than 80 cents on the dollar of transfers to non-recipients because Medicaid also has a moral hazard effect of increasing health care use, so that Medicaid spending on medical care exceeds total (consumer plus third party) spending on medical care for the uninsured).

4.2 Normative Analysis

In this last section we consider the normative implications of the estimates of willingness to pay relative to costs. This requires either that we stipulate to the functional form of the normative utility function (as in Finkelstein, Hendren and Luttmer 2015), or that we interpret demand estimates (as in Finkelstein, Hendren and Shepard 2017) as the welfare-relevant metric. The latter assumption – that “demand reveals value” – is a standard assumption in economics, and has been widely used to estimate the surplus from private (often employer-provided) insurance (e.g., Einav, Levin and Finkelstein 2010).

However, assuming that demand reveals value may be a problematic assumption in general, and for a low-income population in particular. Any type of behavioral bias – such as inattention, misperceived probabilities, or inertia – can drive a wedge between demand and underlying value. Spinnewijn (forthcoming) discusses such issues in detail and Baicker et al. (2015) describe some of the evidence of behavioral biases in the healthcare context. Even without behavioral biases, neoclassical frictions such as liquidity constraints may reduce willingness to pay below the normative value of insurance, since premiums must be paid out of current income. These important caveats notwithstanding, we proceed with normative analysis of subsidies for formal health insurance for low income adults. Interestingly, despite the different caveats, the outcomes-based approach in Finkelstein, Hendren and Luttmer (2015), and the demand-based approach of Finkelstein, Hendren and Shepard (2017) yield very similar results.

4.2.1 Efficiency Analysis

Comparisons of recipient willingness to pay to gross cost ($\frac{\gamma}{G}$) speaks to the question: if low-income adults had to finance the cost of formal insurance through lump sum taxes, would they be better off? Market failures such as adverse selection could in theory allow such budget-neutral subsidies to increase welfare (Einav and Finkelstein, 2011). However, the evidence suggests that although adverse selection exists, it does not, by itself, provide an economic rationale for budget-neutral health insurance subsidies for low-income adults: willingness to pay lies far below *own* costs, so that even if insurers were able to price discriminate on the basis of willingness to pay, individuals would not want to purchase insurance at prices that would allow the insurance company to break even. In other words, average welfare of the uninsured would not increase if the low income uninsured had to finance the cost of their formal insurance through lump sum taxes.

The substantial amount of implicit insurance for the uninsured provides an alternative rationale for subsidizing formal insurance through a mechanism known as the Samaritan’s dilemma (Buchanan 1975). Suppose that an uninsured individual receives a large health shock and requires medical care. Society (the good Samaritan) faces a dilemma: if it chooses to provide charity care to this uninsured individual, it encourages others to forgo health insurance. In other words, because the uninsured do not face the full social cost of being uninsured when uncompensated care exists, they have less incentive to take up formal insurance coverage. This introduces a key potential inefficiency: society is unlikely to choose the optimal level of implicit insurance. Indeed, this op-

timal level likely varies across individuals with different preferences. However, because individuals cannot buy formal insurance that “tops up” what is available as uncompensated care through informal insurance, the existence of this informal insurance creates an implicit tax on formal insurance; formal insurance premiums must cover the infra-marginal care that would have been provided as uncompensated care in the absence of formal insurance. Formal insurance choices may therefore be distorted away from what they would be in the absence of this implicit tax (Coate, 1995).²⁰

One way to overcome this issue is to commit to not providing medical care to the uninsured, but it is hard to imagine a moral society committing to deny care to an uninsured victim of a car accident, for example. Another way to wriggle out of this predicament is to subsidize formal insurance so that individuals face the net cost of formal insurance (i.e., can effectively pay to “top up” the implicit insurance); a health insurance mandate that penalizes individuals for being uninsured can be used to similar effect (Mahoney 2015).

In the presence of substantial transfers to non-recipients, therefore, the question of whether it is efficient for low-income adults to have formal insurance requires a comparison of willingness to pay (γ) to the *net* cost of formal insurance (C), rather than the gross cost (G). Due to moral hazard, it is a priori uncertain whether, for a given risk averse individual, $\frac{\gamma}{C}$ is above or below 1; it will exceed 1 only if the marginal welfare gains from (budget neutral) formal risk protection are greater than the moral hazard costs of this formal insurance (Zeckhauser 1970). Empirically, the results on this question are inconclusive. Finkelstein, Hendren and Luttmer (2015) estimate $\frac{\gamma}{C}$ ranging from 0.5 to 1.2; Finkelstein, Hendren and Shepard’s (2017) similarly estimate that γ hovers in the neighborhood of C . What seems clear is that low income uninsured adult’s willingness to pay is not substantially in excess of the net costs of formal insurance.

