Anti Insurance:
Health Insurance Worsens Risk Exposure

Lee M. Lockwood*

February 21, 2021

Preliminary: Comments Appreciated

Abstract

Health insurance provides valuable insurance against health care cost risk. But there are other risks, and health insurance's overall insurance value depends on the extent to which it insures or worsens those as well. Using a variety of approaches, I find that health insurance systematically worsens other risks, and this worsening is sufficiently costly that on average for U.S. households health insurance worsens risk exposure *on net*. The key driver is the uneven, health care-focused safety net. The safety net—charity care, bad debt, bankruptcy, and means-tested support—provides considerable protection against otherwise-uninsured health care costs, especially in the worst, highest-marginal utility states of the world. Health insurance displaces such support, which undoes without replacing the valuable insurance against other risks that such support otherwise would have provided. This is an important cost of the many major policies that subsidize health insurance or health care.

*University of Virginia and NBER, leelockwood@virginia.edu. I am grateful to David Benson, Karim Chalak, Federico Ciliberto, Kerem Cosar, Manasi Deshpande, Eric French, Leora Friedberg, Adam Leive, Attila Lindner, Benjamin B. Lockwood, Amalia Miller, Matthew Notowidigdo, John Pepper, Juan Carlos Suárez Serrato, and seminar participants at the ASSA meetings, Clemson, Duke, Johns Hopkins, LSE/IFS, Michigan, Texas-Austin, USC, Virginia, and Wharton for helpful comments. I thank Erica Sprott, Lichen Wang, and Yiyu Xing for excellent research assistance. I am grateful to the Bankard Fund for Political Economy for financial support.

1 Introduction

Health insurance is a core aspect of government policy and a major household asset. In the US, government spending on health insurance and health insurance subsidies exceeds $1.6 trillion per year and total health insurance spending exceeds $2.7 trillion per year, over $21,000 per household. Yet despite the vast resources involved, relatively little is known about the benefits of health insurance, including that of its core insurance function.

In this paper, I use a variety of approaches to quantify the insurance value of health insurance: the difference between its ex ante value and its mean ex post value. This insurance value is determined by the extent to which health insurance helps people when they want it most, i.e., when the marginal utility of income is greatest. Each approach points to the same, robust conclusion: Health insurance has *negative* insurance value—it *worsens* risk exposure. My estimates suggest that the associated welfare cost is a few hundred dollars per year per household. Even were it not for moral hazard and other costs, for many households health insurance would be worth less than its mean net benefit.

The key driver is the safety net. The "safety net"—charity care, bad debt, bankruptcy, and means-tested support of all kinds—provides substantial protection against otherwise-uninsured health care costs, especially when times are tight. In so doing, it transforms health spending from a major *risk creator* into a net *risk absorber*, thereby transforming the insurance value of health insurance from highly positive to significantly negative. First, the safety net reduces the value of health insurance protection against health care costs. To a rough approximation, the safety net provides implicit catastrophic health insurance (Mahoney, 2015). The residual health spending remaining for health insurance is therefore mainly the non-catastrophic spending that optimal insurance theory recommends *not* covering, due to the limited insurance benefit of such coverage. Second, the safety net causes health insurance to worsen other risks. The safety net provides more protection against otherwise-uninsured health care costs when non-health care circumstances are worse, such as when income and assets are low following a work-limiting health shock or job loss. This provides valuable insurance against these other risks—insurance that health insurance undoes. This undoing transforms health insurance from being roughly neutral with respect to other risks to worsening them systematically.

I begin by estimating a sufficient statistic approximation to the insurance value of eliminating out-of-pocket health spending, i.e., of going from status quo levels of out-of-pocket health spending to zero out-of-pocket health spending (as, say, providing someone with Medicaid would do). I start with this because of its policy relevance and its lack of need for strong assumptions. Under relatively mild assumptions that encompass a wide variety of models, to

first order the insurance value of eliminating out-of-pocket health spending is proportional to the covariance between marginal utility and out-of-pocket health spending. The key modeling requirement, then, is modeling marginal utility. My main assumption, although I test many others as well, follows the standard approach of modeling marginal utility as a decreasing function of non-health consumption, using data from the Panel Study of Income Dynamics (PSID). Intuitively, this approach is based on the idea that a household's chosen level of consumption reveals the overall tightness of its constraint, bypassing the need to model the constraint in its entirety.

For each of many types of households and states of the world—young and old, insured and uninsured, high- and low-income—and under a wide range of assumptions, the covariance between marginal utility and out-of-pocket health spending is significantly negative: Higher out-of-pocket spending is associated with *lower*, not higher, marginal utility. This is because higher out-of-pocket spending is associated with higher, not lower, consumption. This means that eliminating status quo out-of-pocket health spending would worsen risk exposure. My main estimates imply that for both insured and uninsured non-elderly households, eliminating out-of-pocket health spending would worsen risk exposure by an amount equivalent in utility terms to reducing annual income by about $300.

Why does health insurance worsen risk exposure? A variety of evidence suggests that the key driver is the safety net. Consider income, a large and risky part of household budgets. Across each of many types of households and estimation strategies, income is positively correlated with out-of-pocket health spending. Hence, health insurance pays off more on average when income is higher. This worsens income risk, at a cost, decompositions suggest, of about half of health insurance's overall net anti-insurance cost. This worsening of income risk occurs even though total health care costs and gross health insurance benefits are virtually uncorrelated with income. It is only net health insurance benefits, net of displaced support from the safety net, that are positively related to income. Net health insurance benefits are positively related to income because safety net support is negatively related to income: The safety net provides more protection against otherwise-uninsured health care costs when income is lower. Displacing this valuable insurance against income risk transforms health insurance from being roughly neutral with respect to income risk to worsening it systematically.[1]

The safety net transforms health insurance's targeting so powerfully that it causes health insurance to worsen even the major health risk of experiencing health shocks that lead to hospitalization. Although health insurance's (small) average net transfer from non-hospitalization to hospitalization states provides moderately valuable insurance, health insurance's trans-

[1] An example of an income-related risk that health insurance worsens is unemployment risk. I estimate that health insurance is worth about $230 less in years in which a household experiences unemployment. Health insurance is, in part, negative unemployment insurance.

fers *within* hospitalization states have larger anti-insurance costs. Hospitalization, as Dobkin et al. (2018) document, is associated with significant uninsured earnings losses, losses that are much greater in some states than in others. The safety net provides more protection against otherwise-uninsured health care costs when earnings losses are greater, thereby partially insuring these earnings losses. Health insurance undoes this insurance, at a welfare cost that exceeds the benefit of its insuring of hospitalization-related health care costs. Were it not for the safety net, by contrast, a variety of evidence suggests that health insurance coverage of hospitalization-related health care costs would provide significant net insurance value.

To determine whether the safety net can account for health insurance's negative overall insurance value, to better understand the roles of different features of the safety net and of the underlying risk process more generally, and to quantify the value of inframarginal health insurance coverage, I construct a simple model guided by the key empirical regularities. The model is based on standard models of health spending risk but adds two features motivated by the empirical evidence: non-health care risk and a health care-focused safety net. The key empirical inputs to the model are (i) the joint distribution of health care consumption and income, which I use as a summary measure of the risk households face, and (ii) the relationship between out-of-pocket health spending and income among uninsured households with high levels of health care consumption, which I use to calibrate a simple model of the safety net.

The model matches well a variety of key empirical patterns not directly targeted by the calibration, including the sufficient statistic estimates of health insurance's net insurance value. In this analysis, as in the sufficient statistic estimations, the conclusion that health insurance worsens risk exposure is extremely robust. Here, this conclusion extends to the inframarginal health insurance held by insured households. Health insurance would worsen risk exposure even if income risk and safety net protection against health care costs were considerably smaller than what they appear to be empirically.

Counterfactual analyses of the model highlight the key roles of non-health care cost risk and the safety net, which together transform health insurance's targeting from being extremely valuable to costly. In terms of health care cost risk, the safety net's substantial protection against otherwise-uninsured health care costs reduces the value of health insurance's insurance of health risk by two orders of magnitude, from extremely valuable (about equivalent to a $7,000 increase in annual consumption) to a matter of near indifference ($70 increase in annual consumption). So it is not that insurance against health spending risk is not valuable; such insurance is extremely valuable. It is that *marginal* insurance against health spending risk, on top of that provided by the safety net, is not very valuable. In terms of

other risks, the safety net's contingent nature—in particular, its greater protection against health care costs in states of the world with worse realizations of other risks—and implicit taxation of health insurance benefits transform health insurance from being roughly neutral with respect to other risks to worsening them systematically, at a cost equivalent to a $550 decrease in annual consumption.

That health insurance worsens risk exposure does not mean that it reduces welfare or that households are making mistakes by holding it. There are a variety of potential other, non-insurance benefits of health insurance, which I discuss in Section 5, including reducing efficiency costs associated with safety net support. That health insurance has negative insurance value just means that one important component of the overall welfare effect of health insurance and health insurance-promoting policies is not the benefit previously thought but a cost. It also reverses certain conclusions about the relative insurance value of different types of contracts. For example, the catastrophic contracts recommended by optimal insurance theory would have higher, not lower, insurance value than typical contracts—though still negative. Given the safety net's transformation of health spending risk, it is hard to find contracts in the predominant class—those that cover a portion of the insured's total health care costs—that do not worsen risk exposure. But outside of that class, alternative contracts that limited the displacement of safety net support could provide valuable insurance against health risk while at the same time limiting the worsening of non-health risk by standard contracts. One possibility is cash indemnity contracts that paid a fixed benefit based on the individual's health diagnosis. Whereas the traditional view is that cash indemnity contracts would provide less insurance than the share-of-costs contracts prevailing today (e.g., Zeckhauser, 1970), accounting for other risks reverses this conclusion.

The main contribution of this paper is to quantify, using a variety of approaches, the insurance value of health insurance and the importance of key underlying mechanisms. My findings build on and help reconcile two strands of related literature. The first strand seeks to quantify the overall value of health insurance, including not only its insurance value but also its transfer value and any other aspects that people might value (Dague, 2014; Gallen, 2015; Hackmann, Kolstad and Kowalski, 2015; Finkelstein, Hendren and Shepard, 2019; Finkelstein, Hendren and Luttmer, 2019). A key finding of this literature is that many people are willing to pay less, often much less, for health insurance than the actuarially fair value of gross health insurance benefits. That health insurance worsens risk exposure complements the main previous explanation for this finding, that the safety net reduces health insurance's transfer value and the value of its insurance of health risk (e.g., Mahoney, 2015; Finkelstein, Hendren and Luttmer, 2019). My findings complement and extend these findings by showing the crucial role of other risks, and their interaction with the health care-focused safety net, in not only reducing the insurance value of health insurance but even making it negative.

Negative insurance value can help explain why estimates of willingness to pay are often nearby and even below average *net* benefits, net of displaced safety net support (Finkelstein, Hendren and Shepard, 2019; Finkelstein, Hendren and Luttmer, 2019). More generally, my findings show how the safety net, by insuring a wide variety of risks, can transform the targeting properties of other, marginal sources of support. This could be important in other contexts beyond health insurance as well.

The second strand of related literature seeks to quantify the insurance value of health insurance.[2] To the best of my knowledge, all previous studies have concluded that health insurance has positive insurance value. In fact, the prevailing view is that the insurance value of health insurance is not merely positive but *inefficiently high*. Typical health insurance contracts cover not only catastrophic but also routine health care expenditures. Hence, the thinking goes, they provide more insurance than is worth buying given moral hazard and administrative costs (Feldstein, 1973; Pauly, 1986). My finding that health insurance has negative insurance value fits well with the evidence that its overall value is often low, and my investigation of mechanisms uncovers the key reason previous quantitative studies have concluded that health insurance has positive insurance value: by neglecting the relationship between health care spending and other risks. That seemingly innocuous modeling assumption, often made for simplicity or lack of data, causes the analysis to miss what surprisingly turns out to be health insurance's most consequential effect on risk exposure: its worsening of other risks.

2 Theory

This section defines insurance value, derives a sufficient statistic approximation to insurance value, and shows how non-health care risk and the safety net can combine to decrease health insurance's insurance value, potentially even to the point of making it negative. Here and throughout most of the paper, I focus on insurance value, setting aside other welfare-relevant effects of health insurance such as moral hazard costs and possible non-insurance benefits. I briefly return to these issues in Section 5.

[2]See, for example, Feldstein (1973); Feldman and Dowd (1991); Feldstein and Gruber (1995); Manning and Marquis (1996); Blomqvist (1997); Finkelstein and McKnight (2008); Engelhardt and Gruber (2011); Kowalski (2015); Finkelstein, Hendren and Luttmer (2019).

2.1 Insurance value of a change in constraints: Definition and sufficient statistic

Ex post value of an arbitrary change in constraints: V.— Ex post indirect utility in the realized state of the world θ is

$$v(B(\theta);\theta) \equiv \max_{x \in B(\theta)} u(x;\theta) \equiv u\left[x(B(\theta);\theta);\theta\right] \equiv u\left[c(B(\theta);\theta), a(B(\theta);\theta);\theta\right], \tag{1}$$

where $B(\theta)$ is the budget constraint, x is a vector of all goods and bads, $x(B(\theta);\theta)$ is the demand function, i.e., the choices that maximize utility subject to the constraint, and the final equivalency splits the vector of goods into two parts: a particular good ("consumption") c and "all other goods" a. This formulation allows for the possibility that preferences, budget constraints, or both depend on the realized state of the world.[3]

Consider an arbitrary change in the ex post budget constraint from $B_0(\theta)$ to $B_1(\theta)$. The ex post equivalent variation of this change, measured in terms of consumption c—i.e., the increase in consumption that would make the individual exactly as well off as they would be under the new constraint—is the V defined implicitly by the following equation:[4]

$$u\left[c(B_0(\theta);\theta) + V, a(B_0(\theta);\theta);\theta\right] = u\left[c(B_1(\theta);\theta), a(B_1(\theta);\theta);\theta\right]. \tag{2}$$

Ex ante value of an arbitrary change in constraints: EAV.— Consider an arbitrary change in ex post constraints from $B_0 = \{B_0(\theta)\}_{\theta\in\Theta}$ to $B_1 = \{B_1(\theta)\}_{\theta\in\Theta}$. The ex ante value, in terms of consumption c in all states of the world, of this change is the EAV defined implicitly by the following equation:

$$E\left\{u\left[c(B_0(\theta);\theta) + EAV, a(B_0(\theta);\theta);\theta\right]\right\} = E\left\{u\left[c(B_1(\theta);\theta), a(B_1(\theta);\theta);\theta\right]\right\}, \tag{3}$$

where the expectations are taken over the distribution of possible states of the world, $\theta \sim F(\theta)$. The ex ante value, EAV, is the increment to consumption c in all states relative to its levels under the original constraints such that the individual is exactly as well off ex ante as they would be under the new constraints.

[3]The von Neumann-Morgenstern utility function $u(\cdot;\theta)$ here could be expected continuation utility in a dynamic setting with multiple periods, in which case the state of the world θ is a "state-time" that embeds all relevant information that has been revealed up to that point, as in Chetty (2006). I use the simpler notation and language only for expositional simplicity; the theory applies to dynamic settings as well.

[4]I measure value in terms of consumption rather than income in order to avoid the measure itself having insurance effects. As I explain below, income is implicitly taxed by the health care-related safety net more in some states of the world than in others. As a result, a uniform increase in income in all states helps the individual more in some states than in others and so has insurance effects.

Sufficient statistic approximation to insurance value: $Cov\left(\widehat{\lambda},\ V\right)$.— A first order approximation to the ex ante value of the changes in constraints is

$$EAV \approx \frac{E\left(\lambda \times V\right)}{E\left(\lambda\right)} = \underbrace{E\left(V\right)}_{\text{Transfer value}} + \underbrace{Cov\left(\widehat{\lambda},\ V\right)}_{\text{Insurance value}}, \tag{4}$$

where λ is the marginal utility of consumption, $\widehat{\lambda} \equiv \lambda/E(\lambda)$ is the normalized marginal utility of consumption (normalized to have mean one), and the expectations and the covariance are across states of the world.[5] This follows Finkelstein, Hendren and Luttmer (2019) in decomposing the overall ex ante value into a "transfer" component equal to the mean ex post value and a "pure insurance" component that reflects the excess, if any, of the overall ex ante value over the mean ex post value.

The covariance between normalized marginal utility and the ex post value of the change in constraints,

$$Cov\left(\widehat{\lambda}, V\right) = E\left[\underbrace{\left(\widehat{\lambda} - E\left(\widehat{\lambda}\right)\right)}_{\text{"Marginal utility gap"}} \times \underbrace{\left(V - E(V)\right)}_{\text{"Value gap"}}\right], \tag{5}$$

is a first order approximation to the insurance value of the change.[6] Insurance value is increasing in the extent to which V is an *indicator* of marginal utility, i.e., in the extent to which the change in constraints benefits the individual more when times are tight (i.e., when marginal utility is high) than when they are loose (when marginal utility is low). The change in constraints has positive insurance value, i.e., is worth more ex ante than its mean ex post value, if and only if its ex post value covaries positively with marginal utility, i.e., if its value gaps tend to be the same sign as the associated marginal utility gaps. If instead its

[5]The approximation in equation (4), which would be exact if the marginal utility of consumption in each state were constant over the relevant range of consumption, follows straightforwardly from plugging equation (2) into equation (3) and taking first order approximations around the allocation under the original constraints. This first order approximation to the full ex ante value of the change in constraints is not necessarily the same as the marginal ex ante value of changing any particular parameter of the constraints. Given the hard-to-measure, state-contingent implicit taxation of health insurance benefits by the safety net, it is easier to estimate a first order approximation to the (full) ex ante value of completing health insurance (i.e., of going from the status quo to full health insurance) than to estimate the marginal ex ante value of changing any particular parameter of health insurance coverage (e.g., changing its coinsurance rate).

[6]This covariance, closely related to those in Finkelstein, Hendren and Luttmer (2019) and Lieber and Lockwood (2019), generalizes the familiar "marginal utility gap" from the benefit part of the Baily-Chetty optimal social insurance formula (Baily, 1978; Chetty, 2006) to situations in which the change in constraints potentially have richer effects than the simple binary structure envisioned by the literature on unemployment insurance (in which the ex post value of a change in unemployment insurance takes one of two values: the benefit increase in unemployed states of the world and the opposite of the tax increase in employed states). The (normalized) marginal utility gap for a given state is the markup on transfers from all states of the world to that state, $\left(\widehat{\lambda} - E\left(\widehat{\lambda}\right)\right) = \frac{\lambda - E(\lambda)}{E(\lambda)}$.

ex post value covaries negatively with marginal utility, the change in constraints has negative insurance value, i.e., an anti-insurance cost: It worsens risk exposure. Such a change, which benefits the individual more in states with lower marginal utility on average, increases the welfare cost of the risks the individual faces and so is worth less ex ante than its mean ex post value. A change in constraints with negative insurance value does not necessarily have negative overall ex ante value, just ex ante value that is below the mean ex post value.

