Health Insurance Increases Financial Risk

Lee M. Lockwood*

September 3, 2022

Abstract

The value of health insurance depends on the extent to which it insures all risks, not just health care costs. Using a variety of approaches, I find that for U.S. households, health insurance systematically exacerbates other risks—enough that it increases consumption risk *on net*. The key driver is the implicit health insurance provided by discounts, charity care, and bad debt. This implicit insurance provides significant protection against health care costs, especially when other circumstances are worse. This not only reduces the value of additional insurance against health care costs, it also insures other risks. Health insurance, in displacing this insurance of other risks, exacerbates them. This is an important cost of health insurance and of the unevenness of the safety net.

*University of Virginia and NBER, leelockwood@virginia.edu. I thank, without implicating, David Benson, Karim Chalak, Federico Ciliberto, John Cochrane, Kerem Cosar, Manasi Deshpande, Liran Einav, Amy Finkelstein, Eric French, Leora Friedberg, Tal Gross, Nathaniel Hendren, John Bailey Jones, Daniel Kaliski, Kory Kroft, Tim Layton, Adam Leive, Ethan Lieber, Attila Lindner, Benjamin B. Lockwood, Gideon Magnus, Grace McCormack, Amalia Miller, Casey Mulligan, Matthew Notowidigdo, Edgar Olsen, John Pepper, Mark Shepard, Jonathan Skinner, Isaac Sorkin, Juan Carlos Suárez Serrato, Sandip Sukhtankar, and seminar participants at the AEI, the ASSA meetings, Chicago, Clemson, Columbia, Duke, Johns Hopkins, LSE/IFS, Michigan, the NBER Public Economics meeting, Texas-Austin, USC, Virginia, and Wharton for helpful comments. I thank 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 central component of government policy and a major household asset. In the U.S., government spending on health insurance exceeds \$1.6 trillion per year, and total health insurance benefits exceed \$2.7 trillion per year, over \$21,000 per household (US Centers for Medicare & Medicaid Services, 2019; US Census, 2020). A fundamental motivation for health insurance is insurance: insuring living standards against financial risk. Yet relatively little is known about the insurance effects of health insurance: the extent to which its budget-neutral reallocation of resources across different states of the world targets lower-consumption, higher-marginal utility states.

In this paper, I use a variety of approaches to analyze the insurance effects of health insurance. Each points to the same, robust conclusion: On average for U.S. households, health insurance *increases* financial risk. Although it insures health care costs, it exacerbates other risks enough that it increases consumption risk on net. My estimates suggest that the per-household welfare cost is on the order of \$500 per year, about one-third of the corresponding net resource cost. This means that for many households, health insurance is worth less ex ante than its mean ex post value worth of cash, i.e., it has negative "insurance value."[1]

The key driver is the implicit health insurance provided by discounts, charity care, and bad debt. In the worst, highest-marginal utility states of the world, implicit insurance helps most—so formal insurance helps least. To fix ideas, imagine a household with typical earnings of \$100,000 per year that faces two states of the world in which a spouse suffers a non-fatal heart attack: one where the survivor is able to continue working after a brief, paid leave ("bad") and one where the survivor can no longer work ("worst"). In both states, some costs are not insured, and the household is billed \$10,000. In the bad state, the household pays the \$10,000 bill, leaving it net income of \$90,000. In the worst state, where the loss of the survivor's earnings reduces household income to \$50,000, the household never repays its bill, leaving it net income of \$50,000. Now consider a supplemental health insurance policy that covers all costs. This increases net income by \$10,000 in the bad state but by *\$0* in the worst state; in the worst state it just reduces bad debt. This coverage, in displacing implicit insurance support, increases the gap in net income between the bad and worst states, from \$40k (\$90k − \$50k) to \$50k (\$100k − \$50k). While valuable, this coverage is unlikely to be

[1] This in no way contradicts the financial and health benefits of health insurance coverage (see Finkelstein et al., 2018, for a review). Rather, my findings imply that health insurance coverage, while valuable, is not *as valuable* to the holder ex ante as its mean ex post value worth of cash. Decomposing the ex ante value into two components, the mean ex post value and the budget-neutral reallocation of resources across states of the world, my finding is that health insurance's budget-neutral reallocation of resources across different states of the world—its pure-insurance component—has negative value. Negative insurance value does not imply negative overall value or negative welfare effects; there are many other potential benefits of health insurance. I discuss these and the broader implications of the results in Section 5.4.

as valuable ex ante as its mean ex post value worth of cash, since the cash would help more in the worst state.[2]

Implicit insurance provides significant protection against otherwise-uninsured health care costs, especially in the worst states and even for households in strong financial positions. This reshapes health spending in two ways that transform it from a major *risk creator* into a net *risk absorber*, and thereby transform the insurance effects of health insurance from highly valuable to costly. First, implicit insurance provides considerable protection against catastrophic costs (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 value of such coverage. Second, implicit insurance provides more protection against health care costs when other circumstances are worse, such as when income and assets are low following a job loss or a work-limiting health shock. This provides valuable insurance against other risks—insurance that health insurance undoes. This undoing transforms health insurance from being roughly neutral with respect to other risks to exacerbating them systematically. The exacerbation of other risks dominates the insurance of health care costs because other risks are much less well-insured. This is health insurance in a highly second-best world in which other risks and implicit insurance cause health spending to hedge not only other risks but even overall consumption risk.

I start with a descriptive analysis of health spending, using data from the Medical Expenditure Panel Survey (MEPS) and the Panel Study of Income Dynamics (PSID). This reveals three key determinants of the insurance effects of health insurance. First, health spending risk is limited. Among uninsured households, the standard deviation of out-of-pocket spending is around $3,000, an order of magnitude smaller than that of income. Second, health spending hedges other risks. For example, out-of-pocket spending drops when households experience negative income shocks such as from unemployment, which partially offsets the associated income losses. Third, health spending hedges consumption risk. Out-of-pocket spending is lower when consumption is lower. This suggests that health insurance, which helps more when out-of-pocket spending otherwise would be higher, increases the volatility of consumption across states of the world. Together, these findings suggest that health insurance increases consumption risk because its exacerbation of other risks dominates its insurance of health care costs. Having health insurance means having more insurance against health care costs but less insurance against other risks—and risk as a whole.

All three patterns are driven by implicit insurance. In terms of its effects on out-of-pocket spending, implicit insurance resembles catastrophic health insurance with a *state-contingent*

[2] A similar argument applies to any inframarginal coverage the household might hold. Just substitute $50,000 of uninsured costs for the $10,000 above, and increase the charity care and bad debt accordingly, especially in the worst state.

deductible that is low on average—even for households in strong financial positions—and even lower (more coverage) when circumstances are worse. For example, among uninsured households with health care charges of at least \$20,000, out-of-pocket spending is about \$5,000 on average and even lower when income or wealth is low. This "pre-existing coverage" leaves for health insurance only the residual costs not covered by implicit insurance. Unlike *total* health care costs, which are highly variable, sometimes very large, and exacerbate other risks—a prime candidate for insurance—these *residual* costs are not very variable, seldom large, and hedge other risks—a prime candidate for anti-insurance.

Consider income, a large and risky part of household budgets. For each of many types of households and estimation strategies, income is positively correlated with out-of-pocket spending, both in the cross-section and within households over time. So holding health insurance is, in part, like holding stock in one's employer: It is worth less when income is lower and so increases the welfare cost of income risk. This is true even though total health care costs and, among insured households, gross health insurance benefits are virtually uncorrelated with income. It is only net health insurance benefits, net of displaced support from implicit insurance, that are positively correlated with income. This is because implicit insurance support is negatively correlated with income: Implicit insurance 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 exacerbating it systematically.[3]

I then turn to estimating the insurance value of health insurance. My main approach builds on the sufficient statistics approach from the literature on unemployment insurance, including modeling marginal utility as a decreasing function of observed consumption spending and utilizing panel data specifications that isolate within-household variation over time. Exploiting the long panel nature of the PSID, I estimate the value of coverage from a variety of perspectives, from immediately prior to the coverage, when relatively little risk remains, to "behind the veil of ignorance," when all risk remains. I find that from each perspective and in each of many types of states of the world—insured and uninsured, young and old, high- and low-income—and under a wide range of assumptions, increasing health insurance coverage would have negative insurance value, with the anti-insurance cost growing in the amount of risk that remains to be revealed. My main estimates imply that providing full coverage for one year to uninsured households would have an anti-insurance cost of \$210

[3] An additional force toward health insurance exacerbating income risk is that some types of health care are normal goods, coverage of which is more valuable when income is higher. The exacerbation of income risk is sufficiently strong that health insurance increases the variance of net income (income net of health spending). One income-related risk that health insurance exacerbates is unemployment risk. I estimate that health insurance is worth about \$300 less in years in which a household experiences unemployment. Health insurance is, in part, negative unemployment insurance.

from the perspective of immediately prior to the coverage, $440 from the perspective of ten years prior to the coverage, and $720 from behind the veil. Such coverage, while valuable, would be significantly less valuable to households ex ante than the same mean ex post value worth of cash: 20–70% less valuable, depending on the perspective.[4]

Implicit insurance transforms health insurance targeting so powerfully as to cause health insurance to exacerbate 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, its transfers *within* hospitalization states have larger anti-insurance costs. Hospitalization is associated with significant earnings losses that are much greater in some states than in others (Dobkin et al., 2018). Implicit insurance provides more protection against health care costs when earnings losses are greater. Health insurance displaces this earnings insurance, at a welfare cost that exceeds the benefit of its insurance of the associated health care costs. Were it not for implicit insurance, by contrast, a variety of evidence suggests that coverage of hospitalization-related health care costs would provide considerable insurance value.

To better understand the underlying mechanisms, I construct a simple model guided by the empirical findings. It is an otherwise-standard model of health spending risk with two added features based on the empirical evidence: other risks and implicit insurance. The model matches well the key empirical patterns, including those not targeted directly such as the sufficient statistic estimates. Here, too, the conclusion that health insurance increases financial risk is extremely robust. Health insurance would increase financial risk even if other risks and implicit insurance protection were significantly smaller than they appear to be empirically. Counterfactual analyses suggest that if it were not for implicit insurance and other risks, health insurance would provide highly valuable insurance, with a baseline insurance value of $3,100 per year. But the greater protection from implicit insurance when other circumstances are worse causes health insurance to exacerbate other risks at a welfare cost of $560. And the protection from implicit insurance against catastrophic costs transforms health insurance protection against health care costs from quite valuable (welfare benefit of $3,100) to a matter of near indifference (benefit of $70). So it is not that insurance against health care costs is of little value; it is of considerable value. It is that *marginal* insurance, on top of that provided by implicit insurance, is of little insurance value.

[4]The associated welfare costs are large. An annual cost of $500 per household would translate into an aggregate cost among U.S. households of $64 billion. While this is small relative to gross government spending on major health insurance programs (e.g., spending on Medicaid exceeds $550 billion), it is large in absolute terms, comparable to government spending on food stamps ($70 billion), the Earned Income Tax Credit ($70 billion), Supplemental Security Income ($55 billion), and cash welfare ($30 billion) (all figures from 2015; see US Department of Health and Human Services, 2015, 2016; US Department of Agriculture, 2016; US Internal Revenue Service, 2016; US Social Security Administration, 2016).

That health insurance increases the financial risk of "typical" households in the U.S. does not mean that it does so for all households in the U.S. or for households in other countries, or that it would continue to do so following large changes in the economy. My analysis highlights the importance of key determinants that are known to vary both across and within countries. Neither does the fact that health insurance increases financial risk for many mean that it reduces welfare or that households are making mistakes by holding it. There are a variety of other potential benefits of health insurance, including reducing private and external costs associated with implicit insurance. That health insurance has negative insurance value just means that one important component of its overall welfare effect is not the benefit previously thought but a cost. I discuss these and other issues of interpretation in Section 5.4.

The main contribution of this paper is to quantify, using a variety of approaches, the insurance value of health insurance and the importance of underlying mechanisms. My findings build on and help reconcile two strands of related literature. The first seeks to quantify the insurance value of health insurance.[5] To the best of my knowledge, all previous studies have concluded that health insurance has positive insurance value. In fact, a common view is that its insurance value is not merely positive but inefficiently high (because of over-insurance due to subsidies; e.g., Feldstein, 1973; Pauly, 1986). My investigation of mechanisms reveals a key reason previous studies have reached this conclusion: the omission of other risks or implicit insurance. That seemingly-innocuous omission causes the analysis to miss what surprisingly turns out to be health insurance's most consequential effect on financial risk: its exacerbation of other risks.

The second strand of related literature seeks to quantify the overall value of health insurance.[6] A key finding is that willingness to pay is often quite low, similar to or even below the mean *net* benefit (e.g., Finkelstein et al., 2019a,b). This is puzzling if health insurance has positive insurance value but accords well with my finding of negative insurance value.[7] My findings complement and extend earlier research by showing the crucial role of other risks, and their interaction with implicit insurance, in not only reducing the value of health insurance but even potentially making it smaller than the mean net benefit. My findings also reverse certain conclusions about the welfare effects of different types of coverage. Implicit insurance, despite its well-known disadvantages, has a less-appreciated advantage over formal

[5] 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), French and Jones (2011), Kowalski (2015), and Finkelstein et al. (2019a).

[6] See, for example, French and Jones (2011), Dague (2014), Gallen (2015), Hackmann et al. (2015), Finkelstein et al. (2019a), Finkelstein et al. (2019b), and Mulligan (2021).

[7] The main previous explanation for low overall value is that implicit insurance reduces health insurance's mean net benefit and the value of its protection against health care costs (e.g., Mahoney, 2015; Finkelstein et al., 2019a). As is recognized, this can explain why the overall value would be not much above the mean net benefit but not why it would be below.

insurance: It provides valuable insurance against other risks. This advantage is so great that the displacement of implicit insurance by formal health insurance increases the volatility of consumption. Indemnity insurance that paid fixed benefits based on health diagnoses would have not only lower moral hazard costs but higher, not lower, insurance value than typical contracts. More generally, my findings show how means-tested support can transform the targeting of other sources of support. This could be important in other contexts as well.

2 Insurance Value: Definition and Sufficient Statistic

Ex post value.— The ex post equivalent variation V of an arbitrary change in the ex post budget constraint, measured in terms of consumption, is defined implicitly by

$$u\left(c_0 + V, a_0; \theta\right) = u\left(c_1, a_1; \theta\right), \tag{1}$$

where c is consumption, a is a vector of "all other goods," θ is the state of the world, $u(c, a; \theta)$ is ex post utility, (c_0, a_0) is the allocation under the original constraint, and (c_1, a_1) is the allocation under the new constraint. V is the increase in consumption under the original constraint that would make the individual exactly as well off as they would be under the new constraint.[8]

Insurance value.— The ex ante value EAV, measured in terms of consumption in all states of the world, of an arbitrary change in ex post constraints is defined implicitly by

$$E\left[u\left(c_0 + EAV, a_0; \theta\right)\right] = E\left[u\left(c_1, a_1; \theta\right)\right], \tag{2}$$

where the expectations are taken over the distribution of possible states of the world, $\theta \sim F(\theta)$. EAV is the increment to consumption in all states such that the individual is exactly as well off ex ante as they would be under the new constraints. "Insurance value," what Finkelstein et al. (2019a) call "pure-insurance value," is defined as the amount by which the ex ante value exceeds the mean ex post value:

$$EAV = E(V) + \text{"Insurance value."} \tag{3}$$

[8]Preferences and constraints may depend on the state of the world θ. The von Neumann-Morgenstern utility function $u(c, a; \theta)$ could be expected continuation utility in a dynamic setting with multiple periods, in which case a includes consumption in future periods and 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 for expositional simplicity, but the theory applies to dynamic settings as well. I measure value in terms of consumption rather than income to avoid the measure itself having insurance effects, since income is implicitly taxed by implicit health insurance more in some states than others.

This insurance value is the "pure" insurance value surplus: the ex ante value of the differential targeting of certain states of the world relative to others, holding the mean value fixed. It answers the question: Ex ante, how much more valuable is the change in constraints than the same mean ex post value worth of cash?

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

$$\underbrace{EAV}_{\text{Ex ante value}} \approx \frac{E\left(\lambda \times V\right)}{E\left(\lambda\right)} = \underbrace{E\left(V\right)}_{\text{Mean ex post value}} + \underbrace{Cov\left(\widehat{\lambda},\ V\right)}_{\text{“Insurance value”}}, \tag{4}$$

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

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

is a first order approximation to the insurance value of the change. 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 marginal utility is higher. 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 ex post value covaries negatively with marginal utility, the change in constraints has negative insurance value; it is worth less ex ante than its mean ex post value.[10]

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, and many risks of varying persistence. In a broad class of models, any effects that

[9] This approximation follows from plugging equation (1) into equation (2) and taking first order approximations around the allocation under the original constraints (see Appendix A.1). The insurance value covariance, closely related to that in Finkelstein et al. (2019a) and analogous to a redistribution value covariance in optimal taxation, generalizes the insurance value part of the Baily-Chetty analysis of optimal social insurance (Baily, 1978; Chetty, 2006) to situations in which the ex post value of the change in constraints can take more than two different values (see Appendix A.2).

[10] As is clear from equations (3) and (4), negative insurance value does not imply negative ex ante value, just ex ante value that is smaller than the mean ex post value. If the ex post value is non-negative in all states, $V \geqslant 0$, the ex ante value is necessarily non-negative as well, $EAV \geqslant 0$.

such factors or others might have on the insurance value of a change in constraints manifest themselves through this covariance.[11] 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. 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.[12]

3 Data, Institutions, and Empirical Approach

Data.— *PSID.—* The Panel Study of Income Dynamics (PSID) has many advantages for analyzing the insurance effects of health insurance. It has measures of out-of-pocket spending and health insurance. It has rich measures of income and non-health consumption. And its rich demographic measures and long panel structure are useful for isolating varying amounts of risk that remains to be revealed from risk that has already been revealed. I use data on households interviewed in at least one of the eleven waves from 1999–2019 inclusive. These waves occur every second year. The resulting sample has 85,769 household-wave observations. My baseline measure of non-health consumption is annual expenditure on food (including the value of food stamps received), housing, transportation, clothing, travel, recreation, education, and child care. Standard errors are always clustered by household.

MEPS.— The Medical Expenditure Panel Survey (MEPS) has rich, high-quality information on health care consumption and expenditures, as well as information on household demographics and income. This is especially useful for investigating the role of implicit insurance in shaping health insurance targeting. 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 waves from 1996–2018, which occur annually. The resulting sample has 268,235 family-year observations. Total health care costs are defined as follows. For households with health insurance, total costs are total annual payments, including from the insurer and the household. For households without health insurance, total costs are annual charges scaled by 0.60, the payments-charge ratio among non-elderly households with health insurance. I

[11] 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. Or if persistent health shocks tighten constraints not only by increasing current and future health spending but also by decreasing current or future income, that would tend to increase the insurance value of health insurance. Or if individuals forgo or postpone care when times are tight, that would tend to decrease the insurance value of health insurance by reducing its ex post value in high-marginal utility states.

[12] With optimization, a change in constraints can be valued to first order with knowledge of the constraints and the status quo allocation; behavioral responses have no first-order effect on utility by the envelope theorem.

follow Mahoney (2015) in scaling by this ratio to reflect typical discounts relative to charges.

In both datasets, my baseline measures of out-of-pocket spending are inclusive measures of the types of services typically covered by health insurance, including hospital care, doctor visits, and prescriptions. My measures of income include income from all sources, including from social insurance and means-tested programs, so that they reflect the net risk in income accounting for all sources of income risk and insurance. My measures of hospitalizations are indicators of whether a member of the household was a patient in a hospital overnight or longer at any point in the prior year *and* there is no child in the household young enough that the hospital stay may have been related to childbirth. The aim is to focus on hospitalizations driven by health shocks, as in Dobkin et al. (2018). All monetary variables are converted to real 2020 dollars using the CPI-U-RS. Throughout, I use household weights to ensure that the estimates reflect the experiences of the U.S. population. Appendix B contains details of variable construction, and Appendix Tables A1 and A2 show summary statistics of the main estimation samples.

Institutions.— *Health insurance.—* Throughout, I focus on health insurance benefits, abstracting from how they are funded. I focus on contracts that cover a share of health care costs, which describes the vast majority of contracts in use in the U.S. (Cutler, 2002). The fundamental effect of such contracts is to reduce what the insured is billed for a given amount of health care. While this can lead to over-consumption of health care (moral hazard), such contracts are thought to provide better risk protection than other types.[13]

Implicit health insurance.— Discounts, charity care, and bad debt provide significant protection against otherwise-uninsured health care costs. For example, individuals without formal health insurance pay only about one-fifth to one-third of their health care costs out of pocket.[14] Unlike formal safety net programs, such implicit insurance, while greater for individuals in worse circumstances, is considerable for individuals in strong financial positions as well (see Section 4). In terms of its effect on the value of health insurance, the key feature of implicit insurance is that it is a "secondary payer": It reduces the private cost of

[13] For example, indemnity insurance that paid fixed benefits based on health diagnoses would leave within-diagnosis risk in health care costs uninsured (Zeckhauser, 1970).

[14] See Hadley et al. (2008), Coughlin et al. (2014), and Finkelstein et al. (2019a). Charity care arises from not only charitable motives but legal obligations. For instance, to qualify for certain tax exemptions, nonprofit hospitals (roughly 70% of all hospitals), must provide a "community benefit" in the form of charity care or medical research and teaching (Gov. Account. Off., 2008). Bad debt arises, in part, from the legal obligation that hospitals must provide emergency medical care on credit even when repayment is unlikely. In practice, most hospitals provide non-emergency care on credit as well (see Mahoney, 2015, and the references therein). Much of the care provided on credit is never paid for. For example, uninsured individuals repay only about 10-20% of what they are billed (LeCuyer and Singhal, 2007). Medical debt often is defaulted on implicitly rather than discharged explicitly through bankruptcy. For example, whereas unpaid medical bills affect nearly one-fifth of consumers' credit reports and comprise a majority of all collections lines (CFPB, 2014), in a given year less than 1% of Americans file for personal bankruptcy.

otherwise-uninsured health care costs. Health insurance necessarily displaces such support. This displacement implicitly taxes health insurance, reducing its ex post value by the value of the displaced support. If such implicit taxation is greater in some states of the world than others, it can transform health insurance targeting. This paper analyzes the insurance effects of formal health insurance accounting for its displacement of implicit insurance.[15]

Empirical approach.— *Ideal experiment.—* Many of my analyses seek to characterize the effects of (hypothetical) health insurance coverage expansions: increases in coverage from status quo levels.[16] As discussed in Section 2, in a broad class of models, the insurance value to first order of such a coverage expansion is the covariance, across states of the world in the status quo, between normalized marginal utility and the ex post value of the coverage. This sufficient statistic, a straightforward generalization of the insurance value part of the well-established Baily-Chetty approach, depends only on marginal utility and the ex post value in the status quo. It does not depend on any other outcomes, including counterfactual outcomes away from the status quo (e.g., with the expanded coverage) or causal effects of the contemplated change in coverage. The main empirical challenge is the ever-present challenge for all analyses of risk: modeling the (unobservable) distribution of potential states of the world. The ideal (infeasible) experiment would be to "re-run" an individual's life many times to observe the full distribution of possible states of the world they might experience.[17]

Effects and value of health insurance.— The main effects of health insurance are reduced medical debt, improved health, and reduced out-of-pocket spending (Finkelstein et al., 2018). Reduced medical debt, as Finkelstein et al. (2018) discuss, has clearer benefits to creditors than to the individual.[18] Improved health, from moral hazard effects on health care con-

[15]Specifically, I analyze the insurance effects of having health insurance coverage relative to having cash in the same states. Medicaid, the means-tested public health insurance program, has insurance effects not only from its coverage of health care costs but also from its means tests and individuals' take-up decisions. My analysis applies to the first aspect. The second is an interesting topic for future research.