The large role played by implicit insurance for the uninsured raises another natural efficiency question: Is it relatively more efficient to provide a given amount of coverage formally or informally? This is largely unaddressed by the existing empirical literature and is an important area for further work. A primary efficiency cost of formal health insurance is increased health care expenditures (moral hazard). Of course, implicit insurance may also distort healthcare demand and supply. For example, it may set relative prices differently across different types of care (e.g. making the emergency room effectively more subsidized than primary care). It may also create disincentives for healthcare providers to locate in areas with high rates of uninsured, and particularly to provide emergency services in such areas. To our knowledge, these interesting and potentially important channels have received little empirical attention.

²⁰This same type of choice distortion can arise if individuals cannot apply subsidies for formal insurance to a range of options; historically this was the case for Medicaid (Cutler and Gruber 1996), but more recently public subsidies for health insurance have been accompanied by substantial choice (e.g. Medicare Part D; ACA exchanges). A related issue arises when insured individuals cannot pay on the margin to “top up” care covered by insurance (Einav, Finkelstein and Williams 2016).

4.2.2 Distributional Analysis

In practice, most formal health insurance expansions are partially or fully subsidized: the recipients of formal insurance expansions rarely finance the full costs of this expansion. Insurance expansions therefore have an important distributional component as well. In particular, recent government interventions in health insurance have disproportionately expanded coverage for low income individuals through public subsidies financed by general revenue.

A distributional perspective suggests two key modifications to simply analyzing $\frac{\gamma}{C}$ or $\frac{\gamma}{G}$. First, on the cost side, we have until now considered only formal insurance payments (see equations (5) and (6)).²¹ The costs of a means tested subsidy, however, additionally include any reduced income tax revenue for the government that occurs if individuals decrease labor market earnings in order to meet the eligibility requirements. The resulting fiscal externality from lower income tax revenue would further increase the total cost of the health insurance expansion. Existing evidence suggests that such labor market responses to subsidized health insurance coverage may be substantial (Garthwaite, Gross and Notowidigdo, 2014; Dague et al. 2017), although smaller responses have been found for traditional Medicaid beneficiaries and individuals with potential earnings close to the federal poverty line (Baicker et al. 2014; Leung and Mas forthcoming).

Second, on the benefits side, we should apply a social welfare function to estimates of γ in a low income population, or weight γ using a parameterization of social marginal utilities of income that translate individual willingness to pay into social willingness to pay (Saez and Stantcheva 2016). In other words, even if recipient willingness to pay does not exceed costs, social benefits may exceed cost.

To illustrate this second point, we consider a utilitarian social welfare function over individual utility functions. In this case, social welfare is each recipient's 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 Consumer Expenditure Survey suggests that the median consumption in the recipient population relative to the general population is about 40 percent for the Oregon Health Insurance Experiment and 60 percent for the subsidized low-income health insurance exchange in Massachusetts. To calibrate the social value of insurance expansions for the low income population, suppose individual utility is CRRA with a coefficient of relative risk aversion of σ (i.e., a marginal utility of consumption of $\frac{1}{c^\sigma}$). For $\sigma = 3$, social welfare from insurance coverage would be about 15 times recipient willingness to pay for Medicaid in Oregon, and about 5 times that for the low-income health insurance exchanges in Massachusetts. Even with log utility, social welfare would be 2.5 times recipient willingness to pay for Oregon Medicaid and 1.7 times recipient value for Massachusetts' health insurance. Since γ was estimated to be at least 20 percent of G , these estimates suggest that social welfare would certainly exceed C , and likely exceed G . Naturally, more concave social welfare functions than

²¹These costs included both the inframarginal payments for healthcare the individual consumed when uninsured as well as the increased payments due to the moral hazard response to insurance; the latter represents a fiscal externality from the insurance on the government.

utilitarian would further increase social welfare, as would more concave individual utility functions. Publicly-subsidized health insurance expansions for low income adults therefore seem to pass a social cost-benefit test under a fairly conservative parameterization of the social welfare function.

An alternative benchmark is Hendren’s (2014) suggestion to compare the redistributive “bang for the buck” across alternative redistributive instruments. In this context, one would compare the increase in welfare for low-income recipients per dollar of government expenditures on formal health insurance (i.e. γ/G) to what could be achieved through alternative transfer programs to a similar population. This avoids having to take a stand on the curvature of the social welfare function or the individual utility function, and instead asks: are health insurance subsidies more or less costly as a form of redistribution than other transfer programs?

A natural comparison is to the Earned Income Tax Credit (EITC) – the primary form of tax subsidies to low income individuals in the United States. Hendren (2016b) draws on existing empirical estimates of the impact of the EITC on labor supply to estimate that the EITC generates roughly \$0.90 of welfare to EITC recipients for every dollar of government spending. Welfare gains are less than spending due to the labor supply distortions created by the EITC. The EITC “bang for the buck” of 0.9 thus appears to be substantially above estimates for health insurance subsidies to low income adults which, as discussed, place γ/G at 0.5 or lower.