This general framework nests a wide range of models, including ones with self-insurance, informal insurance, liquidity constraints, investments in health capital, state-dependent utility, many risks of varying persistence, many periods, many goods, etc. In a broad class of models, any effects that such factors and others might have on the insurance value of a change in constraints manifest themselves through this covariance. For instance, better self-insurance or borrowing opportunities would, other things equal, likely lead to smaller variation in marginal utility across states and so a smaller insurance value. As another example, if persistent health shocks tighten constraints not only by increasing current out-of-pocket health spending but also by decreasing current or future income, that would tend to increase the covariance between marginal utility and the ex post value of health insurance. The key advantage of aiming to recover a first order approximation to insurance value rather than its exact value is the significant reduction in the number and strength of assumptions required to do so. Rather than modeling the full data generating process, all one needs to know—exactly what one needs to know—is the covariance between normalized marginal utility and the ex post value. The key assumption is that the individual optimizes with respect to the constraints. With optimization, changes in constraints can be valued using the envelope theorem. The main modeling requirement is modeling marginal utility, since the ex post value is often fairly straightforward to deduce from a combination of economic logic and available data.[7]

2.2 How the safety net and non-health care risk can transform health insurance targeting

Whereas the general framework above is the basis for the key empirical tests in later sections, this section uses a highly stylized special case in order to develop intuition about the key mechanisms. Consider a simple model in which utility is a state-independent, increasing,

[7]This framework recovers to first order the value of changes in constraints, not any effects beyond those such as a direct effect on ex ante utility from "peace of mind." Although such factors might affect the full ex ante value of health insurance, it is less obvious that they would affect health insurance's pure insurance value, the focus of this paper.

concave function of non-health consumption c (hereafter, "consumption"),

$$c(\theta, HI) = y(\theta) - oop(\theta, HI), \tag{6}$$

where y is resources, *oop* is out-of-pocket health spending, and HI summarizes the health insurance contract the individual holds, if any. This simple model captures the key mechanism in most economic models of health insurance: Health insurance reduces out-of-pocket health spending and thereby protects consumption against the financial consequences of health shocks.

The insurance value of going from no health insurance to full health insurance is approximately

$$Cov\left(\widehat{\lambda}, V\right) \approx -\frac{\gamma(\overline{c})}{\overline{c}} Cov\left(c, oop\right) = \underbrace{\frac{\gamma(\overline{c})}{\overline{c}} Var(oop)}_{\text{"Direct effect"}} \underbrace{-\frac{\gamma(\overline{c})}{\overline{c}} Cov\left(y, oop\right)}_{\text{"Portfolio effect"}}, \tag{7}$$

where all variables take their values in the no-health insurance counterfactual (notation suppressed for simplicity), $\overline{c} \equiv E(c)$ is mean consumption, and $\gamma(\overline{c}) \equiv \frac{-\overline{c}u''(\overline{c})}{u'(\overline{c})} > 0$ is the coefficient of relative risk aversion at $\overline{c}$. The approximation is exact if marginal utility is linear in consumption.[8] Setting aside approximation error, health insurance has non-negative insurance value if and only if out-of-pocket health spending covaries negatively with consumption, since in that case eliminating out-of-pocket health spending tends to help the individual more in states in which consumption is lower and so marginal utility higher.

The rightmost part of equation (7) decomposes the covariance between out-of-pocket health spending and consumption into two parts. The "direct effect," $Var(oop)$, reflects that other things equal, greater out-of-pocket health spending reduces consumption. Insuring out-of-pocket health spending reduces consumption risk through this channel by reducing the extent to which health shocks drive down consumption, thereby conferring a valuable insurance benefit. The "portfolio effect," $-Cov(y, oop)$, comes from the correlation, if any, between out-of-pocket health spending and resources. This effect arises in second-best contexts in which people face uninsured risks other than that from out-of-pocket health spending. In such settings, health insurance's insurance value is determined not only by its impact on out-of-pocket health spending but also by its impact on the other risks the individual faces.

[8] In this simple model, the ex post value of going from any given status quo health insurance to full health insurance in a particular state equals status quo out-of-pocket health spending in that state. More generally, the ex post value of an increase in health insurance coverage would include the private values of any moral hazard response, which is second order, and of any reduction in reliance on the safety net (e.g., avoidance of associated stigma or financial costs), which could be first order. I return to these issues in Section 3.

The portfolio effect is valuable, further enhancing health insurance's overall insurance value, if out-of-pocket health spending covaries negatively with resources, since in that case health insurance benefits partially hedge the risk in resources as well as the risk in health care costs. The portfolio effect is costly, detracting from health insurance's overall insurance value, if out-of-pocket health spending covaries positively with resources, since in that case health insurance worsens the risk in resources. If such an effect is strong enough, it can dominate the direct effect. In the second best, insuring one risk does not necessarily improve risk exposure on net.

Common approaches rule out anti insurance by assumption.— Since the direct effect is always non-negative, $Var(oop) \geqslant 0$, equation (7) shows that a sufficient condition for health insurance to have non-negative insurance value is that its portfolio effect be non-negative as well, $-Cov\left(y, oop\right) \geqslant 0$, i.e., that out-of-pocket health spending covary non-positively with resources. Common approaches assume exactly that. One such approach assumes that out-of-pocket health spending is the only risky part of the constraint. In this case, the portfolio effect is zero and health insurance's insurance value is $\frac{\gamma(\bar{c})}{\bar{c}} Var(oop) \geqslant 0$.[9] More generally, approaches that assume that out-of-pocket health spending is exogenous to resources, say because out-of-pocket health spending is fully determined by health alone, rule out that health insurance has negative insurance value.[10]

How health insurance can worsen risk exposure: non-health care risk.— For health insurance to worsen risk exposure on net, (i) health insurance must worsen certain risks, and (ii) the cost of such worsening must exceed the benefit of any improving of other risks. In terms of equation (7), it must be that the portfolio effect is sufficiently negative to outweigh the direct effect, $-Cov(y, oop) < -Var(oop)$ or, equivalently, $Corr(y, oop) > Sd(oop)/Sd(y)$. Given that bad health is associated with lower earnings as well as higher out-of-pocket health spending, for the portfolio effect to be negative at all, i.e., for out-of-pocket health spending to covary positively with resources, it must be that some other

[9] Although non-health care risk is a necessary condition for *completing* health insurance (i.e., going from status quo health insurance to full health insurance) to have negative insurance value, incomplete expansions of health insurance could potentially have negative insurance value even without such risk. If safety net support is non-linear in otherwise-uninsured health care costs (e.g., if safety net support approximates catastrophic coverage above a deductible, as Mahoney's (2015) evidence suggests is the case for bankruptcy and the threat thereof), incomplete expansions of health insurance (e.g., going from no health insurance to a contract that covers, say, 85% of total health care costs, as many private and [non-Medicaid] government contracts do) can worsen risk exposure. Intuitively, the net private ex post value of such expansions, net of implicit taxation by the safety net, are concentrated not in the worst, highest-marginal utility states but in less bad, lower-marginal utility states. This is because, at the margin, health insurance benefits mostly displace safety net support in the worst states.

[10] If out-of-pocket health spending is exogenous to resources and bad health not only increases out-of-pocket health spending but also decreases resources (consistent with a variety of empirical evidence on the income costs of bad health, e.g., Dobkin et al., 2018), the portfolio effect is positive, $-Cov\left(y, oop\right) \geqslant 0$, and so further increases health insurance's insurance value beyond the (positive) direct effect.

factor is a strong enough force toward out-of-pocket health spending covarying positively with resources as to dominate the force in the other direction from health risk. Whatever the drivers, health insurance is more likely to worsen risk exposure on net the more precisely it targets higher-resource states ($Corr(y, oop)$ more negative) and the smaller the variation in out-of-pocket health spending relative to that in resources ($Sd(oop)/Sd(y)$ smaller). But why might health insurance worsen non-health care risk, even to the point of targeting higher-resource states on net? And why might variation in out-of-pocket health spending relative to that in resources be small? The uneven, health care-focused safety net.

How the safety net affects health insurance targeting.— The safety net—charity care, bankruptcy, bad debt, and means-tested programs—provides substantial protection against otherwise-uninsured health care costs. For example, Finkelstein, Hendren and Luttmer (2019) estimate that low-income uninsured adults in the Oregon Health Insurance Experiment paid only about 20% of their medical care costs, with the remaining 80% paid by other parties. Although a growing body of research documents the quantitative importance of such protection (e.g., Garthwaite, Gross and Notowidigdo, 2018) and investigates the implications for the demand for health insurance (e.g., Mahoney, 2015; Finkelstein, Hendren and Shepard, 2019; Finkelstein, Hendren and Luttmer, 2019) (see Finkelstein, Mahoney and Notowidigdo, 2018, for a review), less appreciated is the extent to which such protection fundamentally reshapes health insurance's targeting properties and so its insurance value. As a "secondary payer," the safety net protects against *otherwise-uninsured* health care costs. Such support is necessarily displaced by health insurance.[11] As a result, the safety net implicitly taxes health insurance. Crucially for health insurance's insurance value, the safety net does not implicitly tax health insurance to the same extent in all states of the world. Because the safety net helps most in the worst states, it implicitly taxes health insurance most in the worst states. This decreases health insurance's insurance value—potentially even making it negative.

The safety net decreases health insurance's insurance value in two ways. First, the safety net decreases the value of health insurance's insurance of health care costs. Safety net protection against otherwise-uninsured health care costs decreases both the extent to which consumption varies with health care costs and the extent to which health insurance targets high-cost states. In terms of equation (5), this decreases both the marginal utility and value gaps in the high-cost states that would otherwise drive health insurance's value. In terms of equation (7), this decreases $Var(oop)$ and so the valuable "direct effect."[12] Second, the

[11]For example, Dobkin et al. (2018) find that health insurance reduces households' unpaid medical bills, Garthwaite, Gross and Notowidigdo (2018) find that health insurance reduces the provision of uncompensated care by hospitals, and Finkelstein, Hendren and Luttmer (2019) estimate that 60% of Medicaid spending is a transfer to providers of uncompensated care for the low-income uninsured.

[12]Because of these reinforcing effects on both the marginal utility and value gaps—revealed by the for-

safety net causes health insurance to worsen non-health care risk. The safety net's greater protection against otherwise-uninsured health care costs when non-health care circumstances are worse provides valuable insurance against non-health care risk—insurance that health insurance undoes without replacing. In terms of equation (5), this decreases the value gaps in the high-marginal utility states with bad realizations of non-health care risk. In terms of equation (7), this is a force toward out-of-pocket health spending covarying positively with resources and so a costly portfolio effect of health insurance, $-Cov(y, oop) < 0$. Recall that in terms of equation (7), health insurance's net effect is to worsen risk exposure if $Cov(y, oop) > Var(oop)$, or equivalently if $Corr(y, oop) > Sd(oop)/Sd(y)$. The safety net increases $Corr(y, oop)$ and decreases $Sd(oop)$, both of which are forces toward health insurance worsening risk exposure on net.[13]

Figure 1 shows visually how the combination of non-health care risk and the safety net can affect the insurance value of health insurance. It depicts the results of a stylized simulation in which consumption equals resources less out-of-pocket health spending (as in equation (6)), marginal utility is linear in consumption, and health care and non-health care risks are independent, with $h \sim unif[\$0,\ \$10,000]$ and $y \sim unif[\$40,000,\ \$60,000]$. Without a safety net, there is a strong positive relationship between the ex post value of health insurance, $V = oop = h$, and marginal utility, since without health insurance, for any given level of gross resources y, worse health care shocks leave fewer net resources for consumption. Health insurance has positive insurance value in this case because it insures health care risk (helping more when health care costs are higher, $cov(V, h) = cov(oop, oop) = var(oop) > 0$) and is neutral with respect to other risks ($cov(V, y) = cov(oop, y) = cov(h, y) = 0$).[14]

The safety net can fundamentally transform the relationship between the value of health insurance and marginal utility, potentially even switching the sign of the correlation. The arrows in Figure 1 show the effects of a stylized safety net that provides catastrophic coverage above a resources-dependent deductible, $oop = \min\{h,\ \delta_0 + \delta_1 y\}$, in the realistic case in which

mula's dependence on the *square* of the standard deviation of out-of-pocket health spending: $Var(oop) = [Sd(oop)]^2$—the safety net can decrease the value of health insurance's insurance of health risk considerably. For example, if the safety net covered 80% of the health care costs of individuals without insurance in all states (ignoring that it covers a larger share in states with higher spending but in keeping with the estimates of Finkelstein, Hendren and Luttmer (2019) and the Kaiser Commission on Medicaid and the Uninsured (Coughlin et al., 2014) that out-of-pocket health spending by the uninsured is only about 20% of their total health care costs on average), that would reduce the direct effect by a factor of *25*: $V(tot)/V[(1-0.8)tot] = 1/(1-0.8)^2 = 1/.04 = 25$. Accounting for the safety net's greater protection in higher-cost states would likely make the effect even larger.

[13]To be sure, the safety net as a whole also decreases $Sd(y)$, which decreases the cost of worsening risk in resources. But to the extent that the safety net provides more protection against risk in otherwise-uninsured health care costs than in resources, it tends to decrease $Sd(oop)/Sd(y)$ and so be a force toward health insurance worsening risk exposure on net.

[14]The slope of the population regression of marginal utility on the value of health insurance is $\beta_{\lambda|V} = Cov(\lambda, V)/Var(V)$. So the sign of the population regression slope is the same as the sign of the first order approximation to health insurance's insurance value, $sign(\beta_{\lambda|V}) = sign[Cov(\lambda, V)]$.

the safety net provides greater protection against otherwise-uninsured health care costs when gross resources are lower, $\delta_1 > 0$. Safety net protection against otherwise-uninsured health care costs shifts states southwestward: west (reduction in V) because it reduces the out-of-pocket health spending remaining for health insurance to cover and south (reduction in λ) because, by reducing out-of-pocket health spending, it increases consumption. Crucially, the safety net does not shift each state equally, which would leave the slope and so the sign of health insurance's insurance value unchanged. Instead, it shifts states with worse health care or resources farther. The key effect for health insurance's insurance value is that this shifts the bad-health states that otherwise would have contributed most to its insurance value from the valuable northeast quadrant (positive value and marginal utility gaps) southwestward toward the costly northwest quadrant (negative value gap and positive marginal utility gap). This is a potentially strong force toward health insurance worsening risk exposure on net. Although health insurance would continue to insure (the diminished) health care cost risk ($cov(V, h) = cov(oop, h) > 0$), it would no longer be neutral with respect to other risk but worsen it ($cov(V, y) = cov(oop, y) < 0$). In this case, whether health insurance has positive insurance value depends on whether the value of its insurance of health care cost risk exceeds the cost of its worsening of other risk, which in turn depends on the extent of each of these risks and the sensitivity of safety net support to them.[15,16]

Summary.— Were it not for non-health care risk, health insurance would have a straightforward effect on risk exposure: It would help more when health care costs were higher and thereby provide valuable insurance against health care risk. But individuals face not only health care risk, which health insurance is specifically designed to target well, but also non-health care risk, which it is not. This section showed how the safety net can transform health insurance from being neutral with respect to non-health care risk to systematically worsening it.[17] Because of non-health risk and the safety net, it is as if health insurance were a bundle of two assets: one that pays off more when health care costs are higher, which provides valuable insurance against health care cost risk, and one that pays off more when non-health

[15] If safety net support, and so the size of these shifts, were independent of resources, health insurance would continue to have positive, albeit reduced, insurance value, since it would not worsen other risk. In terms of the figure, an alternative safety net that provided catastrophic coverage above a uniform deductible (independent of resources), $oop = \min\{h,\ \delta\}$, would shift states onto a line that was vertical, not downward-sloping as in the figure. Such shifts would leave the slope, and so the sign (though not the magnitude) of health insurance's insurance value unchanged.

[16] Appendix Section A characterizes analytically the effects on the insurance value of health insurance of a simple model of the effect of the safety net on out-of-pocket health spending.

[17] In principle, other mechanisms could potentially cause health insurance to worsen non-health care risk. One is correlated risks. If bad non-health shocks were sufficiently negatively correlated with bad health shocks, health insurance would worsen non-health risk. Another possibility is responsiveness of the demand for health care to non-health shocks, perhaps due to demand responses to liquidity (as documented, for example, by Gross, Layton and Prinz, 2020). Empirically, collectively these two mechanisms are not strong enough to make health insurance worsen risk exposure on net. See footnote 26 and Appendix Section D for details.

care circumstances are better, which worsens other risks. Intuitively, whereas the portion of the variation in out-of-pocket health spending driven by health shocks is a source of risk and so is valuable to insure, the portion driven by non-health care circumstances absorbs risk—it hedges risk in those circumstances—and so is costly to insure. Health insurance's overall effect on risk exposure therefore reflects competing pro- and anti-insurance effects and so could be to either improve or worsen risk exposure on net, depending on the size of health care and non-health care risks (marginal utility gaps) and the extent to which health insurance insures or worsens them (value gaps). Which effect dominates is an empirical question. The key lesson of the theory for the empirical work is to use flexible approaches that admit both possibilities, so that the result can be determined by the key empirical relationships, not modeling assumptions.

3 Data and Empirical Approaches

3.1 Sufficient statistic approach and the PSID data

As shown in Section 2, in a broad class of models, to first order the insurance value of any arbitrary change in constraints is the covariance, across states of the world, between normalized marginal utility and the ex post value of the change, $Cov\left(\widehat{\lambda}, V\right)$. The key data requirements to estimate this statistic are (i) measures of consumption (to model marginal utility without having to model constraints in their entirety), (ii) measures of health care utilization or spending (to model the ex post value of changes in health insurance), and (iii) rich enough measures of household characteristics to isolate risk from predictable heterogeneity. Although these are demanding requirements that are not met perfectly by any available data, the Panel Study of Income Dynamics (PSID) has many important advantages for these estimations. It includes fairly detailed data on different kinds of consumption spending. It includes out-of-pocket health spending. And its rich demographic measures and long panel structure enable estimation strategies that better isolate risk from predictable heterogeneity. Appendix Table A1 shows summary statistics of the main estimation samples.

My main estimates of insurance value are based on regressions of the form,

$$\widehat{\lambda}_{it} = \alpha + \beta V_{it} + \gamma X_{it} + \varepsilon_{it}, \tag{8}$$

where i is a household, t is calendar time, $\widehat{\lambda} \equiv \frac{\lambda}{E(\lambda|X)}$ is normalized marginal utility of income, V is the ex post value of health insurance, and X_{it} is a vector of controls. The regression results can be converted into the key covariance using the relationship $Cov(\widehat{\lambda}, V) =$

$\beta Var(V|X)$, where $Var(V|X)$ is the residualized variance of V (i.e., the variance of the residuals from a regression of V on the controls X).

Normalized marginal utility, $\widehat{\lambda}_{it} = \lambda_{it}/E(\lambda|X)$.— My main specifications follow the common practice of using measured non-health consumption spending together with an assumed functional form of marginal utility of income to proxy for marginal utility. Intuitively, this approach is based on the idea that a household's chosen level of consumption spending reveals the overall tightness of its constraint, bypassing the need to model the constraint in its entirety. Given the possibility of measurement error and the sensitivity of marginal utility to low consumption levels, I impose a consumption floor of $5,000. I assume constant relative risk aversion utility over consumption with a coefficient of relative risk aversion of three: $\lambda_{it} = c_{it}^{-3}$. I test the robustness of the results to many alternative assumptions about marginal utility. The key advantage of using measured rather than modeled consumption is that it means the key relationship, that between consumption and out-of-pocket health spending, is determined by the data, not modeling assumptions.

Ex post value, V_{it}.— The ex post value of a change in health insurance coverage depends on the change in coverage under consideration. My main analysis considers completing health insurance, i.e., expanding health insurance coverage from its status quo level (which for the uninsured is no health insurance and for the insured is whatever coverage they have) to full health insurance (i.e., eliminating out-of-pocket health spending). As a baseline, I assume that to first order the ex post value in a particular state of the world of such a change is status quo out-of-pocket health spending in that state: $V_{it} = oop_{it}$. This follows the literature on health insurance, which focuses on the impact of health insurance on out-of-pocket health spending, and captures to first order the value of health insurance if there are no private costs to the household of relying on the safety net.[18] I also test the robustness of the results to alternative assumptions about the value to the household of health insurance reducing the household's reliance on the safety net (since there can be private costs to the household of relying at least in part on the safety net, e.g., stigma costs, hassle costs, and financial costs of declaring bankruptcy or worsening future credit access). I also consider other changes in health insurance coverage with different ex post values.