[16]While focusing on coverage expansions has benefits in terms of tractability, transparency, and policy relevance, a cost is that it does not speak directly to the inframarginal coverage that insured households hold in the status quo. To shed light on this interesting, though less policy-relevant, issue, I investigate heterogeneity in the effects of coverage expansions on different types of households and in the effects of different types and levels of coverage, and I combine empirical evidence on heterogeneity and mechanisms with economic logic and modeling. The results suggest that inframarginal coverage has if anything an even higher anti-insurance cost per dollar than marginal coverage in the same states, likely due to greater crowd out of implicit insurance (e.g., see footnotes 35 and 45).

[17]This experiment varies the state of the world, not health insurance. Exogenous variation in health insurance is not only unnecessary, using it to estimate insurance value requires extensive knowledge of the utility function and the causal effects of health insurance (Finkelstein et al., 2019a). See Appendices A.2–A.4 for details about the close relationships between my approach, the widely-used Baily-Chetty approach, and Finkelstein et al.'s (2019a) "optimization approach," and for an overview of an alternative approach based on exogenous variation in health insurance.

[18]Evidence of benefits to individuals from reduced medical debt "remains limited" (Finkelstein et al., 2018, p. 270). The vast majority of unpaid medical bills sent to collections remain unpaid (Avery et al., 2003), and hospital admissions have little effect on credit scores despite increasing medical debt (Dobkin et

sumption, while potentially of considerable ex ante value, is unlikely to significantly increase insurance value and might even decrease it.[19] Perhaps in part from such considerations, the standard approach focuses on reduced out-of-pocket spending alone. This captures the main financial effect of health insurance and, under standard assumptions, is a first order approximation to its ex post value.[20] Hence, the insurance effects of health insurance depend crucially on the distribution of out-of-pocket spending it would cover. This idea is the basis of my analyses (and of virtually all other analyses of the insurance effects of health insurance of which I am aware). Still, where relevant, I test robustness to large private benefits of improved health and reduced medical debt.

Risk and regression specifications.— I follow the common practice of using variation within households over time and in the cross-section to proxy for risk, using a variety of panel data specifications and control variables to isolate varying amounts of risk that remains to be revealed from risk that has already been revealed. I investigate the insurance effects of health insurance from three main perspectives: immediately before the coverage begins ("short run"), ten years before the coverage begins ("medium run"), and "from behind the veil of ignorance" ("long run"). As Hendren (2020) emphasizes, the value of insurance depends critically on what risk has already been revealed when the value is assessed, so analyses based on perspectives when some risk has already been revealed can be misleading about the full ex ante value of insurance. For example, from the perspective of immediately before a spell of coverage begins, an individual already knows their health history up to that point. Neither health insurance nor anything else can insure the already-realized risk of having experienced that history as opposed to others. But from earlier perspectives, the same future coverage could insure not only the risk that remains at the later perspective but also the additional risk of which "later perspective" one will experience.[21]

---

al., 2018). More generally, the extensive reliance on implicit insurance, including by individuals in otherwise-strong financial positions, suggests that when times are tight, many individuals perceive relying on implicit insurance to be less costly than cutting back on consumption to pay their bills in full.

[19] The ex post value of moral hazard effects is second order for optimizing households, and is lower, other things equal, when the marginal utility of consumption is higher. Of course, second order does not imply small. For example, if an individual might benefit from an advanced cancer treatment that is unavailable or unaffordable without health insurance, the ex post value of insurance could be quite high even absent any reduction in medical debt or out-of-pocket spending. For this reason, where relevant I test robustness to large private values of moral hazard in key states of the world.

[20] If the household optimizes and there are no first-order effects on its cost of relying on implicit insurance, the ex post value of health insurance to first order is the "mechanical" reduction in out-of-pocket spending it would cause if behavior were held fixed. This follows from the usual envelope theorem logic that to first order, the value to an optimizing household of a change in constraints is the associated "mechanical effect"—the reduction in net expenditure that would occur if not for behavioral responses—since re-optimization gains are second order. While my sufficient statistic analysis assumes that households optimize, my other analyses, including of the effects of health insurance on net income and consumption risk, do not.

[21] Intuitively, from later perspectives where more risk has been revealed, the set of "lifetime states of the world" one might yet experience is a subset of those one might have experienced from earlier perspectives, as one's realized experience rules out certain states. While the insurance value of insuring risk as a whole tends to decrease with the amount of risk already revealed (Hendren, 2020), the insurance value of health insurance

The short run perspective of immediately before the coverage begins is based on regressions of the within-household change in log consumption or log income from one PSID wave to the next on the within-household change in the log of one plus the ex post value of the coverage, plus year dummies and a cubic in age. The medium run perspective of ten years before the coverage begins is based on regressions that are identical except that they use within-household changes in the key variables from one wave to the fifth wave after that, ten years later. The long run perspective of someone behind the veil is based on regressions of log consumption or log income on the log of one plus the ex post value of the coverage, plus year dummies, a cubic in age, and a quadratic in household size.[22] In a few instances, I consider the perspective of someone who knows their education level but nothing else. This perspective, which is between the medium and long run perspectives, aims to capture the risk within but not across different earning ability groups. The corresponding regressions are the same as the long run regressions except that they add education category dummies to the controls. Finally, I occasionally use household fixed effects regressions as a simple way of isolating within-household variation. These aim to capture risk between the short and medium run perspectives. I also test robustness to many alternatives.

4 Descriptive Analysis of Health Spending

Fact 1: Health spending risk is limited.

Figure 1a shows a histogram and estimated kernel density of the distribution of annual out-of-pocket health spending among non-elderly uninsured households, and Appendix Table A3 shows associated statistics. Health spending is small on average (average of $1,060) and only modestly variable (standard deviation of $2,720 and 99th percentile of $11,460). By comparison, income and consumption are far more variable. For example, among non-elderly households in the PSID, the *within-household* standard deviations of annual income and consumption are $34,910 and $15,620, respectively. Even if health spending were purely a risk creator, it would be a small risk relative to others that households face.

Health spending risk is limited mainly because of implicit insurance. Among non-elderly households, total health care costs are large on average (average of $9,610) and highly variable

could in principle increase or decrease due to its opposing pro- and anti-insurance effects, depending on the relative persistence of the underlying risks.

[22]This follows the common "steady state" assumption that the cross-sectional distribution approximates the distribution of states of the world faced by someone behind the veil. The controls for time, age, and household size are not meant to exclude risk but to reduce the impact of any aggregate risk or misspecification of price indices (time dummies) and of any misspecification of household equivalence scales or age-dependent utility (age and household size controls).

Figure 1: Health spending risk is limited

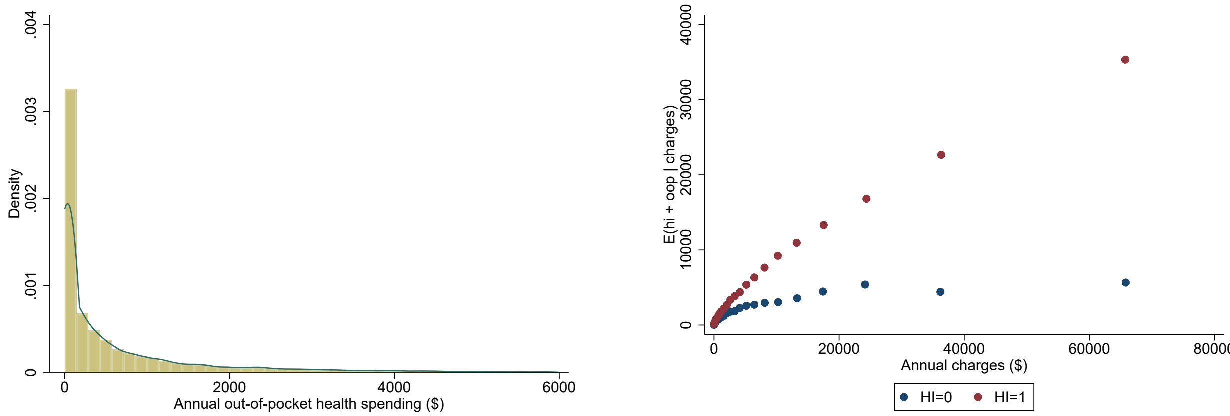

(a) OOP spending among uninsured HHs (b) Implicit catastrophic health insurance

Notes: Left panel: Histogram and estimated kernel density function of annual out-of-pocket health spending among non-elderly uninsured households. The average is $1,060, the standard deviation is $2,720, and the 99th percentile is $11,460. This figure cuts off at $6,000 for legibility. Right panel: Conditional mean of total combined 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 health insurance (higher, red dots) and without health insurance (lower, blue dots). This is a binned scatter plot. This figure excludes households with charges in excess of $100,000 for legibility. Both panels are based on MEPS data and include all outliers, without any trimming or winsorizing.

(standard deviation of $23,030), albeit significantly less variable than income (see Appendix Table A3).[23] It is only net health care costs, net of implicit insurance support, that are small on average and not that variable. Figure 1b, analogous to Figure 1A in Mahoney (2015), shows a non-parametric estimate of the conditional mean of total combined payments by health insurers (health insurance benefits) and households (out-of-pocket spending) as a function of charges, a rough measure of health care utilization. At low charges, total payments 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 $5,000, even among households with tens of thousands of dollars of charges. The difference, presumably but in keeping with direct measures of implicit insurance support, arises from greater reliance on implicit insurance among uninsured households. The nominally uninsured, though lacking formal health insurance, have substantial implicit health insurance from discounts, charity care, and bad debt. This implicit insurance resembles catastrophic coverage above a modest deductible (Mahoney, 2015).

[23] Even among elderly households, the within-household standard deviation of income exceeds the overall standard deviation of total health care costs ($30,650 versus $24,590). Although income from Social Security and defined benefit pensions is fairly stable, the elderly face significant risk in earnings (among those still working) and asset income (see, e.g., Blundell et al., 2020). Beyond income, the elderly face major risks in long-term care costs, the prices of housing and other assets, and household composition.

Although implicit insurance provides more protection to households in worse financial positions, it provides considerable protection to households in strong financial positions as well. Appendix Figure A1 shows that among households with at least $20,000 of charges, mean out-of-pocket spending among uninsured households with a college degree is $7,210, not much above that of uninsured households with less than a high school degree ($4,430) and far below total payments among insured households ($33,870). Similarly, Mahoney (2015) finds that among households in the top ventile of their respective charges distributions, households with financial costs of bankruptcy of at least $50,000 have mean out-of-pocket spending of about $7,000, not much above that of households with lower costs of bankruptcy (about $3,500) and far below total payments among insured households (about $28,000) (see Figure 1B in Mahoney, 2015).[24] Appendix Figure A2 shows the shares of different groups of households that report having had problems paying or having been unable to pay their medical bills in the past 12 months. Among uninsured households with a college degree, this share is 15%. Even among households with health insurance, this share is 9%. Whereas formal safety net programs restrict eligibility to individuals of limited means, implicit insurance helps a much broader set of people, including anyone who, at least in some states of the world, would receive a discount on their care or would not pay a large medical bill in full.

Implications.— Health insurance protection against health care cost risk is likely of limited insurance value, i.e., its ex ante value is unlikely to be much greater than its mean ex post value. Implicit insurance provides considerable protection against otherwise-uninsured health care costs, which to a rough approximation is the type of catastrophic coverage recommended by optimal insurance theory. This leaves for health insurance mainly the non-catastrophic costs that optimal insurance theory recommends *not* covering (since the insurance value would be outweighed by administrative and moral hazard costs). So while *total* insurance against health care costs, including from implicit insurance, likely generates significant insurance value, additional, *marginal* insurance on top of that provided by implicit insurance is unlikely to generate much insurance value.

Fact 2: Health spending hedges other risks.

The most important risk for many households is income risk. Figure 2a shows, for non-elderly households, a non-parametric estimate of the conditional mean of income as a function of out-of-pocket health spending, controlling for household fixed effects, year dummies, and a

[24] While $7,000 is a lot to spend on health care, even such an extreme realization of out-of-pocket spending (above the 95th percentile) is small in comparison to these households' tens of thousands of dollars worth of seizable assets and to many other risks they face, such as income losses from unemployment. It is on the order of the *average* cost of common home repair projects, such as to HVAC systems ($4,950), septic tanks ($4,530), and roofs ($8,370) (statistics from the American Housing Survey, 2019).

Figure 2: Health spending hedges other risks

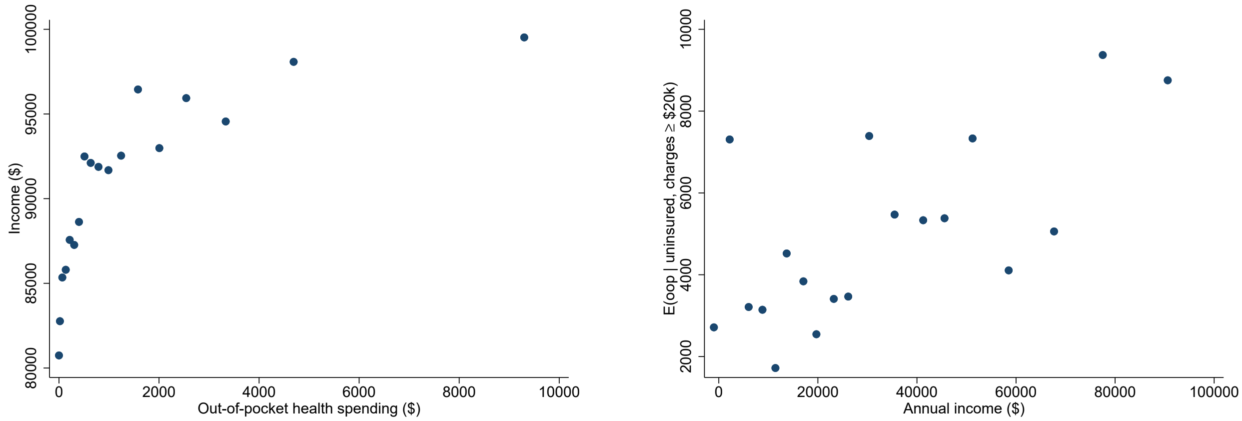

(a) Health spending hedges income risk (b) Implicit HI support decreases in income

Notes: Left panel: Conditional mean of income as a function of out-of-pocket health spending among non-elderly households in the PSID, controlling for household fixed effects, year dummies, and a cubic in age. This is a binned scatter plot using the methods of Cattaneo et al. (2019). Right panel: Conditional mean of out-of-pocket health spending as a function of income among uninsured households in the MEPS with annual health care charges of at least $20,000. In this sample, average out-of-pocket spending is $5,210, average charges are $63,960, and the conditional mean of charges is decreasing in income, so the greater out-of-pocket spending among higher-income households in the figure is not due to higher charges. This is a binned scatter plot. This panel uses raw variables, including all outliers without any winsorizing or trimming. For better legibility, this figure excludes households with income above $100,000 (the 80th percentile of the income distribution among uninsured households with charges of at least $20,000).

cubic in age. Associated heterogeneity and robustness results are in Appendix Tables A4 and A5. For each type of state and each perspective, health spending and income covary positively (though only weakly for the elderly in the short and medium runs), despite that health insurance coverage and generosity also covary positively with income. Health spending tends to be lower when income is lower and higher when income is higher, partially offsetting income shocks. Health spending is at least in part a risk absorber: It hedges income risk.

That health spending hedges income risk is striking given that health risk is a force toward health spending exacerbating income risk. Health shocks increase health spending and decrease income, which is a force toward health spending being higher when income is lower, which would exacerbate income risk. So the positive relationship between health spending and income must reflect stronger countervailing forces. The key one is implicit insurance. Total health care costs are if anything slightly *negatively* correlated with income (baseline correlation of -0.02; see Appendix Table A6). It is only net health care costs, net of implicit insurance support, that are positively correlated with income and so hedge income risk.[25]

[25] Another force toward health spending hedging income and other risks is that certain types of health care are normal goods (see Appendix C and Acemoglu et al., 2013; Gross et al., 2020). This is in keeping with Grossman (1972)-type models of optimal investment in health capital and with models of optimal investment in durable goods more generally (Browning and Crossley, 2009).

Figure 2b shows a non-parametric estimate of the conditional mean of out-of-pocket health spending as a function of income among uninsured households with annual charges of at least $20,000. This figure reveals two key findings. First, implicit insurance provides significant protection across all income levels: Average out-of-pocket spending is far below charges across all income levels, including the highest.[26] Second, implicit insurance 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 spending and income is presumably driven by lower-income households receiving more support from implicit insurance, not differences in total health care costs. Similarly, Mahoney (2015) finds that for a given level of charges, out-of-pocket spending is positively related to (seizable) net assets. Implicit insurance helps more when circumstances are worse. In this way, implicit insurance is not standard catastrophic insurance that "covers" costs above a fixed deductible but more like special catastrophic insurance with a state-contingent deductible that is lower (greater coverage) when circumstances are worse. As a result, it implicitly insures risk in income, assets, and non-health care circumstances more generally.

Implications.— Health insurance exacerbates other risks: Having health insurance means having more exposure to other risks. For example, that out-of-pocket spending hedges income risk means that health insurance, by reducing out-of-pocket spending, exacerbates income risk. In this way, holding health insurance is like holding stock in one's employer: It is worth less when income is lower and so increases the welfare cost of income risk. The implicit health insurance provided by discounts, charity care, and bad debt provides valuable insurance against a wide variety of risks—insurance that health insurance undoes. This undoing transforms health insurance from being roughly neutral with respect to other risks to exacerbating them systematically.

Case study: Unemployment.— Health insurance exacerbates unemployment risk. Out-of-pocket spending drops when households experience unemployment (see Appendix Table A7), despite health insurance coverage dropping as well. Health spending therefore hedges unemployment risk, partially offsetting the effects of the associated income losses. Health insurance, by decreasing health spending, displaces this hedge. So health insurance is, in part, negative unemployment insurance. My main estimate is that for non-elderly households health insurance effectively undoes about $300 worth of unemployment insurance benefits per unemployment spell.

[26] Of course, high-income households without health insurance are an unusual population. But that they receive significant protection from implicit insurance is consistent with other evidence that even households in strong financial positions enjoy such protection. One manifestation of this is the finding in the previous paragraph: Implicit insurance is pivotal in causing health spending to hedge income risk for all kinds of households, including those with health insurance and high income.

Fact 3: Health spending hedges consumption risk.

Figure 3: Health spending hedges consumption risk

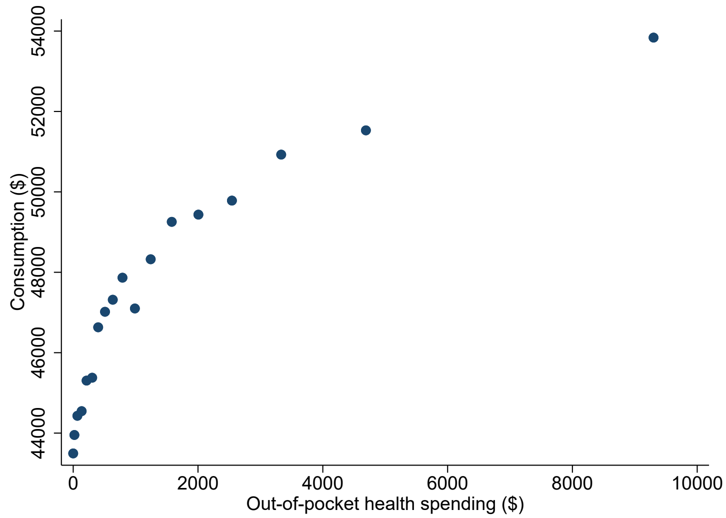

Notes: Conditional mean of consumption as a function of out-of-pocket health spending among non-elderly households in the PSID, controlling for household fixed effects, year dummies, and a cubic in age. This is a binned scatter plot using the methods of Cattaneo et al. (2019).

Figure 3 shows, for non-elderly households, a non-parametric estimate of the conditional mean of non-health consumption as a function of out-of-pocket health spending, controlling for household fixed effects, year dummies, and a cubic in age. Associated heterogeneity and robustness results are in Appendix Tables A8 and A9. For each type of state and each perspective, health spending and consumption covary positively. Health spending tends to be lower when consumption is lower and higher when consumption is higher, thereby mitigating consumption shocks. This suggests that if health spending were to disappear, consumption would increase least in the states of the world in which consumption is lowest and the volatility of consumption would increase. In a wide variety of models, consumption is a revealed-preference measure of the overall tightness of the constraint, so that health spending decreases consumption risk suggests that health spending is more risk absorber than risk creator *on net*.

How could health spending decrease consumption risk? Consider the variance of net income,

$$Var(y - oop) = Var(y) + \left[\underbrace{Var(oop)}_{\text{“Partial effect”}} - \underbrace{2Cov(y, oop)}_{\text{“Portfolio effect”}} \right]. \tag{6}$$

The term in brackets is the total effect of health spending. The “partial effect” reflects that, other things equal, greater health spending reduces net income. This is a force toward health spending increasing net income risk. The “portfolio effect,” from the interaction with income risk, could increase or decrease net income risk depending on the sign of $Cov(y, oop)$. As discussed in Fact 2, health spending hedges income risk: $Cov(y, oop) > 0$. This hedge is so strong, in fact, that for most combinations of states and perspectives, including each type of state from the long run perspective, the portfolio effect dominates the partial effect and health spending reduces the variance of net income: $2Cov(y, oop) > Var(oop)$, which is equivalent to the slope of the regression of income on health spending exceeding one-half, $\beta_{y|oop} \equiv \frac{Cov(y,oop)}{Var(oop)} > 1/2$. In non-elderly states, for example, \$1 higher health spending is associated with income that is higher by \$1.03 from the short run perspective, \$3.04 from the medium run perspective, and \$8.38 from the long run perspective (and so with net income that is higher by \$0.03, \$2.04, and \$7.38, respectively). Beyond net income, the results from Fact 3 suggest that for all combinations of states and perspectives, health spending decreases the volatility of consumption itself.[27] In second-best contexts in which people face uninsured risks beyond health care costs, the net effect of exposure to health care costs reflects not only its partial effect but also its interaction with other risks. In the second best, greater exposure to one risk does not imply greater exposure to risk on net.