Again, however, this estimate of γ/G considers only the value of the insurance to recipients and gives zero weight to the transfers to external parties ΔI ; recall that $\Delta I \approx 0.6G$. A critical question, therefore, is the ultimate economic incidence of these transfers. Consider for example the estimates in Finkelstein, Hendren and Luttmer (2015) of γ/G ranging from 0.2 to 0.5. If the government bears the incidence of uncompensated care payments – as would be the case if public funds were used to reimburse hospitals for uncompensated care – then the cost to the government of providing the insurance would decrease to $0.4G$, and the relevant comparison to the 0.9 estimate for the EITC would be to $\gamma/0.4G$, which is 0.5 to 1.25. If the low-income individuals bear the ultimate incidence of the transfers ΔI – as would be the case if reductions in uncompensated care costs for the newly insured allowed provided to provide better care to the remaining low-income uninsured – then the relevant comparison of the 0.9 estimate for the EITC is to $(\gamma + 0.6G)/G$, which ranges from 0.8 to 1.1. This suggests that subsidizing health insurance for low income adults is in the same ballpark as cash transfers through the EITC if either the government or the low income beneficiaries bear the incidence of uncompensated care provision to the low-income uninsured. On the other hand, if the ultimate incidence of ΔI is on the high end of the income distribution – such as hospital owners or the privately insured – then the relevant EITC comparison could be to $(\gamma + 0.3G)/G$ or 0.5 to 0.8, at which point the “bang for the buck” looks lower for health insurance subsidies than the EITC.²² Finkelstein, Hendren and Shepard (2017) undertake a comparable set of exercises, and the results are similar in their context. Whether or not subsidized health insurance is more or less efficient than the EITC as a way to transfer resources to low income individuals depends critically on the

²²In this calculation we assume that half of the \$0.6 transfer can be brought back to the bottom of the income distribution, using Hendren’s (2014) estimate that \$1 falling near the top of the income distribution can be turned into \$0.50 to an EITC beneficiary through modifications in the income tax schedule.

economic incidence of the large transfers that subsidized health insurance generates to providers of uncompensated care.²³

The normative analysis suggests the efficiency case for subsidized health insurance is marginal. In particular, the evidence indicates that average welfare of the uninsured would be lower if low-income uninsured had to finance the full cost of formal insurance through lump sum taxes. Instead, willingness to pay seems to be in the ballpark of the net cost of formal insurance, implying that any welfare gains from budget-neutral subsidies are minimal. In practice, of course, health insurance subsidies have a substantial distributional component. Given the higher marginal utility of consumption of lower income adults, the social value of formal health insurance to low income adults likely exceeds its (gross) cost and certainly its net cost. However, whether subsidies for health insurance for low income adults are an efficient form of redistribution compared to cash transfers, such as the Earned Income Tax Credit, depends critically on where in the income distribution the ultimate economic incidence of the transfers from formal health insurance expansions to non-recipients falls.

4.3 The Economic Incidence of ΔI

Determining the ultimate economic incidence of the large transfers that health insurance subsidies generate to external parties is a complicated, challenging, and not-yet-resolved issue. We begin by describing evidence on the direct financiers of the implicit insurance for the uninsured. Then we turn to the ultimate economic incidence; we briefly sketch some of the conceptual issues and describe the limited evidence to date.

Direct financiers of implicit insurance payments

Implicit insurance payments for the uninsured (I_u) appear to be directly financed primarily by hospitals and the public sector. Garthwaite et al. (forthcoming) estimate that each uninsured individual costs hospitals approximately \$800 per year in uncompensated care costs. Early evidence from the ACA also finds evidence of financial benefits to hospitals from expanding formal health insurance, with Medicaid expansion states experiencing a decline in hospital uncompensated care costs relative to non-expansion states (Dranove et al., 2016). Direct benefits to hospitals are consistent with the lobbying efforts by hospitals around health insurance expansions and repeats that we described in the Introduction.

²³Means-tested subsidized health insurance would also look worse as a redistributive tool than the estimates presented here if these means-tested subsidies encourages substantial declines in labor market earnings. As noted earlier, this additional fiscal externality would increase costs per dollar of transfer to recipients. The estimates discussed here do not incorporate any such labor market responses. Alternatively, under the assumption that the labor supply response per dollar of health insurance transfer is the same as the labor supply response per dollar of cash transfer like the EITC, one could safely ignore any labor market responses to subsidized health insurance and simply compare the estimates of willingness to pay per dollar of health insurance subsidy against the benchmark of 1 (instead of 0.9 which accounts for the EITC's labor market distortions).