Controls, X_{it}.— Given my focus on insurance, the key covariance is across states of the world for a given household.[19] I follow the common practices of using cross-sectional heterogeneity

[18]This follows from the usual envelope theorem logic that to first order the value to an optimizing household of a reduction in the price of a good is the associated "mechanical" reduction in the household's expenditure on that good—the reduction in expenditure that would occur if the household did not adjust its consumption of the good in response to the change. Although the household might in fact adjust its consumption of this or other goods in response to the price change, any re-optimization gains from doing so are second order.

[19]Although I focus on insurance given a household's type, it is worth mentioning that unlike many

as a proxy for the risk facing any particular household and of using control variables and panel data specifications to try to isolate risk from ex ante predictable heterogeneity. As a baseline, I control for a cubic in age, a quadratic in household size, and year dummies (as in Finkelstein, Hendren and Luttmer, 2019), plus education category dummies (to try to limit to within-earning ability group comparisons). I also test robustness to many alternatives. Standard errors are always clustered by household.

3.2 Other approaches and the MEPS data

In addition to estimating sufficient statistic approximations to the insurance value of various changes in health insurance coverage, I also use a variety of complementary analyses to quantify and understand health insurance's insurance value. I investigate health insurance's targeting properties with respect to income, unemployment, hospitalization, and other health and non-health risks. I investigate how the safety net transforms health spending risk. And I use a structural model to quantify the value of changes in coverage beyond those that are well-suited to the sufficient statistics approach and to investigate the importance of different mechanisms.

For many of these analyses, it is useful to have richer measures of health care consumption and spending than those available in the PSID. I therefore supplement the PSID data with data from the Medical Expenditure Panel Survey (MEPS). The main advantage of the MEPS is its rich, high-quality information on health care consumption and expenditures. This enables me to investigate the role of the safety net in shaping health insurance's targeting properties and insurance value. Appendix Table A2 shows summary statistics of the main estimation samples.

applications of differential commodity taxes or social insurance programs, even a government with a nonlinear income tax would generally wish to consider at least some of the effects of health insurance on redistribution across types (which can be viewed as insurance against risk from behind the veil) in addition to insurance for each type. Some health insurance-induced redistribution occurs *within* earnings groups (e.g., from those in better to worse health), across which the income tax cannot redistribute. Hence, the usual result about the dominance of the income tax over differential commodity taxes (Atkinson and Stiglitz, 1976) does not apply due to welfare-relevant preference heterogeneity. This is an interesting issue for future research.

4 The Insurance Value of Health Insurance: Estimates and Mechanisms

4.1 Health insurance worsens risk exposure

I begin by estimating first-order approximations to the insurance value of completing health insurance, i.e., of going from status quo health insurance to full health insurance as, for example, Medicaid expansions do. In a broad class of models, to first order this insurance value is the covariance, across states of the world, between normalized marginal utility and out-of-pocket health spending. I estimate this covariance using regressions as in equation (8). This covariance is roughly proportional to the opposite of the covariance between non-health consumption and out-of-pocket health spending (as shown in equation (7)). This covariance is the key empirical relationship driving the results.

Table 1 presents estimates of the insurance value of completing health insurance in each of four sets of state-times: those in which the household head is non-elderly, non-elderly and without health insurance, non-elderly and with health insurance, and elderly and with health insurance. In all cases, the estimated insurance value is negative. Appendix Table A3 shows that the estimated insurance value is negative for each education group and for each of several different types of state-times: ones in which lagged liquidity is low versus high, age is younger versus older, and lagged health is good versus bad. Appendix Table A4 shows that the results are robust to a wide range of changes in the specification, including one that aims to overstate the value of reducing reliance on the safety net. Appendix Table A15 shows estimates of the insurance value of completing coverage of each of the three types of health care spending in the PSID (hospital/nursing, doctor visits, and prescriptions), which are also all negative.[20]

The estimates imply that the welfare cost of the resulting worsening of risk exposure is economically significant. For non-elderly uninsured households, my baseline estimates imply that the worsening of their risk exposure from holding full health insurance like Medicaid would have a welfare cost equivalent to decreasing their annual income by $310. This anti-insurance cost is 33% of the mean ex post value (the transfer value in equation (4)). So to first order, full health insurance coverage is worth 67 cents per dollar of expected net benefits ex ante ($1 of transfer value less 33 cents of anti-insurance cost), or $610 per year.

[20] Finkelstein, Hendren and Luttmer (2019) estimate the covariance between marginal utility and out-of-pocket health spending in both the PSID and the Consumer Expenditure Survey using a variety of specifications. They, too, find that the estimated insurance value is robustly negative. Given the striking inconsistency of a negative insurance value with leading theories of health insurance (and of insurance more generally), they concluded that these unexpected results may have been driven by measurement error.

For non-elderly insured households, the worsening of risk exposure from completing their health insurance (filling in the gaps in their current coverage) would have a welfare cost equivalent to decreasing their annual income by $320, 23% of the associated transfer value. For elderly insured households, the worsening of risk exposure from completing their health insurance (filling in the gaps in their Medicare and other coverage) would have a welfare cost equivalent to decreasing their annual income by $130, 7% of the associated transfer value.[21]

Proximate cause: The strongest indicator of high marginal utility is unusually low*, not high, out-of-pocket health spending.—* Figure 2, an empirical analogue of the theoretical Figure 1, shows a non-parametric estimate of the conditional mean of normalized marginal utility as a function of out-of-pocket health spending. This figure reveals the strong negative relationship between marginal utility and out-of-pocket spending that drives the negative insurance value estimates. On average, higher out-of-pocket health spending is associated with lower, not higher, marginal utility. Among states of the world with low out-of-pocket health spending, those in which marginal utility is high dominate, and among states of the world with high out-of-pocket health spending, those in which marginal utility is low dominate. As a result, exposure to status quo out-of-pocket health spending provides a valuable hedge on net: Out-of-pocket health spending is more "risk absorber" than "risk creator."

4.2 Health insurance worsens income risk

The sufficient statistic estimates indicate that health insurance worsens risk exposure on net.[22] In terms of equation (7), this suggests that the valuable "direct effect" of insuring health care costs is dominated by a costly "portfolio effect" of worsening other risks. This section investigates health insurance's targeting properties with respect to income risk. Income risk is for many households the most important risk they face. It reflects the combined effects of many specific underlying risks, health and non-health. That bad health decreases income and increases out-of-pocket health spending is a force toward out-of-pocket health spending covarying negatively with income, which is a force toward health insurance improving income risk. So absent a stronger countervailing effect, one would expect health insurance to improve income risk.

[21] To be clear, these are estimates of the insurance value of going from status quo health insurance to full health insurance, as a beneficiary of a Medicaid expansion does. As such, they do not reveal the insurance value of any (inframarginal) health insurance that households hold in the status quo. It could be that status quo health insurance has higher insurance value than full health insurance, or that health insurance has higher insurance value in states of the world in which households hold health insurance than in states in which they do not. I return to these issues in Section 4.4.

[22] The key advantage of the sufficient statistics approach—that under certain assumptions, all one needs to know to answer (to first order) the question of interest is the value of the sufficient statistic—has an associated drawback: that the statistics themselves reveal little about the underlying mechanisms. Much of the remainder of the paper is devoted to understanding the underlying mechanisms and their implications.

Table 2 shows the results of regressions of income on out-of-pocket health spending and the baseline controls, with associated heterogeneity and robustness analyses in Appendix Tables A5 and A6, respectively. For each of several household types and each of a wide variety of specifications, out-of-pocket health spending and income covary strongly *positively*. For non-elderly households, the correlation is 0.16, and on average a $1 higher realization of out-of-pocket health spending is associated with almost $5 higher income. That out-of-pocket health spending covaries positively with income means that health insurance worsens income risk on net. Holding health insurance is, in part, like holding stock in one's employer: The payoffs are positively correlated with income, thereby increasing the welfare cost of income risk.[23]

How costly is health insurance's worsening of income risk? One rough but revealing measure, based on equation (7), is

$$\frac{\text{Cost of worsening income risk}}{\text{Ben of improving health care risk}} = \frac{\frac{\gamma(\bar{c})}{\bar{c}} Cov\left(y, oop\right)}{\frac{\gamma(\bar{c})}{\bar{c}} Var(oop)} = \beta_{y|oop} = \frac{Sd(y)}{Sd(oop)} Corr(y, oop), \quad (9)$$

where $\beta_{y|oop}$ is the coefficient on out-of-pocket health spending in a regression of income on out-of-pocket health spending and controls. This links the regression in Table 2 to rough measures of the welfare cost of worsening income risk and welfare benefit of improving health care cost risk. The results, viewed through the lens of this simple model and across different types of households, imply that the worsening of income risk is about three to five times more costly than the improving of health care cost risk is beneficial.[24]

Why is the worsening of income risk so much more costly than the improving of health care cost risk is beneficial? The final equality of equation (9) suggests two key factors. First, health insurance payoffs are fairly strongly positively correlated with income ($Corr(y, oop) \approx 0.16$) and so anti-target income risk fairly precisely. Second, variation in income is far greater than that in out-of-pocket health spending ($Sd(y)/Sd(oop) \approx 31$), so a worsening of income risk is much more costly than a same-sized improvement of health care cost risk is beneficial.[25]

[23] I also find that health insurance worsens unemployment risk. On average, health insurance is worth about $230 less in years in which the household experiences unemployment than in other years (see Appendix Table A8). This suggests that health insurance is, in part, negative UI, effectively undoing about 60% of one week's worth of average UI benefits.

[24] This measure based on equation (7) is rough because it assumes that marginal utility is linear in consumption and that consumption equals income minus out-of-pocket health spending. A more general, though less revealing, measure, derived in Appendix Section C, suggests that health insurance's worsening of income risk can account for about half of its overall net anti-insurance cost (see Appendix Table A7).

[25] This analysis also shows the bias in common approaches that take out-of-pocket health spending as an exogenous source of risk, independent of resources. If out-of-pocket health spending were exogenous, it would be purely a risk-creator, not a risk-absorber. The estimates in Table 2, when viewed through the lens of equation (9), suggest that out-of-pocket health spending's risk-absorber role is about 3–5 times more consequential for welfare than its risk-creator role, so the common approach of neglecting the relationship between out-of-pocket health spending and other exposures leaves out more than it captures.

4.3 The safety net transforms health insurance targeting

Why does health insurance worsen income risk, and why does this worsening appear to be so much more costly than health insurance's improving of health care cost risk is beneficial? A variety of evidence suggests that the key driver is the safety net. The safety net—charity care, bad debt, bankruptcy, and means-tested support—provides significant protection against otherwise-uninsured health care costs. Garthwaite, Gross and Notowidigdo (2018) estimate that each additional uninsured person costs hospitals approximately $800 each year. The Consumer Financial Protection Bureau estimates that unpaid medical bills comprise a majority of all collections lines on consumers' credit reports and affect nearly one-fifth of credit reports (CFPB, 2014). Dranove, Garthwaite and Ody (2015) estimate that new unpaid hospital debt totals about $20 billion per year. Perhaps the most powerful, albeit indirect, measure of safety net support is "missing" out-of-pocket health spending:

$$oop = tot - hi - sn, \tag{10}$$

where tot is total health care costs, hi is (gross) health insurance benefits, and sn is inferred safety net support (e.g., uncompensated care). This attributes the gap, if any, between net-of-health insurance health care bills and out-of-pocket health spending to the safety net. For households without health insurance, this gap is large. The Kaiser Commission on Medicaid and the Uninsured estimates that the uninsured pay for only about 20% of their health care costs out-of-pocket (Coughlin et al., 2014), and Finkelstein, Hendren and Luttmer (2019) estimate that low-income uninsured adults in the Oregon Health Insurance Experiment also paid only about 20% of their medical care costs out-of-pocket.

Implicit catastrophic health insurance with a state-contingent deductible.— The safety net helps more when health care costs are higher and non-health care circumstances are worse. In terms of health care costs, Figure 3a, analogous to Figure 1A in Mahoney (2015), shows a non-parametric estimate of the conditional mean of the sum total of payments by health insurers (health insurance benefits) and households (out-of-pocket health spending) as a function of charges, a rough measure of health care utilization. At low charges, total payments by insurers and households are similar for uninsured and insured households. But as charges increase, a gap opens up, with total payments for insured households increasing roughly linearly in charges whereas total payments for uninsured households level off around $3,000–$5,000. The difference presumably, in keeping with the evidence on direct measures of safety net support, arises from greater reliance on the safety net among uninsured households. As a rough approximation, the safety net's effect on out-of-pocket health spending matches that of full catastrophic health insurance above a relatively low deductible (Mahoney, 2015).

In terms of non-health care circumstances, Figure 3b shows a non-parametric estimate of the conditional mean of out-of-pocket health spending as a function of income among non-elderly households without health insurance with annual health care charges of at least $20,000. This figure, like the previous one, reveals the large extent of implicit health insurance from the safety net: Mean out-of-pocket health spending is far below charges across all income levels. But the main point is that safety net support is decreasing in the realization of income. In this sample, average charges are *negatively* related to income, so the positive relationship between out-of-pocket health spending and income is presumably driven by lower-income households receiving more support from the safety net. As a result, safety net protection against otherwise-uninsured health care costs implicitly insures income risk.

Implications for health insurance targeting.— These patterns of safety net protection against otherwise-uninsured health care costs undermine the insurance value of health insurance in two ways. Key to both is the "secondary payer" nature of the safety net. The safety net protects against *otherwise uninsured* health care costs. Health insurance reduces otherwise-uninsured health care costs—that is its core function—and thereby displaces support that the safety net otherwise would have provided. In this way, the safety net implicitly taxes health insurance benefits, reducing the ex post value of health insurance in states in which health insurance displaces safety net support.

Health insurance's insurance value can be written

$$Cov\left(\widehat{\lambda}, V\right) = \underbrace{Cov\left(\widehat{\lambda}, hi\right)}_{\text{Gross HI benefits}} - \underbrace{Cov\left(\widehat{\lambda}, \tau hi\right)}_{\text{Displaced safety net support}}, \tag{11}$$

where $V = (1-\tau)hi$ is the ex post value of health insurance, hi is the gross health insurance benefit, and τ is the (possibly state-contingent) implicit average tax rate—the proportional reduction in the ex post value of health insurance to the household due to displaced support from the safety net. Implicit taxation of health insurance benefits by the safety net essentially "carves out" a chunk of health insurance benefits from health insurance's pattern of (net) ex post value. Were implicit taxation uniform across states (constant τ), it would scale down health insurance's insurance value but not switch its sign: $Cov\left(\widehat{\lambda}, (1-\tau)hi\right) = (1-\tau)Cov\left(\widehat{\lambda}, hi\right)$. But the evidence above suggests that the safety net taxes health insurance most in the worst, highest-marginal utility states. Health insurance's insurance value is positive only if its implicitly-taxed-away benefits, τhi, would not have provided greater insurance value (the second term on the right-hand-side) than its benefits as a whole would have if, counterfactually, they were not implicitly taxed (the first term on the right-hand-side). The better targeted is safety net protection against otherwise-uninsured health care

costs in terms of marginal utility, the more the safety net reduces health insurance's insurance value, potentially even to the point of making it negative.

The first way in which the safety net undermines the insurance value of health insurance is by reducing the value of health insurance's insurance against health care costs. In Figure 3a, mean out-of-pocket spending by the uninsured never rises above about $5,000, even among households with tens of thousands of dollars of charges. Hence, the nominally uninsured, though lacking formal health insurance, hold substantial implicit health insurance from the safety net, which, at least in terms of its impact on out-of-pocket health spending, closely resembles the type of catastrophic coverage above a deductible recommended by optimal insurance theory (Mahoney, 2015; Finkelstein, Mahoney and Notowidigdo, 2018). In other words, the safety net has "already taken" the most valuable chunk of insurance that health insurance otherwise could have provided, leaving mainly the relatively minor risk that optimal insurance theory recommends *not* covering (since the insurance value would be outweighed by administrative and moral hazard costs). This reduces both marginal utility and the ex post value of health insurance in the high-cost states that would otherwise drive health insurance's value.

As a rough measure of this effect, recall that in equation (7) the "direct effect" benefit of insuring health care costs is proportional to the variance of out-of-pocket health spending. In the MEPS, I estimate that among uninsured households the variance of out-of-pocket health spending is only about 2% of the variance of total health care costs, suggesting that the safety net reduces the direct effect by about two orders of magnitude. So while *total* insurance against health care costs, including that provided by the safety net, is likely extremely valuable given the possibility of catastrophically high costs, the additional, *marginal* insurance that health insurance can provide on top of that already provided by the safety net is unlikely to be very valuable. In this way, the safety net "steals health insurance's thunder."

The second way in which the safety net undermines the insurance value of health insurance is by causing health insurance to worsen other risks, beyond health care costs. Consider income. The positive correlation between income and out-of-pocket health spending arises despite income being virtually uncorrelated with total health care costs and gross health insurance benefits.[26] This suggests that the positive correlation between income and out-

[26]Estimates of the correlation between income and out-of-pocket health spending are reported in Tables 2, A5, A6, and A9. Estimates of the correlations between income and total health care costs and between income and gross health insurance benefits are reported in Appendix Table A9. That income and out-of-pocket health spending are positively correlated despite income being virtually uncorrelated with total health care costs and gross health insurance benefits is true for the insured as well as the uninsured, since, as Dobkin et al. (2018) and others have shown, the insured have some of their otherwise-uninsured health care costs paid for by other parties as well. That income is virtually uncorrelated with total health care

of-pocket health spending is driven by a negative correlation between income and safety net support, such as that suggested by Figure 3b. So even though "health insurance itself" is (i.e., gross health insurance benefits are) roughly neutral with respect to other risks, by displacing safety net support health insurance systematically worsens other risks. In other words, the safety net transforms health insurance from being roughly neutral with respect to other risks to worsening them systematically.[27]

In these two reinforcing ways, the safety net could potentially transform health spending from major risk creator to net risk absorber and so cause health insurance to worsen risk exposure on net. The safety net's substantial protection against catastrophic health care costs greatly reduces the extent to which health care costs create risk. The safety net's greater protection against health care costs when non-health shocks are worse makes health spending into an absorber of non-health risk. In other words, exposure to the residual, net-of-safety-net health care cost risk hedges non-health risk.

Consider hospitalization. Using a variety of definitions and estimation strategies, I find that health insurance worsens hospitalization-related risk (see Appendix Tables A10 and A11). To understand why, note that hospitalization is associated with not only high health care costs but also heterogeneous, sometimes substantial income losses (Dobkin et al., 2018). So the value of health insurance's insurance of hospitalization risk depends on its insurance or worsening of the associated income losses, not just the health care costs. Figure 4a decomposes the markup on (insurance value per dollar of transfer value of) health insurance's hospitalization-related targeting into the sum of three terms (see Appendix Table A12 for related statistics). The first is health insurance's average net transfer from non-hospitalization to hospitalization states. On average, health insurance is worth \$1,150 more in hospitalization than non-hospitalization states, which, given the higher average marginal utility in hospitalization states, provides moderately valuable insurance. This "hospitalization indemnity" component contributes 0.20 to the overall markup on health insurance's hospitalization-related targeting. The second term is health insurance's average net transfer from hospitalization states with lower- to higher- income realizations, defined as below versus above the 25th percentile of residualized income, residualized with the baseline controls. On average, health insurance is worth \$960 more in hospitalization states with above- than below-25th percentile income, which worsens income risk. This component contributes -0.21 to the overall markup on

costs casts doubt on the ability of two alternative mechanisms, other than the safety net, to explain why health insurance worsens other risks: a negative correlation between health care costs and other risks and a sensitivity of demand for health care to other risks. See Appendix Section D for details.