The valuable hedge of other risks dominates the cost of health care cost risk largely because other risks are far greater than health care cost risk, increasingly so from earlier perspectives where more risk remains. Income alone has a standard deviation on the order of 20–30 times that of health spending (as shown in Appendix Table A3).[28] This connects the three main findings of the analysis of health spending risk. Health spending decreases risk on net (Fact 3) because its hedge of other risks (Fact 2) dominates its partial effect of being risky itself, because health spending risk is limited relative to other risks (Fact 1). As one rough measure of magnitudes, simple calculations suggest that eliminating health spending would increase the within-household standard deviation of consumption among non-elderly households by roughly twice as much as eliminating unemployment insurance would and the overall standard deviation by about four times as much (see Appendix Figure A3).

Implications.— Health insurance increases consumption risk. It is likely worth less ex ante than its mean ex post value. Its insurance of health care costs is dominated by its exacerbation of other risks. To a rough approximation, implicit insurance is catastrophic

[27] If the marginal propensity to consume were $\alpha > 0$ in all states, then other things equal the variance of consumption if health spending were eliminated would be $Var(c + \alpha oop) = Var(c) + \alpha^2 Var(oop) - 2\alpha Cov(c, oop)$, where c and oop are their values in the status quo. Eliminating health spending would increase the variance of consumption if $2\alpha Cov(c, oop) > \alpha^2 Var(oop)$, i.e., if $\beta_{c|oop} > -\alpha/2$. All estimates of $\beta_{c|oop}$ are positive and significant and so well above $-\alpha/2$ for all $\alpha > 0$. In non-elderly states, for example, \$1 higher health spending is associated with consumption that is higher by 32 cents from the short run perspective, 95 cents from the medium run perspective, and \$2.36 from the long run perspective.

[28] The ratio of the portfolio effect to the partial effect is proportional to the relative magnitude of the risks: $\frac{2Cov(y,oop)}{Var(oop)} = 2Corr(y, oop) \times \frac{Sd(y)}{Sd(oop)}$. Not only do many non-health shocks have large and long-lasting effects on income and assets (e.g, job loss, wage and asset price shocks), even many health shocks have larger and longer-lasting effects on income than on health care costs (e.g., see Dobkin et al., 2018, on hospitalization).

insurance with a state-contingent deductible; it helps more when health care costs are higher and non-health care circumstances are worse. This reshapes health spending in two ways that transform it from a major risk creator to a net risk absorber. First, by providing significant protection against catastrophic costs, it greatly reduces health care cost risk. Second, by providing greater protection against health care costs when other circumstances are worse, it implicitly insures other risks. The former steals health insurance's thunder from insuring health care costs. The latter causes health insurance to exacerbate other risks. Other things equal, more health insurance means less insurance against other risks—and risk as a whole.

Case study: Hospitalization.— A priori, one might expect hospitalization risk to be a major contributor to the value of health insurance. Hospitalization is associated with not only high health care costs but also substantial income losses (Dobkin et al., 2018), so transfers into the average hospitalization state potentially could provide highly-valuable insurance. But the associated income losses are highly variable (Dobkin et al., 2018), and health insurance helps more in some hospitalization states than others. Appendix Table A10 shows results related to health insurance targeting of three sets of states: non-hospitalization, hospitalization with better income realizations ("better"), and hospitalization with worse income realizations ("worse"), where worse is defined as being in the bottom quartile among hospitalization states.[29] Among uninsured households, hospitalization is associated with a $1,230 greater change in out-of-pocket spending from the previous wave and a $470 smaller change in consumption. Health insurance would therefore increase net income modestly in hospitalization states relative to non-hospitalization states on average and thereby reduce the small consumption gap between them. But health insurance would not help equally in all hospitalization states. The estimates suggest that full coverage would help almost $1,500 more in "better" than "worse" states, which would increase the already-large consumption gap between them ($7,420). So the overall effect of health insurance on hospitalization-related risk reflects two opposing factors: transfers from non-hospitalization to hospitalization states insure the associated health care costs, but transfers within hospitalization states exacerbate the associated income risk. The net effect is not only unlikely to be as valuable as might have been expected a priori, it could even have negative insurance value.

[29]These are the results of regressions of the first differences (short run) and levels (long run) of out-of-pocket spending, consumption, and income on indicators for different types of hospitalization states and controls. I focus on the first-differences results for uninsured households here, but the key patterns are similar in levels (with larger magnitudes) and for all non-elderly households (see Appendix Table A10).

5 Insurance Value of Health Insurance

5.1 Sufficient statistic estimates

Implementation.— Recall from Section 2 that the sufficient statistic approximation to the insurance value of a change in constraints—the excess of its ex ante value over its mean ex post value, both measured in terms of resources in the same states of the world—is the covariance, across states of the world in the status quo, between normalized marginal utility and the ex post value of the change, $Cov\left(\widehat{\lambda}, V\right)$ (see equation (4)).[30] As discussed in Section 3, I aim to estimate insurance value from three main perspectives. The short run perspective of immediately before the coverage begins is based on the following regression:

$$\Delta\log(c_{it}) = \alpha + \beta\Delta\log(1+V_{it}) + \gamma X_{it} + \varepsilon_{it}, \tag{7}$$

where i is a household, t is calendar time, $\Delta\log(c_{it}) \equiv \log(c_{it}) - \log(c_{it-1})$ is the within-household change in log consumption from one wave to the next, $\Delta\log(1+V_{it}) \equiv \log(1+V_{it}) - \log(1+V_{it-1})$ is the within-household change in the log of one plus the ex post value of the coverage, and the controls X_{it} are year dummies and a cubic in age. With state-independent utility with constant coefficient of relative risk aversion $\gamma > 0$, the desired covariance is approximately,

$$Cov\left(\widehat{\lambda}, V\right) \approx -\gamma \times \beta \times \frac{Var(V)}{E(V)}, \tag{8}$$

where the key assumption, analogous to that in much of the unemployment insurance literature, is that the slope across states of the world is equal to the slope of the respective within-household changes (see Appendix A.5).[31] The medium run perspective of ten years before the coverage begins is based on a regression that is identical except that it uses within-household changes in the key variables from one wave to the fifth wave after that, ten years later. The long run perspective of someone behind the veil is based on a regression of log consumption on the log of one plus the ex post value of the coverage, plus year dummies, a

[30]Economic logic and quantitative results of the structural model of Section 5.3 both suggest that the approximation error of this first-order approximation tends to make the sufficient statistic overstate the insurance value of health insurance (i.e., to be biased against anti-insurance). Intuitively, it overstates the benefit of insuring health care costs and understates the cost of exacerbating other risks by ignoring that the marginal benefit of decreasing a distortion decreases as the size of the distortion decreases and vice-versa.

[31]My regressions based on equation (7), including the specific control variables, are the health insurance analogue of a common specification in the literature on unemployment insurance (e.g., Hendren, 2017). As discussed in Section 3, the goal is to estimate a covariance across states of the world, not a causal effect of health insurance. With this approach, exogenous variation in health insurance is of no immediate use.

cubic in age, and a quadratic in household size.[32]

Consumption and marginal utility, c_{it}, λ_{it}.— My main specifications follow the common practice of modeling marginal utility as a decreasing function of measured consumption spending. My main measure of consumption is total annual expenditure on food, housing, transportation, clothing, travel, recreation, education, and child care, as measured in the PSID. Given the possibility of measurement error and the sensitivity of marginal utility to low consumption levels, I impose an annual consumption floor of $5,000. This affects less than one percent of observations, and the results are quite similar if I use half or twice this amount. As discussed above, as a baseline I assume state-independent, constant relative risk aversion utility, $\lambda_{it} = c_{it}^{-\gamma}$, with a baseline coefficient of relative risk aversion of three: $\gamma = 3$. I test the robustness of the results to many alternative assumptions about marginal utility, including different risk aversion, using food rather than total consumption, proxying consumption as income minus health spending, and different assumptions about state-dependent utility. Using measured rather than modeled consumption ensures that the key relationship, between consumption and health spending, is determined by the data. Intuitively, this approach is based on the idea that a household's consumption reveals the tightness of its constraint, bypassing the need to model the constraint in its entirety.

Ex post value, V_{it}.— I estimate the value of (hypothetically) supplementing status quo health insurance coverage with full coverage above various stop-loss thresholds. As a baseline, I assume that to first order the ex post value of full coverage above stop-loss $d \geqslant 0$ is $V_{it} = \max\{0, \ oop_{it} - d\}$. Full coverage is the special case with $d = 0$: $V_{it} = oop_{it}$.[33] As discussed in Section 3, this follows the standard approach of focusing on health spending, and, with optimization, it captures to first order the value of changes in coverage if there are no first-order effects on the household's cost of relying on implicit insurance. I also test robustness to large private benefits of improved health and reduced medical debt.

Results.— Table 1 presents estimates of the insurance value of going from the status quo to full health insurance coverage in three sets of states: non-elderly uninsured, non-elderly insured, and elderly insured. In all cases, the estimated insurance value is significantly negative, and it becomes increasingly negative for earlier perspectives from which more risk remains to be revealed. Providing full coverage in uninsured states has an anti-insurance cost of $210 from the perspective of immediately before the coverage begins, $440 from ten years before the coverage begins, and $720 from behind the veil. Such coverage is worth less

[32] These long run regressions are the same as the regressions used by Finkelstein et al. (2019a) in their most closely-related analysis of the value of Medicaid. I use log specifications to reduce the influence of outliers. Levels specifications tend to produce results that are similar but less precise.

[33] I focus on full coverage above a threshold in part because its effect on out-of-pocket spending, unlike that of other changes in health insurance, is straightforward to infer even with unobserved, nonlinear implicit taxation of health insurance by implicit insurance.

Table 1: Insurance value of completing health insurance: Sufficient statistic estimates

	Non-elderly uninsured			Non-elderly insured			Elderly insured		
	Short run (1)	Medium run (2)	Long run (3)	Short run (4)	Medium run (5)	Long run (6)	Short run (7)	Medium run (8)	Long run (9)
Corr(log(c),log(oop))	0.09	0.17	0.25	0.05	0.13	0.29	0.03	0.07	0.23
(*se*)	(0.017)	(0.027)	(0.014)	(0.007)	(0.011)	(0.008)	(0.015)	(0.018)	(0.014)
Insurance value	-205	-439	-721	-89	-289	-758	-82	-199	-785
(*se*)	(38)	(70)	(42)	(13)	(24)	(22)	(38)	(49)	(48)
Mean ex post value	1,016	1,016	1,016	1,505	1,505	1,505	2,086	2,086	2,086
Markup	-0.20	-0.43	-0.71	-0.06	-0.19	-0.50	-0.04	-0.10	-0.38

Notes: Statistics related to the value of completing health insurance in each of three sets of states: non-elderly uninsured, non-elderly insured, and elderly insured. Short run and medium run columns are based on regressions of within-household changes in log consumption on within-household changes in the log of one plus out-of-pocket spending, plus year dummies and a cubic in age, where the changes are from one wave to the next (short run) or from one wave to five waves later (medium run) (see equation (7)). Long run is based on regressions of log consumption on the log of one plus out-of-pocket spending, plus year dummies, a cubic in age, and a quadratic in household size. Short run aims to capture the value of coverage from the perspective of immediately before the coverage begins, medium run from ten years before the coverage begins, and long run from behind the veil of ignorance. Corr(log(c), log(oop)) is the correlation between the relevant changes in (short and medium run) or levels of (long run) log consumption and the log of one plus out-of-pocket spending, both residualized with the corresponding controls. "Insurance value," $Cov\left(\widehat{\lambda}, V\right)$, is $-\gamma \times \beta \times \frac{Var(V)}{E(V)}$, where $\gamma = 3$ is the coefficient of relative risk aversion, β is the regression coefficient on the out-of-pocket spending term, and $V = oop$ (see equation (8)). "Markup" is insurance value per dollar of mean ex post value, $Cov\left(\widehat{\lambda}, V\right)/E\left(V\right)$. Standard errors, which are clustered at the household level, reflect sampling uncertainty in β but treat $E(V)$ and $Var(V)$ as non-stochastic. Data are from the PSID.

ex ante than the mean ex post value worth of cash in the same states by 20%, 43%, and 71%, respectively. Figure 4 shows that the estimated anti-insurance cost increases roughly linearly in the time until coverage begins, as more and more risk remains to be revealed. The longer until the coverage begins, the greater the increase in average consumption associated with a given increase in health spending and so the greater the extent to which the value of health insurance is concentrated in higher-consumption, lower-marginal utility states. Filling in the gaps in coverage in insured states has an anti-insurance cost that is about half as large as that of providing full coverage in uninsured states in the short and medium runs but that is similar in the long run. These estimates imply that in both insured and uninsured states, increases in coverage, say from expanding Medicaid, are worth significantly less ex ante than the mean ex post value worth of cash.

The conclusion that health insurance has negative insurance value is highly robust.[34] The estimated insurance value is negative for coverage of each type of health care in the PSID (hospital, doctor/outpatient/dental, prescriptions/other; see Appendix Table A11). Insur-

[34] Because the estimated insurance value becomes more negative as more risk is included, to be conservative the robustness and heterogeneity analyses focus on the short run perspective.

Figure 4: Insurance value as a function of the amount of risk that remains to be realized

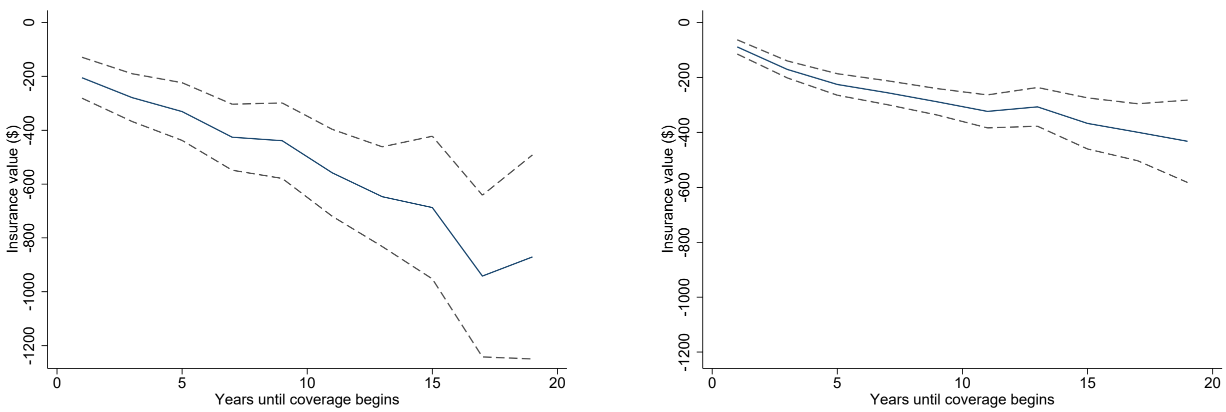

(a) Non-elderly uninsured

(b) Non-elderly insured

Notes: Insurance value of completing health insurance in a particular year as a function of the length of time until the coverage begins. A longer time means more risk remains to be realized. The result for "y years until coverage begins" is based on a regression of the $(y+1)$-year change in log consumption on the $(y+1)$-year change in $\log(1+V)$ (i.e., from one wave to $\frac{y+1}{2}$ waves later for $y \in \{1, 3, 5, \ldots 19\}$), plus year dummies and a cubic in age. "Insurance value," $Cov\left(\hat{\lambda}, V\right)$, is $-\gamma \times \beta \times \frac{Var(V)}{E(V)}$, where $\gamma = 3$ is the coefficient of relative risk aversion, β is the regression coefficient on the out-of-pocket spending term, and $V = oop$ (see equation (8)). Dashed lines are two standard errors above and below the estimated insurance value. Standard errors, which are clustered at the household level, reflect sampling uncertainty in β but treat $E(V)$ and $Var(V)$ as non-stochastic. The corresponding "long run" insurance values to someone behind the veil are $-\$720$ and $-\$760$ for the uninsured and insured, respectively (see Table 1). Data are from the PSID.

ance value is negative for each education group and within each of several subsets of the state space: high and low lagged liquidity, younger and older, and good and bad lagged health (Appendix Table A12). The results are robust to a wide range of assumptions about marginal utility, including the level of risk aversion, the measure or model of consumption, and state-dependent utility (Appendix Table A13). The results are robust to a wide range of changes in the regression specification and control variables (Appendix Table A14). The results are robust to incorporating large private benefits of improved health and reduced medical debt (Appendix A.6 and Appendix Table A15). The markup is significantly negative across widely varying coverage levels, from catastrophic coverage above high deductibles to full coverage (Appendix Figure A4).[35] The key empirical relationship driving the results is that depicted in Figure 3: On average, higher health spending is strongly associated with

[35]The markups on less comprehensive coverage are even more negative than those on more comprehensive coverage. This pattern is suggestive that the "inframarginal" health insurance that many households hold in the status quo also has negative insurance value. So too are the findings of more negative markups from providing coverage to the uninsured than completing the coverage of the insured and of similar markups for coverage of different types of health care and across households with more versus less education. That even catastrophic coverage of high costs has negative insurance value is unsurprising in light of the finding that health spending and consumption covary strongly positively throughout the distribution of health spending, even at the highest levels (see Figure 3). Intuitively, high health spending signals a good realization of other risks, since it reflects not only high health care costs but also limited support from implicit insurance.

higher consumption.[36]

5.2 Insurance of health care costs vs. exacerbation of other risks

Consider a simple model in which consumption equals resources minus health spending, $c = y - oop$, and marginal utility is linear in consumption, $u'(c) = u'(\bar{c}) + u''(\bar{c})(c - \bar{c})$. In that case, to first order the insurance value of going from the status quo to full health insurance coverage is

$$Cov\left(\widehat{\lambda}, V\right) = -\frac{\gamma(\bar{c})}{\bar{c}} Cov\left(c, oop\right) = \frac{\gamma(\bar{c})}{\bar{c}} \left[\underbrace{Var(oop)}_{\text{“Partial effect”}} - \underbrace{Cov\left(y, oop\right)}_{\text{“Portfolio effect”}} \right], \quad (9)$$

where all variables take their values in the status quo, $\bar{c} \equiv E(c)$ is mean consumption, and $\gamma(\bar{c}) \equiv \frac{-\bar{c}u''(\bar{c})}{u'(\bar{c})} > 0$ is the coefficient of relative risk aversion at $\bar{c}$ (see Appendix A.7 for a derivation). The first equality connects the sufficient statistic to the relationship between health spending and consumption. The second equality, which is related to equation (6), connects the sufficient statistic to the extent of health spending risk and the relationship between health spending and other risks.

Interpreting the empirical results through the lens of this model suggests the following. Health insurance increases financial risk on net (Fact 3 and the sufficient statistic results) because the valuable partial effect of insuring health care costs is dominated by a costly portfolio effect of exacerbating other risks. The partial effect is small because health care cost risk is limited (Fact 1), which in turn is because implicit insurance provides considerable protection against health care costs.[37] The portfolio effect is costly because health spending hedges other risks (Fact 2), which in turn is mainly because of the greater protection from implicit insurance when other circumstances are worse (and secondarily because certain types of health care are normal goods). Such protection provides valuable insurance against

[36] Finkelstein et al. (2019a) also estimate sufficient statistic approximations to the insurance value of health insurance. Using a variety of specifications in both the PSID and the Consumer Expenditure Survey, they too find that the estimated insurance value is robustly negative. They concluded that these unexpected results may have been driven by measurement error. While certain types of measurement error could bias the sufficient statistic toward negative insurance value, Appendix D presents several considerations why measurement error is unlikely to explain the sufficient statistic results or the wide variety of corroborating evidence I find in this paper.

[37] Implicit insurance decreases both the extent to which health care costs reduce consumption and the ex post value of health insurance in 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. Because of these reinforcing effects—revealed by the partial effect being the *square* of the standard deviation of health spending—implicit insurance can decrease the partial effect significantly. For example, among uninsured households the variance of health spending is less than one-twentieth that of total health care costs, so implicit insurance reduces the partial effect in this simple framework by roughly a factor of 20.

other risks—insurance that health insurance undoes. The portfolio effect cost is significant because households face substantial risk in income, assets, and living expenses, so even a modest exacerbation of these risks can have a large welfare cost. In terms of the economics of the second best, the exacerbation of other risks is more costly than the protection against health care cost risk is beneficial largely because other risks are far greater than health care cost risk. In terms of equation (5), the partial effect is effectively a small reduction in a small wedge (small value and marginal utility gaps), whereas the portfolio effect is effectively a small increase in a large wedge (small value gaps and large marginal utility gaps). The anti-insurance cost is increasing in the amount of risk that remains to be revealed because the costly portfolio effect grows more rapidly in the amount of risk than the valuable partial effect, since other risks are more persistent than health care costs.[38]

Figure 5: Implicit health insurance can transform the targeting of formal health insurance

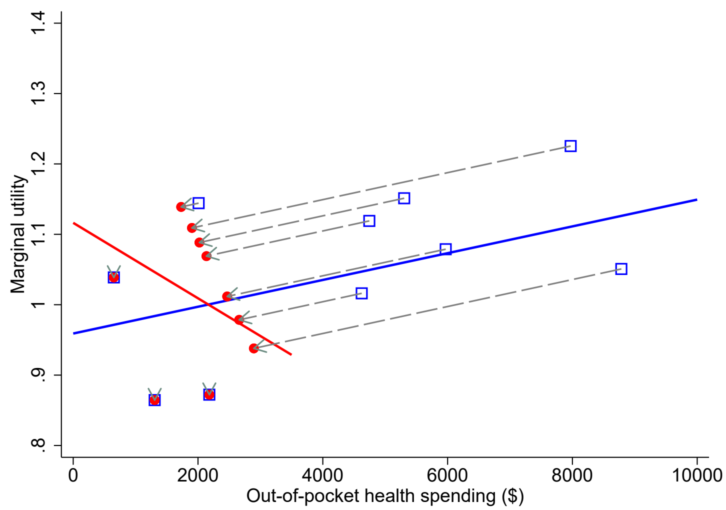

Notes: Scatter plot of simulations of two versions of a stylized model: without implicit health insurance (blue squares and upward-sloping regression line) and with it (red circles and downward-sloping regression line). Risks in total health care costs h and resources y are independent: $h \sim unif[\$0,\ \$10,000]$ and $y \sim unif[\$40,000,\ \$60,000]$. Marginal utility is linear in consumption. Consumption is resources minus health spending: $c = y - oop$. Without implicit insurance, $oop = h$. The arrows show the effects of implicit insurance that provides full coverage above a resources-dependent deductible, $oop = \min\{h,\ d(y)\}$, where $d'(y) > 0$ (less coverage when resources are higher). The sign of the regression slope is the same as the sign of the first order approximation to the insurance value of full health insurance: $\beta_{\widehat{\lambda}|oop} = Cov\left(\widehat{\lambda}, oop\right)/Var(oop)$. Insurance value is the amount by which the ex ante value exceeds the mean ex post value (see equation (3)). For legibility, the figure shows a (random) subset of the simulated data.