In addition to financing by hospitals, there are a number of ways that the public sector pays for implicit insurance (Hadley et al. 2008). Federal Disproportionate Share Hospital (DSH) payments are intended to partially compensate hospitals for providing uncompensated care. Several states have designed “uncompensated care pools” that provide funding to hospitals that face bad debt from unpaid medical bills and funding shortfalls due to providing uncompensated care to the uninsured seeking emergency medical treatment. Formal health insurance expansions are often accompanied by a decline in uncompensated care, suggesting that policymakers perceive a relationship between formal insurance coverage and hospital finances. For instance, because the Medicaid expansions under the ACA were expected to reduce the costs of uncompensated care, they were accompanied by reductions in DSH payments (Rudowitz 2013). Likewise, the financing of the Massachusetts health insurance expansion in 2006 was partially funded by dissolving the state’s uncompensated care pool (Kolstad and Kowalski 2012).

Ultimate economic incidence of implicit insurance payments

The ultimate economic incidence of changes in implicit insurance payments due to formal health insurance is conceptually complicated and empirically elusive. The list of potentially affected parties is long; it includes share-holders (at for-profit hospitals), buyers (insurance companies and patients), suppliers (e.g., employees, prescription drug and medical device manufactures), hospital competitors (community health clinics), and local, state and federal government. As an example of the complicated interrelationships involved, consider Federal Qualified Health Clinics (FQHCs). Federal, state, and local governments provide grants to FQHCs to provide free or “pay as you can” care to the uninsured population. In 2015, FQHCs served roughly 6 million uninsured patients (out of roughly 25 million patients total) and received roughly \$7 billion in grants from federal, state, and local governments (<https://bphc.hrsa.gov/uds/datacenter.aspx>). Against this backdrop, changes in implicit insurance payments could indirectly affect both the clinics themselves and the governments that fund them through grants. The ACA expanded funding for FQHCs, perhaps in expectation of an increased demand for FQHCs in response to the ACA Medicaid expansions.

Economic theory, unfortunately, provides little guidance on who will bear the incidence. One critical conceptual question is whether ΔI should be thought of as a (negative) shock to a medical provider’s fixed costs or variable costs. If it primarily affects variable costs, theories of pass-through (e.g., Weyl and Fabinger 2013) predict that the incidence will be borne by the inelastic parties, which simply transforms the problem into one of estimating elasticities. Naturally these elasticities may vary depending on the time horizon considered.

If ΔI primarily affects fixed costs, then in the simplest model, ΔI only affects the provider’s exit / entry decisions, as well as their threat point in their negotiations with buyers (insurance companies) and suppliers (labor, prescription drug and medical device manufactures). If the exit / entry threat point is not relevant to the bargaining game, then providers may have little ability to pass-through ΔI to other parties. Any effect on exit / entry - such as the decision to locate or leave an area - may have impacts on all the patients in that area. In a similar vein, ΔI could

affect hospital investments in fixed cost activities - e.g. the adoption of technology - which could likewise broadly affect patients. Consistent with this, Finkelstein (2007) provides evidence that the introduction of Medicare in 1965 affected hospital technology adoption and entry decisions.

Interestingly, many policymakers either implicitly or explicitly assume that hospitals simply pass on uncompensated care costs to privately insured patients. For example, the text of the ACA states, “[t]o pay for [uncompensated care], health care providers pass on the cost to private insurers, which pass on the cost to families. This cost-shifting increases family premiums by on average over \$1,000 a year. By significantly reducing the number of the uninsured, the requirement, together with the other provisions of this Act, will lower health insurance premiums” (42 USC 18091). Cost shifting was also cited by Chief Justice Roberts in the Supreme Court decision upholding the ACA’s constitutionality (*National Federation of Independent Business v. Sebelius*, 567 U.S., 2012).

Despite this widespread assumption of substantial “cost shifting,” there is limited empirical evidence of the ability of hospitals to actually pass on uncompensated care costs. Garthwaite et al. (forthcoming) show that increases in the un-insurance rate are associated with a decline in hospital operating (profit) margins. This indicates that hospitals are not able to *fully* pass on increased uncompensated care costs, at least over the relatively short time horizon they examine. This implies that hospitals are not able to fully evade the ultimate economic incidence of implicit insurance payments, but leaves the precise breakdown across different parties largely uncertain.

This finding of limited cost shifting in the case of uncompensated care costs lines up with several related studies which document limited evidence of cost shifting in other health care contexts (Dranove 1988; Morrissey 1994; Timmins 2014; Dranove, Garthwaite, and Ody 2014; and Clemens and Gottlieb, 2017). Indeed, a recent review of evidence on hospital cost shifting concludes that it occurs “at a relatively low rate” (Frakt 2011). While the ultimate incidence of hospital uncompensated care costs remains unsettled, these studies are consistent with arguments made by hospital trade groups that repeals of public health insurance expansions will cause hospitals to suffer significant financial losses (e.g. Goldstein 2016). Indeed, the initial collapse of the American Health Care Act bill to repeal the ACA in March 2017 appeared to lead to sharp increases in the stock prices of for-profit hospital systems.²⁴

5 Conclusion

This essay has considered the positive and normative implications of publicly subsidized expansions of health insurance coverage. We described three core positive findings. First, publicly-subsized health insurance conveys benefits to the previously uninsured in the form of improvements in health measures, increases in consumption proxies, and reductions in medical debt. Measurement challenges to date suggest that more work is needed to estimate non-mortality health impacts of health insurance, to directly measure consumption impacts, and to examine benefits to recipients from reductions in medical debt.