[27]The unevenness of the safety net—its much greater protection against health care costs than income and other risks—can help explain why even a relatively modest worsening of other risks can be more costly than larger improvements of health care cost risk are beneficial. It also means that safety net protection against otherwise-uninsured health care costs could comprise a non-negligible share of insurance against other risks—insurance that health insurance undoes.

health insurance's hospitalization-related targeting, roughly cancelling out the value of the hospitalization indemnity component. The third term is health insurance's transfers within hospitalization-by-income categories, from states in which health insurance is less to more valuable. On average, these transfers flow from higher- to lower-marginal utility states, which worsens within-category risk. This component contributes −0.65 to the overall markup on health insurance's hospitalization-related targeting, resulting in a total markup from these three components of −0.66 and an overall insurance value of −$75 per year or about −$770 per hospitalization state. Health insurance worsens hospitalization risk because although it provides moderately-valuable insurance against the health care costs associated with hospitalization, it also undoes some of the implicit insurance against the associated earnings losses provided by the health care safety net, at significantly greater welfare cost.

What is the role of the safety net in driving these results? Although I postpone the main counterfactual analysis to the next section, it is useful to consider a simple analysis here. Figure 4b, an empirical analogue of the theoretical Figure 1, plots average marginal utility against the average ex post value of health insurance in three sets of states: non-hospitalization, hospitalization and higher income, and hospitalization and lower income, where lower versus higher income is defined as whether residualized income is below or above the 25th percentile of hospitalization states, as before (see Appendix Table A12 for related statistics). The red markers are based on observed consumption and out-of-pocket health spending. The blue markers are based on a simple model of counterfactual consumption and out-of-pocket health spending were there no health care-related safety net.[28] Whereas in reality hospitalization is associated with a relatively small increase in out-of-pocket health spending and decrease in consumption even among households without health insurance, were it not for the safety net hospitalization would be associated with a large increase in out-of-pocket health spending for households without health insurance and likely a large decrease in consumption. Whereas in reality health insurance's coverage of hospitalization-related risk provides only modestly valuable insurance against health care costs while worsening income risk at significant cost, in this simple no-safety net counterfactual, health insurance's coverage of hospitalization-related risk would provide highly valuable insurance against health care costs while slightly improving income risk as well, resulting in a markup of 2.5 (i.e., the ex ante value is 3.5 times the transfer value) and an overall insurance value of $1,420

[28]Counterfactual no-safety net out-of-pocket health spending is $oop_{\omega}^{no} = oop_{\omega} + E(sn|hosp_{\omega})$, where oop_{ω} is observed out-of-pocket health spending and $E(sn|hosp_{\omega}, y_{\omega})$ is such that (i) average counterfactual no-safety net out-of-pocket health spending is the same in both sets of hospitalization states and (ii) average safety net support in hospitalization states as a whole is $6,000 per year, Dobkin et al.'s (2018) estimate of the effect of hospitalization on unpaid bills among the uninsured. This is likely a lower bound on the costs paid by external parties for the average hospital admission among the non-elderly uninsured. Counterfactual no-safety net consumption is $c_{\omega}^{no} = \max\{\underline{c},\ c - E(sn|hosp_{\omega})\}$. To be clear, the aim is to give a rough sense of the likely effect of the safety net in a transparent, intuitive way, albeit at the expense of the analysis being highly stylized.

per year—for hospitalization-related coverage alone. This simple analysis is suggestive that the safety net transforms health insurance's insurance of hospitalization risk from extremely valuable to somewhat costly.

The observed patterns of safety net support are qualitatively consistent with the safety net driving the key aspects of health insurance targeting uncovered in the empirical analysis. In terms of the relationship between out-of-pocket health spending and marginal utility, shown in Figure 2, the safety net could explain why on average marginal utility is much greater at lower than higher levels of out-of-pocket health spending. States in which marginal utility is highest tend to have the lowest out-of-pocket health spending because the safety net provides the most protection in those states. The underlying health and non-health shocks that drive up marginal utility also trigger greater protection against otherwise-uninsured health care costs from the safety net, which reduces out-of-pocket health spending and the ex post value of health insurance. States in which out-of-pocket health spending is high have below-average marginal utility because, given the safety net, high out-of-pocket health spending signals not only bad health but good enough non-health circumstances that the household is willing and able to pay high health care bills. If income and other resources are not sufficiently high, the safety net keeps out-of-pocket health spending relatively low even if health care costs are high. In this way, the safety net can transform health care spending from risk creator to risk absorber.

The implicit health insurance provided by the safety net has obvious disadvantages relative to formal health insurance, for both the beneficiary and external parties.[29] But the implicit health insurance provided by the safety net also has a less obvious but potentially important advantage over formal health insurance: Implicit health insurance provides valuable all-purpose insurance against other risks that is not otherwise being provided. Formal health insurance displaces without replacing such insurance. So having formal health insurance, other things equal, means being less well-insured against other risks.

[29]These disadvantages include, for example, the private and external costs of relying on charity care or of implicitly or explicitly defaulting on medical debt. The private costs may include stigma, hassle, and worse future credit access. Unfortunately, evidence of these private costs is sparse. In their review article Finkelstein, Mahoney and Notowidigdo (2018) conclude that empirical evidence of the benefits to individuals of reductions in their medical debt "remains limited." The vast majority of medical debt remains unpaid, and Dobkin et al. (2018) find that hospital admissions have little effect on credit scores despite increasing medical debt. From a revealed-preference perspective, people's behavior suggests that when times are tight, many perceive relying on the safety net as being less costly than paying their health care bills in full and cutting back along other dimensions, such as non-health consumption. Of course, as I discuss briefly in Section 5, safety net protection against health care costs has external costs and likely net efficiency costs, though these are irrelevant to the (private) insurance value of health insurance to the individual.

4.4 The safety net causes health insurance to worsen risk exposure

The evidence so far suggests that health insurance worsens risk exposure, because it worsens other, non-health care risks and because this worsening is more costly than its insurance of health care costs is valuable. The safety net's substantial protection against otherwise-uninsured health care costs, especially when health and non-health shocks are worst, appears qualitatively consistent with the key evidence. The first goal of this section is to assess whether this mechanism is quantitatively consistent with the key evidence. To that end, I develop and analyze a simple model of the insurance value of health insurance guided by the key empirical evidence. Finding that the model matches well the key empirical patterns related to health insurance's insurance value, I turn to using the model to investigate key issues that require the additional structure of the model, beyond the general assumptions underlying the sufficient statistic approach. These include quantifying the insurance value of the (inframarginal) health insurance coverage that households actually hold (as opposed to the value of increasing their health insurance coverage), quantifying the respective roles of different mechanisms in determining health insurance's insurance value, and assessing the robustness, heterogeneity, and generalizability of the key results.

Model.— A household draws health care consumption h (potentially a vector of consumption of different types of health care) and resources y from the joint distribution, $(h, y) \sim F(h, y)$. (Non-health) consumption is given by the constraint

$$c(h, y; HI) = \max\{\underline{c},\ y - [tot(h) - hi(h; HI) - ihi(h, y; HI)]\}, \tag{12}$$

where $\underline{c}$ is the consumption floor, $tot(h)$ is the total cost of the household's health care consumption, $hi(h; HI)$ is the health insurance benefit, if any, and $ihi(h, y; HI)$ is the "implicit health insurance" support, if any, from charity care, bad debt, bankruptcy, or means-tested support, over and above any support from the consumption floor. Utility is a state-independent constant relative risk aversion function of consumption,

$$v(h, n; HI) = \frac{c(h, n; HI)^{1-\gamma}}{1-\gamma}, \tag{13}$$

where γ is the coefficient of relative risk aversion.

Baseline calibration.— The key ingredients of the model are the joint distribution of health care consumption and resources, $F(h, y)$, and the model of implicit health insurance, $ihi(h, y; HI)$. For $F(h, y)$, I use the empirical joint distribution of total health care costs and income among households in which the head is between 25 and 64 years old and has a high school degree but no college degree. Income is the maximum of total annual income

and an income floor of $15,000 (somewhat above the tenth percentile). The purpose of the income floor is to be conservative about the income risk that households face. Total health care costs for households with health insurance are total annual payments. For households without health insurance, total health care costs are annual charges scaled by both (i) the payments-charge ratio among non-elderly households with health insurance (which is roughly one-half), to reflect discounts relative to charges, and (ii) 1.25, to reflect moral hazard based on evidence from the Oregon Health Insurance Experiment (Finkelstein et al., 2012).[30] This follows the standard practice of approximating risk (heterogeneity across states of the world) with heterogeneity across households in the cross-section and ensures that the model matches the joint distribution of what are the two most important elements of the budget constraint in this context (health spending and income), including income losses from bad health and income support from unemployment insurance and similar institutions.

I approximate implicit health insurance as providing full catastrophic coverage above an income-dependent deductible,

$$ihi(h, y; HI) = \max\{0,\ h - hi(h; HI) - d_{ihi}(y)\}, \tag{14}$$

where the deductible function $d_{ihi}(y)$ is based on the predicted values from a regression of out-of-pocket health spending on a cubic in income among households *without* health insurance, with annual health care charges of at least $20,000, and in which the head is between 25 and 64 years old and has a high school degree but no college degree. The idea is to estimate the typical amount of health care costs that is *not* covered by the safety net. That there is implicit health insurance on top of the consumption floor captures in a simple way the observed unevenness of the safety net.

I consider a full, comprehensive health insurance contract that covers 100% of health care costs. As for the other parameters, the consumption floor is $\underline{c} = \$5,000$/year, and the coefficient of relative risk aversion is three, $\gamma = 3$.

Remarks.— The aim is to maximize simplicity and transparency by making minimal adjustments to standard approaches. As in the simplest standard approach, everything is driven by the budget constraint and consumption equals income less out-of-pocket health spending. The key difference relative to the standard approach is that here, out-of-pocket

[30]These definitions of health care consumption follow Mahoney (2015). The moral hazard factor of 1.25 is the change in utilization found in the Oregon Health Insurance Experiment (Finkelstein et al., 2012). I scale up the health care consumption of households without health insurance to the level that would be expected if they had full health insurance to err on the side of overstating health risk and so the insurance value of health insurance. The analysis otherwise ignores moral hazard in order to focus on insurance value. The results are similar under an alternative, simple definition in which total health care costs are one-half of charges for all households.

health spending is not exogenous with respect to other "right-hand-side" (non-consumption) elements of the budget constraint. Out-of-pocket health spending is potentially correlated with income both "directly" through the joint distribution of health care consumption and income and "indirectly" through the safety net.[31] Unlike standard approaches, this model admits both possibilities: health insurance could be pro- or anti-insurance. Whereas standard approaches focus exclusively on how out-of-pocket health spending affects the tightness of the constraint, this model also accounts for causality in the opposite direction: The tightness of the constraint potentially affects out-of-pocket health spending, because of the safety net. As income risk approaches zero, the model approaches the standard model in which health insurance is necessarily pro-insurance. But with non-zero income risk, health insurance's insurance value potentially is the net of competing pro- and anti-insurance effects: Health insurance improves health risk but potentially worsens income risk.[32]

Results.— *Overall insurance value.—* The first column of Table 3 shows that with the baseline calibration, health insurance has negative insurance value, equivalent in expected utility terms to reducing consumption in all states of the world by $480 per year. This is about 7% of gross health insurance benefits and 20% of net health insurance benefits (i.e., of the mean ex post value of health insurance). The minimal changes to the standard approach of accounting for non-health risk and implicit health insurance can explain why health insurance would worsen risk exposure on net, to an extent broadly similar to that implied by the corresponding sufficient statistics estimates. The model matches well the key empirical patterns, and the conclusion that health insurance worsens risk exposure on net is highly robust. As shown in Appendix Table A14, health insurance worsens risk exposure on net even if implicit health insurance is much less generous than it appears to be empirically (e.g., by doubling its effective deductibles) and even if income risk is much smaller than it appears to be empirically (e.g., imposing a $50,000 income floor or cutting income risk by up to a factor of nine).[33]

Pro-insurance benefit from insuring health care cost risk.— The second column of Table 3 reports the results for a counterfactual world without income risk: In all states, income is the median of its empirical distribution, $y \equiv y_{med}$, and health risk remains as in the data, $h \sim$

[31] In states in which income is at least the consumption floor ($y \geqslant \underline{c}$), out-of-pocket health spending is $oop(h, y; HI) = \min\{tot(h) - hi(h; HI) - ihi(h, y; HI),\ y - \underline{c}\}$.

[32] Health insurance does not necessarily worsen income risk in this setup, since I allow for any relationship between health and income shocks in the (joint) risk process, $F(h, y)$. Whether health insurance improves or worsens income risk depends on whether any pro-insurance effect from a positive correlation between health and income shocks is larger or smaller than the anti-insurance effect from health insurance displacing support from the safety net.

[33] One reason the results are so robust is that many of the changes to the model that would decrease health insurance's anti-insurance cost of worsening income risk would also decrease its pro-insurance benefit of insuring health care cost risk (e.g., reducing risk aversion, allowing for consumption smoothing, increasing the amount of the observed heterogeneity that is "soaked out" toward isolating risk, etc.).

$F(h)$. In this case, out-of-pocket health spending is purely a risk-creator, not a risk-absorber. Yet even in this case, health insurance's insurance value is small: equivalent to increasing consumption in all states by \$70. That health insurance's insurance of health risk is a matter of near indifference is in keeping with the results in Table A7 that health insurance's insurance of hospitalization-related risk provides little value. The third column of Table 3 also has no income risk but turns off implicit health insurance as well ($ihi(h, y; HI) \equiv 0$). In this case, by contrast, health insurance is extremely valuable, providing insurance value surplus equivalent to increasing consumption in all states by \$7,220. These results suggest that implicit health insurance greatly reduces the value of health insurance's insurance of health risk, cutting it in this simple model by two orders of magnitude. Implicit health insurance provides significant protection against health risk, which markedly decreases both the marginal utility and value gaps in the bad-health states that would otherwise make health insurance highly valuable. These results show that the value of health insurance's insurance of health risk is small not because insurance of health care costs is of little value—it is of great value—but because implicit health insurance provides so much protection that the residual risk that remains for health insurance is of little consequence. Given the significant, generally well-targeted support from the health care-related safety net, health insurance's insurance value almost could not possibly be large.[34]

Anti-insurance cost from worsening income risk.— Comparing the results with versus without income risk (first versus second columns of Table 3) reveals the anti-insurance cost of health insurance's worsening of income risk. Whereas without income risk, health insurance has a small positive insurance value (\$70), with income risk health insurance has a significantly negative insurance value (−\$480), implying a gross anti-insurance cost of \$550 per year. Additional results reveal the key role of the safety net in causing health insurance to worsen non-health risk. Whereas gross health insurance benefits are virtually uncorrelated with income ($corr(hi, y) = corr(h, y) = 0.04$), net health insurance benefits, accounting for implicit taxation by the safety net, are strongly positively correlated with income ($corr(V, y) = corr(oop, y) = 0.35$).[35] As a result, health insurance worsens income risk: In addition to transferring resources from states of the world with better to worse health (good targeting), health insurance also transfers resources from states with lower to higher income (bad targeting). This arises because on average the safety net helps more when income is lower, which means that on average health insurance helps less when income is lower. In

[34]That the residual health care cost risk remaining after the implicit catastrophic coverage from the safety net has little welfare cost is the basis of the key result in optimal insurance theory that such risk should optimally remain uninsured, since the minimal insurance gain from insuring such risk would be dominated by the associated moral hazard and administrative costs.

[35]That out-of-pocket health spending and income are more strongly correlated in this simple model than in the data (0.35 vs. 0.16, respectively) is likely due at least in part to the fact that this simple model does not include any other, non-income determinants of the extent of safety net support of a given level of health care costs, whereas in reality seizable assets (Mahoney, 2015) and other factors matter as well.

other words, the safety net transforms health insurance from being roughly neutral with respect to income risk to worsening it significantly. Health insurance's worsening of income risk is quite costly because of the substantial income risk that households face.

Summary.— The safety net fundamentally reshapes—and, from an insurance perspective, destroys—health insurance's targeting properties. The safety net (i) steals health insurance's thunder, transforming health insurance's insurance against health care cost risk from being hugely valuable to being a matter of near indifference, and (ii) adds insult to injury, transforming health insurance from being roughly neutral with respect to other risks to worsening them systematically, at significant cost. Health insurance displaces support from the safety net and thereby reduces the extent to which the safety net insures all manner of risk, health and non-health. While health insurance more than makes up for the insurance of health care cost risk that it displaces and so provides a (small) positive amount of insurance against health care cost risk on net, it does not make up for the safety net's insurance of other risks that it displaces and so it decreases insurance against other risks on net. Quantitatively, the safety net's effect on the value of health insurance's insurance against health care cost risk is much larger than its effect on the cost of health insurance's worsening of other risks. That safety net support tends to be much greater for health care costs than other risks means that health insurance's displacement of such support might "undo" a meaningful share of the safety net's insurance of other risks. Non-health care risks like income are far greater than health care cost risk, in part because of the unevenness of the safety net, which makes health insurance's worsening of other risks more costly than its improving of health care cost risk is beneficial. The uneven, health care-focused nature of the safety net is the key factor making health insurance worsen risk exposure on net.

5 Implications

The economic analysis of health insurance has focused on the tradeoff between insurance benefit and moral hazard cost. Although there are few quantitative estimates of the insurance benefit, the prevailing view has been that typical health insurance contracts have an insurance benefit that is inefficiently *high* (e.g., Feldstein, 1973; Pauly, 1986). This view is based on the observation that whereas the optimal contract for balancing insurance benefit against moral hazard cost is catastrophic coverage above a high deductible (Arrow, 1963), typical health insurance contracts cover all kinds of expenditures, not just catastrophic ones. Hence, the thinking goes, people with typical health insurance contracts are over-insured: They enjoy more risk protection than would be efficient to buy given moral hazard and administrative costs.

My findings suggest that rather than providing too much of a good thing, typical health insurance contracts provide a negative amount of that particular good thing: They worsen risk exposure. Not only do typical contracts involve the substantial moral hazard and administrative costs that have been studied extensively, they also involve a previously unappreciated anti-insurance cost. For most U.S. households, the benefit of health insurance's insurance of health care cost risk is dominated by the cost of its worsening of other risks. This raises several questions.

Do existing health insurance contracts decrease welfare? Are people making mistakes by holding them? *If* the key costs and benefits of health insurance were those at the center of the economic approach to insurance—risk protection and moral hazard—health insurance would be decreasing welfare on a large scale, worsening risk exposure at substantial moral hazard cost. A rough, back-of-the-envelope calculation combining my estimates of insurance value with Finkelstein, Hendren and Luttmer's (2019) estimates of moral hazard cost from the Oregon Health Insurance Experiment suggests that the combined anti-insurance and moral hazard costs of full health insurance coverage for a typical household is roughly on the order of $1,800 per year (about $300 from worsening risk and about $1,500 from moral hazard).[36] Many feasible alternative contracts would be less bad on both dimensions, including the "null contract" of no health insurance. But risk protection is not the only possible benefit of health insurance. There are a variety of potential other, non-insurance benefits as well.

Coverage of certain health goods and services might have paternalistic benefits if it corrects people's mistakes (e.g., Baicker, Mullainathan and Schwartzstein, 2015) or merit good benefits if people value other people's health consumption per se, above and beyond its effect on the recipient's welfare. It might help internalize real externalities (e.g., from infectious disease) or fiscal externalities (e.g., Miller and Wherry, 2019; Brown, Kowalski and Lurie, 2020; Hendren and Sprung-Keyser, 2020). It avoids efficiency costs associated with the safety net. It might provide peace of mind. A variety of evidence suggests that some of these potential benefits of health insurance could be large. But as Currie and Gahvari (2008) emphasize in their review of the literature on in-kind transfers, of which the implicit subsidies on health care from health insurance can be viewed as a special case, relatively little is known about the magnitude of such benefits.[37]

[36] Finkelstein, Hendren and Luttmer (2019) estimate an average per-person moral hazard cost of around $750 per year. I multiply this by two to get a rough estimate of the average per-household cost.