[38]Because of the opposing partial and portfolio effects, in theory increasing the amount of risk yet to be revealed could increase or decrease the net insurance value, depending on the relative persistence of the underlying risks. In practice, a variety of evidence suggests that certain important other risks tend to be more persistent than health care costs. For example, mass layoffs are associated with long-term earnings losses (e.g., Jacobson et al., 1993), and even hospitalizations have longer-lasting effects on earnings than on health spending (Dobkin et al., 2018). That the anti-insurance cost is increasing in the amount of risk means that willingness to pay for health insurance overstates the full ex ante value, since it excludes the cost of exacerbating already-revealed risk.

Figure 5 shows how a simple approximation to implicit health insurance based on the empirical evidence—full coverage above a resources-dependent deductible—transforms health insurance targeting in a simple model with independent risks in total health care costs and resources. Without implicit insurance, states of the world with higher health spending have higher marginal utility on average, since health spending is uncorrelated with resources. With implicit insurance, however, states with higher health spending may well have lower marginal utility on average, since health spending is positively correlated with resources. Intuitively, because implicit insurance helps more when circumstances are worse, higher health spending signals not only a worse realization of health care costs but also (and possibly even mainly) a better realization of other risks. Such implicit insurance greatly limits health insurance's valuable targeting across states of the world with different health care costs while causing health insurance to systematically exacerbate risk within the key, high-marginal states in which the individual would otherwise, if not for health insurance, rely on implicit insurance. In those states, health insurance's ex post value is negatively related to marginal utility since differences in its ex post value are driven by differences in implicit insurance support—and so in resources—not in total health care costs.

A stark example of the contributions of implicit insurance to the opposing pro- and anti-insurance effects of health insurance is provided by hospitalization risk. I find that health insurance tends to exacerbate hospitalization-related risk, especially beyond the very short run (see Appendix Figure A5). Decomposing the markup to someone behind the veil on health insurance's hospitalization-related targeting, I find that the average net transfer from non-hospitalization to hospitalization states provides moderately valuable insurance, contributing 0.24 to the markup, but that the within-hospitalization transfers exacerbate risk at a cost more than four times as large, for a total markup of −0.87 (see Appendix Figure A6a and Appendix Table A16). To investigate the potential role of implicit insurance, I contrast health insurance's actual targeting with its predicted targeting in the absence of implicit insurance based on a simple model of the effect of implicit insurance on out-of-pocket spending and consumption (see Appendix Figure A6b, an empirical analogue of Figure 5).[39] Whereas in reality hospitalization is associated with a small increase in out-of-pocket spending and a small decrease in consumption even among the uninsured (as shown in Appendix Table A10), without implicit insurance hospitalization is predicted to be associated with

[39] Counterfactual out-of-pocket spending without implicit health insurance is $oop_\omega^{no} = oop_\omega + ihi(hosp_\omega, y_\omega)$, where oop_ω is observed out-of-pocket spending and $ihi(hosp_\omega, y_\omega)$ is such that (i) average counterfactual out-of-pocket spending is the same in hospitalization states with income realizations that are above versus below the 25th percentile realization in hospitalization states and (ii) average implicit insurance support in hospitalization states as a whole is \$6,000 per year, roughly Dobkin et al.'s (2018) estimate of the effect of hospitalization on unpaid bills among the uninsured. This likely understates the support from implicit insurance for the average hospital admission among the non-elderly uninsured. Counterfactual consumption is $c_\omega^{no} = \max\{\underline{c},\ c - ihi(hosp_\omega, y_\omega)\}$. To be clear, this highly-stylized analysis aims only to give a rough sense of the likely effect of implicit insurance in a transparent, intuitive way.

a large increase in out-of-pocket spending and a large decrease in consumption. Whereas in reality coverage of hospitalization-related health care costs provides moderately-valuable insurance against health care costs while exacerbating income risk at greater cost, without implicit insurance such coverage is predicted to provide highly valuable insurance on net (insurance value of $1,630, over twice the associated mean ex post value). These results are suggestive that, as in the simple model depicted in Figure 5, health insurance targeting within the key, high-total cost (hospitalization) states exacerbates risk because it is driven mainly by differences in implicit insurance support rather than total health care costs.

Were it not for other risks, health insurance would have a straightforward effect on financial risk: It would provide valuable insurance against health care costs. But because of other risks and implicit insurance, health insurance effectively bundles two assets: one that pays off more when health care costs are higher, which provides valuable insurance against health care costs, and one that pays off more when non-health care circumstances are better, which exacerbates other risks. The overall effect therefore reflects opposing pro- and anti-insurance effects: Having health insurance means having more insurance of health care costs but less insurance of other risks. In principle, either could dominate. In practice, the exacerbation of other risks dominates and health insurance increases financial risk on net.

5.3 Structural analysis of mechanisms

To better understand the underlying mechanisms and to assess the generalizability of the results, this section develops and analyzes a simple model guided by the key empirical regularities. The model is based on standard models of health spending risk but adds two features based on the empirical evidence: other risk and implicit insurance.[40]

Model.— A household draws health care consumption h and resources y from the joint distribution $F(h, y)$. (Non-health) consumption is determined by the constraint

$$c(h, y; HI) = \max\{\underline{c},\ y - [tot(h) - hi(h; HI) - ihi(h, y; HI)]\}, \tag{10}$$

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 "implicit health insurance" support over and above any support from the consumption floor. Ex ante expected utility is the expected value of a state-independent constant relative risk aversion

[40]The model is kept as simple as possible toward the goal of understanding, not estimating, the value of health insurance. The sufficient statistic is my preferred approach to estimation given its considerable advantages in terms of flexibility and robustness.

function of consumption,

$$v(h, n; HI) = \frac{c(h, n; HI)^{1-\gamma}}{1-\gamma}, \tag{11}$$

where γ is the coefficient of relative risk aversion.

Empirical inputs.— The key ingredients are the joint distribution of health care consumption and resources, $F(h, y)$, and implicit health insurance support, $ihi(h, y; HI)$. For $F(h, y)$, I use the joint distribution of residualized total health care costs and residualized income among non-elderly households in the MEPS, residualized with year dummies, a cubic in age, a quadratic in household size, and education category dummies. The aim is to approximate relatively long-run risk where the household knows its permanent skill or ability level, as captured by its education, but all other risk remains to be revealed. Income is the maximum of residualized total annual income and an income floor of \$15,000 (about the tenth percentile of income). Total health care costs are as before (see Section 3) except that I inflate those of the uninsured by 25% to reflect moral hazard. This follows the common practice of proxying for risk with cross-sectional heterogeneity and ensures that the model matches the joint distribution of what likely 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 other sources.[41]

Implicit health insurance provides full coverage above an income-dependent deductible,

$$ihi(h, y; HI) = \max\{0,\ h - hi(h; HI) - d_{ihi}(y)\}. \tag{12}$$

The deductible function, $d_{ihi}(y)$, is based on the predicted values from a regression of out-of-pocket spending on a cubic in income and year dummies, a cubic in age, a quadratic in household size, and education category dummies among non-elderly households in the MEPS without health insurance and with annual health care charges of at least \$20,000 (a regression version of Figure 2b). The idea is to estimate the typical amount of health care costs that is *not* covered by implicit insurance (i.e., that is below the effective deductible). That there is implicit health insurance on top of the consumption floor captures in a simple way the observed unevenness of the safety net.

The consumption floor is $\underline{c} = \$5,000$/year. The coefficient of relative risk aversion is $\gamma = 3$. I focus mainly on comprehensive health insurance that covers 100% of health care costs.

[41]The purpose of the income floor is to be conservative about the income risk that households face. The moral hazard factor of 25% is the change in utilization from health insurance in the Oregon Health Insurance Experiment (Finkelstein et al., 2012). My aim in scaling up the health care consumption of uninsured households to the predicted level with full health insurance is to err on the side of overstating health care cost risk and so the insurance value of health insurance. The analysis otherwise ignores moral hazard in order to focus on insurance value. As discussed in Appendix A.6, my sufficient statistic estimates suggest that moral hazard if anything increases health insurance's anti-insurance cost.

Remarks.— As in the simplest standard approach, everything is driven by the budget constraint and consumption equals income minus health spending. The key difference is that here, health spending is not exogenous with respect to other non-consumption elements of the constraint. Health spending is potentially correlated with income both "directly," through the joint distribution of health care consumption and income, and "indirectly," through implicit insurance. Unlike standard approaches, this model admits both possibilities: Health insurance could be pro- or anti-insurance. 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 may have opposing pro- and anti-insurance effects: Health insurance insures health care cost risk but may exacerbate income risk.[42]

Results.— *Insurance value.—* The insurance value of health insurance is significantly, robustly negative: Health insurance is worth considerably less ex ante than its mean ex post value. With the baseline parameters, the insurance value is equivalent to reducing consumption in all states of the world by \$490 per year (see Appendix Table A17). This is about 10% of gross health insurance benefits and 19% of net benefits (mean ex post value). The small change to the standard approach of accounting for the link between resources and health spending through implicit insurance can explain why health insurance would increase financial risk, to an extent broadly similar to that implied by the corresponding sufficient statistic estimates.[43] The model matches well the key empirical patterns, and the conclusion that health insurance increases financial risk is highly robust. Health insurance increases financial risk for all six age-by-education groups (Appendix Table A18), for households with less income risk and less implicit insurance coverage than the typical household appears to have empirically (Appendix Table A19), and for many other changes to the model (Appendix Table A20).[44] As in the sufficient statistic analysis, here too I find that health insurance exacerbates hospitalization-related risk and that there is no tendency for less extensive coverage to have a higher markup (Appendix Table A20).[45]

[42]Whereas standard approaches focus exclusively on how health spending affects the tightness of the constraint, this model allows the tightness of the constraint to affect health spending as well. Whether health insurance insures or exacerbates income risk depends on whether any pro-insurance effect from a positive correlation between income and health outweighs the anti-insurance effects from health insurance displacing implicit insurance and from certain types of health care being normal goods.

[43]The structural analysis, which aims to capture all risk within education groups, is between the medium run (ten years) and long run (behind the veil) perspectives in the sufficient statistic analysis. The corresponding estimates are −\$440 and −\$720, respectively (versus −\$490 in the structural model).

[44]One reason the results are so robust is that many of the changes to the model that would decrease the cost of exacerbating income risk would also decrease the benefit of insuring health care costs (e.g., reducing risk aversion, allowing for consumption smoothing, and decreasing the amount of risk by adopting a later perspective). An analogous point applies to the sufficient statistic analysis as well.

[45]In the model, not only is insurance value negative under a wide range of parameter values, the markups associated with going from no coverage to typical levels of coverage are more negative than those associated with going from typical levels of coverage to full coverage. This is further suggestive evidence that the inframarginal health insurance held by many households in the status quo increases financial risk.

Figure 6: Other risks and implicit health insurance reduce the insurance value of health insurance

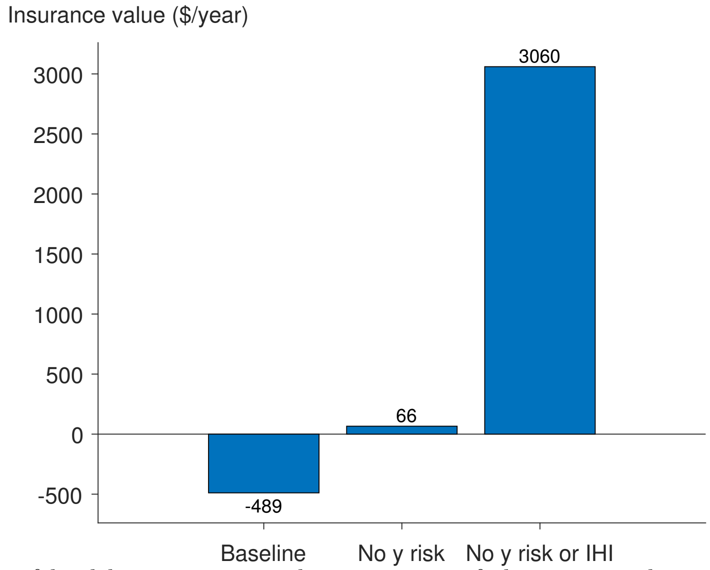

Notes: Insurance value of health insurance in three versions of the structural model: the baseline model ("Baseline"), no income risk ("No y risk"), and neither income risk nor implicit health insurance ("No y risk or IHI"). "No income risk" has health risk as in the data, $h \sim F(h)$, but median income in all states, $y \equiv y_{med}$. "Neither income risk nor implicit health insurance" has $y \equiv y_{med}$ and $ihi(h, y; HI) \equiv 0$. Insurance value is the amount by which the ex ante equivalent variation of health insurance exceeds its mean ex post value (see equation (3)), using consumption-based equivalent variation. See Appendix Table A17 for related statistics. The baseline risk process aims to approximate relatively long run risk: all risk within education groups but not the risk of being in one education group as opposed to another.

Mechanisms.— Figure 6 and Appendix Table A17 show results for three counterfactuals: the baseline model, no income risk, and neither income risk nor implicit health insurance. With neither income risk nor implicit insurance, health insurance generates insurance value surplus worth \$3,060 per year (versus −\$490 in the baseline). The large impact of income risk and implicit insurance reflects two reinforcing effects. First, implicit insurance causes health insurance to exacerbate income risk, at a welfare cost of about \$560 per year (insurance value of −\$490 in the baseline versus \$70 without income risk). This exacerbation is entirely due to implicit insurance. Whereas gross health insurance benefits are virtually uncorrelated with income (correlation of −0.002), net health insurance benefits, accounting for implicit taxation by implicit insurance, are strongly positively correlated with income (correlation of 0.17). Second, implicit insurance transforms health insurance's gross pro-insurance benefit from insuring health care costs from highly valuable (worth \$3,060 per year) to a matter of near indifference (\$70 per year). So health insurance's insurance of health care costs is of little value not because insurance of health care costs is of little value—it is of considerable value—but because implicit insurance provides so much protection that the residual risk remaining for health insurance is of little consequence.[46]

[46]Although other risks have a smaller effect than implicit insurance, their effect is large both relative to the effect of health care costs and in absolute terms. Relative to protection against health care costs, the exacerbation of other risks has a welfare impact eight times as large (\$560/\$70 = 8). In absolute terms,

These results, in keeping with the descriptive and sufficient statistic evidence, suggest that implicit insurance fundamentally transforms—and, from an insurance perspective, ruins—health insurance targeting. Implicit insurance steals the thunder of formal health insurance by transforming its insurance against health care costs from highly valuable to a matter of near indifference, and adds insult to injury by transforming it from being roughly neutral with respect to other risks to exacerbating them at significant cost.

5.4 Implications

Does health insurance decrease welfare? *If* the key costs and benefits of health insurance were those at the center of the economic approach—risk protection and moral hazard—then *relative to equally valuable cash transfers*, health insurance would be decreasing welfare on a large scale. Health insurance has not only a substantial moral hazard cost but also a previously unappreciated anti-insurance cost. A rough, back-of-the-envelope calculation combining existing estimates of moral hazard with my estimates of insurance value suggests that for a typical household, the combined moral hazard and anti-insurance cost of full coverage is on the order of $2,000 per year.[47] Many feasible alternative contracts would be less bad on both dimensions, including the "null contract" of a cash transfer.

But cash transfers are not always feasible, and risk protection is not the only benefit of health insurance. Political economy factors might favor health insurance over cash transfers (see Currie and Gahvari, 2008, for a review). Providing health insurance in kind could improve targeting relative to cash (Nichols and Zeckhauser, 1982). Coverage of certain health goods and services can have a variety of benefits. It has paternalistic benefits if people would otherwise under-consume care (e.g., Baicker et al., 2015). It has merit good benefits if people value others' health care consumption or health insurance coverage per se, beyond the effect on the recipient's welfare (Musgrave, 1959). It can help internalize externalities from infectious disease or health care-related innovation. It can create positive "fiscal externalities" from increased productivity and net tax payments.[48,49] It reduces reliance on

$560 per household corresponds to an aggregate welfare cost to U.S. households on the order of $70 billion per year, comparable to government spending on food stamps ($70 billion), the EITC ($70 billion), SSI ($55 billion), and cash welfare ($30 billion) (see footnote 4).

[47]Finkelstein et al. (2019a) estimate an average per-person moral hazard cost of around $750 per year. I multiply this by two to get a rough estimate of average per-household cost ($1,500) and add to it an anti-insurance cost of $500 as a rough summary measure between the sufficient statistic estimates for uninsured households in the medium and long runs ($440 and $720, respectively).

[48]Such fiscal externalities can be so large as to make certain health insurance expansions *more than pay for themselves* (e.g., Miller and Wherry, 2019; Brown et al., 2020; Hendren and Sprung-Keyser, 2020; Goodman-Bacon, 2021). Such expansions, which help the direct recipients at negative cost to the government, are desirable under virtually all social welfare functions, regardless of any anti-insurance cost.

[49]Combining these reasons why moral hazard responses can be beneficial (mistakes, externalities, and

implicit insurance (e.g., Garthwaite et al., 2018) and the associated costs. It could provide peace of mind (despite increasing financial risk). It can help people access care. A variety of evidence suggests that some of these benefits are large. Yet much remains unknown about the overall magnitude of such benefits. This is an important priority for future research.[50] That typical contracts increase financial risk just means that one key element of their overall welfare effect is for many households not the benefit previously thought but a cost.

Beyond ex ante welfare, my results highlight a stark tradeoff, in cases in which cash benefits are feasible, between the benefits of health insurance and ex post welfare in the worst states. This tradeoff arises because health insurance tends to help least in the worst states. How could that be, given the well-established benefits of health insurance, including—indeed, especially—in the worst states? As Finkelstein et al. (2018) discuss, the benefits of health insurance to individuals take three main forms. One is reduced out-of-pocket spending. As shown throughout this paper, out-of-pocket spending tends to be lowest in the worst states, mainly from greater implicit insurance support but also from individuals postponing care when times are tight. Another benefit is reduced medical debt. Such reductions, as Finkelstein et al. (2018) discuss, have clearer benefits to creditors than to the individual (especially in the worst states: what would someone pay out of their limited resources in the worst states to reduce their medical debt?), and they are an example of the displacement of implicit insurance that causes health insurance to help least in the worst states. The third benefit is improved health from greater health care consumption, especially in the worst states. Counter-intuitively, this benefit is, in part, evidence that health insurance helps *less* than cash in the worst states. It suggests that there are states of the world with such urgent demands beyond health that individuals postpone even high-value care. While the ex ante value of health insurance in such states is high, individuals' behavior suggests that their perceived ex post value is low, i.e., that ex post in such states they would trade their health insurance coverage for even small amounts of cash to meet their other urgent demands.[51]

Could health insurance provide better risk protection? Within the class of contracts that pay a share of health care costs, which comprises the vast majority of contracts currently

merit good considerations) with the result that health insurance increases financial risk 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 provide a moral hazard benefit at the expense of an insurance cost.

[50] A related priority is to incorporate the reshaping of health care costs and health insurance targeting by implicit insurance into the large bodies of research on household choices when facing risk, including consumption, labor supply, and, especially, health insurance, as such analyses typically maintain assumptions that rule out that health spending decreases risk and that health insurance increases financial risk.

[51] Of course, that individuals sometimes prioritize other things over health does not mean that doing so is in their interest, so revealed-preferences logic should be applied with care. But in light of the major risks that individuals face, it is natural that there would be high-marginal utility states in which individuals would be better off ex post with a little more cash than with health insurance coverage.

in use in the U.S. (Cutler, 2002), it is not easy to identify contracts that would decrease financial risk. Such contracts' targeting is undermined by implicit insurance. This is health insurance in a highly second-best world in which the residual health spending that remains after implicit insurance is more risk absorber than risk creator. That makes it hard for such contracts, whose main effect is to decrease health spending, to decrease financial risk.[52] Yet there would appear to be scope to better insure health risk. Even people with health insurance are exposed to consumption risk from health shocks (e.g., French and Jones, 2004; Dobkin et al., 2018; Meyer and Mok, 2019). Could a well-designed contract insure health risk while not exacerbating other risks so much? One candidate is indemnity insurance that paid fixed benefits based on health diagnoses.[53] The prevailing view is that indemnity insurance is better than typical contracts in terms of moral hazard but worse in terms of risk protection, since it provides no within-diagnosis insurance (e.g., some heart attacks are costlier than others; see Zeckhauser, 1970). Although the lack of within-diagnosis insurance is an important limitation, my results suggest that indemnity insurance would be better, not worse, in terms of risk protection.[54]

How generalizable are the results? The analysis highlights the importance of two factors that likely vary significantly across households and economies: exposure to other risks and the supply of and demand for implicit insurance protection against health care costs.[55] The smaller the exposure to other risks, such as those in income, assets, and living expenses, the smaller the anti-insurance cost from exacerbating them. Still, my results suggest that, other things equal, exposure to other risks would have to be significantly smaller than that of typical U.S. households to make health insurance decrease financial risk, and that even then its insurance value would be small. Even among elderly households in the U.S., who face less income risk and more health care cost risk than non-elderly households, income alone is much more variable than total health care costs. In such contexts in which exposure to other risks far exceeds that to health care costs, even a modest exacerbation of other risks

[52]Even catastrophic coverage increases financial risk (see Appendix Figure A4 and footnote 35). Its anti-insurance cost is smaller than that of more extensive coverage, however, so its advantages include not only a smaller moral hazard cost but also a larger insurance benefit (smaller anti-insurance cost).

[53]For example, it might pay $10,000 in the event of a heart attack. Though indemnity health insurance is rare today, it was common in the past (see, e.g., Cutler, 2002), and indemnity insurance is common in other contexts, such as life insurance, annuities, and, increasingly, long-term care insurance. For long-term care insurance, see Doty et al. (2010) for an overview of such policies and Lieber and Lockwood (2019) for an analysis of their targeting effects.

[54]The idea is that indemnity insurance would displace implicit insurance less than typical contracts do, since someone with indemnity insurance still has health care bills on which they could potentially receive support from implicit insurance. I estimate that a per-hospital-day indemnity for uninsured non-elderly households that did not displace implicit insurance would generate 59 cents of insurance value per dollar of expected value to someone behind the veil (see Appendix Table A21).

[55]For example, households in Denmark are well-insured against severe nonfatal health shocks but not against the death of a spouse (Fadlon and Nielsen, 2021). Households in Canada have roughly half of their persistent earnings losses from layoffs and hospital stays replaced by taxes and transfers (Stepner, 2019).

can dominate a larger reduction in health care cost risk.