²⁴See, e.g., <https://twitter.com/ddiamond/status/845391255953723393>.

Second, publicly-subsidized health insurance also conveys substantial benefits to the non-recipients who previously provided informal insurance payments for the uninsured. Both empirical evidence and lobbying campaigns by hospital groups for health insurance expansions (and against contractions) suggest that hospitals “foot the bill” for much of the uninsured’s care on the margin. A key unanswered – and challenging – question going forward is who bears the ultimate economic incidence of these informal insurance payments.

Third, low income adults’ willingness to pay for publicly-subsidized insurance is substantially below the costs to insurers of providing this coverage. This helps explain limited take-up of even heavily subsidized formal health insurance. The substantial benefits from formal insurance to non-recipients (i.e., the presence of large amounts of uncompensated care for the low income uninsured) are likely important in explaining this low willingness-to-pay for formal insurance relative to cost; willingness to pay is much closer to costs net of transfers to non-recipients.

In addition to these positive findings, we also described some of their normative implications. The fact that willingness to pay is substantially below gross costs indicates that low-income adults would be worse off if they had to finance the cost of formal insurance through lump sum taxes. These subsidies are therefore difficult to rationalize as welfare-increasing, budget-neutral subsidies that combat market failures like adverse selection. Subsidies can potentially be rationalized as a way to combat the implicit tax that uncompensated care for the uninsured imposes on formal insurance, but willingness to pay appears similar to net costs, suggesting any welfare gains from budget-neutral subsidies are minimal.

In practice, of course, these subsidies are not financed by low income adults but have a substantial distributional component. Given the higher marginal utility of consumption of lower income adults, it seems likely that the social value of formal health insurance to low income adults exceeds its (gross) cost and certainly its net cost. However, whether subsidies for health insurance for low income adults are an efficient form of redistribution compared to cash transfers such as the Earned Income Tax Credit is unclear. The answer depends critically on the open and challenging question of where in the income distribution the ultimate economic incidence of the transfers from formal health insurance expansions to non-recipients falls.

At a broad level, the work highlighted here suggests the importance of uncompensated care in understanding both the positive and normative implications of public subsidies for health insurance, and underscores the need for more work on this topic. The existence of uncompensated care and the underlying Samaritan’s dilemma (Buchanan 1975, Coate 1995) is one of the textbook rationales for government intervention in health insurance markets. Yet in recent decades it has received comparatively less empirical attention than other such textbook rationales, such as adverse selection. On the descriptive front, it would be valuable to have more evidence on the nature and form of uncompensated care, how willingness to pay for formal insurance varies with the availability of uncompensated care, and how provider behavior is affected by their uncompensated care obligations. On the normative front, key open questions concern the relative efficiency of implicit insurance provided through uncompensated care and formal insurance coverage, and the

ultimate economic incidence of implicit insurance provision.

6 References

- Anderson M, Dobkin C, Gross T. 2012. The Effect of Health Insurance Coverage on the Use of Medical Services. *American Economic Journal: Economic Policy*, 4(1): 1-27
- Angrist JD, Pischke J. 2009. *Mostly harmless econometrics: an empiricist's companion*. Princeton, New Jersey: Princeton University Press
- Aron-Dine A, Finkelstein A, Einav L. 2013. The RAND Health Insurance Experiment, Three Decades Later. *Journal of Economic Perspectives* 27(1): 197-222
- Avalere Health. 2016. The state of exchange: a review of trends and opportunities to grow and stabilize the market
- Avery RB, Calem PS, Canner GB. 2004. *Consumer credit scoring: do situation circumstances matter?*. BIS Working Papers 146
- Baicker K, Mullainathan S, Schwartzstein J. 2015. Behavioral hazard in health insurance. *Quarterly Journal of Economics* 130(5): 1623-1667
- Baicker K, Finkelstein A. 2013. Effects of Medicaid on clinical outcomes - reply to letters. *New England Journal of Medicine* 369(6): 581-583
- Baicker K, Taubman S, Allen H, Bernstein M, Gruber J, Newhouse J, Schneider E, Wright B, Zaslabsky A, Finkelstein A. 2013. The Oregon experiment - effects of Medicaid on clinical outcomes. *New England Journal of Medicine* 366(18):1713-1722
- Baicker K, Finkelstein A, Song J, Taubman S. 2014. The Impact of Medicaid on Labor Market Activity and Program Participation: Evidence from the Oregon Health Insurance Experiment. *American Economic Review Papers and Proceedings*, 104(5): 322-28
- Barcellos S, Jacobson M. 2015. The effects of Medicare on medical expenditure risk and financial strain. *American Economic Journal: Economic Policy* 7(4):41-70
- Cardon J, Hendel I. 2001. Asymmetric information in health insurance: evidence from the National Medical Expenditure Survey. *RAND Journal of Economics* 32(3):408-427
- Card D, Dobkin C, Maestas N. 2008. The impact of nearly universal insurance coverage on health care utilization: evidence from Medicare. *American Economic Review* 98(5):2242-58
- Card D, Dobkin C, Maestas. 2009. Does Medicare save lives?. *Quarterly Journal of Economics* 124(2):597-636
- Chan D, Gruber J. 2010. How sensitive are low income families to health plan prices?. *American Economic Review* 100(2):292-96
- Chandra A, Gruber J, McKnight R. 2011. The importance of the individual mandate - evidence from Massachusetts. *New England Journal of Medicine* 364(4):293-5
- Chetty R, Stepner M, Abraham S, Lin S, Scuderi B, Turner N, Bergeron A, Cutler D. 2016. The association between income and life expectancy in the United States, 2001-2014. *Journal of the American Medical Association* 315(16):1750-1766