[37] Combining the result that health insurance worsens risk exposure with the result that moral hazard responses can be valuable, not costly (whether because of mistakes, externalities, or merit good considerations) would, *if* moral hazard's benefits outweigh its costs, turn the traditional welfare analysis of health insurance upside down. Rather than providing an insurance benefit at the expense of a moral hazard cost, health insurance might instead provide a moral hazard benefit at the expense of an insurance cost.

A partial exception to this negative conclusion about the state of knowledge is the externality associated with uncompensated care. A growing literature has documented the large role of such care (e.g., Mahoney, 2015; Garthwaite, Gross and Notowidigdo, 2018). Uncompensated care is so important, in fact, that for many individuals the vast majority of the incidence of their health insurance decision falls on other parties: Its effect on uncompensated care far exceeds the transfer value to the individual (Finkelstein, Hendren and Luttmer, 2019). This creates a Samaritan's dilemma rationale for mandating or encouraging health insurance coverage, which is further enhanced by any associated efficiency costs beyond just the pure transfers from other parties. Given such considerations, might health insurance increase welfare despite its anti-insurance cost? My results reveal an important tradeoff. On the one hand, efficiency costs associated with uncompensated care are a force toward the efficient contract displacing more uncompensated care. On the other hand, displacing uncompensated care undoes without replacing some of the valuable insurance that such care otherwise would have provided. This is a force toward the efficient contract being less comprehensive and displacing less uncompensated care.

In summary, that typical contracts have substantial anti-insurance costs just means that one key element of their overall welfare effect is for many households not the benefit it has been widely thought to be but a cost. The overall welfare effect of such contracts, and of the policies that increase their prevalence, depends on the size of their other, non-insurance benefits, about which much remains unknown despite notable recent progress. Quantifying such benefits is an important priority for future research.[38]

Which health insurance contracts would provide the best risk protection? Although risk protection is just one determinant of the overall value of health insurance, set non-insurance factors aside for the moment to consider: What health insurance contracts would provide the best risk protection? Within the class of contracts that subsidizes health care, which comprises the vast majority of contracts currently in use in the U.S., it is not easy to identify contracts that would provide valuable insurance on net. Such contracts' targeting of marginal utility is undermined by the safety net. This is health insurance in a highly second-best world in which the residual out-of-pocket health spending that remains after the safety net is more risk absorber than risk creator. That makes it hard for health care subsidy contracts, whose main effect is to decrease out-of-pocket health spending, to

[38] A related priority is to incorporate the safety net's reshaping of health care cost risk and health insurance's targeting properties into the large bodies of research on people's choices in the face of health risk, including consumption, labor supply, and, especially, health insurance choices, as such analyses typically maintain assumptions that rule out health insurance worsening risk exposure.

improve risk exposure.[39,40]

Yet there would appear to be scope for well-designed contracts to provide valuable insurance against health risk. Even people with health insurance are exposed to consumption risk from health shocks, mainly from the associated earnings losses (e.g., Dobkin et al., 2018; Meyer and Mok, 2019). Could a well-designed contract insure health risk while at the same time avoid worsening non-health risk? One candidate is contracts that pay cash indemnities based on health, e.g., $10,000 if the insured has a heart attack.[41] The prevailing view is that cash indemnity contracts are better than subsidy contracts in terms of moral hazard but worse in terms of risk protection, since they provide no within-diagnosis insurance (e.g., some heart attacks involve much higher costs than others; see, e.g., Zeckhauser (1970)). Although the lack of within-diagnosis insurance is an important limitation, my results suggest that cash indemnity contracts would be better, not worse, in terms of risk protection. Most important, cash indemnity benefits would likely displace safety net support to a lesser extent than health care subsidies do and thereby worsen other risks to a lesser extent.[42] Moreover, unlike health care subsidy contracts, cash indemnity contracts could partially insure the costs of health shocks beyond health care costs, such as associated earnings losses. Whereas health care subsidies are limited to reducing health care costs, assuming that covering more than 100% of health care costs would often create significant ex post moral hazard, cash indemnity

[39]Unless, that is, one is willing to consider negative subsidies, i.e., taxes on health care. That health care subsidies help less on average when marginal utility is high means that health care taxes would hurt less on average when marginal utility is high, thereby providing valuable insurance. My results imply that from non-negative subsidy rates, the cost of cutting the subsidy rate in terms of worsening health care cost risk is significantly smaller than the benefit in terms of insuring other risks. To be clear, I discuss health care taxes not as a policy recommendation but to help build intuition about the main results and what they imply.

[40]Even catastrophic health insurance contracts tend to worsen risk exposure. This is because high out-of-pocket health spending signals not only high health care costs but also a good realization of non-health care risks, since it reveals a high willingness and ability to pay health care bills. In other words, although higher out-of-pocket health spending means a tighter constraint *other things equal*, other things are systematically not equal. In the current equilibrium, the differences in these other things dominate the "main" effect on average, so even catastrophic health insurance tends to worsen risk exposure. But to be clear, catastrophic contracts have smaller anti-insurance costs than typical contracts do. So in the context of health insurance, the advantage of catastrophic contracts is not just their lower moral hazard costs but also their higher insurance benefits (lower anti-insurance cost).

[41]Though cash indemnity health insurance is rare today, it was common in the past (see, e.g., Cutler, 2002), and cash indemnity policies are common in many other contexts, such as life insurance, annuities, and, increasingly, long-term care insurance. For long-term care insurance, see Doty, Mahoney and Sciegaj (2010) for an overview of such policies and Lieber and Lockwood (2019) for an analysis of their targeting effects.

[42]Although a full analysis of the extent to which cash indemnity health insurance would displace safety net support is beyond the scope of this paper, the uneven nature of the safety net—its much greater protection against otherwise-uninsured health care costs than against losses of income and assets—makes it likely that cash indemnity health insurance would displace less support than health care subsidies do. Whereas someone with full traditional health insurance has no otherwise-uninsured health care costs on which they can receive safety net support, someone with cash indemnity health insurance potentially *does* have large health care bills on which they can potentially receive safety net support (e.g., through charity care or by not paying the bills), despite their cash indemnity.

benefits need not be limited to covering (typical levels of) health care costs but could cover (typical levels of) other costs of bad health as well. I estimate, for example, that a fixed cash indemnity benefit per day spent in the hospital that did not displace safety net support would generate 65 cents worth of insurance value surplus per dollar of expected value (see Appendix Table A15).[43]

The safety net significantly decreases the value of health insurance. Does that make it bad? That safety net protection against otherwise-uninsured health care costs decreases the value of health insurance so much is in part a sign of the safety net's *success*. It reflects that the health care-related safety net provides valuable all-purpose insurance against a wide range of risks, health and non-health alike. But it is also a sign—and unappreciated cost—of the growing *unevenness* of the safety net. The safety net provides much more protection against otherwise-uninsured health care costs than against losses of income and assets and other risks. Such unevenness has likely become more pronounced in recent decades with the growth of health care costs and the shift from traditional cash assistance to in-work benefits (Hoynes and Schanzenbach, 2018). More uniform protection against different risks would not only decrease the value of health insurance less but also likely provide more valuable insurance. Whether such insurance would be cost-effective is an important question for future research. In any event, given the current configuration of the safety net, health insurance is complementary with reducing other risk exposures, whether through government policy changes such as expanding unemployment insurance or household choices such as the riskiness of financial portfolios.

6 Conclusion

For many U.S. households, health insurance worsens risk exposure—it has negative insurance value. This worsening arises because health insurance worsens non-health care cost risk, and because this worsening is more costly than health insurance's improving of health care risk is beneficial. The key driver is the safety net. The safety net provides significant protection against otherwise-uninsured health care costs, especially when health and non-health shocks are worst. Health insurance displaces such support, thereby undoing the valuable insurance it otherwise would have provided while only replacing the insurance against health care

[43]Appendix Table A15 shows the value of other hypothetical contracts as well. It shows that each of several types of coverage of health spending worsens risk exposure, whereas each of three types of fixed indemnity benefits improve it. The types of coverage of health spending are spending on: health care as a whole, hospital care, doctor visits, and prescriptions. The types of fixed indemnity benefits are for: whether the household head or spouse experienced a hospitalization, the number of days the head or spouse spent in the hospital, and whether the household head reported that their health was bad (presumably not a viable contract due to verifiability or moral hazard issues, reported for illustrative purposes only).

cost risk, not other risks. This is an important cost of the roughly $1.6 trillion per year of government spending on health insurance provision and health insurance subsidies.

My findings point to at least two priorities for future research. One is to better understand heterogeneity in health insurance's insurance value across different types of households and different institutional contexts. Given the importance of other risks and the safety net to health insurance's value, there is likely important heterogeneity both across households within a society and across societies with different institutions. But the more important priority for future research is to quantify other major components of the overall welfare effect of health insurance, especially potential non-insurance benefits. That for many households the insurance effect of typical contracts is not the benefit it has been thought to be but a cost raises the urgency of quantifying other benefits such contracts might have and determining whether they could be realized in a less costly way.

References

Arrow, Kenneth J. 1963. "Uncertainty and the Welfare Economics of Medical Care." *The American Economic Review*, 53(5): 941–973.

Atkinson, Anthony Barnes, and Joseph E Stiglitz. 1976. "The design of tax structure: direct versus indirect taxation." *Journal of Public Economics*, 6(1–2): 55–75.

Baicker, Katherine, Sendhil Mullainathan, and Joshua Schwartzstein. 2015. "Behavioral Hazard in Health Insurance." *The Quarterly Journal of Economics*, 130(4): 1623–1667.

Baily, Martin Neil. 1978. "Some aspects of optimal unemployment insurance." *Journal of Public Economics*, 10(3): 379–402.

Blomqvist, Ake. 1997. "Optimal non-linear health insurance." *Journal of Health Economics*, 16(3): 303–321.

Brown, David W, Amanda E Kowalski, and Ithai Z Lurie. 2020. "Long-Term Impacts of Childhood Medicaid Expansions on Outcomes in Adulthood." *Review of Economic Studies*, 87(2): 792–821.

CFPB. 2014. "Consumer credit reports: a study of medical and non-medical collections." Rep., Consumer Financial Protection Bureau, Washington, DC.

Chetty, Raj. 2006. "A general formula for the optimal level of social insurance." *Journal of Public Economics*, 90(10–11): 1879–1901.

Coughlin, TA, J Holahan, K Caswell, and M McGrath. 2014. "Uncompensated Care for the Uninsured in 2013: A Detailed Examination." Rep., Kaiser Fam. Found., Menlo Park, CA.

Currie, Janet, and Firouz Gahvari. 2008. "Transfers in Cash and In-Kind: Theory Meets the Data." *Journal of Economic Literature*, 46(2): 333–383.

Cutler, David M. 2002. "Health care and the public sector." In *Handbook of Public Economics*. Vol. 4 of *Handbook of Public Economics*, , ed. A. J. Auerbach and M. Feldstein, Chapter 31, 2143–2243. Elsevier.

Dague, Laura. 2014. "The effect of Medicaid premiums on enrollment: A regression discontinuity approach." *Journal of Health Economics*, 37(C): 1–12.

Deshpande, Manasi, and Lee M. Lockwood. 2020. "Beyond Health: Non-Health Risk and the Value of Disability Insurance." University of Chicago Working Paper.

Dobkin, Carlos, Amy Finkelstein, Raymond Kluender, and Matthew J. Notowidigdo. 2018. "The Economic Consequences of Hospital Admissions." *American Economic Review*, 108(2): 308–52.

Doty, Pamela, Kevin J Mahoney, and Mark Sciegaj. 2010. "New state strategies to meet long-term care needs." *Health affairs*, 29(1): 49–56.

Dranove, David, Craig Garthwaite, and Christopher Ody. 2015. "A floor-and-trade proposal to improve the delivery of charity-care services by US nonprofit hospitals." Discussion Paper 2015-07, Hamilton Project, Washington, DC.

Engelhardt, Gary V, and Jonathan Gruber. 2011. "Medicare Part D and the financial protection of the elderly." *American Economic Journal: Economic Policy*, 3(4): 77–102.

Feldman, Roger, and Bryan Dowd. 1991. "A New Estimate of the Welfare Loss of Excess Health Insurance." *American Economic Review*, 81(1): 297–301.

Feldstein, Martin, and Jonathan Gruber. 1995. "A Major Risk Approach to Health Insurance Reform." In *Tax Policy and the Economy, Volume 9. NBER Chapters*, 103–130. National Bureau of Economic Research, Inc.

Feldstein, Martin S. 1973. "The welfare loss of excess health insurance." *Journal of Political Economy*, 81(2, Part 1): 251–280.

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

Finkelstein, Amy, Nathaniel Hendren, and Erzo F. P. Luttmer. 2019. "The Value of Medicaid: Interpreting Results from the Oregon Health Insurance Experiment." *Journal of Political Economy*, 127(6): 2836–2874.

Finkelstein, Amy, Nathaniel Hendren, and Mark Shepard. 2019. "Subsidizing Health Insurance for Low-Income Adults: Evidence from Massachusetts." *American Economic Review*, 109(4): 1530–1567.

Finkelstein, Amy, Neale Mahoney, and Matthew J. Notowidigdo. 2018. "What Does (Formal) Health Insurance Do, and for Whom?" *Annual Review of Economics*, 10(1): 261–286.

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

Gallen, Trevor S. 2015. "Using Participant Behavior to Measure the Value of Social Programs: The Case of Medicaid." Working paper, Dept. Econ., Purdue Univ.

Garthwaite, Craig, Tal Gross, and Matthew J. Notowidigdo. 2018. "Hospitals as Insurers of Last Resort." *American Economic Journal: Applied Economics*, 10(1): 1–39.

Gross, Tal, Timothy Layton, and Daniel Prinz. 2020. "The Liquidity Sensitivity of Healthcare Consumption: Evidence from Social Security Payments." National Bureau of Economic Research, Inc NBER Working Papers 27977.

Hackmann, Martin B., Jonathan T. Kolstad, and Amanda E. Kowalski. 2015. "Adverse Selection and an Individual Mandate: When Theory Meets Practice." *American Economic Review*, 105(3): 1030–1066.

Hendren, Nathaniel, and Ben Sprung-Keyser. 2020. "A Unified Welfare Analysis of Government Policies." *The Quarterly Journal of Economics*, 135(3): 1209–1318.

Hoynes, Hilary W., and Diane Whitmore Schanzenbach. 2018. "Safety Nets Investments in Children." *Brookings Papers on Economic Activity*, 49(1 (Spring): 89–150.

Kowalski, Amanda E. 2015. "Estimating the tradeoff between risk protection and moral hazard with a nonlinear budget set model of health insurance." *International Journal of Industrial Organization*, 43(C): 122–135.

Lieber, Ethan M. J., and Lee M. Lockwood. 2019. "Targeting with In-Kind Transfers: Evidence from Medicaid Home Care." *American Economic Review*, 109(4): 1461–1485.

Mahoney, Neale. 2015. "Bankruptcy as implicit health insurance." *American Economic Review*, 105(2): 710–46.

Manning, Willard G., and M. Susan Marquis. 1996. "Health insurance: The tradeoff between risk pooling and moral hazard." *Journal of Health Economics*, 15(5): 609–639.

Meyer, Bruce D., and Wallace K.C. Mok. 2019. "Disability, earnings, income and consumption." *Journal of Public Economics*, 171(C): 51–69.

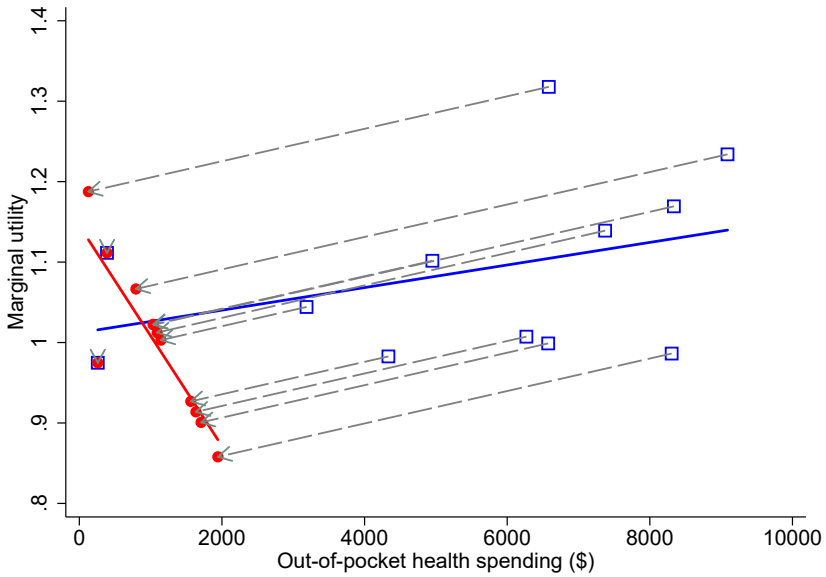
Miller, Sarah, and Laura R. Wherry. 2019. "The Long-Term Effects of Early Life Medicaid Coverage." *Journal of Human Resources*, 54(3): 785–824.

Pauly, Mark V. 1986. "Taxation, health insurance, and market failure in the medical economy." *Journal of Economic Literature*, 24(2): 629–675.

Zeckhauser, Richard. 1970. "Medical insurance: A case study of the tradeoff between risk spreading and appropriate incentives." *Journal of Economic Theory*, 2(1): 10–26.

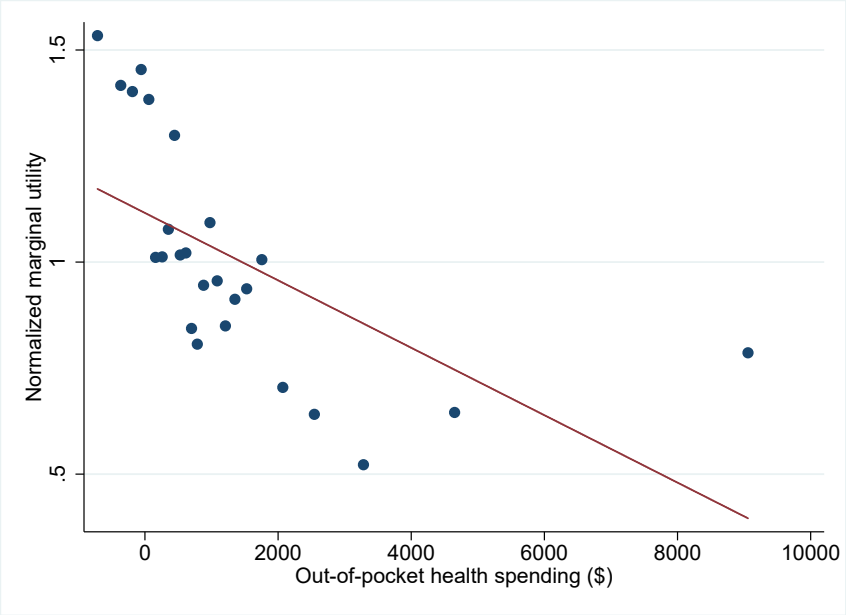
Figures and Tables

Figure 1: How the safety net can transform health insurance targeting



Notes: Figure shows scatter plot of simulated data without a safety net (blue squares and blue best-fit line) and with a safety net (red circles and red best-fit line). In the simulation, consumption equals resources less out-of-pocket health spending ($c = y - oop$), marginal utility is linear in consumption, and health and non-health risks are independent, with $h \sim unif[\$0,\ \$10,000]$ and $y \sim unif[\$40,000,\ \$60,000]$. Without a safety net, $oop = h$, so bad health shocks decrease consumption and so increase marginal utility. The arrows show the effects of a stylized safety net that provides catastrophic coverage above a resources-dependent deductible, $oop = \min\{h,\ \delta_0 + \delta_1 y\}$, in the realistic case in which the safety net provides greater protection against otherwise-uninsured health care costs when gross resources are lower, $\delta_1 > 0$. The safety net reduces both out-of-pocket health spending (westward movement in figure) and marginal utility (southward), more in states in which income is lower. The differential reduction of out-of-pocket health spending in states in which income is low means that high out-of-pocket health spending signals high income, and it transforms the relationship between out-of-pocket health spending and marginal utility from one in which they covary positively to one in which they covary negatively. Because of the safety net, out-of-pocket health spending is determined not only (and, in this simple simulation, not mainly) by health care consumption but also (mainly) by ability to pay (income)—and is thus positively related to the realization of income and negatively related to marginal utility. In all cases, the sign of the regression slope is the same as the sign of the first order approximation to health insurance's insurance value, $sign(\beta_{\hat{\lambda}|V}) = sign\left[Cov\left(\hat{\lambda}, V\right)\right]$, since $\beta_{\hat{\lambda}|V} = Cov\left(\hat{\lambda}, V\right)/Var(V)$. The slope of the "effects of the safety net" arrows does not exactly match the slope of the regression line only because of sampling error due to the small number of simulation draws.