The smaller the supply of and demand for implicit insurance, the smaller the anti-insurance cost. I find that in the structural model, less implicit insurance (upward-shifted $d_{ihi}(y)$ schedule) can lead to health insurance having positive insurance value. But the required shift is sufficiently large that it causes predicted health spending to exceed even the observed spending of uninsured households with high (⩾\$50,000) financial costs of bankruptcy. Moreover, to the extent that less implicit insurance protection is associated with having more resources, those same resources may help households smooth consumption in the face of shocks, a factor excluded from the simple hand-to-mouth model. So health insurance is unlikely to have high insurance value except for households with an unusual combination of characteristics: low ability to smooth consumption and (yet) unusually low willingness and ability to rely on implicit insurance, plus unusually low exposure to other risks.

Does this mean implicit health insurance is bad? That implicit insurance decreases health insurance's value so much is in part a sign of its *success*. It reflects the valuable all-purpose insurance that implicit health insurance provides against a variety of risks, health and non-health alike. In this particular sense, implicit insurance is "the best health insurance that money can't buy."[56] But it is also a sign, and an unappreciated cost, of the *unevenness* of the safety net. The safety net provides much more protection against health care costs than against other risks. More uniform protection might provide better insurance. In any event, with the current safety net health insurance is complementary with reducing other risk exposures.

6 Conclusion

For many U.S. households, health insurance increases financial risk and is worth less ex ante than its mean ex post value worth of cash. Health insurance exacerbates risks other than health care costs, at a cost that exceeds the benefit from insuring health care costs. The key driver is the implicit health insurance provided by discounts, charity care, and bad debt. Implicit insurance provides significant protection against otherwise-uninsured health

[56]To re-purpose Gruber's (2008) memorable remark about Medicaid. To be clear, this is a narrow statement about insurance effects that ignores the many important costs of implicit insurance, including the uneven, incomplete nature of its coverage, its efficiency costs, and that it violates certain notions of fairness. The insurance-promoting targeting of implicit insurance likely reflects a combination of supply- and demand-side factors. Suppliers of implicit insurance presumably aim to provide greater support to individuals in worse circumstances (or to recover a greater share of costs from individuals in better circumstances). Demanders of implicit insurance support presumably request greater discounts and repay less of their health care bills when their circumstances are worse. Regardless of the mechanisms, the net effect is that implicit insurance provides valuable all-purpose insurance that is apparently difficult to replicate by other means.

care costs, especially in the worst states. Health insurance displaces such support, thereby undoing the insurance it would have provided against other risks. This is an important cost of health insurance and related policies.

My findings point to at least two priorities for future research. One is to better understand heterogeneity in health insurance's value. Given the importance of other risks and implicit insurance to health insurance's value, there is likely considerable heterogeneity both across households within a society and across societies with different institutions. Another priority is to quantify other major components of the overall welfare effect of health insurance, especially its non-insurance benefits. That for many households the insurance effect is not the benefit previously thought but a cost raises the urgency of quantifying other benefits and determining if they could be realized in a less costly way.

References

Acemoglu, Daron, Amy Finkelstein, and Matthew J Notowidigdo, "Income and health spending: Evidence from oil price shocks," *Review of Economics and Statistics*, 2013, *95* (4), 1079–1095.

Avery, Robert B., Paul S. Calem, and Glenn B. Canner, "An overview of consumer data and credit reporting," *Federal Reserve Bulletin*, February 2003, *89* (Feb), 47–73.

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

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

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

Blundell, Richard, Margherita Borella, Jeanne Commault, and Mariacristina De Nardi, "Why Does Consumption Fluctuate in Old Age and How Should the Government Insure It?," NBER Working Papers 27348, National Bureau of Economic Research, Inc June 2020.

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

Browning, Martin and Thomas F Crossley, "Shocks, stocks, and socks: Smoothing consumption over a temporary income loss," *Journal of the European Economic Association*, 2009, *7* (6), 1169–1192.

Cattaneo, Matias D., Richard K. Crump, Max H. Farrell, and Yingjie Feng, "On Binscatter," Papers 1902.09608, arXiv.org February 2019.

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

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

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

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

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

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

Deshpande, Manasi and Lee M. Lockwood, "Beyond Health: Non-Health Risk and the Value of Disability Insurance," *Econometrica*, forthcoming.

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

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

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

Fadlon, Itzik and Torben Heien Nielsen, "Family Labor Supply Responses to Severe Health Shocks: Evidence from Danish Administrative Records," *American Economic Journal: Applied Economics*, July 2021, *13* (3), 1–30.

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

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

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

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

_ , Erzo F. P. Luttmer, and Matthew J. Notowidigdo, "What Good Is Wealth Without Health? The Effect Of Health On The Marginal Utility Of Consumption," *Journal of the European Economic Association*, January 2013, *11*, 221–258.

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

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

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

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

French, Eric and John Bailey Jones, "On the distribution and dynamics of health care costs," *Journal of Applied Econometrics*, 2004, *19* (6), 705–721.

_ **and** _ , "The Effects of Health Insurance and Self-Insurance on Retirement Behavior," *Econometrica*, May 2011, *79* (3), 693–732.

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

Ganong, Peter and Pascal Noel, "Consumer Spending during Unemployment: Positive and Normative Implications," *American Economic Review*, July 2019, *109* (7), 2383–2424.

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

Goodman-Bacon, Andrew, "The Long-Run Effects of Childhood Insurance Coverage: Medicaid Implementation, Adult Health, and Labor Market Outcomes," *American Economic Review*, 2021, *111* (8), 2550–93.

Gov. Account. Off., "Nonprofit hospitals: variation in standards and guidance limits comparison of how hospitals meet community benefit requirements," 2008. Rep. GAO-08-880, Gov. Account. Off., Washington, DC.

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

Grossman, Michael, "On the Concept of Health Capital and the Demand for Health," *The Journal of Political Economy*, 1972, *80* (2), 223–255.

Gruber, Jonathan, "Covering the Uninsured in the United States," *Journal of Economic Literature*, September 2008, *46* (3), 571–606.

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

Hadley, Jack, John Holahan, Teresa Coughlin, and Dawn Miller, "Covering the Uninsured in 2008: Current costs, Sources of Payment, and Incremental Costs," *Health Affairs*, 2008, *27* (5), 399–415.

Hendren, Nathaniel, "Knowledge of future job loss and implications for unemployment insurance," *American Economic Review*, 2017, *107* (7), 1778–1823.

_ , "Measuring Ex Ante Welfare in Insurance Markets," *The Review of Economic Studies*, 06 2020, *88* (3), 1193–1223.

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

Jacobson, Louis S., Robert J. LaLonde, and Daniel G. Sullivan, "Earnings Losses of Displaced Workers," *The American Economic Review*, 1993, pp. 685–709.

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

Kroft, Kory, "Takeup, social multipliers and optimal social insurance," *Journal of Public Economics*, April 2008, *92* (3–4), 722–737.

LeCuyer, Nick A and Shubham Singhal, "Overhauling the US health care payment system," *The McKinsey Quarterly*, June 2007.

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

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

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

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

_, Wallace K. C. Mok, and James X. Sullivan, "Household Surveys in Crisis," *Journal of Economic Perspectives*, November 2015, *29* (4), 199–226.

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

Mulligan, Casey B, "The Value of Employer-Sponsored Health Insurance," Working Paper 28590, National Bureau of Economic Research March 2021.

Musgrave, Richard A, *The Theory of Public Finance*, New York and London: McGraw-Hill, 1959.

Nichols, Albert L and Richard J Zeckhauser, "Targeting transfers through restrictions on recipients," *American Economic Review*, 1982, *72* (2), 372–377.

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

Smith, James P, "Healthy Bodies and Thick Wallets: The Dual Relation between Health and Economic Status," *Journal of Economic Perspectives*, 1999, *13* (2), 145–166.

Stepner, Michael, "The Insurance Value of Redistributive Taxes and Transfers," 2019. mimeo, University of Toronto.

US Census, "America's Families and Living Arrangements: 2020," 2020. `https://www.census.gov/data/tables/2020/demo/families/cps-2020.html`, Table H1.

US Centers for Medicare & Medicaid Services, "National Health Expenditure Accounts," 2019. Office of the Actuary, National Health Statistics Group: https://www.cms.gov/Research-Statistics-Data-and-Systems/Statistics-Trends-and-Reports/NationalHealthExpendData/NationalHealthAccountsHistorical, NHE Tables, Table 03.

US Department of Agriculture, "The Food Assistance Landscape: FY 2015 Annual Report," 2016. Econ. Information Bull. EIB-150. https://www.ers.usda.gov/publications/pub-details/?pubid544062.

US Department of Health and Human Services, "2015 Actuarial Report on the Financial Outlook for Medicaid," 2015. https://www.medicaid.gov/medicaid/finance/downloads/medicaid-actuarial-report-2015.pdf.

_ , "FY 2015 Federal TANF & State MOE Financial Data," 2016. https://www.acf.hhs.gov/sites/default/files/ofa/tanf_financial_data_fy_2015.pdf.

US Internal Revenue Service, "Statistics of Income Bulletin Historical Table 1," 2016. https://www.irs.gov/statistics/soi-tax-stats-historical-table-1.

US Social Security Administration, "Annual Report of the Supplemental Security Income Program," 2016. https://www.ssa.gov/OACT/ssir/SSI16/ssi2016.pdf.

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

Zeldes, Stephen P, "Consumption and liquidity constraints: An empirical investigation," *Journal of Political Economy*, 1989, *97* (2), 305–346.

Online Appendix

Health Insurance Increases Financial Risk

Lee M. Lockwood

A Sufficient Statistic Approximation to Insurance Value

A.1 Derivation of equation (4)

Recall equation (4),

$$\underbrace{EAV}_{\text{Ex ante value}} \approx \frac{E(\lambda \times V)}{E(\lambda)} = \underbrace{E(V)}_{\text{Mean ex post value}} + \underbrace{Cov\left(\widehat{\lambda},\ V\right)}_{\text{Insurance value}}. \tag{4}$$

This is a first order approximation to the ex ante value of a change in ex post constraints whose ex post values in different states are V (which may vary across states).

Start from equation (2),

$$E\left[u\left(c_0 + EAV, a_0; \theta\right)\right] = E\left[u\left(c_1, a_1; \theta\right)\right]. \tag{2}$$

Use equation (1),

$$u\left(c_0 + V, a_0; \theta\right) = u\left(c_1, a_1; \theta\right), \tag{1}$$

to write the right-hand-side of equation (2) in terms of V:

$$E\left[u\left(c_0 + EAV, a_0; \theta\right)\right] = E\left[u\left(c_0 + V, a_0; \theta\right)\right]. \tag{13}$$

Take first-order approximations to the utility levels inside the expectations on both sides of equation (13) around the allocation under the original constraint, (c_0, a_0):

$$E\left[u\left(c_0, a_0; \theta\right) + u_c\left(c_0, a_0; \theta\right) EAV\right] \approx E\left[u\left(c_0, a_0; \theta\right) + u_c\left(c_0, a_0; \theta\right) V\right], \tag{14}$$

where $u_c(c, a; \theta)$ is the marginal utility of consumption c when the allocation is (c, a) and the state is θ. Subtracting $E\left[u\left(c_0, a_0; \theta\right)\right]$ from both sides and moving the constant EAV

outside the expectation yields

$$E\left[u_c\left(c_0,a_0;\theta\right)\right]EAV\approx E\left[u_c\left(c_0,a_0;\theta\right)V\right]. \tag{15}$$

Solving for EAV yields

$$EAV\approx\frac{E\left[u_c\left(c_0,a_0;\theta\right)V\right]}{E\left[u_c\left(c_0,a_0;\theta\right)\right]}. \tag{16}$$

Using that $E(XY)=E(X)E(Y)+Cov(X,Y)$ yields

$$EAV\approx E(V)+\frac{Cov\left[u_c\left(c_0,a_0;\theta\right),V\right]}{E\left[u_c\left(c_0,a_0;\theta\right)\right]}. \tag{17}$$

Passing $E\left[u_c\left(c_0,a_0;\theta\right)\right]$ into the covariance and using that $\widehat{\lambda}\equiv\frac{u_c(c_0,a_0;\theta)}{E[u_c(c_0,a_0;\theta)]}$ yields

$$EAV\approx E(V)+Cov\left(\widehat{\lambda},V\right), \tag{18}$$

which is equation (4), as was to be shown.

A.2 Relationship to Baily-Chetty

The "insurance value" covariance, $Cov\left(\widehat{\lambda},V\right)$, generalizes the insurance value part of the Baily-Chetty analysis of optimal social insurance (Baily, 1978; Chetty, 2006) to situations in which the ex post value of the change in constraints, V, can take more than two different values. To see the connection to the familiar Baily-Chetty analysis, consider the special case in which V takes one of two values, V_H with probability p and V_L with probability $(1-p)$. Then the insurance value covariance can be written,

$$\begin{aligned}Cov\left(\widehat{\lambda},V\right)&=E\left[\left(\widehat{\lambda}-E\left(\widehat{\lambda}\right)\right)\left(V-E(V)\right)\right]\\&=p\left(V_H-E(V)\right)\left[E\left(\widehat{\lambda}|V=V_H\right)-E\left(\widehat{\lambda}\right)\right]\\&\quad+\left(1-p\right)\left(V_L-E(V)\right)\left[E\left(\widehat{\lambda}|V=V_L\right)-E\left(\widehat{\lambda}\right)\right].\end{aligned} \tag{19}$$

Noting that $E(V)=pV_H+(1-p)V_L$, and so $(V_H-E(V)=(1-p)(V_H-V_L)$ and $(V_L-E(V)=p(V_L-V_H)$, this can be simplified to

$$Cov\left(\widehat{\lambda},V\right)=p(1-p)\left(V_H-V_L\right)\left[E\left(\widehat{\lambda}|V=V_H\right)-E\left(\widehat{\lambda}|V=V_L\right)\right]. \tag{20}$$

The term in brackets is the familiar "marginal utility gap" from the Baily-Chetty analysis. Typical implementations of this analysis to unemployment insurance consider the following two sets of states of the world: unemployed states, in which the individual is assumed to receive an unemployment insurance benefit, and employed states, in which the individual is assumed to pay unemployment insurance taxes. In this case, the marginal utility gap is that between states of the world in which the individual is unemployed ($V = V_H$) versus employed ($V = V_L$).

This sufficient statistic depends only on marginal utility in the status quo and the ex post value of the contemplated change in constraints. It does not depend on any other outcomes, including counterfactual outcomes away from the status quo or causal effects of the contemplated change in constraints. So estimating it does not require estimating causal effects of the contemplated change in constraints.

The reason that many implementations of the Baily-Chetty approach and related approaches require causal effects of the contemplated change in constraints is that they aim to characterize optimal benefits or, more generally, account for costs as well as value. Costs depend on behavioral responses to the change in constraints. Value, by contrast, does not to first order with optimization, because with optimization, behavioral responses have no first order impact on value by the envelope theorem. The insurance value covariance is about the value of the change in constraints, not the cost, so causal effects of the change are not necessary.

A.3 Relationship to Finkelstein, Hendren, and Luttmer (2019)

My sufficient statistic approach to estimating insurance value is very similar to Finkelstein et al.'s (2019a) "consumption-based optimization approach." The difference is that I estimate a first-order approximation to insurance value, whereas Finkelstein et al. (2019a) make two assumptions to go beyond a first order approximation.

In both cases, the key statistic is the "insurance value covariance" of equation (4): the covariance across states of the world of normalized marginal utility and the ex post value of the contemplated change in health insurance coverage. In both cases, the ex post value of the contemplated change in health insurance is assumed to be the mechanical reduction in out-of-pocket spending (though I test robustness to other assumptions). In both cases, this statistic is estimated using standard strategies for approximating the (unobservable) distribution of states of the world using observable variation across households or over time within households. In neither case is exogenous variation in health insurance or other factors used to estimate this statistic.[57]

[57]Finkelstein et al. (2019a) use exogenous variation in health insurance, generated by the Oregon Health Insurance Experiment, for several purposes, just not for estimating insurance value in their optimization approaches. For example, they use it to estimate the cost of providing the coverage, the value of the

The difference is that Finkelstein et al. (2019a) make two assumptions to go beyond a first order approximation. They assume that (i) the marginal insurance value of hypothetically increasing the extent of health insurance coverage from a baseline of full coverage (i.e., of hypothetically reducing the financial cost to the individual of consuming health care from zero to a negative value, i.e., paying the individual to consume care) is zero and (ii) the marginal insurance value of increasing the extent of coverage is linear in the extent of coverage between no coverage and full coverage (a simple statistical extrapolation). Together, these assumptions imply that the full insurance value of going from no coverage to full coverage—which is the integral of the marginal insurance value over that range of coverage—is one-half the marginal insurance value from a baseline of no coverage.

There are at least two options for going beyond a first order approximation in this context. One would be to follow Finkelstein et al. (2019a) in combining an assumed value of the marginal insurance value at an unobserved counterfactual coverage level with an assumed functional form of the marginal insurance value. The main challenge for this option is that, because of the opposing pro- and anti-insurance effects of health insurance, in theory even the sign of the marginal insurance value at any given coverage level is ambiguous.[58] Another option would be to combine my estimates of the marginal insurance value of increasing coverage from its status quo level among households with versus without health insurance, plus an assumed functional form of the marginal insurance value between those coverage levels. The main challenge for this option is that insured and uninsured households differ in many important ways beyond their health insurance coverage. Moreover, both of these options face the additional challenge that the opposing pro- and anti-insurance effects of health insurance make it considerably more difficult to use theory to guide the choice of the functional form of the marginal insurance value between no coverage and typical coverage. Given these challenges, adopting either of these approaches risks diminishing the key strength of the sufficient statistic approach: its validity under a wide range of assumptions.

Moreover, the approximation error in the first order approximation likely works against

coverage in their "complete-information approach" (described in the next section), and the private value of the moral hazard response in their optimization approaches.

[58]This is true even at full coverage. Although the marginal insurance value of hypothetically increasing health insurance coverage from a baseline of full coverage is zero in a simple model in which health care costs are the only risk (since in that case, there would be no variation in marginal utility across states of the world with full coverage), in richer models with other risks, there is no clear prediction of even the sign of this marginal insurance value. This marginal insurance value is positive if, in the counterfactual with full coverage, greater health care consumption is positively related to marginal utility (e.g., if this covariance mainly reflects the realization of health risk: that people in worse health consume more care and have higher marginal utility, say, due to earning less). But this marginal insurance value is negative if, in the counterfactual with full coverage, greater health care consumption is negatively related to marginal utility (e.g., if this covariance mainly reflects the realization of non-health risk: that people with worse non-health shocks have higher marginal utility and consume less care, say, due to having lower demand for care and facing time or utility costs of consuming care).

my key conclusions. Economic logic and quantitative results of the structural model of Section 5.3 both suggest that the approximation error of this first-order approximation tends to make the sufficient statistic overstate the insurance value of health insurance (i.e., to be biased against anti-insurance). Intuitively, it overstates the benefit of insuring health care costs by ignoring that the marginal benefit of decreasing a distortion decreases as the size of the distortion decreases, and it understates the cost of exacerbating other risks by ignoring that the marginal cost of increasing a distortion increases as the size of the distortion increases. This suggests that my sufficient statistic estimates are lower bounds on health insurance's anti-insurance cost.

Finkelstein et al. (2019a) implement their consumption-based optimization approach with three different sources of information about consumption: two datasets with direct measures of consumption and one simple model of simulated consumption.[59] Their analyses based on direct measures of consumption, which use the PSID and the Consumer Expenditure Survey, reveal robust negative relationships between marginal utility and out-of-pocket spending, consistent with my findings. Their "consumption proxy" analysis of simulated consumption assumes that consumption is equal to the difference between average consumption and the per capita net excess of out-of-pocket spending over its average,

$$c = \overline{c} - \frac{oop - \overline{oop}}{n},$$

where $\overline{c}$ is average consumption expenditure among the low-income uninsured, $\overline{oop}$ is average out-of-pocket spending among untreated compliers in the Oregon Health Insurance Experiment, and n is family size. This is a simple hand-to-mouth model of consumption in which the only risk is in health spending (an instance of a common class of models in the literature on health spending risk and health insurance). It necessarily implies that consumption is negatively correlated with out-of-pocket spending across states of the world and so that health insurance has positive insurance value.

A.4 Alternative approach based on the causal effects of health insurance

Finkelstein et al. (2019a) discuss two types of approaches to estimating the value of health insurance, which they term "optimization approaches" and a "complete-information approach." As discussed in the preceding section, their main optimization approach is closely related to my sufficient statistic approach. This section briefly describes their complete-information approach, which to my knowledge is the approach to valuing health insurance

[59] The Oregon Health Insurance Experiment did not collect information about consumption.

based most closely on the causal effects of health insurance.

The idea of the complete-information approach is to quantify the value of health insurance to the individual by combining (i) a completely-specified utility function and (ii) the causal effect of health insurance on the distribution of all of the arguments of utility (consumption, health, health care, peace of mind, etc.). With these ingredients, it is straightforward to quantify the value of health insurance. For example, to calculate the ex ante equivalent variation of health insurance coverage (the increment to wealth in all states of the world that would make someone without health insurance as well off ex ante as they would be with health insurance), first use the causal effects of health insurance and the utility function to calculate the causal effect of health insurance on ex ante utility, then use the utility function to calculate the increment to wealth that would cause the same increase in ex ante utility.

As Finkelstein et al. (2019a) discuss, although this approach has certain advantages, it is quite demanding in terms of its information requirements. Finkelstein et al. (2019a) emphasize the detailed knowledge about the utility function and the causal effects of health insurance that is required. Another requirement that they do not emphasize is that one needs complete information on the counterfactual outcomes with and without health insurance in *all* states of the world. This could be a considerable challenge in practice, as it requires either that compliance with the experimental or quasi-experimental variation in health insurance be representative of all states of the world or that the analyst make assumptions about the distribution of unobserved counterfactual outcomes in "non-compliant" states (never takers and always takers).[60] These considerations are why I focus on an alternative optimization approach instead.

A.5 Derivation of equation (8)

The goal is to estimate the covariance across states of the world of normalized marginal utility and the ex post value of health insurance, $Cov\left(\widehat{\lambda}, V\right)$. In order to use regressions of the log of (changes in) consumption and ex post value instead of levels to try to reduce the effects of sampling and measurement error, I use two approximations. The first is a

[60] Representative compliance requires that an individual be equally likely to be a "complier," i.e., to have his or her health insurance status shifted by the instrument, in all states of the world. This would be violated if, for instance, in states in which the ex post value is large, the individual is more likely to obtain health insurance regardless of the treatment assignment (an always taker). Or if in states in which the ex post value is small, the individual is less likely to take up health insurance when it is offered (a never taker). These particular patterns of unrepresentative compliance would cause complier states to exhibit less variation in the ex post value of health insurance, and likely in marginal utility as well, than exists across all states.

log-linearization of marginal utility:

$$\log\left(\widehat{\lambda}\right) \approx \widehat{\lambda} - 1, \tag{21}$$

which is a first-order Taylor approximation to $\log\left(\widehat{\lambda}\right)$ around $\widehat{\lambda} = E\left(\widehat{\lambda}\right) = 1$. Rearranging, using the definition of normalized marginal utility ($\widehat{\lambda} \equiv \lambda/E(\lambda)$), and assuming state-independent, constant relative risk aversion utility over consumption ($\lambda = c^{-\gamma}$) yields

$$\widehat{\lambda} \approx 1 + \log\left(\widehat{\lambda}\right) = (1 - \log[E(\lambda)]) + \log(\lambda) = (1 - \log[E(\lambda)]) - \gamma \log(c), \tag{22}$$

where γ is the coefficient of relative risk aversion. Hence, this approximation to normalized marginal utility is linearly decreasing in log consumption, with slope equal to the coefficient of relative risk aversion.