- Clemens J, Gottlieb JD. 2017. In the shadow of a giant: Medicare's influence on private physician payments. *Journal of Political Economy* 125(1):1-39
- Cutler, D and J Gruber, 1996. "Does Public Insurance Crowd Out Private Insurance?" *Quarterly Journal of Economics* 111(2): 391-430.
- Coate S. 1995. Altruism, the Samaritan's dilemma, and government transfer policy. *American Economic Review* 85(1):46-57
- Congressional Budget Office (CBO). 2013. The distribution of major tax expenditures in the individual income tax system
- Consumer Financial Protection Bureau (CFPB). 2014. Consumer credit reports: a study of medical and non-medical collections
- Coughlin TA, Holahan J, Caswell K, McGrath M. 2014. Uncompensated care for the uninsured in 2013: a detailed examination. The Henry J Kaiser Family Foundation
- Cunningham R, Cunningham RM. 1997. *The blues: a history of the Blue Cross and Blue Shield system*. Illinois: Northern Illinois University Press
- Currie J, Gruber J. 1996. Saving babies: the efficacy and cost of recent changes in the Medicaid eligibility of pregnant woman. *Journal of Political Economy* 104(6):1263-96
- Dague L. 2014. The effect of Medicaid premiums on enrollment: a regression discontinuity approach. *Journal of Health Economics* 37:1-12
- Dague, L, DeLeire T, Leininger L. 2017. The Effect of Public Insurance Coverage for Childless Adults on Labor Supply. *American Economic Journal: Economic Policy*, 9(2): 124-54
- Diamond D. 2017. Twitter. <https://twitter.com/ddiamond/status/845391255953723393>
- Dobkin C, Finkelstein A, Kluender R, Notowidigdo MJ. 2016. *The economic consequences of hospital admissions*. NBER Working Paper 22288
- Dranove D. 1988. Pricing by non-profit institutions: the case of hospital cost-shifting. *Journal of Health Economics* 7(1):47-57
- Dranove D, Garthwaite C, Ody C. 2014. *How do hospitals respond to negative financial shocks? The impact of the 2008 stock market crash*. NBER Working Paper 18853
- Dranove D, Garthwaite C, Ody C. 2015. A floor-and-trade proposal to improve the delivery of charity-care services by U.S. nonprofit hospitals. The Hamilton Project Discussion Paper 2015-07
- Dranove D, Garthwaite C, Ody C. 2016. Uncompensated care decreased at hospitals in Medicaid expansion states but not at hospitals in nonexpansion states. *Health Affairs* 35(8):1471-9
- Einav L, Finkelstein A. 2011. Selection in insurance markets: theory and empirics in pictures. *Journal of Economic Perspectives* 25(1):115-138
- Einav L, Finkelstein A, and Levin J. 2010. Beyond testing: empirical models of insurance markets. *Annual Review of Economics* 2:311-36
- Einav L, Finkelstein A, Schimpf P, Ryan S, Cullen M. 2013. Selection on moral hazard in health insurance. *American Economic Review* 103(1):178-219
- Einav, L., Finkelstein A., and Williams, H. 2016. Paying on the margin for medical care: evidence from breast cancer treatments. *American Economic Journal: Economic Policy* 8(1)