Figure 2: Marginal utility and out-of-pocket health spending



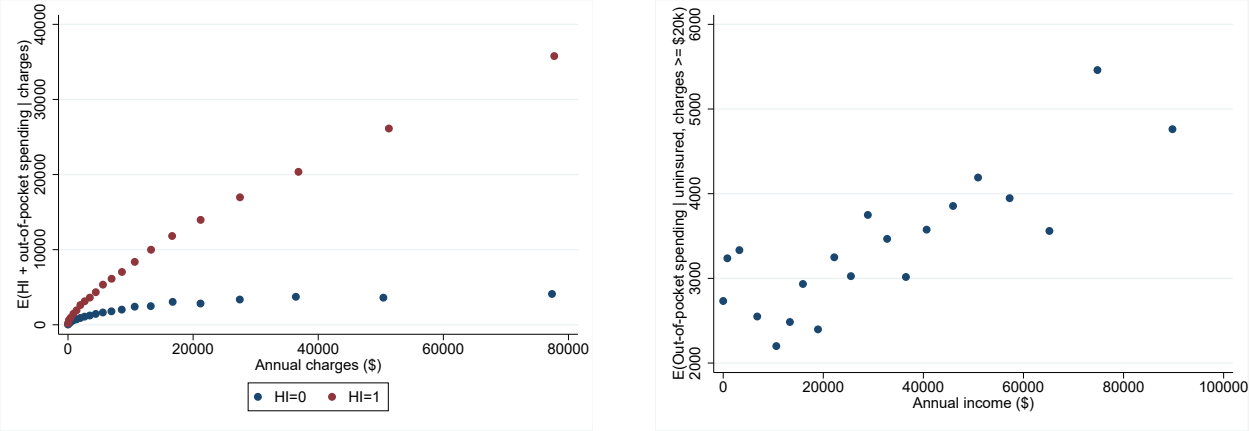
Notes: Figure shows the conditional mean of normalized marginal utility as a function of out-of-pocket spending and the corresponding regression line for non-elderly households. This is an empirical analogue of theoretical Figure 1. This is a bin scatter plot that groups the data into equal-sized bins based on out-of-pocket health spending and plots one point for each bin, whose coordinates are average out-of-pocket health spending and average normalized marginal utility among households in that bin. Marginal utility is normalized to have a mean of one and is residualized using the baseline controls (year dummies, a cubic in age, a quadratic in household size, and education category dummies).

Table 1: Sufficient statistic estimates

	Non-elderly (1)	Non-elderly uninsured (2)	Non-elderly insured (3)	Elderly insured (4)
$\widehat{\beta} \times \left(sd(oop)/sd\left(\widehat{\lambda}\right)\right) = Corr\left(\widehat{\lambda}, oop\right)$	-0.04	-0.06	-0.04	-0.02
$se\left(\widehat{\beta}\right) \times \left(sd(oop)/sd\left(\widehat{\lambda}\right)\right)$	(0.006)	(0.011)	(0.006)	(0.014)
Insurance value ($), $Cov\left(\widehat{\lambda}, oop\right)$	-326	-307	-318	-131
se(Insurance value) ($) $= se\left(\widehat{\beta}\right) \times var(oop)$	(43)	(55)	(53)	(77)
Transfer value ($), $E\left(oop\right)$	1,298	917	1,359	1,906
Markup, $M = Cov\left(\widehat{\lambda}, oop\right)/E\left(oop\right)$	-0.25	-0.33	-0.23	-0.07

Notes: Table presents estimates related to the insurance value of eliminating out-of-pocket health spending. $\widehat{\beta}$ is the estimate of the coefficient on out-of-pocket health spending in a regression of normalized marginal utility on out-of-pocket health spending and controls (equation 8), scaled to be in standard deviation units (so a coefficient of -0.04 means that a one-standard deviation increase in out-of-pocket health spending is associated with a 0.04 decrease in marginal utility. The standard error of this coefficient, shown on the next line, is expressed in the same units. The controls are year dummies, a cubic in age, a quadratic in household size, and education category dummies. The markup is the insurance value per dollar of transfer value, $M = Cov\left(\widehat{\lambda}, oop\right)/E\left(oop\right)$. Monetary amounts are in dollars per household per year.

Figure 3: Implicit health insurance from the safety net



(a) Safety net support increasing in health costs

(b) Safety net support decreasing in income

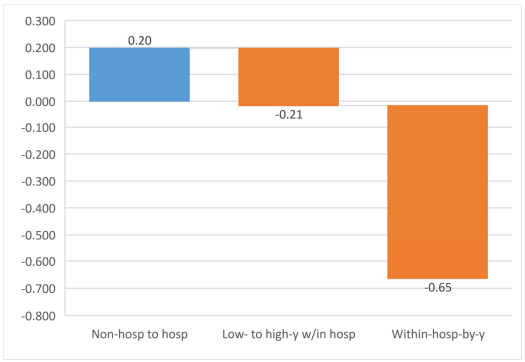
Notes: Left panel: Conditional mean of the sum of total payments by health insurers (health insurance benefits) and households (out-of-pocket health spending) as a function of charges (a rough measure of health care utilization) for households with versus without health insurance. The higher, red dots are for households with health insurance. The lower, blue dots are for households without health insurance.
Right panel: Conditional mean of out-of-pocket health spending as a function of income among households without health insurance with annual health care charges of at least $20,000. For reference, in this sample mean health care charges is $72,400 and the conditional mean of health care charges is a decreasing function of income, so the greater out-of-pocket health spending among higher-income households does not reflect higher charges.
Both panels are bin scatter plots, which group the data into equal-sized bins based on the x-axis variable and plots one point for each bin, whose coordinates are the average of the x- and y-axis variables among households in that bin.

Table 2: Health insurance and income risk

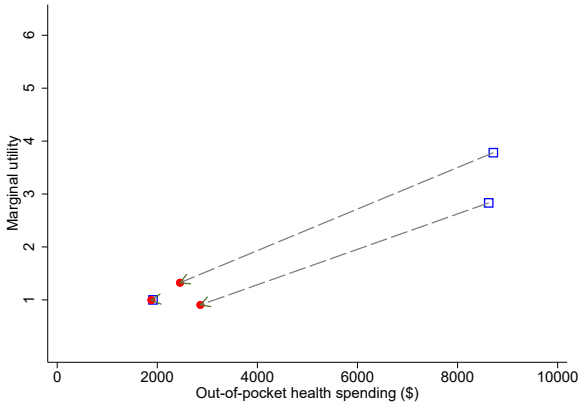
	Non-elderly (1)	Non-elderly uninsured (2)	Non-elderly insured (3)	Elderly insured (4)
$\widehat{\beta}_{y\|oop}$	4.91	3.38	4.86	2.77
(se)	(0.32)	(0.51)	(0.34)	(0.43)
$Corr(y, oop)$	0.16	0.17	0.16	0.14
$Sd(oop)$	2,038	1,952	2,047	2,601
$Sd(y)$	62,316	38,416	64,119	52,605

Notes: The first row shows the coefficient estimate on out-of-pocket health spending from a regression of income on out-of-pocket health spending and the baseline controls (year dummies, a cubic in age, a quadratic in household size, and education category dummies). In a simple model in which marginal utility is linear in consumption and consumption equals income minus out-of-pocket health spending, this regression coefficient equals the ratio of the "portfolio effect" cost from worsening income risk to the "direct effect" benefit from improving health risk (see equations (7) and (9)). The second row is the corresponding standard error. The third row is the correlation between income and out-of-pocket health spending, residualized with the baseline controls. The fourth and fifth rows are the standard deviations of out-of-pocket health spending and income, respectively, each residualized with the baseline controls.

Figure 4: Health insurance and hospitalization risk



(a) HI worsens hospitalization risk



(b) Safety net transforms HI targeting

Notes: Left panel: Waterfall figure of the decomposition of the markup on (insurance value per dollar of transfer value of) health insurance transfers into hospitalization states into the sum of three terms. The first is health insurance's average net transfer from non-hospitalization to hospitalization states. The second is health insurance's average net transfer from hospitalization states with lower- to higher- income realizations, defined as below versus above the 25th percentile of residualized income, residualized with the baseline controls. The third is health insurance's transfers within hospitalization-by-income categories, from states in which health insurance is less to more valuable. See Appendix Table A12 for related statistics.
Right panel: An empirical analogue of the theoretical Figure 1, plots average marginal utility against the average ex post value of health insurance in three sets of states: non-hospitalization, hospitalization and higher income, and hospitalization and lower income, where lower versus higher income is defined as whether residualized income is below or above the 25th percentile of hospitalization states, as in the left panel. The red markers are based on observed consumption and out-of-pocket health spending. The blue markers are based on a simple model of counterfactual consumption and out-of-pocket health spending were there no health care-related safety net. Counterfactual no-safety net out-of-pocket health spending is $oop_{\omega}^{no} = oop_{\omega} + E(sn|hosp_{\omega})$, where oop_{ω} is observed out-of-pocket health spending and $E(sn|hosp_{\omega}, y_{\omega})$ is such that (i) average counterfactual no-safety net out-of-pocket health spending is the same in both sets of hospitalization states and (ii) average safety net support in hospitalization states as a whole is \$6,000 per year, Dobkin et al.'s (2018) estimate of the effect of hospitalization on unpaid bills among the uninsured, likely a lower bound on the costs paid by external parties for the average hospital admission among the non-elderly uninsured. Counterfactual no-safety net consumption is $c_{\omega}^{no} = \max\{\underline{c},\ c - E(sn|hosp_{\omega})\}$.

Table 3: Structural model results

	Baseline	No income risk	No income risk and no implicit HI
	(1)	(2)	(3)
Insurance value	-479	68	7,220
Transfer value	2,380	2,320	6,397
Markup	-0.20	0.03	1.13
$Corr\left(\widehat{\lambda}, V\right)$	-0.09	1.00	0.68
$Corr\left(\widehat{\lambda}, hi\right)$	0.07	0.50	0.78
$Corr(hi, y)$	0.04	0.00	0.00
$Corr(V, y)$	0.35	0.00	0.00
$Sd(oop)$	1,733	1,583	10,040

Notes: Table shows results from the structural model. The first column is the baseline specification. The second column is a counterfactual world without income risk: Health risk is as in the data, $h \sim F(h)$, but income is the median of its empirical distribution in all states, $y \equiv y_{med}$. The third column also has no income risk and turns off implicit health insurance as well: $ihi(h, y; HI) \equiv 0$. The insurance value is the amount by which the equivalent variation of health insurance exceeds its mean ex post value ("transfer value"). The equivalent variation is a consumption-based measure: It is the amount by which the *consumption* of a household without health insurance must be increased in all states of the world in order to make the household as well off as it would be with health insurance. The markup is the ratio of insurance value to transfer value. All results are for households whose heads are non-elderly high school graduates without college degrees.

Online Appendix

Anti-Insurance: Health Insurance Worsens Risk Exposure

Lee M. Lockwood

A Theory Appendix

This appendix section characterizes analytically the effects of the safety net on the insurance value of health insurance in a highly stylized model. Assume, as in Section 2.2, that utility is a state-independent, increasing, concave function of non-health consumption c (hereafter, "consumption"),

$$c(\theta, HI) = y(\theta) - oop(\theta, HI), \tag{6}$$

where y is resources, *oop* is out-of-pocket health spending, and HI summarizes the health insurance contract the individual holds, if any. In this case, as shown in Section 2.2, to first order the insurance value of going from no health insurance to full health insurance is approximately

$$Cov\left(\widehat{\lambda}, V\right) \approx -\frac{\gamma(\overline{c})}{\overline{c}} Cov\left(c, oop\right), \tag{15}$$

where all variables take their values in the no-health insurance counterfactual (notation suppressed for simplicity), $\overline{c} \equiv E(c)$ is mean consumption, and $\gamma(\overline{c}) \equiv \frac{-\overline{c}u''(\overline{c})}{u'(\overline{c})} > 0$ is the coefficient of relative risk aversion at $\overline{c}$. The approximation is exact if marginal utility is linear in consumption.

Suppose that the individual's out-of-pocket health spending without health insurance would be

$$oop = h - sn = h - (s_h h - s_y y) = (1 - s_h)h + s_y y, \tag{16}$$

where h is the total cost of the individual's health care consumption, sn is total (explicit or implicit) support from the health care-related safety net, and the values of the two safety net parameters, $s_h \in [0, 1]$ and $s_y \in [0, 1]$, are such that $oop \in [0, h]$ in all states. Assuming for simplicity that health care costs and resources are uncorrelated, $Cov(y, h) = 0$, then to

first order health insurance's insurance value is approximately

$$\underbrace{Cov\left(\widehat{\lambda}, V\right)}_{\text{Overall net insurance value}} \approx \underbrace{\frac{\gamma(\overline{c})}{\overline{c}}(1-s_h)^2 Var(h)}_{\text{Gross pro-insurance benefit}} - \underbrace{\frac{\gamma(\overline{c})}{\overline{c}}s_y(1-s_y)Var(y)}_{\text{Gross anti-insurance cost}} . \quad (17)$$

This equation shows that if safety net support is decreasing in resources ($s_y > 0$), then the overall net insurance value of health insurance is the net of two competing factors: a pro-insurance benefit from insuring health care cost risk h and an anti-insurance cost from worsening non-health care cost risk y.[44] Intuitively, whereas the portion of the variation in out-of-pocket health spending driven by health shocks is a source of risk and so is valuable to insure, the portion of the variation in out-of-pocket health spending driven by non-health shocks absorbs risk—it hedges non-health risk—and so is costly to insure.

Equation (17) shows that the safety net decreases the insurance value of health insurance in two ways. First, it decreases the gross pro-insurance benefit of insuring health risk. It does so by decreasing both the extent to which consumption varies with health ("marginal utility gaps") and the extent to health insurance targets bad-health states ("value gaps").[45] Because of these reinforcing effects, safety net coverage of otherwise-uninsured health care costs can decrease health insurance's gross pro-insurance benefit significantly. For example, if the safety net would cover 80% of total health care costs ($s_h = 0.80$) if the individual did not have health insurance, then the safety net decreases health insurance's gross pro-insurance benefit by a factor of 25 ($1/[1-0.80]^2 = 1/.04 = 25$).[46]

The second way in which the safety net decreases the insurance value of health insurance is by causing it to worsen non-health risk. In the simple case with $Cov(h, y) = 0$, gross health insurance benefits are neutral with respect to risk in resources y, neither improving nor worsening it. But *net* health insurance benefits, accounting for implicit taxation by the safety net, worsen risk in resources: $Cov(oop, y) = s_y Var(y) \geqslant 0$. This is because health

[44]This decomposition splits risks into two categories: health care cost risk and all other risks. This is not necessarily the same as splitting risk into health versus non-health risk, since some of the variation in resources y might be driven by health risk (e.g., earnings losses from bad health) and some of the variation in health care consumption h might be driven by non-health risk (e.g., demand responses to liquidity shocks).

[45]That both the marginal utility gaps and the value gaps are proportional to $(1 - s_h)$ is why the gross pro-insurance benefit is proportional to the square of $(1 - s_h)$.

[46]The safety net covering 80% of total health care costs is consistent with the estimate by the Kaiser Commission on Medicaid and the Uninsured that the uninsured pay for only about 20% of their health care costs out-of-pocket (Coughlin et al., 2014) and the estimate by Finkelstein, Hendren and Luttmer (2019) that low-income uninsured adults in the Oregon Health Insurance Experiment also paid only about 20% of their medical care costs out-of-pocket. The effect of the safety net on health insurance's gross pro-insurance benefit highlights an important distinction. The insurance value of health insurance's additional, marginal insurance of health risk on top of that provided by the safety net might be low even if the insurance value of total insurance against health risk, including that provided by the safety net, is high.

insurance displaces safety net support against otherwise-uninsured health care costs to a greater extent in states in which resources are lower. As a result, health insurance reduces the extent to which the safety net insures risk in resources and thereby increases the welfare cost of that risk.[47] Equivalently, the safety net causes health insurance to worsen risk in resources, by implicitly taxing health insurance benefits to a greater extent when resources are lower.

Equation (17) shows that the safety net-induced gross anti-insurance cost of health insurance's worsening of risk in resources is proportional to the variance of resources, since more variable resources increases both the extent to which consumption varies with resources ("marginal utility gaps") and the extent to health insurance worsens risk in resources by targeting high-resource states ("value gaps"). The gross anti-insurance cost of health insurance's worsening of risk in resources is non-monotonic in s_y (proportional to $s_y(1-s_y)$). On the one hand, the greater is s_y, the greater the implicit insurance against risk in resources from the health care-related safety net that is displaced by health insurance. On the other hand, the greater is s_y, the smaller the effective resources risk that the individual faces accounting for the implicit insurance against risk in resources from the safety net.

B Data Appendix

B.1 Panel Study of Income Dynamics (PSID)

The PSID sample covers households interviewed in at least one of the ten waves between 1999 and 2017 inclusive. 1999 is the first year that certain key variables, such as many of the measures of consumption and health spending, were collected. 2017 is the last year of data available as of this writing. The final sample contains 83,831 household-wave observations.

The baseline measure of non-health consumption is annual expenditure on food, housing, transportation, clothing, travel, recreation, education, and child care. I add the value of food stamps to the measure of food and so overall consumption. Given the possibility of measurement error and the sensitivity of marginal utility to low consumption levels, I impose a consumption floor of $5,000.

The baseline measure of out-of-pocket health spending includes spending on hospital care, surgeries, dental bills, and prescriptions. It also includes spending on nursing home care,

[47] That safety net support is decreasing in resources means that safety net protection against health care costs implicitly insures risk in resources as well as health spending risk. Even though the form of support is coverage of health costs, that such support is negatively related to resources means that it partially insures that risk.

but I limit the sample to individuals not in nursing homes.

"Unemployment" is an indicator of whether the head or spouse was unemployed at any point in the past year. "Hospital" is an indicator of whether head or spouse was hospitalized at any point in the past year. For the analyses of hospitalization risk, I exclude hospitalizations in waves in which the household includes a child under two years old. I do this to exclude hospitalizations related to child-birth, in order to better focus on hospitalizations driven by health shocks, as in Dobkin et al. (2018).

Liquid assets are defined as holdings of checking or savings accounts, money market funds, certificates of deposit, government bonds, and Treasury bills, excluding those in employer-based pensions or IRAs.

Variables judged a priori likely to have large outliers—in this case, the monetary variables consumption, out-of-pocket health spending, income, liquidity—are winsorized at their (weighted) first and 99th percentiles; that is, values below the first percentile are set equal to the first percentile and values above the 99th percentile are set equal to the 99th percentile.

Summary statistics on the main estimation samples are in Appendix Table A1.

B.2 Medical Expenditure Panel Survey (MEPS)

I use the Household Component of the MEPS, which is a nationally representative survey of the U.S. civilian non-institutionalized population. I use all of the waves from 1996–2017. 2017 is the last wave available as of this writing. The resulting sample has 293,040 family-year observations. Summary statistics on the main estimation samples are in Appendix Table A2.