The second approximation is a log-linearization of the ex post value around its mean:

$$\log(V) \approx \log(E(V)) + \frac{1}{E(V)}(V - E(V)), \tag{23}$$

which is a first-order Taylor approximation to $\log(V)$ around $V = E(V)$. Rearranging yields

$$V \approx E(V)(1 - \log(E(V))) + E(V)\log(V). \tag{24}$$

So this approximation to V is linearly increasing in $\log(V)$ with slope $E(V)$.

With these in hand, the covariance of normalized marginal utility and the ex post value of health insurance can be written:

$$Cov\left(\widehat{\lambda}, V\right) \approx -\gamma Cov(\log(c), \log(V))\, E(V) = -\gamma\beta Var(\log(V))E(V) \approx -\gamma\beta\frac{Var(V)}{E(V)}, \tag{25}$$

where β is the slope of the regression of log consumption (or the change therein) on the log of the ex post value (or the change therein). This is the approximation in equation (8), as was to be shown.

A.6 Robustness to large private benefits of improved health and reduced medical debt

The main effects of health insurance are reduced medical debt, improved health, and reduced out-of-pocket spending (Finkelstein et al., 2018). As discussed in Section 3, evidence of benefits to individuals from reduced medical debt "remains limited" (Finkelstein et al., 2018,

p. 270), and improved health, from moral hazard effects on health care consumption, while potentially of considerable ex ante value, is unlikely to significantly increase insurance value and might even decrease it (since the ex post value of moral hazard effects is lower, other things equal, when the marginal utility of consumption is higher). Reduced out-of-pocket spending is the main financial effect of health insurance and, under standard assumptions, a first order approximation to its ex post value. Perhaps in part from such considerations, virtually all analyses of the insurance value of health insurance focus on its impact on out-of-pocket spending.[61] Still, where relevant I test robustness to large private benefits of improved health and reduced medical debt.

Appendix Table A15 reports the results of several tests of the potential effects of large private benefits of improved health and reduced medical debt on the insurance value of health insurance. Columns (2)–(7) increase the ex post value of health insurance by $20,000 in states of the world in which the private benefit of improved health could plausibly be large, including states in which the household head or spouse receives a new cancer diagnosis, states in which the head or spouse has ever received a cancer diagnosis, states in which the head's health recently declined, states in which the head's health is bad, and 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 reduced out-of-pocket spending, from improved health (from moral hazard).[62] The estimated insurance values are always significantly negative, and they remain so even when the ex post value of health insurance is increased by $100,000 in these states (not shown).

The main reason that the results are so robust to even high values of improved health is that bad health is not a strong indicator of marginal utility. A key reason for this, in turn, is presumably the considerable protection against health care costs provided by implicit insurance. Such protection significantly reduces the extent to which bad health shocks increase marginal utility by greatly limiting a key channel by which they otherwise would: increased health spending. Regardless of the underlying mechanisms, bad health is a much weaker indicator of marginal utility than unemployment, for example. That, in turn, is another manifestation of the key proximate reason that health insurance increases financial

[61] The main exception I know of is Finkelstein et al.'s (2019a) "complete-information approach," which I discuss in Appendix A.4.

[62] This seems likely to be conservative not only because $20,000 is a large value of the ex post surplus to the household from moral hazard in a single state, but also because using a uniform value within a given health category ignores the within-category anti-insurance effect from the fact that, other things equal, the ex post value to the household (in terms of resources in that state of the world) of a given health improvement is lower when the marginal utility of consumption is higher. In other words, although moral hazard effects are like inframarginal health care consumption in that they have opposing pro- and anti-insurance effects—a pro-insurance effect from being more valuable when health is worse and an anti-insurance effect from being more valuable when non-health circumstances are better—these tests account for only the pro-insurance effect.

risk on net: Other risks are much less well-insured than health care costs, even among households without formal health insurance. As a result, health insurance's exacerbation of other risks dominates its insurance of health care costs.[63]

Columns (8) and (9) test the potential effects of large private benefits of reduced medical debt on the insurance value of health insurance. Column (8) adds the full amount of the household's outstanding medical bills to the ex post value of health insurance. Column (9) adds the lesser of this amount and $10,000. Both aim to overstate any additional ex post value of health insurance to the household, over and above that from reduced out-of-pocket spending, from reduced medical debt.[64] In both cases, the estimated insurance value remains negative.

These results suggest that even large private benefits of health insurance from improved health and reduced medical debt do not overturn—and perhaps even strengthen—the conclusion that health insurance has negative insurance value, i.e., that it is worth less ex ante than its mean ex post value.

[63]The estimated insurance values in these alternative specifications that increase the ex post value of health insurance in bad-health states are not just negative but more negative than the corresponding baseline estimate. In addition to bad health not being a strong indicator of marginal utility, another contributing factor is that adding a large value to the ex post value of health insurance in certain relatively rare states increases variation in the ex post value, which is a force toward the insurance value increasing in absolute value. Recall equation (8),

$$Cov\left(\widehat{\lambda}, V\right) \approx -\gamma \times \beta \times \frac{Var(V)}{E(V)}. \quad (8)$$

So other things equal, an increase in $Var(V)/E(V)$ increases $\left|Cov\left(\widehat{\lambda}, V\right)\right|$. This is why some of the alternative specifications that increase the ex post value of health insurance in bad-health states increase the estimated anti-insurance cost despite decreasing $Corr(\log(c), \log(1+V))$. The more important result, however, is that even large values of improved health do not change the sign of this correlation, i.e., that on average the ex post value of health insurance is positively related to consumption (and so negatively related to marginal utility).

[64]In theory, reducing debt by $X should be worth at most $X to the household, since it could simply repay $X to achieve that. Other options include not repaying—the most common choice—or discharging through bankruptcy. See footnote 18. Another sense in which this is conservative is that it ignores the anti-insurance effect from the fact that, other things equal, the ex post value to the household (in terms of resources in that state of the world) of a given reduction in medical debt is lower when the marginal utility of consumption is higher (e.g., due to a lower willingness to pay for a given reduction in any hassle or stigma cost of medical debt).

A.7 Derivation of equation (9)

Recall equation (9),

$$Cov\left(\widehat{\lambda}, V\right) = -\frac{\gamma(\overline{c})}{\overline{c}} Cov\left(c, oop\right) = \frac{\gamma(\overline{c})}{\overline{c}} \left[\underbrace{Var(oop)}_{\text{“Direct effect”}} - \underbrace{Cov\left(y, oop\right)}_{\text{“Portfolio effect”}} \right]. \tag{9}$$

This holds in the simple special case in which consumption equals resources less out-of-pocket spending, $c = y - oop$, marginal utility is linear in consumption, $u'(c) = u'(\overline{c}) + u''(\overline{c})(c - \overline{c})$, and $V = oop$. To see this, first note that when marginal utility is linear in consumption, normalized marginal utility can be written

$$\widehat{\lambda} \equiv \frac{u'(c)}{E[u'(c)]} = \frac{u'(\overline{c}) + u''(\overline{c})(c - \overline{c})}{u'(\overline{c})} = \alpha + \frac{u''(\overline{c})}{u'(\overline{c})} c = \alpha - \frac{\gamma(\overline{c})}{\overline{c}} c, \tag{26}$$

where α is a constant and $\gamma(\overline{c}) \equiv -\frac{u''(\overline{c})c}{u'(\overline{c})}$ is the coefficient of relative risk aversion at $c = \overline{c}$. Plugging equation (26) and $V = oop$ into $Cov\left(\widehat{\lambda}, V\right)$ yields

$$Cov\left(\widehat{\lambda}, V\right) = Cov\left(\alpha - \frac{\gamma(\overline{c})}{\overline{c}} c, oop\right) = Cov\left(-\frac{\gamma(\overline{c})}{\overline{c}} c, oop\right) = -\frac{\gamma(\overline{c})}{\overline{c}} Cov\left(c, oop\right), \tag{27}$$

which is the first equality of equation (9). For the second equality, plug $c = y - oop$ into the right-hand-side of equation (27) and rearrange to find

$$-\frac{\gamma(\overline{c})}{\overline{c}} Cov\left(c, oop\right) = -\frac{\gamma(\overline{c})}{\overline{c}} Cov\left(y - oop, oop\right) = \frac{\gamma(\overline{c})}{\overline{c}} \left[Var\left(oop\right) - Cov\left(y, oop\right)\right], \tag{28}$$

which is the second equality of equation (9), as was to be shown.

B Data Appendix

B.1 Panel Study of Income Dynamics (PSID)

Sample.— My PSID sample covers households interviewed in at least one of the eleven waves between 1999 and 2019 inclusive. 1999 is the first year that certain key variables, such as many of the measures of consumption and health spending, were collected. 2019 is the last year of data available as of this writing. I exclude households living in nursing homes and households whose head (or reference person) is younger than 25 years old. The final sample contains 85,769 household-wave observations.

Out-of-pocket health spending ("health spending").— The baseline measure of out-of-pocket spending includes annualized spending on hospital care, doctor visits, outpatient surgery, dental bills, prescriptions, in-home medical care, special facilities, and other services. The underlying variables also include spending on nursing home care, but in practice this is unlikely to have much effect given that I exclude households in nursing homes. I also occasionally use the underlying disaggregated measures, which are (i) hospital bills (and nursing home expenses, though that part is largely removed by my sample restriction); (ii) doctor visits, outpatient surgery, and dental bills; and (iii) prescriptions, in-home medical care, special facilities, and other services. The underlying questions ask about spending during the past calendar year in the 2013–2019 survey waves and during the past two calendar years combined in the 1999–2011 survey waves. I divide the latter measures by two to annualize them. When I restrict the sample to those waves that use the annual measure, which, as I discuss below, aligns better with the timing of the consumption and income variables, the results are similar but stronger, implying a larger anti-insurance cost.

Consumption spending.— The baseline measure of non-health consumption is total annualized expenditure on food, housing, transportation, clothing, travel, recreation, education, and child care. Spending on food includes spending on food at home, away from home, and deliveries. I add to it the annualized value of food stamps in order to better measure food consumption rather than spending (since the conceptual object of interest is consumption, not spending). Given the possibility of measurement error and the sensitivity of marginal utility to low consumption levels, I impose an annual consumption floor of $5,000 on total consumption and, separately for the analyses based on food consumption only, of $1,000 on food consumption. The total consumption floor affects less than one percent of observations. The food consumption floor affects just over one percent of observations. The results are quite similar if I use half or twice the baseline consumption floor amount. The underlying questions about consumption spending allow respondents to choose whether to report their spending per month, per year, or per other unit of time. As Zeldes (1989) discusses, the calendar time period in which respondents are recalling their consumption spending is ambiguous. Zeldes (1989) argues that these questions aim to measure the rate of spending at the time of the interview rather than spending during a particular time period. If so, the corresponding conceptual experiment would be closer to health insurance that reimbursed out-of-pocket spending at the end of the year than to health insurance that covered health care costs as they were incurred throughout the year. As mentioned above, when I restrict the sample to those waves that have better-aligned measures of consumption and out-of-pocket spending, the results are similar but stronger, implying a larger anti-insurance cost. I also test the robustness of the results to using consumption proxies based on income and out-of-pocket spending, whose time reference periods coincide exactly in several waves, in place of measured consumption and find results that are broadly similar to the main results (see, e.g., Appendix Table A13). Appendix D lays out several considerations why measurement error from this or other sources is unlikely to explain the broad pattern of results in practice.

Health insurance.— There is no single ideal way to classify households as being insured or uninsured in a particular wave. One issue is that many households have multiple individuals, who may have different insurance status. Another is that a given individual might have health insurance during some but not all of the time period of interest. As a baseline, I classify households as being insured or uninsured using the main household-level health insurance coverage measures in the PSID. What could be a complicating factor that turns out to be useful is that this main household-level health insurance coverage measure switches during the sample period from being an indicator of whether anyone in the household *had* health insurance coverage at any time since the last wave (1999–2011 waves) to being an indicator of whether anyone in the household *did not have* health insurance coverage at any time since the last wave (2013–2019 waves). The former measure can be used to identify "pure-uninsured" households (those in which no one in the household had health insurance at any time since the last wave) but can only identify a relatively loose definition of insured households (by this measure, a household is insured if anyone in the household had health insurance at any time, even if only briefly and even if others in the household did not). The latter measure can be used to identify "pure-insured" households (those in which everyone in the household had health insurance at all times since the last wave) but can only identify a relatively loose definition of uninsured households (by this measure, a household is uninsured if anyone in the household did not have health insurance at any time, even if only briefly and even if others in the household did have health insurance). I adopt these "impure" measures of insured or uninsured households as my baseline measures because I find that the resulting estimates are very similar to those based on the "pure" measures—itself a manifestation of the finding that the sufficient statistic estimates for the insured and uninsured are similar. It is because of this similarity that I view the benefit of using the "impure" measures in terms of greater sample size as exceeding the cost in terms of classifying certain households as insured or uninsured despite not everyone in the household having that insurance status during the entire time period of interest.

Hospitalization.— My measure of hospitalizations is an indicator of whether the head or spouse was a patient in a hospital overnight or longer at any point in the prior year *and* there is no child under two years old in the household. I limit to hospitalizations in which there is no child under two years old in the household to exclude hospitalizations related to childbirth, in order to better focus on hospitalizations driven by health shocks, as in Dobkin et al. (2018).

Other variables.— My measure of income includes income from all sources, including from social insurance and means-tested programs, so that it reflects the net risk in income accounting for all sources of income risk and insurance. This measure refers to income received in the previous calendar year. My measure of unemployment is an indicator of whether the head or spouse was unemployed at any point in the past year. The education categories that I use to create the education category dummy variables are no degree or GED only, high school degree, some college (including an associate's degree), and college degree or above (bachelor's, master's, or doctorate, including in law [J.D.] or medicine [M.D.]). 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.

Outliers.— Variables expected to have large outliers—consumption, out-of-pocket spending, and income—plus one other variable that turned out to have an extreme outlier—the value of food stamps—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. I do this to avoid having the estimates unduly affected by outlier values that may be errors. If I instead use raw rather than winsorized measures of consumption and out-of-pocket spending, my main sufficient statistic estimate of the short run anti-insurance cost of health insurance for non-elderly uninsured households increases from $210 to $504.

Converting to real dollars.— All monetary variables are converted to real 2020 dollars using the CPI-U-RS.

Survey weights.— Throughout, I use family weights to ensure that the estimates reflect the experiences of the U.S. population.

Standard errors.— Throughout, I cluster standard errors at the household level.

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

B.2 Medical Expenditure Panel Survey (MEPS)

Sample.— 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 waves from 1996–2018. I exclude families whose reference person is younger than 25 years old. The resulting sample has 268,235 family-year observations.

Out-of-pocket health spending ("health spending").— The baseline measure of out-of-pocket spending includes annual spending on office-based visits, hospital outpatient visits, emergency room visits, inpatient hospital stays, prescription medicines, dental visits, home health care, and other medical expenses.

Total health care costs.— Total health care costs are defined as follows. For households with health insurance, total costs are total annual payments, including from the insurer and the household. For households without health insurance, total costs are annual charges scaled by the payments-charge ratio among non-elderly households with health insurance. I follow Mahoney (2015) in scaling by this ratio, which is 0.60, to reflect typical discounts relative to charges.

Health insurance.— As discussed in Section B.1, there is no single ideal way to classify households as being insured or uninsured in a particular wave. For many of my MEPS-based analyses, it is important to have a "pure" measure of uninsured households, since the goal is to understand the average level and variability of out-of-pocket spending, and its relationship to total health care costs, of households without any health insurance. To that end, my baseline measure of health insurance status in the MEPS is an indicator of whether anyone in the family *had* health insurance coverage at any time in the last year. A correct response of "No" to this question implies that no one in the family had health insurance at any time in the last year—a "pure" uninsured household. As mentioned in Section B.1, my health insurance measure in the PSID is sometimes different, so statistics on the insured and uninsured based on these different measures are not directly comparable. This is not an issue for my analyses.

Hospitalization.— My measure of hospitalizations is an indicator of whether anyone in the family was a patient in a hospital overnight or longer at any point in the prior year *and* there is no child under one year old in the family at the time of the interview. I limit to hospitalizations in which there is no child under one year old in the family to exclude hospitalizations related to childbirth, in order to better focus on hospitalizations driven by health shocks, as in Dobkin et al. (2018). This definition conditions on a slightly different age range of any children than that in the PSID because of the different time frequencies of the PSID (every two years during my sample period) and the MEPS (every year).

Other variables.— My measure of income is a broad measure of income received in the previous calendar year, including income from social insurance and means-tested programs, so that it reflects the net risk in income accounting for all sources of income risk and insurance. The education categories that I use to create the education category dummy variables are no degree or GED only, high school degree, and college degree or above (bachelor's, doctorate, or other degree).

Outliers.— For variables judged a priori likely to have large outliers—measures of health care consumption, health care costs, health care expenditure, health care charges, and income—I use raw versions, including all outliers, when it works against me (e.g., in analyses whose key results are that out-of-pocket spending is low on average and not so variable) and winsorized versions in analyses when the goal is to estimate a relationship between different variables. In the latter case, these variables 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. I do this to avoid having the estimates unduly affected by outlier values that may be errors. I report each instance where I use the raw, unwinsorized measures in the corresponding table or figure notes.

Converting to real dollars.— All monetary variables are converted to real 2020 dollars using the CPI-U-RS.

Survey weights.— Throughout, I use MEPS family weights to ensure that the estimates reflect the experiences of the U.S. non-institutionalized population.

Summary statistics on the main estimation samples are reported in Appendix Table A2.

C Income Effects of Demand for Health Care: An Additional Force Toward Health Insurance Exacerbating Other Risks

In addition to displacing implicit health insurance, another force toward health insurance exacerbating other risks is income effects of demand for health care. To the extent that the demand for certain types of health care is greater when income is greater, or more generally when the realization of other, non-health care risks are more favorable, that is a force toward health insurance exacerbating other risks. Such demand responses, which arise naturally if certain types of health care are normal goods, are a force toward the ex post value of health insurance being greater when the realization of other, non-health care risks are more favorable. 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.

Using detailed data from the Medical Expenditure Panel Survey on health care costs and health care consumption, I find that office visits covary slightly positively with income (correlation of 0.03 among the non-elderly), consistent with office visits being a normal good, but other types of health care, such as inpatient care and prescriptions, tend to covary negatively with income (see Appendix Table A6). That certain types of health care covary positively with income is consistent with the responsiveness of health care consumption to non-health driven changes in income or liquidity found by Acemoglu et al. (2013) and Gross et al. (2020). It is also consistent with the theoretical prediction of models of optimal investment in durable goods, such as health (Grossman, 1972), that such investments tend to be more sensitive to circumstances than other forms of consumption spending (Browning and Crossley, 2009). Intuitively, utility depends largely on the stock of such a durable rather than the investment flow, so temporarily postponing investment can be a low-cost way of making ends meet when times are tight.

That other types of health care covary slightly negatively with income is consistent with there being important costs of bad health beyond health care costs, such as earnings reductions. This is consistent with a variety of evidence on the non-health care costs of bad health (e.g., see Smith (1999) for a review and Dobkin et al. (2018) for an analysis of hospitalization). Because of such costs, demand responses, although a force toward health insurance exacerbating other risks, are not strong enough alone to make health insurance exacerbate income risk on net. For that, the displacement of implicit insurance is pivotal.

D Measurement Error

Measurement error in key variables is an obvious concern for any analysis. The concern is amplified in situations like this in which the results are contrary to strong priors. While classical measurement error would tend to attenuate the results rather than bias them toward health insurance having negative insurance value, this section considers the possibility that non-classical measurement error in certain key variables might bias the results toward health insurance having negative insurance value, i.e., being worth less ex ante than its mean ex post value.

A key result in both the descriptive analysis of health spending risk and the sufficient statistic analysis of the insurance value of increases in health insurance coverage is that the correlation between out-of-pocket health spending and consumption is strongly, robustly positive. I find this result in the PSID, and Finkelstein et al. (2019a) find related results in both the PSID and the Consumer Expenditure Survey.[65] If measurement error in out-of-pocket spending and consumption were positively correlated—i.e., if positive (negative) errors in out-of-pocket

[65]Specifically, Finkelstein et al. (2019a) estimate the correlation between marginal utility and out-of-pocket spending, or the logs thereof, using a variety of specifications. In all cases the estimated correlation is negative.

spending tended to be matched to positive (negative) errors in consumption—that would be a force toward measured out-of-pocket spending and measured consumption being positively correlated.

One mechanism that could potentially generate positively-correlated measurement errors is a type of recall bias in which different respondents base their responses on different recall windows (e.g., some report how much they spent in the past month and others in the past year), *and* these recall windows are not recorded in the data. Several considerations suggest that this particular bias is not a major concern for the analysis. Most directly, the key survey questions appear to be well-protected against such a problem. In the PSID, the questions about out-of-pocket spending ask about spending during an explicit time period, either the past calendar year (in later survey waves) or the past two calendar years combined (in earlier survey waves). For example, in the 2017 wave respondents were asked, "About how much did you (and your family) pay out-of-pocket for doctor, outpatient surgery, and dental bills in 2016?" A respondent answering correctly has no scope for choosing a recall window. The questions about consumption spending are somewhat different in that they allow respondents to choose whether to report their spending per month, per year, or per other unit of time. For example, in the 2017 wave respondents were asked, "How much did you [and your family living there] spend altogether in 2016 on trips and vacations, including transportation, accommodations, and recreational expenses on trips?" Respondents can choose to report their spending per month, per year, or per other unit of time, and their chosen time unit is recorded in the data. These two sets of questions are not only designed to avoid problems from respondent choices of recall windows, they are also structured differently enough that it is hard to see how such correlated recall window bias might occur. Moreover, taking advantage of the fact that the explicit recall window for the out-of-pocket spending questions in the PSID changed during my sample time period, I find that the sufficient statistic estimates are similar across the two recall windows, with somewhat stronger results (insurance value more negative) with the one-year recall window (as would be expected given that such a window better aligns the timing of the out-of-pocket spending and consumption measures).