- Einav L, Finkelstein A. 2017. Moral hazard in health insurance: what do we know and how do we know it?. Unpublished mimeo
- Ehrlich I, Becker G. 1972. Market insurance, self-insurance, and self-protection. *Journal of Political Economy* 80(4):623-48
- Engelhard G, Gruber J. 2011. Medicare Part D and the financial protection of the elderly. *American Economic Journal: Economic Policy* 3(4):77-102
- Federal Trade Commission (FTC). 2012. Report to Congress under Section 319 of the Fair and Accurate Credit Transactions Act of 2003
- Finkelstein A. 2007. The aggregate effects of health insurance: evidence from the introduction of Medicare. *Quarterly Journal of Economics* 122(1):1-37
- Finkelstein A, Hendren N, Luttmer E. 2015. *The value of Medicaid: interpreting the results of the Oregon Health Insurance Experiment*. NBER Working Paper 21308
- Finkelstein A, Hendren N, Shepard M. 2017. Subsidizing health insurance for low income adults: evidence from Massachusetts. Unpublished mimeo
- Finkelstein A, McKnight R. 2008. What did Medicare do? The initial impact of Medicare on mortality and out of pocket spending. *Journal of Public Economics* 92(7):1644-68
- Finkelstein A, Taubman S, Wright B, Bernstein M, Gruber J, Newhouse J, Allen H, Baicker K. 2012. The Oregon Health Insurance Experiment: evidence from the first year. *Quarterly Journal of Economics* 127(3):1057-106
- Finkelstein A, Taubman S, Allen H, Wright B, Baicker K. 2016. Effect of Medicaid coverage on ED use - further evidence from Oregon's experiment. *New England Journal of Medicine* 375(16):1505-1507
- Frakt A. 2011. How much do hospitals cost shift? A review of the evidence. *The Milbank Quarterly* 89(1):90-130
- Fuchs, Victor. 1974. *Who shall live? Health, economics, and social choice*. Singapore: World Scientific Publishing
- Garthwaite C, Gross T, Notowidigdo MJ. Forthcoming. Hospitals as insurers of last resort. *American Economic Journal: Applied Economics*
- Garthwaite C, Gross T, Notowidigdo MJ, Graves J. 2017. Insurance expansion and hospital emergency department access: evidence from the Affordable Care Act. *Annals of Internal Medicine* 166(3):172-179
- Garthwaite, C., Gross T, Notowidigdo MJ. 2014. "Public Health Insurance, Labor Supply, and Employment Lock." *Quarterly Journal of Economics* 129(2): 653-696
- Goldstein A. 2016. Hospitals warn Trump, Congress of massive losses with Affordable Care Act repeal. *Washington Post*, December 6
- Goodman-Bacon A. Public Insurance and Mortality: Evidence from Medicaid Implementation. *Journal of Political Economy*, forthcoming.
- Government Accountability Office (GAO). 2008. Nonprofit hospitals: variation in standards and guidance limits comparison of how hospitals meet community benefit requirements. GAO-08-880

- Gross T, Notowidigdo MJ. 2011. Health insurance and the consumer bankruptcy decision: evidence from expansions of Medicaid. *Journal of Public Economics* 95(7-8):767-778
- Grossman M. 1972. On the concept of health capital and the demand for health. *Journal of Political Economy* 80(2):223-225
- Hadley J, Holahan J, Coughlin T, Miller D. 2008. Covering the uninsured in 2008: current costs, sources of payment, and incremental costs. *Health Affairs* 27(5):399-415
- Hackman M, Kolstad J, Kowalski A. Adverse selection and an individual mandate: when theory meets practice. *American Economic Review* 105(3):1030-66
- Hendren N. 2014. *The inequality deflator: interpersonal comparison without a social welfare function*. NBER Working Paper 201351
- Hendren, N. 2016a. *Measuring ex-ante welfare in insurance markets*. Working Paper
- Hendren, N. 2016b. The policy elasticity. *Tax Policy and the Economy* 30(1):51-89
- Herring B. 2005. The effect of the availability of charity care to the uninsured on the demand for private health insurance. *Journal of Health Economics* 24(2):225-52
- Hu L, Kaestner R, Mazumder B, Miller S, Wong A. 2016. *The effect of the patient protection and Affordable Care Act Medicaid expansions on financial well-being*. NBER Working Paper 22170
- Internal Revenue Service (IRS). 2007. Hospital compliance project: interim report
- Johnson L. 1965. Remarks with President Truman at the signing in independence of the Medicare bill. LBJ Presidential Library. <http://www.lbjlibrary.net/collections/selected-speeches/1965/07-30-1965.html>
- Kolstad J, Kowalski A. 2012. The impact of health care reform on hospital and preventive care: evidence of Massachusetts. *Journal of Public Economics* 96(11-12):909-29
- Krueger AB, Kuziemko I. 2013. The demand for health insurance among uninsured Americans: results of a survey experiment and implications for policy. *Journal of Health Economics* 325(5):780-93
- Lakdawalla D, Sood N. 2009. Innovation and the welfare effects of public drug insurance. *Journal of Public Economics* 93(3-4):541-548
- LeCuyer NA, Singhal S. 2007. Overhauling the US health care payment system. *The McKinsey Quarterly*
- Leung P, Mas A. Forthcoming. Employment Effects of the ACA Medicaid Expansions. *Industrial Relations*
- LoSasso A, Buchmueller T. 2004. The effect of the State Children's Health Insurance Program on Health Insurance Coverage. *Journal of Health Economics* 23(5): 1059-82
- Mahoney N. 2015. Bankruptcy as implicit health insurance. *American Economic Review* 105(2):710-46
- Mann RJ, Porter K. 2010. Saving up for bankruptcy. *Georgetown Law Journal* 98(2):289-339
- Mazumder B, Miller S. 2016. The effects of the Massachusetts health reform on household financial distress. *American Economic Journal: Economic Policy* 8(3):284-313
- Michalopoulos C, Wittenburg D, Israel D, Schore J, Warren A, Zutshi A, Freedman S, Schwartz