Variables judged a priori likely to have large outliers—in this case, measures of health care consumption, health care costs, health care expenditure, health care charges, and income—are winsorized at their (weighted) first and 99th percentiles; that is, values below the first percentile are set equal to the first percentile and values above the 99th percentile are set equal to the 99th percentile.

C Decomposing insurance value

To quantify the welfare benefit or cost of health insurance's improving or worsening of a particular risk, decompose the first order approximation to insurance value as

$$Cov\left(\widehat{\lambda}, V\right) = \underbrace{Cov_g\left[E\left(\widehat{\lambda}|g\right), E\left(V|g\right)\right]}_{\text{Across-}g\text{ targeting}} + \underbrace{E_g\left[Cov_g\left(\widehat{\lambda}, V|g\right)\right]}_{\text{Within-}g\text{ targeting}}, \tag{18}$$

where $g \in G$ is a partition of the state space, i.e., a division of the state space into mutually exclusive, exhaustive subsets g.[48] For example, when the state space is partitioned based on the realization of income, the first term is a measure of the insurance value of health insurance's worsening of income risk. It is the contribution to health insurance's overall insurance value from its average transfers across groups of states with similar income realizations. It reflects the insurance value of health insurance's differential targeting of higher- versus lower-income states on average. If in higher-income states the ex post value of health insurance tends to be greater and marginal utility lower, health insurance's differential targeting of higher-income states has negative insurance value.The second term in equation (18) is the contribution to health insurance's overall insurance value from its transfers within each group of states g. If within sets of states with similar income realizations health insurance tends to help more when marginal utility is lower, such differential targeting has negative insurance value.

This decomposition can be used to quantify health insurance's worsening of income risk under much more general assumptions than those required for the "direct versus portfolio effect" analysis in Section 4.2. The results are reported in the first column of Appendix Table A7. Dividing the state space into quintiles of realized annual income, I find that health insurance's average net transfers across income quintiles, which shift resources from lower- to higher-income states, account for about half of health insurance's net anti-insurance cost. The remaining half comes from its worsening of within-income quintile risk, which arises because it transfers from higher- to lower-marginal utility states within income quintiles on average. These results suggest that health insurance's worsening of income risk is a key driver of its anti-insurance cost.[49]

[48] This decomposition builds on that in Deshpande and Lockwood (2020), which decomposes the markup on (binary) disability insurance benefits into across- versus within-group components.

[49] To be clear, this decomposition measures the contribution to health insurance's marginal insurance value of its average net transfers across *quintiles* of realized annual income versus its net transfers within those quintiles. The "across" component does not capture exactly the cost of health insurance's worsening of annual income risk for two reasons. On the one hand, it leaves out within-quintile income risk, which is a force toward making it understate the cost of health insurance's worsening of income risk. On the other hand, it includes any other, non-income risks that are correlated with realized income quintile (e.g., wealth), which, to the extent that such correlated factors tend to be positively correlated with income, tends to make

D Other mechanisms that could cause health insurance to worsen non-health risk

Why does health insurance worsen non-health risk? Perhaps the most obvious possibility is correlated risks. If non-health risk were sufficiently negatively correlated with health risk, then insuring health risk, as health insurance does, would worsen non-health risk. But using detailed data from the Medical Expenditure Panel Survey on health care consumption, I find that although certain types of health care consumption, especially office visits, covary slightly positively with income, health care consumption as a whole, as measured by total charges, covaries slightly negatively with income (see Appendix Table A9). Of course, this does not mean that health and non-health shocks are negatively correlated since income reflects the combined effects of health and non-health shocks. But it does suggest that health and non-health shocks are not sufficiently negatively correlated to cause health insurance to worsen income risk. The (slight) negative relationship between heath care consumption and income is a force toward making health insurance partially hedge income risk, which tends to increase health insurance's overall insurance value. The reason that health insurance worsens risk exposure must lie elsewhere.

A second possibility is demand responses. To the extent that the demand for health care is sensitive to non-health shocks, that could cause health insurance to worsen non-health risk. For example, if in the absence of health insurance people would cut back on or postpone health care consumption in the event of unemployment, health insurance would be worth less in unemployed states of the world and thereby exacerbate that risk. Such a mechanism is consistent with both the responsiveness of health care consumption to non-health driven changes in liquidity found by Gross, Layton and Prinz (2020) and my finding, reported in Appendix Table A9, that office visits covary positively with income. But while this mechanism is a force toward health insurance worsening non-health risk, it is not strong enough to make health insurance worsen income risk on net (since, as just discussed, total charges are negatively correlated with income).

it overstate the cost of health insurance's worsening of income risk. Analogous considerations apply to the other decompositions based on equation (18) as well.

Appendix Figures and Tables

Table A1: Summary statistics of the main estimation samples in the PSID

	Non-elderly			Elderly
	All	Uninsured	Insured	
Age	44.6	42.6	44.9	75.2
Family size	2.5	2.1	2.6	1.6
Income	86,096	42,651	92,754	59,462
Consumption	43,312	30,222	45,314	28,817
Out-of-pocket health spending	1,298	917	1,359	1,906
Hospital	0.13	0.10	0.14	0.24
Unemployment	0.11	0.20	0.09	0.02

Notes: Table shows summary statistics of the main estimation samples in the Panel Study of Income Dynamics (PSID). These are family-level averages using family weights. Monetary variables are in dollars per year. Non-elderly are families whose head is between 25 and 64 years old, inclusive. Elderly are families whose head is 65 and older. Hospital and Unemployment are indicators of whether the head or spouse was hospitalized or unemployed in the last year. I use the 1999–2017 waves (which occur in odd-numbered years), with a total sample size of 77,806 household-wave observations.

Table A2: Summary statistics of the main estimation samples in the MEPS

	Non-elderly			Elderly
	All	Uninsured	Insured	
Age	43.8	42.4	44.2	74.6
Family size	2.6	1.7	2.5	1.6
Income	85,203	42,018	92,060	58,943
Out-of-pocket health spending	1,520	968	1,594	2,265
Health care charges	25,581	7,970	26,729	50,513
Health care total expenditure	12,796	3,354	13,806	22,067

Notes: Table shows summary statistics of the main estimation samples in the Medical Expenditure Panel Survey (MEPS). These are family-level averages using family weights. Monetary variables are in dollars per year. Non-elderly are families whose head is between 25 and 64 years old, inclusive. Elderly are families whose head is 65 and older. The MEPS top-codes age (at 90 from 1996–2000 and at 85 from 2001–2018), so the reported average age of the elderly sample in this table is affected by that. All analyses of the MEPS data in this paper that control for age include an indicator of whether the age is the top-coded value. I use the 1996–2018 waves, with a total sample size of 293,040 household-wave observations.

Table A3: Sufficient statistic estimates: Heterogeneity

	All	Education				Liquidity (lagged)		Age		Health (lagged)	
		<HS or GED	High school	Some college	College+	⩽$500	>$500	25–39	40–64	Good	Bad
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
$\widehat{\beta} \times \left(sd(oop)/sd\left(\widehat{\lambda}\right)\right) = Corr\left(\widehat{\lambda}, oop\right)$	-0.04	-0.07	-0.06	-0.05	-0.03	-0.07	-0.02	-0.02	-0.05	-0.04	-0.08
Insurance value ($), $Cov\left(\widehat{\lambda}, oop\right)$	-326	-323	-361	-363	-385	-338	-203	-147	-406	-369	-536
se(Insurance value) ($) $= se\left(\widehat{\beta}\right) \times var(oop)$	(43)	(62)	(58)	(88)	(135)	(49)	(68)	(45)	(59)	(63)	(92)
Transfer value ($), $E\left(oop\right)$	1,298	909	1,201	1,171	1,604	839	1,457	881	1,539	1,375	1,464
Markup, $M = Cov\left(\widehat{\lambda}, oop\right)/E\left(oop\right)$	-0.25	-0.36	-0.30	-0.31	-0.24	-0.40	-0.14	-0.17	-0.26	-0.27	-0.37

Notes: Table presents estimates related to the insurance value of eliminating out-of-pocket health spending. $\widehat{\beta}$ is the estimate of the coefficient on out-of-pocket health spending in a regression of normalized marginal utility on out-of-pocket health spending and controls (equation 8), scaled to be in standard deviation units (so a coefficient of -0.04 means that a one-standard deviation increase in out-of-pocket health spending is associated with a 0.04 decrease in marginal utility. The standard error of this coefficient, shown on the next line, is expressed in the same units. Column (1) reproduces the main results for all non-elderly state-times. Columns (2)–(5) present results for each of the four different education groups. These results show estimates of the value of completing health insurance (eliminating out-of-pocket health spending) in all non-elderly state-times separately for households in each education category. These results show heterogeneity across different "ex ante risk types" (households with different education levels). Columns (6)–(11), by contrast, show heterogeneity across different state-times for a given risk type, i.e., heterogeneity in the value of completing health insurance in different state-times. Columns (6) and (7) split the set of state-times into two groups: those in which lagged liquidity is smaller or greater than $500. Liquidity is defined as holdings of checking or savings accounts, money market funds, certificates of deposit, government bonds, and Treasury bills, excluding those in employer-based pensions or IRAs. Its median is about $1,400. Lagged liquidity is liquidity in the preceding wave. Columns (8) and (9) split the set of state-times into two groups: those in which the household head's age is 25–39 or 40–64. Columns (10) and (11) split the set of state-times into two groups: those in which the household head's self-reported health status is (i) X, Y, or Z ("Good") or X or Y ("Bad"). The markup is the insurance value per dollar of transfer value, $M = Cov\left(\widehat{\lambda}, oop\right)/E\left(oop\right)$. Monetary amounts are in dollars per household per year.

Table A4: Sufficient statistic estimates: Robustness

	Baseline	No ed controls	Avg income control	Health, income controls	Log utility	Food consumption	C proxy $c = y - oop$	Costly safety net
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\widehat{\beta} \times \left(sd(oop)/sd\left(\widehat{\lambda}\right)\right) = Corr\left(\widehat{\lambda}, oop\right)$	-0.04	-0.06	-0.03	-0.04	-0.11	-0.03	-0.02	-0.03
$se\left(\widehat{\beta}\right) \times \left(sd(oop)/sd\left(\widehat{\lambda}\right)\right)$	(0.006)	(0.005)	(0.005)	(0.006)	(0.006)	(0.006)	(0.007)	(0.006)
Insurance value (\$), $Cov\left(\widehat{\lambda}, oop\right)$	-326	-456	-197	-277	-168	-253	-179	-288
se(Insurance value) (\$) $= se\left(\widehat{\beta}\right) \times var(v)$	(43)	(43)	(37)	(42)	(9)	(47)	(58)	(48)
Transfer value (\$), $E\left(oop\right)$	1,298	1,298	1,298	1,298	1,298	1,298	1,298	1,507
Markup, $M = Cov\left(\widehat{\lambda}, oop\right)/E\left(oop\right)$	-0.25	-0.35	-0.15	-0.21	-0.13	-0.20	-0.14	-0.19

Notes: Table presents estimates related to the insurance value of going from status quo health insurance to full health insurance (no out-of-pocket health spending). $\widehat{\beta}$ is the estimate of the coefficient on V in a regression of normalized marginal utility on V and controls (equation 8), scaled to be in standard deviation units (so a coefficient of -0.04 means that a one-standard deviation increase in V is associated with a 0.04 decrease in marginal utility. The standard error of this coefficient, shown on the next line, is expressed in the same units. In all but the last column, column (8), the ex post value of going from status quo health insurance to full health insurance is assumed to equal out-of-pocket health spending in that state, $V = oop$. Column (1) reproduces the main results. Columns (2)–(4) use different sets of control variables. The baseline controls (column (1)) are year dummies, a cubic in age, a quadratic in household size, and education category dummies. Column (2) drops the education category dummies, in which case the controls match those in Finkelstein, Hendren and Luttmer (2019). Column (3) substitutes average income (of each household, across survey waves) for the education category dummies as an alternative control for earning-ability type. This is meant to be conservative, i.e., to bias the specification against health insurance worsening risk exposure, in that it over-controls for differences in ex ante, long-term earning ability by controlling for ex post realized income as well. Column (4) adds lagged income and dummies for lagged self-reported health status to the baseline controls. The aim is to isolate short-term risk. Column (5) uses a coefficient of relative risk aversion of one (log utility) rather than the baseline value of three. Column (6) assumes that marginal utility is a function of food consumption rather than total non-health consumption. Column (7) assumes that marginal utility is a function of the "consumption proxy" of income less out-of-pocket health spending with a floor, $c = \max\{\$5,000,\ y - oop\}$. Column (8) adds \$2,000 to the ex post value of health insurance in states in which the household experiences a hospitalization. The aim is to overstate any additional ex post value of health insurance to the household, over and above that from reducing its out-of-pocket health spending, from reducing its usage of the safety net and so any associated private costs (e.g., hassle or stigma costs, or financial costs of declaring bankruptcy or otherwise having bad debt). The markup is the insurance value per dollar of transfer value, $M = Cov\left(\widehat{\lambda}, V\right)/E\left(V\right)$. Monetary amounts are in dollars per household per year.

Table A5: Health insurance and income risk: Heterogeneity

	All	Education				Liquidity (lagged)		Age		Health (lagged)	
		<HS or GED	High school	Some college	College+	⩽$500	>$500	25–39	40–64	Good	Bad
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
$\widehat{\beta}_{y\|oop}$	4.91	4.54	4.20	4.81	5.14	4.40	4.14	5.43	4.50	5.08	4.55
(se)	(0.32)	(0.62)	(0.43)	(0.46)	(0.59)	(0.41)	(0.36)	(0.46)	(0.38)	(0.37)	(0.54)
$Corr(y, oop)$	0.16	0.23	0.18	0.17	0.15	0.19	0.13	0.17	0.15	0.16	0.24
$Sd(oop)$	2,038	1,869	1,970	1,892	2,200	1,706	2,119	1,660	2,228	2,064	2,366
$Sd(y)$	62,316	37,154	46,871	55,072	77,616	39,095	65,770	52,434	66,755	65,194	45,194

Notes: Table presents estimates related to health insurance's worsening of income risk. It is a supporting table to Table 2. All columns use the non-elderly sample. Column (1) reproduces the main results for all non-elderly state-times (as in column (1) of Table 2). Columns (2)–(5) present results for each of the four different education groups. These results show estimates in all non-elderly state-times separately for households in each education category. These results show heterogeneity across different "ex ante risk types" (households with different education levels). Columns (6)–(11), by contrast, show heterogeneity across different state-times for a given risk type. Columns (6) and (7) split the set of state-times into two groups: those in which lagged liquidity is smaller or greater than $500. Liquidity is defined as holdings of checking or savings accounts, money market funds, certificates of deposit, government bonds, and Treasury bills, excluding those in employer-based pensions or IRAs. Its median is about $1,400. Lagged liquidity is liquidity in the preceding wave. Columns (8) and (9) split the set of state-times into two groups: those in which the household head's age is 25–39 or 40–64. Columns (10) and (11) split the set of state-times into two groups: those in which the household head's self-reported health status is (i) X, Y, or Z ("Good") or X or Y ("Bad"). The first row shows the coefficient estimate on out-of-pocket health spending from a regression of income on out-of-pocket health spending and the baseline controls (year dummies, a cubic in age, a quadratic in household size, and education category dummies). In a simple model in which marginal utility is linear in consumption and consumption equals income minus out-of-pocket health spending, this regression coefficient equals the ratio of the "portfolio effect" cost from worsening income risk to the "direct effect" benefit from improving health risk (see equations (7) and (9)). The second row is the corresponding standard error. The third row is the correlation between income and out-of-pocket health spending, residualized with the baseline controls. The fourth and fifth rows are the standard deviations of out-of-pocket health spending and income, respectively, each residualized with the baseline controls.

Table A6: Health insurance and income risk: Robustness

	Baseline	No ed controls	Avg income control	Health, income controls	Costly safety net
	(1)	(2)	(3)	(4)	(5)
$\widehat{\beta}_{y\|oop}$	4.91	6.64	0.93	1.25	3.86
(se)	(0.32)	(0.35)	(0.12)	(0.15)	(0.29)
$Corr(y, oop)$	0.16	0.20	0.05	0.06	0.14
$Sd(oop)$	2,038	2,041	1,999	2,078	2,038
$Sd(y)$	62,316	67,471	35,123	40,589	62,316

Notes: Table presents estimates related to health insurance's worsening of income risk. It is a supporting table to Table 2. All columns use the non-elderly sample. Column (1) reproduces the main results (as in column (1) of Table 2). Columns (2)–(4) use different sets of control variables. The baseline controls (column (1)) are year dummies, a cubic in age, a quadratic in household size, and education category dummies. Column (2) drops the education category dummies, in which case the controls match those in Finkelstein, Hendren and Luttmer (2019). Column (3) substitutes average income (of each household, across survey waves) for the education category dummies as an alternative control for earning-ability type. This is meant to be conservative, i.e., to bias the specification against health insurance worsening risk exposure, in that it over-controls for differences in ex ante, long-term earning ability by controlling for ex post realized income as well. Column (4) adds lagged income and dummies for lagged self-reported health status to the baseline controls. The aim is to isolate short-term risk. Column (5) adds \$2,000 to the ex post value of health insurance in states in which the household experiences a hospitalization. The aim is to overstate any additional ex post value of health insurance to the household, over and above that from reducing its out-of-pocket health spending, from reducing its usage of the safety net and so any associated private costs (e.g., hassle or stigma costs, or financial costs of declaring bankruptcy or otherwise having bad debt). The first row shows the coefficient estimate on out-of-pocket health spending from a regression of income on out-of-pocket health spending and the baseline controls (year dummies, a cubic in age, a quadratic in household size, and education category dummies). In a simple model in which marginal utility is linear in consumption and consumption equals income minus out-of-pocket health spending, this regression coefficient equals the ratio of the "portfolio effect" cost from worsening income risk to the "direct effect" benefit from improving health risk (see equations (7) and (9)). The second row is the corresponding standard error. The third row is the correlation between income and out-of-pocket health spending, residualized with the controls. The fourth and fifth rows are the standard deviations of out-of-pocket health spending and income, respectively, each residualized with the controls.

Table A7: Decompositions of health insurance's insurance value

	Income	Unemployment	Hospitalization	Self-reported health
Across share	0.47	0.05	-0.06	-0.12
$Corr_g\left[E\left(\widehat{\lambda}\middle\vert g\right), E\left(V\middle\vert g\right)\right]$	-0.61	-1.00	1.00	0.90
$Sd_g\left[E\left(V\middle\vert g\right)\right]$	252	69	345	119
$Sd_g\left[E\left(\widehat{\lambda}\middle\vert g\right)\right]$	0.99	0.26	0.08	0.36
Within share	0.53	0.95	1.06	1.12
$E_g\left[Corr_g\left(\widehat{\lambda}, V\middle\vert g\right)\right]$	-0.04	-0.04	-0.05	-0.04
$E_g\left[Sd_g\left(V\middle\vert g\right)\right]$	1,985	2,042	2,100	2,028
$E_g\left[Sd_g\left(\widehat{\lambda}\middle\vert g\right)\right]$	2.76	3.86	4.29	3.64

Notes: Table shows statistics related to decompositions of health insurance's insurance value based on equation (18) for non-elderly households. These divide the state space into quintiles of realized income (first column), whether the head or spouse experienced unemployment (second column), whether the head or spouse experienced a hospitalization (third column), and the head's self-reported health status (fourth column). "Across share" is the share of health insurance's insurance value accounted for by health insurance's average net transfers across the groups of states defined by the column header, i.e., the ratio of the "across-g targeting" term in equation (18) to the overall insurance value: $Cov_g\left[E\left(\widehat{\lambda}\middle\vert g\right), E\left(V\middle\vert g\right)\right]/Cov\left(\widehat{\lambda}, V\right)$. "Within share" is the share of health insurance's insurance value accounted for by its transfers within each group of states defined by the column header, i.e., the ratio of the "within-g targeting" term in equation (18) to the overall insurance value: $E_g\left[Cov_g\left(\widehat{\lambda}, V\middle\vert g\right)\right]/Cov\left(\widehat{\lambda}, V\right)$.