Several additional considerations are reassuring not only about that particular type of recall bias but also about the possible role of measurement error in the key evidence and conclusions more generally. First, the key finding that the correlation between out-of-pocket spending and consumption is positive is robust across a wide range of specifications and measures of consumption and out-of-pocket spending, in both the PSID and the Consumer Expenditure Survey. Second, a corroborating key finding, also replicated in multiple datasets, is that out-of-pocket spending and income are strongly positively correlated—enough in most cases as to make net income covary positively with out-of-pocket spending (see Appendix Tables A4 and A5 and the discussion on page 18). Because of this, even setting aside the consumption measures in the PSID and Consumer Expenditure Survey, I estimate negative insurance value based on simple consumption-proxy measures, for example using a consumption proxy of income minus out-of-pocket spending (see Appendix Table A13).[66] Third, I find that PSID measures of out-of-pocket spending match quite well the corresponding measures in MEPS, which are widely thought to be of high quality. This is true in terms of means and standard deviations (see Appendix Table A3) as well as correlations with income (see Appendix Tables A4 and A6). Fourth, making two small changes to the workhorse model of health spending risk—adding other, non-health care risk and implicit health insurance, both based on empirical evidence—causes model-predicted correlations between out-of-pocket spending and consumption, and between out-of-pocket spending and income, to be strongly positive, to an extent similar to that observed in the various datasets.

In terms of external validation beyond the PSID, MEPS, and Consumer Expenditure Survey, Ganong and Noel (2019) find that in bank account data with measures of monthly income and spending based on the universe of Chase consumer checking and credit card accounts, out-of-pocket spending on medical copays drops 17% from three months prior to receiving unemployment insurance (UI) benefits to one month before UI benefit exhaustion and a further 14% one month after exhaustion (their Table 2 on page 2400). Consumption spending and income are dropping at the same time as well. So the parts of the covariation between out-of-pocket spending and consumption, and between out-of-pocket spending and income, associated with unemployment shocks and UI benefit exhaustion exhibit positive correlations. I find the same qualitative patterns in the PSID based on unemployment and other non-health care shocks. Of course, Ganong and Noel's (2019) findings on out-of-pocket spending around unemployment and UI benefit exhaustion do not imply that the *overall* covariance across states of the world between out-of-pocket spending and consumption (or income) is positive, but they accord well with my findings in the PSID.

More generally, measurement error is unlikely to explain why such a wide range of evidence based on a variety of approaches—from descriptive evidence about health spending risk (including not only its marginal distribution but how it relates to consumption, income, and assets), to sufficient statistic estimates based on different measures of consumption and proxies of consumption, to structural analyses based on key features of the data—all points to the same, robust conclusions.

[66]On the issue of the particular type of recall bias discussed above, the key income variables in the PSID have an explicit one-year recall window, which would seem to leave little scope for that type of recall bias to affect the income-based results.

E Decomposition of Hospitalization-Related Health Insurance Targeting

Appendix Figure A6 and Appendix Table A16 report the results of a decomposition of hospitalization-related health insurance targeting. This section shows the underlying theory. Decompose the ex post value of an increase in health insurance coverage into the following components:

$$\begin{aligned} V = E(V|h=0) + (1-h)\left[V - E(V|h=0)\right] \\ + h\left[E(V|h=1) - E(V|h=0)\right] \\ + g\left[E(V|g=1) - E(V|h=1)\right] + b\left[E(V|b=1) - E(V|h=1)\right] \\ + g\left[V - E(V|g=1)\right] + b\left[V - E(V|b=1)\right], \end{aligned} \tag{29}$$

where h is an indicator of whether the head or spouse experienced a hospitalization, g is an indicator of whether the head or spouse experienced a hospitalization and household income was "good" (above a threshold value), and b is an indicator of whether the head or spouse experienced a hospitalization and household income was "bad" (below the same threshold value). The second line, which corresponds to the average net transfer from non-hospitalization to hospitalization states, is the "Non-hosp to hosp" component in Appendix Figure A6a and Appendix Table A16. The average net transfer from non-hospitalization to hospitalization states, $[E(V|h=1) - E(V|h=0)]$, is what is reported in the corresponding "Benefit" column of Appendix Table A16. The third line, which corresponds to the average net transfer from hospitalization states with bad to good income realizations, is the "Low- to high-y w/in hosp" component of Appendix Figure A6a and Appendix Table A16. The average net transfer from hospitalization states with bad to good income realizations, $[E(V|g=1) - E(V|b=1)]$, is what is reported in the corresponding "Benefit" column of Appendix Table A16. The fourth line, which corresponds to the transfers within hospitalization-by-income categories from states in which the increase in health insurance is less to more valuable, is the "Within-hosp-by-y" component of Appendix Figure A6a and Appendix Table A16. The first line, which corresponds to the value in non-hospitalization states of the change in health insurance, does not enter the analysis of hospitalization-related targeting except that the mean ex post value in non-hospitalization states helps determine the average net transfer from non-hospitalization to hospitalization states. This decomposition is similar in spirit to that of Deshpande and Lockwood (forthcoming).

Appendix Figures and Tables

Figure A1: Implicit health insurance support by education

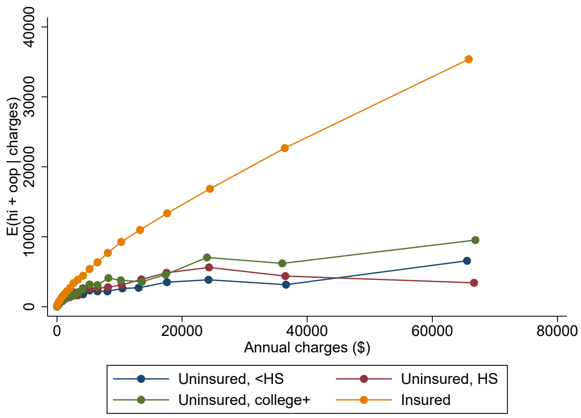

(a) Payments as a function of charges

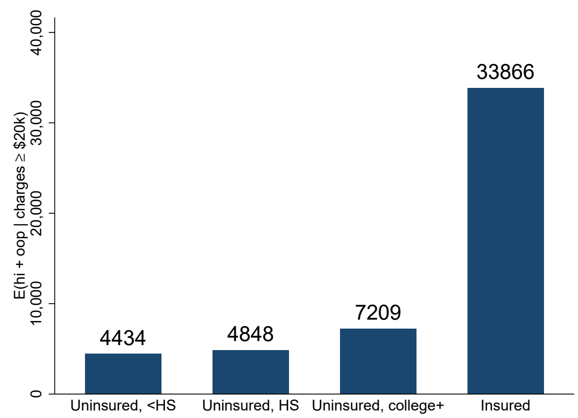

(b) Payments among HHs with charges ⩾$20k

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 health insurance (highest curve) and without health insurance by education category. This is a binned scatter plot. This figure excludes households with charges in excess of $100,000 for legibility.
Right panel: Mean of the sum of total payments by health insurers (health insurance benefits) and households (out-of-pocket health spending) among households with charges of at least $20,000.
Both panels are based on MEPS data and include all outliers, without any trimming or winsorizing.

Figure A2: Medical debt

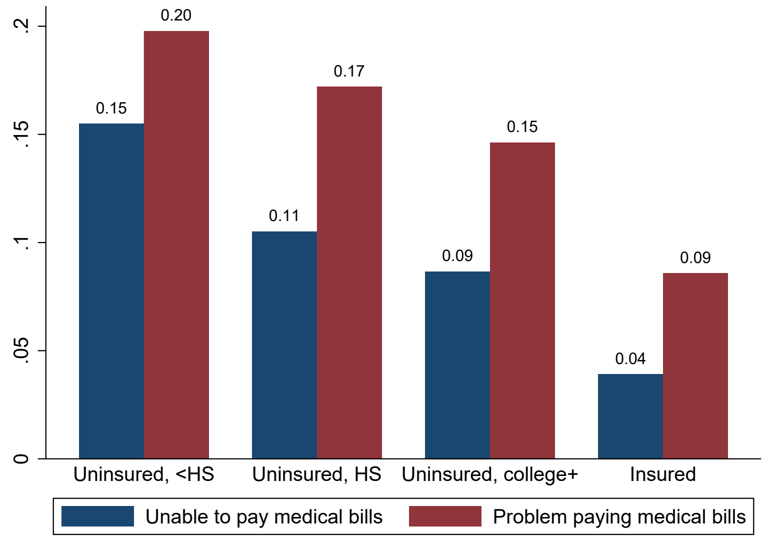

Notes: Figure shows the share of each group of non-elderly households who respond "Yes" to (i) "Does anyone in your family currently have any medical bills that you are unable to pay at all?" (which I label "Unable to pay medical bills") and, separately, (ii) "In the past 12 months did anyone in the family have problems paying or were unable to pay any medical bills?" (which I label "Problem paying medical bills"). These are based on MEPS data from 2014 on (as these variables were added to the survey in 2014). For this figure, a family is classified as "insured" only if everyone in the family had health insurance in every month of the year (in order to be a "pure" measure of being insured). (The uninsured are the usual pure measure: no one in the family had health insurance at any point during the year.)

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	75.0
Family size	2.5	2.2	2.6	1.6
Income	95,762	47,250	103,653	66,380
Consumption	47,563	32,968	49,934	31,833
Out-of-pocket health spending	1,436	1,016	1,505	2,086
Hospital	0.10	0.09	0.10	0.24
Unemployment	0.11	0.20	0.09	0.02
Sample size	73,874	11,108	62,409	11,895

Notes: Summary statistics from the Panel Study of Income Dynamics (PSID). These are family-level averages using family weights. Monetary variables are in real 2020 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 overnight or unemployed in the last year, excluding hospitalizations in which there is a child under two years old in the household (to avoid hospitalizations associated with childbirth). I use the 1999–2019 waves, which occur in every odd-numbered year. Sample size is the number of household-year observations. Note that the measure of health insurance status in the PSID differs from that in the MEPS, so the insured and uninsured groups are not directly comparable across datasets (see Appendix Section B).

Figure A3: Simulated effects of eliminating out-of-pocket health spending versus eliminating unemployment insurance on the volatility of consumption and net income

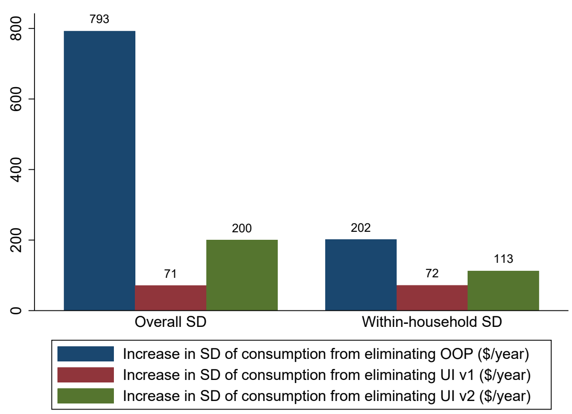

(a) Consumption

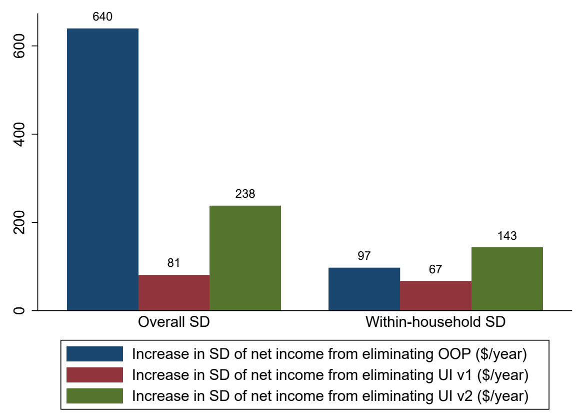

(b) Net income

Notes: Simulated effects on the overall standard deviation and the within-household standard deviation of consumption (panel (a)) and net income (income minus out-of-pocket spending) (panel (b)) of eliminating out-of-pocket spending versus eliminating unemployment insurance (UI). The simulated effect of eliminating out-of-pocket spending is to increase consumption and net income by the status quo amount of out-of-pocket spending, e.g., $c_{it}(oop = 0) = c_{it} + oop_{it}$, where $c_{it}(oop = 0)$ is counterfactual consumption if out-of-pocket spending were eliminated and c_{it} and oop_{it} are actual, observed consumption and out-of-pocket spending, respectively. The simulated effect of eliminating UI is to decrease consumption and net income by the status quo UI benefit amount, e.g., $c_{it}(UI = 0) = c_{it} - b_{it}$, where b_{it} is the UI benefit received by household i in period t under the status quo. The model of net income assumes that gross income is unchanged in response to eliminating out-of-pocket spending or UI. The model of consumption assumes that changes in out-of-pocket spending and UI benefits affect consumption one-for-one in each state. While this hand-to-mouth assumption likely overstates the *absolute* effects of these changes on consumption, the goal of this analysis is to get a sense of the *relative* effects of reducing out-of-pocket spending versus reducing UI. “UI v1” sets b_{it} to reported UI benefits received in the PSID. Unfortunately, this measure understates UI receipt by about one-third (though does not understate benefits conditional on receipt; see Meyer et al., 2015). So I also consider an alternative measure, “UI v2,” which assumes that every household in which the head or spouse was unemployed at any time during the previous year receives the average UI benefit among non-elderly households who report positive benefits: $b_{it} = unemp_{it} \times \bar{b}$, where $unemp_{it}$ is an indicator of whether the head or spouse was unemployed at any time during the previous year and the average benefit $\bar{b}$ is $4,990. Given that limitations on eligibility and incomplete take up among the eligible mean that only a minority of the unemployed receive benefits (e.g., see Kroft, 2008, on take up), “UI v2” likely overstates UI benefits significantly. Data are from the PSID. The sample is non-elderly households.

Figure A4: Markup as a function of the level of coverage

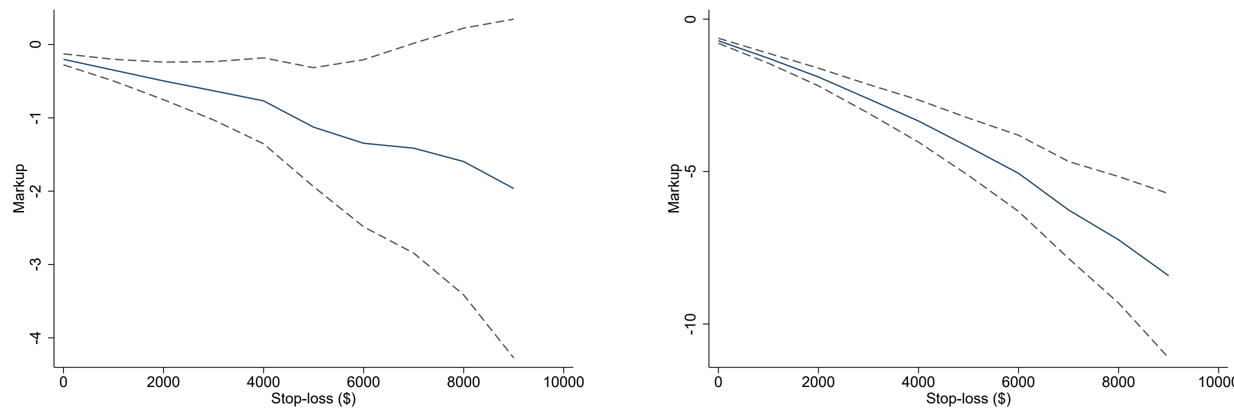

(a) Short run (b) Long run

Notes: Sufficient statistic estimates of the markup (insurance value per dollar of mean ex post value, $M = Cov\left(\widehat{\lambda}, V\right)/E(V)$) on different levels of health insurance coverage for the non-elderly uninsured. The coverage takes the form of full coverage above a stop-loss. The stop-loss amount is the x-axis. The ex post value of coverage with a stop-loss of $d \geqslant 0$ is the excess, if any, of out-of-pocket spending over d: $V = \max\{0,\ oop - d\}$. A stop-loss of \$0 corresponds to full coverage of all costs. Panel (a) (Short run) is based on regressions of within-household changes in log consumption on within-household changes in $\log(1 + V)$, plus year dummies and a cubic in age, where the changes are from one wave to the next. This aims to capture the value of coverage from the perspective of immediately before the coverage begins. Panel (b) (Long run) is based on regressions of log consumption on $\log(1 + V)$, plus year dummies, a cubic in age, and a quadratic in household size. The aim is to capture the value of coverage from behind the veil of ignorance. Neither specification enforces that the overall ex ante value be non-negative, which must be true of an expansion of health insurance coverage and which is equivalent to the markup being no less than negative one. The goal of this analysis is not to estimate the level of the markup but to understand how the markup on less extensive coverage (higher stop-loss) compares to that on more extensive coverage (lower stop-loss). An alternative specification that does enforce this restriction (based on "levels" regressions of normalized marginal utility on the ex post value and controls, which are otherwise not as well-behaved) similarly shows no tendency for less extensive coverage to have a less-negative markup than more extensive coverage. "Insurance value," $Cov\left(\widehat{\lambda}, V\right)$, is $-\gamma \times \beta \times \frac{Var(V)}{E(V)}$, where $\gamma = 3$ is the coefficient of relative risk aversion and β is the regression coefficient on the V term (see equation (8)). Dashed lines are two standard errors above and below the estimated markup. Standard errors, which are clustered at the household level, reflect sampling uncertainty in β but treat $E(V)$ and $Var(V)$ as non-stochastic. Data are from the PSID. Monetary amounts are in real 2020 dollars per household per year. Non-elderly are households whose heads are 25–64.

Figure A5: Insurance value of health insurance's hospitalization-related targeting as a function of the amount of risk that remains to be realized

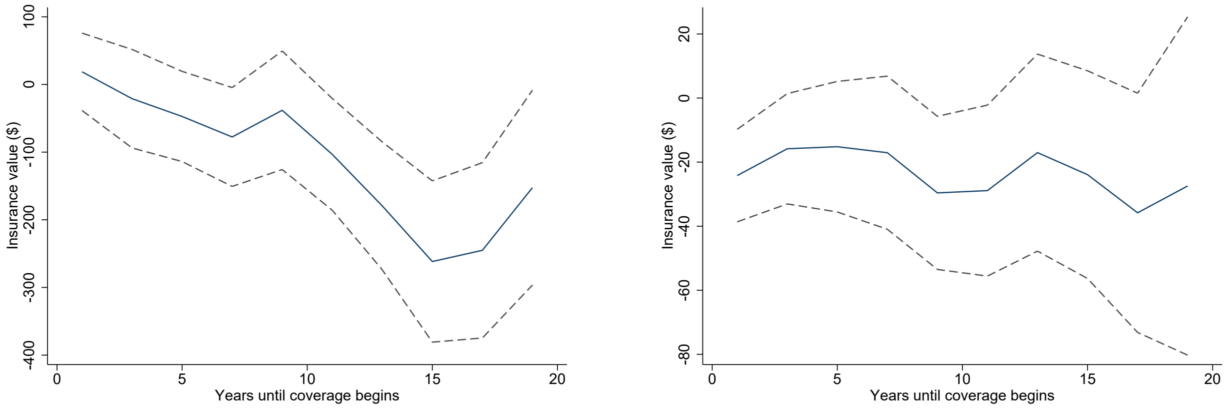

(a) Non-elderly uninsured

(b) Non-elderly insured

Notes: Insurance value of completing health insurance coverage of hospitalization risk in a particular year as a function of the length of time until the coverage begins. To isolate the targeting of hospitalization risk, I "shut down" targeting in non-hospitalization states by setting the ex post value to its average value (average out-of-pocket health spending) in non-hospitalization states (whereas in each hospitalization state it equals out-of-pocket 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 child under two years old 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. A longer time means more risk remains to be realized. The result for "y years until coverage begins" is based on a regression of the $(y+1)$-year change in log consumption on the $(y+1)$-year change in $\log(1+V)$ (i.e., from one wave to $\frac{y+1}{2}$ waves later for $y \in \{1, 3, 5, \ldots 19\}$), plus year dummies and a cubic in age. "Insurance value," $Cov\left(\widehat{\lambda}, V\right)$, is $-\gamma \times \beta \times \frac{Var(V)}{E(V)}$, where $\gamma = 3$ is the coefficient of relative risk aversion and β is the regression coefficient on the V term (see equation (8)). Dashed lines are two standard errors above and below the estimated insurance value. Standard errors, which are clustered at the household level, reflect sampling uncertainty in β but treat $E(V)$ and $Var(V)$ as non-stochastic. The corresponding "long run" insurance values to someone behind the veil are $-\$210$ and $-\$220$ for the uninsured and insured, respectively. Data are from the PSID. Monetary amounts are in real 2020 dollars per household per year. Non-elderly are households whose heads are 25–64.

Figure A6: Health insurance and hospitalization risk

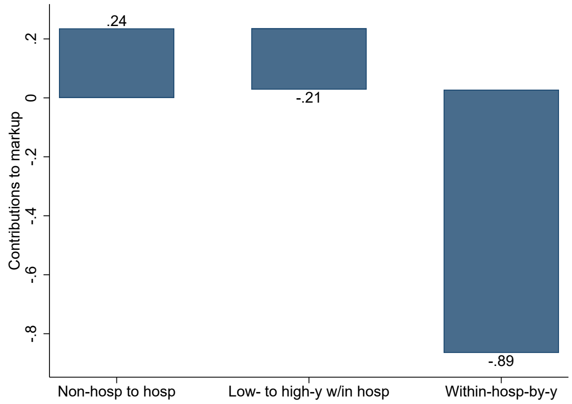

(a) HI exacerbates hospitalization risk

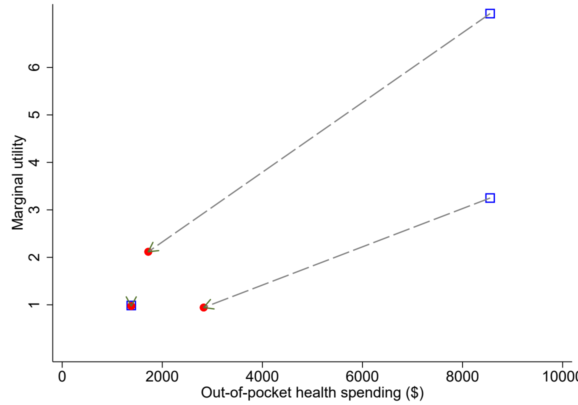

(b) Implicit HI undermines HI targeting

Notes: Left panel: Waterfall chart of the decomposition of the markup on the net transfers into hospitalization states from completing health insurance, from the perspective of an individual who knows their education but nothing else about the realization of risk, into the sum of three terms. The first ("Non-hosp to hosp") is the average net transfer from non-hospitalization to hospitalization states. The second ("Low- to high-y w/in hosp") is the 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 year dummies, a cubic in age, a quadratic in household size, and education category dummies. The third ("Within-hosp-by-y") is the transfers within hospitalization-by-income categories, from states in which completing health insurance is less to more valuable. The markup is insurance value per dollar of mean ex post value. See Appendix Section E for the underlying theory and Appendix Table A16 for related statistics. The underlying regressions are of normalized marginal utility on V (not log consumption on $\log(1+V)$), plus year dummies, a cubic in age, a quadratic in household size, and education category dummies. This "exact" specification, which aims to approximate relatively long run risk (all risk within education groups but not the risk of being in one education group as opposed to another), guarantees that the sum of the values of each component sum to the overall value, whereas specifications based on approximations to normalized marginal utility or V do not. Right panel: Average marginal utility and average ex post value of completing health insurance in three sets of states: non-hospitalization (leftmost markers), hospitalization and higher income, and hospitalization and lower income (highest set of markers), where lower versus higher income is defined as in the left panel. The red circle markers are based on observed consumption and out-of-pocket spending. The blue square markers are based on a simple model of counterfactual consumption and out-of-pocket spending were there no implicit health insurance. Counterfactual out-of-pocket spending is $oop_{\omega}^{no} = oop_{\omega} + ihi(hosp_{\omega}, y_{\omega})$, where oop_{ω} is observed out-of-pocket spending and $ihi(hosp_{\omega}, y_{\omega})$ is such that (i) average counterfactual out-of-pocket spending is the same in both sets of hospitalization states and (ii) average implicit health insurance 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 consumption is $c_{\omega}^{no} = \max\{\underline{c},\ c - ihi(hosp_{\omega}, y_{\omega})\}$. This is a hospitalization-related empirical analogue of the theoretical Figure 5, except that it is not the full scatter plot of the joint distribution of marginal utility and out-of-pocket spending (and so the slope of a regression line through these points is not necessarily proportional to insurance value). Both panels: The sample is non-elderly households in the PSID. The results are similar for the uninsured alone (see Appendix Table A16).