L. 2011. The accelerated benefits demonstration and evaluation project: impacts on health and employment at 12 months. New York, New York: mdrc

Miller S, Wherry L. 2016. Early coverage, access, utilization, and health effects of the Affordable Care Act Medicaid Expansions: a quasi-experimental study. *Annals of Internal Medicine* 164(12):795-803

Miller S, Wherry L. 2017. Health and access to care during the first 2 years of the ACA Medicaid expansions. *New England Journal of Medicine* 376(10):947-56

Morrissey M. 1994. *Cost shifting in health care: separating evidence from rhetoric*. Washington DC: AEI Press

National Center for Health Statistics. Health, United States, 2015: In Brief

National Federation of Independent Business v. Sebelius. 2012. 567 U.S. 1

Newhouse JP. 1993. *Free for all?*. Cambridge: Harvard University Press

Nicholson S, Pauly MV, Burns LR, Baumritter A, Asch DA. 2000. Measuring community benefits provided by for-profit and nonprofit hospitals. *Health Affairs* 19(6):168-177

Obama, Barack. 2010. Remarks by the President and Vice President at signing of the health insurance reform bill. The White House Office of the Press Secretary. <https://obamawhitehouse.archives.gov/photos-and-video/video/president-obama-signs-health-reform-law#transcript>

Ollove, Michael. 2013. Hospitals press states to expand Medicaid. Kaiser Health News

Patient Protection and Affordable Care Act (ACA). 2010. 42 U.S.C. § 18001

Panhans M. 2017. Adverse selection in ACA exchange markets: evidence from Colorado. Working Paper

Pauly M. 1968. The economics of moral hazard: comment. *American Economic Review* 58(3):531-537

Rudowitz, R. 2013. How do Medicaid Disproportionate Share Hospital (DSH) payments change under the ACA?. The Kaiser Commission on Medicaid and the Uninsured

Saez E, Stantcheva S. 2016. Generalized social marginal welfare weights for optimal tax theory. *American Economic Review* 106(1):24-45

Simon K. 2005. Adverse selection in health insurance markets? Evidence from state small-group health insurance reforms. *Journal of Public Economics* 89(9-10):1865-77

Spinnewijn J. Forthcoming. Heterogeneity, demand for insurance, and adverse selection. *American Economic Journal: Economic Policy*

Taubman S, Allen H, Wright B, Baicker K, Finkelstein A. 2014. Medicaid increases emergency department use: evidence from Oregon's health insurance experiment. *Science* 343(6168):263-268

Tebaldi P. 2016. *Estimating equilibrium in health insurance exchanges: price competition and subsidy design under the ACA*. Working Paper

Timmins LL. 2014. *How do hospitals respond to financial pain? Evidence from hospital markets in Texas*. Working Paper

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

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

U.S. Internal Revenue Services. 2015. EITC calendar year report.

Weyl G, Fabinger M. 2013. Pass-through as an economic tool: principles of incidence under imperfect competition. *Journal of Political Economy* 121(3):528-83

Wherry LR, Meyer BD. 2016. Saving Teens: Using a Policy Discontinuity to Estimate the Effects of Medicaid Eligibility. *Journal of Human Resources* 51(3): 556-588.

Wherry, LR., Miller S, Kaestner R, and Meyer B. "Childhood Medicaid Coverage and Later Life Health Care Utilization", Review of Economics and Statistics, forthcoming.

Zeckhauser R. 1970. Medical insurance: a case study of the tradeoff between risk spreading and appropriate incentives. *Journal of Economic Theory* 2(1):10-26

Figure 1: Payments vs. Charges, For Privately Insured and Uninsured

This figure is a reproduction of Panel A of Figure 1 in Mahoney (2015). Figure plots payments against charges for privately insured and uninsured households. Payments are the sum of out-of-pocket payments and payments from private insurance providers. (Payments by the uninsured are therefore simply out-of-pocket payments.) Charges are the list price of medical care and proxy for the level of medical utilization. The plot was created by averaging payments and charges at twentieths of the charge distribution. Pooled 1996–2005 MEPS, excluding households with public insurance or a member age 65 or older, inflation-adjusted to 2005 USD using the CPI-U.