Table A8: Hospitalization and unemployment

	Pr(event)	$E(V$|event$)$	$E(V$|not-event$)$	Value gap	Markup to event	Markup to not-event	Markup gap	Insurance value per event state
Hospitalization	0.13	2,209	1,168	1,041	0.17	-0.02	0.19	176
Unemployment	0.11	1,094	1,320	-225	0.63	-0.15	0.78	-157
Ratio: |Ue|/|Hosp|	0.82			0.22			4.02	0.89

Notes: Table shows statistics related to hospitalization and unemployment for non-elderly households. Pr(event) is the share of households experiencing the event. $E(V$|event$)$ and $E(V$|not-event$)$ are average out-of-pocket health spending when the household does versus does not experience the event. The value gap is the difference between the two. "Markup to (not-)event" is the markup on transfers from all states to states in which the household does (not) experience the event. The markup gap is the difference between the two. "Insurance value per event state" is the insurance value surplus generated by the *average* differential payment into event states relative to not-event states per event state, which equals the product of the value gap, the markup gap, and $(1 - Pr(\text{event}))$. This is the insurance value of a hypothetical binary cash indemnity benefit that paid off the amount of the value gap in the event state and zero in the non-event state. Health insurance, of course, is *not* a binary benefit; it targets within event and non-event states as well as across event and non-event states. So these insurance values are *not* the value of the insurance that health insurance provides against these risks. For more on the value of the insurance that health insurance provides against hospitalization risk, see Section 4.3. Value gaps and markups are estimated after residualizing V and normalized marginal utility with the baseline controls (year dummies, a cubic in age, a quadratic in household size, and education category dummies). Results are similar for households with versus without health insurance.

Table A9: Correlation between income and various measures of health care utilization and spending

	Non-elderly			Elderly
	All	Uninsured	Insured	
Charges				
Total	-0.02	0.00	-0.03	-0.01
Office visits	0.00	0.02	0.00	0.01
Outpatient hospital	0.01	0.01	0.00	0.00
Outpatient doctor	0.01	0.02	0.00	0.00
Inpatient	-0.03	-0.02	-0.03	-0.02
Quantities				
Office visits	0.04	0.07	0.01	0.05
Outpatient hospital	-0.01	0.01	-0.02	-0.01
Outpatient doctor	-0.01	0.00	-0.02	-0.01
Inpatient discharge	-0.05	-0.03	-0.06	-0.03
Inpatient night	-0.06	-0.03	-0.06	-0.03
Prescriptions	-0.06	-0.02	-0.08	-0.08
Out-of-pocket spending	0.13	0.10	0.13	0.10

Notes: This table shows the correlation between income and various measures of health care utilization and spending (row names) for each of four samples (column names). All variables are residualized using the baseline controls (year dummies, a cubic in age, a quadratic in household size, and education category dummies). Data are from the Medical Expenditure Panel Survey. In principle, the positive correlation between income and out-of-pocket health spending could arise from a positive correlation between income and health care costs (which in turn could arise from a negative correlation between health and non-health shocks ["correlated risks"] or responsiveness of the demand for health care to non-health shocks ["demand responses"]). In practice, health care consumption and total health care costs are only weakly correlated with income and, other than the exception of office visits, tend to be negatively correlated, and so do not account for the positive correlation between income and out-of-pocket health spending.

Table A10: Health insurance and hospitalization risk: Sufficient statistic estimates

	Non-elderly (1)	Non-elderly uninsured (2)	Non-elderly insured (3)	Elderly insured (4)
$\widehat{\beta} \times \left(sd(V)/sd\left(\widehat{\lambda}\right)\right) = Corr\left(\widehat{\lambda}, V\right)$	-0.02	-0.02	-0.02	-0.02
$se\left(\widehat{\beta}\right) \times \left(sd(V)/sd\left(\widehat{\lambda}\right)\right)$	(0.006)	(0.014)	(0.008)	(0.013)
Insurance value (\$), $Cov\left(\widehat{\lambda}, V\right)$	-75	-74	-80	-91
se(Insurance value) (\$) $= se\left(\widehat{\beta}\right) \times var(V)$	(27)	(41)	(36)	(56)
Transfer value (\$), $E\left(V\right)$	1,435	949	1,517	2,083
Markup, $M = Cov\left(\widehat{\lambda}, V\right)/E\left(V\right)$	-0.05	-0.08	-0.05	-0.04

Notes: Table presents estimates related to the insurance value of health insurance's targeting of hospitalization risk. To isolate health insurance's targeting of hospitalization risk, I "shut down" the targeting in non-hospitalization states by setting the ex post value of health insurance to its average value (average out-of-pocket health spending) in non-hospitalization states (whereas in each hospitalization state it equals out-of-pocket health spending in that state): $V = hosp \times oop + (1 - hosp) \times E(oop|hosp = 0)$, where $hosp$ equals one if the household head or spouse experienced a hospitalization in the past year *and* there is no newborn present in the household (to exclude hospitalizations related to childbirth) and zero otherwise. This means there is no within-non-hospitalization-states targeting, only targeting from non-hospitalization states to hospitalization states and within hospitalization states. $\widehat{\beta}$ is the estimate of the coefficient on the ex post value of health insurance V in a regression of normalized marginal utility on V and controls (equation 8), scaled to be in standard deviation units (so a coefficient of -0.02 means that a one-standard deviation increase in V is associated with a 0.02 decrease in marginal utility. The standard error of this coefficient, shown on the next line, is expressed in the same units. The controls are year dummies, a cubic in age, a quadratic in household size, and education category dummies. The markup is the insurance value per dollar of transfer value, $M = Cov\left(\widehat{\lambda}, oop\right)/E\left(oop\right)$. Monetary amounts are in dollars per household per year.

Table A11: Health insurance and hospitalization risk: Sufficient statistic estimates robustness

	Baseline	Costly safety net	$V=0$ if non-hosp	Log utility	Food consumption	C proxy $c=y-oop$	Avg income control	Health, income controls
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\widehat{\beta}\times\left(sd(V)/sd\left(\widehat{\lambda}\right)\right)=Corr\left(\widehat{\lambda},V\right)$	-0.02	-0.01	-0.01	-0.05	-0.02	-0.01	-0.01	-0.01
$se\left(\widehat{\beta}\right)\times\left(sd(V)/sd\left(\widehat{\lambda}\right)\right)$	(0.006)	(0.006)	(0.006)	(0.007)	(0.004)	(0.007)	(0.006)	(0.007)
Insurance value (\$), $Cov\left(\widehat{\lambda},V\right)$	-75	-36	-49	-35	-100	-36	-43	-61
se(Insurance value) (\$) $=se\left(\widehat{\beta}\right)\times var(v)$	(27)	(36)	(32)	(5)	(19)	(31)	(27)	(30)
Transfer value (\$), $E\left(V\right)$	1,435	1,630	241	1,435	1,435	1,435	1,435	1,435
Markup, $M=Cov\left(\widehat{\lambda},V\right)/E\left(V\right)$	-0.05	-0.02	-0.20	-0.02	-0.07	-0.03	-0.03	-0.04

Notes: Table presents estimates related to the insurance value of health insurance's targeting of hospitalization risk. $\widehat{\beta}$ is the estimate of the coefficient on the ex post value of health insurance V in a regression of normalized marginal utility on V and controls (equation 8), scaled to be in standard deviation units (so a coefficient of -0.02 means that a one-standard deviation increase in V is associated with a 0.02 decrease in marginal utility. The standard error of this coefficient, shown on the next line, is expressed in the same units. The markup is the insurance value per dollar of transfer value, $M=Cov\left(\widehat{\lambda},oop\right)/E\left(oop\right)$. Monetary amounts are in dollars per household per year. Column (1) reproduces the main results, reported in Table A10. These set $V=hosp\times oop+(1-hosp)\times E(oop|hosp=0)$, where *hosp* equals one if the household head or spouse experienced a hospitalization in the past year *and* there is no newborn present in the household (to exclude hospitalizations related to childbirth) and zero otherwise. This isolates health insurance's targeting of hospitalization risk by "shutting down" the targeting in non-hospitalization states. The baseline controls are year dummies, a cubic in age, a quadratic in household size, and education category dummies. Column (2) adds \$2,000 to the ex post value of health insurance in all states in which the household experiences a hospitalization. The aim is to overstate any additional ex post value of health insurance to the household, over and above that from reducing its out-of-pocket health spending, from reducing its usage of the safety net and so any associated private costs (e.g., hassle or stigma costs, or financial costs of declaring bankruptcy or otherwise having bad debt). Column (3) uses $V=hosp\times oop$. This overstates the extent to which health insurance targets hospitalization states by ignoring health insurance transfers to non-hospitalization states and so tends to overstate health insurance's insurance value. Column (4) uses a coefficient of relative risk aversion of one (log utility) rather than the baseline value of three. Column (5) assumes that marginal utility is a function of food consumption rather than total non-health consumption. Column (6) assumes that marginal utility is a function of the "consumption proxy" of income less out-of-pocket health spending with a floor, $c=\max\{\$5,000,\ y-oop\}$. Column (7) substitutes average income (of each household, across survey waves) for the education category dummies as an alternative control for earning-ability type. This is meant to be conservative, i.e., to bias the specification against health insurance worsening risk exposure, in that it over-controls for differences in ex ante, long-term earning ability by controlling for ex post realized income as well. Column (8) adds lagged income and dummies for lagged self-reported health status to the baseline controls. The aim is to isolate short-term risk.

Table A12: Health insurance and hospitalization risk: Decomposition estimates

	Baseline				No safety net			
	Insurance value	Contrib to markup	Transfer value	Benefit	Insurance value	Contrib to markup	Transfer value	Benefit
Non-elderly								
Total	-75	-0.66	112		1,417	2.07	684	
Non-hosp to hosp	22	0.20	112	1,153	1,714	2.51	684	7,029
Low- to high-y within hosp	-24	-0.21	0	-401	17	0.02	0	98
Within hosp-by-y	-73	-0.65	0		-313	-0.46	0	
Non-elderly uninsured								
Total	-74	-0.59	125		868	1.41	616	
Non-hosp to hosp	26	0.21	125	1,488	978	1.59	616	7,334
Low- to high-y within hosp	-15	-0.12	0	-570	-1	0.00	0	-10
Within hosp-by-y	-86	-0.69	0		-109	-0.18	0	

Notes: Estimates related to the decomposition of health insurance's targeting of hospitalization risk, as depicted in Figure 4a. This decomposes health insurance transfers into hospitalization states into the sum of three components. The first, "Non-hosp to hosp," is health insurance's average net transfer from non-hospitalization to hospitalization states, which equals the difference in average out-of-pocket health spending between hospitalization and non-hospitalization states. The second, "Low- to high-y within hosp," is health insurance's average net transfer from hospitalization states with lower- to higher-income realizations, defined as below versus above the 25th percentile of residualized income, residualized with the baseline controls. This equals the difference in average out-of-pocket health spending between hospitalization-and-high-income and hospitalization-and-low-income states. The third, "Within hosp-by-y," is health insurance's transfers within the two hospitalization-and-income-is-high/low categories, from states in which health insurance is less to more valuable. In each hospitalization state, this transfer equals the difference between out-of-pocket health spending in that state and average out-of-pocket health spending in that state's hospitalization-by-income group (hospitalization and income high versus low). Because this varies across states within a category, I do not report a benefit value in the "Benefit" column. The "Baseline" columns are based on observed out-of-pocket health spending and consumption (for marginal utility). The "No safety net" columns use a simple model of counterfactual consumption and out-of-pocket health spending were there no health care-related safety net. Counterfactual no-safety net out-of-pocket health spending is $oop_{\omega}^{no} = oop_{\omega} + E(sn|hosp_{\omega})$, where oop_{ω} is observed out-of-pocket health spending and $E(sn|hosp_{\omega}, y_{\omega})$ is such that (i) average counterfactual no-safety net out-of-pocket health spending is the same in both sets of hospitalization states and (ii) average safety net support in hospitalization states as a whole is \$6,000 per year, Dobkin et al.'s (2018) estimate of the effect of hospitalization on unpaid bills among the uninsured, likely a lower bound on the costs paid by external parties for the average hospital admission among the non-elderly uninsured. Counterfactual no-safety net consumption is $c_{\omega}^{no} = \max\{\underline{c},\ c - E(sn|hosp_{\omega})\}$. This coarse model of the safety net, which assumes that safety net support is uniform within each hospitalization-by-low/high-income groups, very likely understates safety net support in lower- relative to higher-income states within these groups and so likely leaves much of the safety net's transformation of health spending risk intact. This would help explain why the "No safety net" results still have health insurance worsening within-group risk. This is a force toward these results understating the value of health insurance in the "No safety net" counterfactual. The markup is the insurance value per dollar of transfer value. The transfer value here is the ex ante expected value of these three components as a whole, which comes entirely from the first component, the average net transfer from non-hospitalization to hospitalization states, since the second and third components have zero ex ante expected value given that they are transfers from some hospitalization states to others. A hospitalization state is defined as one in which the household head or spouse experienced a hospitalization in the past year *and* there is no newborn present in the household (to exclude hospitalizations related to childbirth).

Table A13: Heterogeneity results for the structural analysis

	Non-elderly			Elderly
	$<$ High school	High school	$\geqslant$ College	
	(1)	(2)	(3)	(4)
Insurance value	-102	-479	-909	-523
Transfer value	2,014	2,380	2,810	2,952
Markup	-0.05	-0.20	-0.32	-0.18
$Corr\left(\widehat{\lambda}, V\right)$	0.03	-0.09	-0.13	-0.22
$Corr\left(\widehat{\lambda}, hi\right)$	0.11	0.07	0.04	-0.01
$Corr(hi, y)$	0.01	0.04	0.07	0.07
$Corr(V, y)$	0.24	0.35	0.36	0.43
$Sd(oop)$	1,541	1,733	1,916	1,387

Notes: Table shows heterogeneity results for the structural analysis. The insurance value is the amount by which the equivalent variation of health insurance exceeds its mean ex post value ("transfer value"). The equivalent variation is a consumption-based measure: It is the amount by which the *consumption* of a household without health insurance must be increased in all states of the world in order to make the household as well off as it would be with health insurance. The markup is the ratio of insurance value to transfer value.

Table A14: Robustness tests of the structural analysis

	Baseline	Double implicit HI deductible	Residualize income	Income floor $50k	One-fourth income risk	One-tenth income risk	HI deductible $3k	Log utility
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Insurance value	-479	-214	-307	-263	-96	0	-438	-335
Transfer value	2,380	3,555	2,438	2,477	2,355	2,336	471	2,380
Markup	-0.20	-0.06	-0.13	-0.11	-0.04	0.00	-0.93	-0.14
$Corr\left(\widehat{\lambda}, V\right)$	-0.09	0.08	-0.04	-0.16	-0.03	0.19	-0.29	-0.22
$Corr\left(\widehat{\lambda}, hi\right)$	0.07	0.16	0.08	0.04	0.08	0.16	0.06	0.03
$Corr(hi, y)$	0.04	0.04	0.00	0.04	0.04	0.04	0.02	0.04
$Corr(V, y)$	0.35	0.25	0.25	0.28	0.23	0.20	0.52	0.35
$Sd(oop)$	1,733	3,227	1,759	1,791	1,629	1,601	1,733	1,733

Notes: Table shows robustness tests of the structural analysis. Column (1) is the baseline specification. Column (2) doubles the implicit deductible of implicit health insurance: $\widetilde{d}_{ihi}(y) = 2 \times d_{ihi}(y)$. Columns (3)–(6) reduce non-health risk. Column (3) residualizes income using the baseline controls (year dummies, a cubic in age, a quadratic in household size, and education category dummies). Column (4) increases the income floor from $15,000 to $50,000. Column (5) cuts income risk by a factor of four: It sets income in each state of the world to a weighted average of its observed value and median income across states, with three-fourths of the weight on median income: $\widetilde{y} = 0.25 \times y + 0.75 \times y_{median}$. Column (6) cuts income risk by a factor of ten: It sets income in each state of the world to a weighted average of its observed value and median income across states, with nine-tenths of the weight on median income: $\widetilde{y} = 0.1 \times y + 0.9 \times y_{median}$. Column (7) increases the deductible of the health insurance contract from $0 (i.e., full coverage in the baseline) to $3,000 per year for the household. Column (8) uses log utility (a coefficient of relative risk aversion of one). The insurance value is the amount by which the equivalent variation of health insurance exceeds its mean ex post value ("transfer value"). The equivalent variation is a consumption-based measure: It is the amount by which the *consumption* of a household without health insurance must be increased in all states of the world in order to make the household as well off as it would be with health insurance. The markup is the ratio of insurance value to transfer value. All results are for households whose heads are non-elderly high school graduates without college degrees.

Table A15: Sufficient statistic estimates: Different types of coverage and alternative types of contracts

	Coverage of health spending				Fixed indemnity benefits		
	Total (1)	Hospital (2)	Doctor (3)	Rx (4)	Hospitalized (5)	Hosp days (6)	Bad health (7)
$\widehat{\beta} \times \left(sd(V)/sd\left(\widehat{\lambda}\right)\right) = Corr\left(\widehat{\lambda}, V\right)$	-0.04	-0.02	-0.03	-0.05	0.02	0.03	0.09
$se\left(\widehat{\beta}\right) \times \left(sd(V)/sd\left(\widehat{\lambda}\right)\right)$	(0.006)	(0.004)	(0.006)	(0.006)	(0.007)	(0.011)	(0.013)
Insurance value (\$), $Cov\left(\widehat{\lambda}, V\right)$	-326	-71	-140	-111	0.02	0.56	0.11
se(Insurance value) (\$) $= se\left(\widehat{\beta}\right) \times var(v)$	(43)	(12)	(25)	(14)	(0.008)	(0.196)	(0.016)
Transfer value (\$), $E(V)$	1,298	273	641	354	0.13	0.86	0.15
Total value (\$), $E(V) + Cov\left(\widehat{\lambda}, V\right)$	972	202	501	243	0.15	1.42	0.26
Markup, $M = Cov\left(\widehat{\lambda}, V\right)/E(V)$	-0.25	-0.26	-0.22	-0.31	0.15	0.65	0.76
Total value per dollar of transfer value (\$), $1 + M$	0.75	0.74	0.78	0.69	1.15	1.65	1.76

Notes: Table presents sufficient statistic estimates of the value of different types of health insurance coverage and contracts. $\widehat{\beta}$ is the estimate of the coefficient on the ex post value of the coverage in question, V, in a regression of normalized marginal utility on V and controls (equation 8), scaled to be in standard deviation units (so a coefficient of -0.04 means that a one-standard deviation increase in V is associated with a 0.04 decrease in marginal utility. The standard error of this coefficient, shown on the next line, is expressed in the same units. Column (1) reproduces the main results, which summarize the value of full coverage of all health care spending: $V = oop$. Columns (2)–(4) show the value of coverage of three types of health care spending: hospital care ("Hospital"), doctor visits ("Doctor"), and prescriptions ("Rx"), where in each case V is out-of-pocket health spending on that type of care. Columns (5)–(7) show the value of an alternative type of contract that pays a fixed cash indemnity benefit based on either health care consumption or health status. Column (5) shows the value of a contract that paid out \$1 if the household head or spouse experienced at least one hospitalization in the past year and \$0 otherwise. Column (6) shows the value of a contract that paid out \$1 for each day the household head or spouse was hospitalized in the past year. Column (7) shows the value of a contract that paid out \$1 if the household head reported that their health was X or Y ("Bad health") as opposed to X, Y, and Z. The best way to compare the results across columns is to focus on the markup, i.e., the insurance value per dollar of transfer value, $M = Cov\left(\widehat{\lambda}, V\right)/E(V)$. Monetary amounts are in dollars per household per year.