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	81,420	40,061	88,051	55,885
Out-of-pocket health spending	1,491	986	1,554	2,215
Health care charges	14,308	4,275	15,051	28,779
Health care payments	8,767	2,243	9,489	15,207
Hospital	0.10	0.04	0.10	0.22
Problem paying medical bills	0.11	0.17	0.10	0.07
Unable to pay medical bills	0.06	0.11	0.05	0.03
Sample size	214,083	18,144	156,371	54,152

Notes: Summary statistics from the Medical Expenditure Panel Survey (MEPS). These are family-level averages using family weights. Monetary variables are in real 2020 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. In all analyses of MEPS that control for age, I include an indicator of whether the age is the top-coded value. "Health care payments" are total annual payments, including from the insurer and the household. Hospital is an indicator of whether anyone in the household was hospitalized in the prior year and there is no child under one year old in the household (to avoid hospitalizations associated with childbirth). I use the 1996–2018 waves. Sample size is the number of household-year observations. Note that the measure of health insurance status in the MEPS differs from that in the PSID, so the insured and uninsured groups are not directly comparable across datasets (see Appendix Section B).

Table A3: Out-of-pocket health spending, total health care costs, income, and consumption

	Non-elderly			Non-elderly uninsured			Non-elderly insured			Elderly		
Panel A: MEPS including outliers	Oop	Tot	Oop/Tot	Oop	Tot	Oop/Tot	Oop	Tot	Oop/Tot	Oop	Tot	Oop/Tot
Mean	1,560	9,612	0.16	1,058	2,696	0.39	1,619	10,126	0.16	2,379	16,034	0.15
Standard deviation	2,866	23,033	0.12	2,724	11,963	0.23	2,828	21,916	0.13	4,007	24,593	0.16
95th percentile	5,853	36,524	0.16	4,696	10,510	0.45	5,914	37,115	0.16	7,749	57,708	0.13
99th percentile	11,641	92,476	0.13	11,455	46,787	0.24	11,443	90,078	0.13	16,177	114,182	0.14
Panel B: PSID including outliers	Oop	Income	Consump	Oop	Income	Consump	Oop	Income	Consump	Oop	Income	Consump
Mean	1,506	95,762	48,486	1,126	47,250	33,235	1,570	103,653	50,961	2,378	66,380	32,493
Standard deviation	3,214	134,877	37,491	3,626	59,760	29,540	3,139	141,874	38,051	6,602	87,589	33,568
Within standard deviation	2,541	66,439	20,358	2,027	32,442	9,085	2,490	69,743	20,563	4,933	46,835	27,784
Within standard deviation, 2-wave	1,869	28,524	9,969	1,197	28,560	5,584	1,665	26,382	10,119	2,271	28,504	9,104
Panel C: PSID winsorized	Oop	Income	Consump	Oop	Income	Consump	Oop	Income	Consump	Oop	Income	Consump
Mean	1,436	90,908	47,563	1,016	46,538	32,968	1,505	98,123	49,934	2,086	64,137	31,833
Standard deviation	2,327	79,501	30,908	2,214	44,832	21,546	2,339	81,484	31,533	2,905	64,924	23,902
Within standard deviation	1,597	34,914	15,619	1,174	18,880	9,039	1,587	35,860	15,991	1,881	30,654	12,375
Within standard deviation, 2-wave	1,225	19,164	8,708	813	11,212	5,510	1,168	18,653	8,538	1,400	21,575	8,446

Notes: Statistics on out-of-pocket health spending (Oop), total health care costs (Tot), income, and consumption among different types of households. All variables are measured in real 2020 dollars per year. Panel A uses MEPS data and includes all outliers, without any trimming or winsorizing. Total health care costs are defined as follows. For households with health insurance, total costs are total annual payments, including from the insurer and the household. For households without health insurance, total costs are annual charges scaled by the payments-charge ratio among non-elderly households with health insurance. Panel B uses PSID data and includes all outliers, without any trimming or winsorizing. Panel C uses PSID data and winsorizes each variable at its (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. "Within standard deviation" is the within-household standard deviation among households appearing in any of the eleven waves of the PSID from 1999–2019. The average number of waves in which a non-elderly household appears (as a non-elderly household) is 5.0. The average number of waves in which an elderly household appears (as an elderly household) is 4.0. "Within standard deviation, 2-wave" is the within-household standard deviation in the two waves of the PSID from 2017–2019 among households appearing in both of those waves. Note that the measure of health insurance status in the MEPS differs from that in the PSID, so the insured and uninsured groups are not directly comparable across datasets. See Appendix Section B.

Table A4: Out-of-pocket health spending hedges income risk

	Non-elderly	Non-elderly uninsured	Non-elderly insured	Elderly insured
	(1)	(2)	(3)	(4)
$\widehat{\beta}_{\log(y)\mid\log(oop)}$	0.036	0.032	0.033	0.005
(se)	(0.003)	(0.008)	(0.003)	(0.006)
Corr(log(y), log(oop))	0.12	0.09	0.12	0.02
Implied $\widehat{\beta}_{y\mid oop}$	2.26	1.48	2.12	0.15
(se)	(0.21)	(0.35)	(0.22)	(0.19)
Implied Corr(y, oop)	0.10	0.10	0.10	0.01

Notes: Results from regressions of the log of income on the log of one plus out-of-pocket spending and household fixed effects, year dummies, and a cubic in age for each of four sets of states: non-elderly, non-elderly uninsured, non-elderly insured, and elderly insured. Given the coverage of the panel, these fixed effects regressions capture risk between the short run (one year) and medium run (ten year) perspectives discussed in Section 3 and reported in Appendix Table A5. The first row shows the coefficient estimate on the log of one plus out-of-pocket spending. The second row is the corresponding standard error, which is clustered at the household level. The third row is the correlation between the log of income and the log of one plus out-of-pocket spending, both residualized with household fixed effects, year dummies, and a cubic in age. The fourth row is the implied slope of income with respect to out-of-pocket spending, evaluated at the means of income and out-of-pocket spending: $\widehat{\beta}_{y\mid oop} \equiv \widehat{\beta}_{\log(y)\mid\log(oop)} \times \frac{E(y)}{E(oop)}$. The fifth row is its standard error, which is the product of the standard error in the second row and $\frac{E(y)}{E(oop)}$. The sixth row is the implied correlation between income and out-of-pocket spending, defined as the product of the implied slope of income with respect to out-of-pocket spending, $\widehat{\beta}_{y\mid oop}$, and the ratio of the standard deviation of out-of-pocket spending to the standard deviation of income, each residualized with household fixed effects, year dummies, and a cubic in age. Data are from the PSID.
If $\widehat{\beta}_{y\mid oop} > 0.5$, the variance of net income (income net of out-of-pocket spending) is smaller than that of gross income (see page 18). If $\widehat{\beta}_{y\mid oop} > 1$, out-of-pocket spending covaries positively not only with income but even with net income. Footnote 23 on page 13 discusses the considerable income risk faced by the elderly.

Table A5: Out-of-pocket health spending hedges income risk: Heterogeneity and robustness

	Non-elderly			Non-elderly uninsured			Non-elderly insured			Elderly insured		
	Short run (1)	Medium run (2)	Long run (3)	Short run (4)	Medium run (5)	Long run (6)	Short run (7)	Medium run (8)	Long run (9)	Short run (10)	Medium run (11)	Long run (12)
$\widehat{\beta}_{\log(y)\mid\log(oop)}$	0.016	0.048	0.132	0.021	0.049	0.088	0.014	0.044	0.127	0.006	0.010	0.104
(se)	(0.002)	(0.004)	(0.003)	(0.005)	(0.009)	(0.006)	(0.002)	(0.004)	(0.004)	(0.004)	(0.006)	(0.006)
Corr(log(y), log(oop))	0.06	0.16	0.36	0.06	0.15	0.25	0.05	0.15	0.36	0.02	0.03	0.30
Implied $\widehat{\beta}_{y\mid oop}$	1.03	3.04	8.38	0.94	2.26	4.03	0.94	2.90	8.27	0.17	0.29	3.21
(se)	(0.14)	(0.26)	(0.22)	(0.25)	(0.42)	(0.27)	(0.14)	(0.28)	(0.23)	(0.13)	(0.18)	(0.18)
Implied Corr(y, oop)	0.03	0.09	0.25	0.05	0.11	0.21	0.03	0.08	0.24	0.01	0.01	0.15

Notes: Results from regressions of income variables on out-of-pocket spending variables for each of four sets of states: non-elderly, non-elderly uninsured, non-elderly insured, and elderly insured. This is a supporting table to Appendix Table A4. Short run and medium run columns are based on regressions of within-household changes in log income on within-household changes in the log of one plus out-of-pocket spending, plus year dummies and a cubic in age, where the changes are from one wave to the next (short run) or from one wave to five waves later (medium run). Long run is based on regressions of log income on the log of one plus out-of-pocket spending, plus year dummies, a cubic in age, and a quadratic in household size. Short run aims to capture the income-out-of-pocket spending relationship from the perspective of immediately before the coverage begins, medium run from ten years before the coverage begins, and long run from behind the veil of ignorance. The first row shows the coefficient estimate on the log of one plus out-of-pocket spending. The second row is the corresponding standard error, which is clustered at the household level. The third row is the correlation between the log of income and the log of one plus out-of-pocket spending, both residualized with household fixed effects, year dummies, and a cubic in age. The fourth row is the implied slope of income with respect to out-of-pocket spending, evaluated at the means of income and out-of-pocket spending: $\widehat{\beta}_{y|oop} \equiv \widehat{\beta}_{\log(y)|\log(oop)} \times \frac{E(y)}{E(oop)}$. The fifth row is its standard error, which is the product of the standard error in the second row and $\frac{E(y)}{E(oop)}$. The sixth row is the implied correlation between income and out-of-pocket spending, defined as the product of the implied slope of income with respect to out-of-pocket spending, $\widehat{\beta}_{y|oop}$, and the ratio of the standard deviation of out-of-pocket spending to the standard deviation of income, each residualized with household fixed effects, year dummies, and a cubic in age. Data are from the PSID.

If $\widehat{\beta}_{y|oop} > 0.5$, the variance of net income (income net of out-of-pocket spending) is smaller than that of gross income (see page 18). If $\widehat{\beta}_{y|oop} > 1$, out-of-pocket spending covaries positively not only with income but even with net income. Footnote 23 on page 13 discusses the considerable income risk faced by the elderly.

Table A6: Correlation between income and various measures of health care utilization and spending

	Non-elderly			Elderly
	All	Uninsured	Insured	
Charges				
Total	-0.02	-0.01	-0.03	-0.01
Office visits	0.00	0.01	0.00	0.01
Outpatient hospital	0.00	0.01	0.00	0.00
Outpatient doctor	0.00	0.02	0.00	0.00
Inpatient	-0.03	-0.02	-0.03	-0.02
Quantities				
Office visits	0.03	0.06	0.01	0.04
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	0.12	0.10	0.12	0.08

Notes: Correlation between income and various measures of health care utilization and spending for each of four samples: non-elderly, non-elderly uninsured, non-elderly insured, and elderly insured. All variables are residualized with year dummies, a cubic in age, a quadratic in household size, education category dummies, and an indicator of whether age is at the top code. The aim of the residualization is to approximate relatively long run risk: all risk within education groups but not the risk of being in one education group versus another. Data are from the MEPS.

Table A7: Out-of-pocket health spending hedges unemployment risk

	All non-elderly			Uninsured			Insured		
	Short run (1)	Medium run (2)	Long run (3)	Short run (4)	Medium run (5)	Long run (6)	Short run (7)	Medium run (8)	Long run (9)
$\widehat{\beta}_{oop\|ue}$	-112	-294	-331	-76	-177	-232	-112	-249	-303
(se)	(39)	(70)	(31)	(74)	(132)	(61)	(47)	(85)	(36)
$\widehat{\beta}_{c\|ue}$	-2,385	-6,124	-9,685	-1,134	-2,544	-5,513	-2,596	-5,970	-8,983
(se)	(315)	(682)	(380)	(538)	(1195)	(549)	(386)	(829)	(468)

Notes: Results from regressions of out-of-pocket spending (first two rows) and consumption (last two rows) on an indicator of unemployment (ue) and controls in each of three sets of states: non-elderly, non-elderly uninsured, and non-elderly insured. Each coefficient estimate is from a separate regression, with the corresponding standard error in parentheses below. Short run and medium run columns are based on regressions of within-household changes in the dependent variable on within-household changes in ue, plus year dummies and a cubic in age, where the changes are from one wave to the next (short run) or from one wave to five waves later (medium run). Long run is based on regressions of the dependent variable on ue, plus year dummies, a cubic in age, and a quadratic in household size. Short run aims to capture the relationship from the perspective of immediately before the coverage begins, medium run from ten years before the coverage begins, and long run from behind the veil of ignorance. The indicator of unemployment is a dummy variable equal to one if the household head or spouse experienced an unemployment spell in the previous year and zero otherwise. Data are from the PSID. Standard errors are clustered at the household level.

Table A8: Out-of-pocket health spending hedges consumption risk

	Non-elderly	Non-elderly uninsured	Non-elderly insured	Elderly insured
	(1)	(2)	(3)	(4)
$\widehat{\beta}_{\log(c)\|\log(oop)}$	0.021	0.017	0.020	0.013
(se)	(0.002)	(0.004)	(0.002)	(0.004)
Corr(log(c), log(oop))	0.12	0.10	0.11	0.06
Implied $\widehat{\beta}_{c\|oop}$	0.70	0.54	0.66	0.20
(se)	(0.05)	(0.12)	(0.06)	(0.07)
Implied Corr(c, oop)	0.08	0.08	0.07	0.03

Notes: Results from regressions of the log of consumption on the log of one plus out-of-pocket spending and household fixed effects, year dummies, and a cubic in age for each of four sets of states: non-elderly, non-elderly uninsured, non-elderly insured, and elderly insured. Given the coverage of the panel, these fixed effects regressions capture risk between the short run (one year) and medium run (ten year) perspectives discussed in Section 3 and reported in Appendix Table A9. The first row shows the coefficient estimate on the log of one plus out-of-pocket spending. The second row is the corresponding standard error, which is clustered at the household level. The third row is the correlation between the log of consumption and the log of one plus out-of-pocket spending, both residualized with household fixed effects, year dummies, and a cubic in age. The fourth row is the implied slope of consumption with respect to out-of-pocket spending, evaluated at the means of consumption and out-of-pocket spending: $\widehat{\beta}_{c|oop} \equiv \widehat{\beta}_{\log(c)|\log(oop)} \times \frac{E(c)}{E(oop)}$. The fifth row is its standard error, which is the product of the standard error in the second row and $\frac{E(c)}{E(oop)}$. The sixth row is the implied correlation between consumption and out-of-pocket spending, defined as the product of the implied slope of consumption with respect to out-of-pocket spending, $\widehat{\beta}_{c|oop}$, and the ratio of the standard deviation of out-of-pocket spending to the standard deviation of consumption, each residualized with household fixed effects, year dummies, and a cubic in age. Data are from the PSID.

Table A9: Out-of-pocket health spending hedges consumption risk: Heterogeneity and robustness

	Non-elderly			Non-elderly uninsured			Non-elderly insured			Elderly insured		
	Short run	Medium run	Long run	Short run	Medium run	Long run	Short run	Medium run	Long run	Short run	Medium run	Long run
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
$\widehat{\beta}_{\log(c)\mid\log(oop)}$	0.010	0.029	0.071	0.014	0.030	0.050	0.008	0.026	0.070	0.007	0.016	0.065
(se)	(0.001)	(0.002)	(0.002)	(0.003)	(0.005)	(0.003)	(0.001)	(0.002)	(0.002)	(0.003)	(0.004)	(0.004)
Corr(log(c), log(oop))	0.06	0.16	0.34	0.09	0.17	0.28	0.05	0.14	0.32	0.03	0.08	0.27
Implied $\widehat{\beta}_{c\mid oop}$	0.32	0.95	2.36	0.46	0.98	1.62	0.27	0.88	2.31	0.10	0.25	0.99
(se)	(0.04)	(0.07)	(0.06)	(0.09)	(0.16)	(0.09)	(0.04)	(0.07)	(0.07)	(0.05)	(0.06)	(0.06)
Implied Corr(c, oop)	0.02	0.07	0.19	0.05	0.10	0.19	0.02	0.07	0.18	0.01	0.03	0.13

Notes: Results from regressions of consumption variables on out-of-pocket spending variables for each of four sets of states: non-elderly, non-elderly uninsured, non-elderly insured, and elderly insured. This is a supporting table to Appendix Table A8. Short run and medium run columns are based on regressions of within-household changes in log consumption on within-household changes in the log of one plus out-of-pocket spending, plus year dummies and a cubic in age, where the changes are from one wave to the next (short run) or from one wave to five waves later (medium run). Long run is based on regressions of log consumption on the log of one plus out-of-pocket spending, plus year dummies, a cubic in age, and a quadratic in household size. Short run aims to capture the consumption-out-of-pocket spending relationship from the perspective of immediately before the coverage begins, medium run from ten years before the coverage begins, and long run from behind the veil of ignorance. The first row shows the coefficient estimate on the log of one plus out-of-pocket spending. The second row is the corresponding standard error, which is clustered at the household level. The third row is the correlation between the log of consumption and the log of one plus out-of-pocket spending, both residualized with household fixed effects, year dummies, and a cubic in age. The fourth row is the implied slope of consumption with respect to out-of-pocket spending, evaluated at the means of consumption and out-of-pocket spending: $\widehat{\beta}_{c|oop} \equiv \widehat{\beta}_{\log(c)|\log(oop)} \times \frac{E(c)}{E(oop)}$. The fifth row is its standard error, which is the product of the standard error in the second row and $\frac{E(c)}{E(oop)}$. The sixth row is the implied correlation between consumption and out-of-pocket spending, defined as the product of the implied slope of consumption with respect to out-of-pocket spending, $\widehat{\beta}_{c|oop}$, and the ratio of the standard deviation of out-of-pocket spending to the standard deviation of consumption, each residualized with household fixed effects, year dummies, and a cubic in age. Data are from the PSID.

Table A10: Hospitalization risk

	Non-elderly						Non-elderly uninsured					
	Short run			Long run			Short run			Long run		
	Any	Better	Worse	Any	Better	Worse	Any	Better	Worse	Any	Better	Worse
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
$\widehat{\beta}_{oop\|hosp}$	970	1,079	637	1,104	1,577	-328	1,233	1,591	97	1,447	1,951	-59
(se)	(79)	(89)	(165)	(59)	(71)	(87)	(246)	(293)	(371)	(183)	(223)	(229)
$\widehat{\beta}_{c\|hosp}$	-266	1,049	-4,214	-5,496	800	-24,491	-468	1,388	-6,036	-2,666	1,112	-14,152
(se)	(451)	(501)	(921)	(475)	(541)	(549)	(926)	(980)	(1,959)	(853)	(935)	(1,318)
$\widehat{\beta}_{y\|hosp}$	-1,217	12,684	-42,945	-18,183	2,105	-79,391	68	11,453	-34,086	-8,414	2,977	-43,037
(se)	(938)	(808)	(1,873)	(1,228)	(1,387)	(725)	(2,171)	(2,325)	(2,400)	(2,011)	(2,250)	(1,336)

Notes: Results from regressions of out-of-pocket health spending, consumption, and income (*oop*, c, and y in the row names) (long run columns) or their within-household first differences (short run columns) on either (i) an indicator for any hospitalization (columns labeled "Any") or (ii) two indicators, one for the subset of hospitalizations in which residualized income (long run columns) or its first difference (short run) is above its 25th percentile in hospitalization states ("Better") and one for the subset in which it is below ("Worse"). The hospitalization indicator equals one if the head or spouse experienced at least one hospitalization and there are no children under two years old (to exclude hospitalizations related to childbirth). The reported coefficient in the "Any" column is the coefficient on the hospitalization indicator. The reported coefficient in the "Better" column is the coefficient on the indicator for the subset of hospitalizations in which residualized income (or its first difference) is at least the 25th percentile in hospitalization states. The reported coefficient in the "Worse" column is that on the indicator for the subset of hospitalizations in which residualized income (or its first difference) is below its 25th percentile in hospitalization states. So for a given specification (short run or long run) and population (non-elderly or uninsured) and outcome variable (row names), the estimate in the "Any" column comes from one regression and the estimates in the "Better" and "Worse" columns come from a second regression. Short run specifications restrict the sample to household-waves in which the household did not experience a hospitalization in the previous wave. Data are from the PSID. Standard errors are clustered at the household level. Consumption and income are higher in "Better" states than in non-hospitalization states because conditioning on not having a bottom-quartile income realization within hospitalization states conditions out not only the worst income losses from hospitalization but also the worst income losses from other risks that happen to occur in a hospitalization state. The positive effect on average consumption and income from conditioning out these other risks dominates the negative effect from conditioning on experiencing a ("Better") hospitalization. This is what would be expected as long as other risks can have big effects on consumption and income and are not too positively correlated with hospitalization.

Table A11: Sufficient statistic estimates: Heterogeneity across health care types

	Total (1)	Hospital (2)	Doctor (3)	Rx (4)
Corr(log(c),log(oop))	0.09	0.06	0.07	0.07
(se)	(0.017)	(0.014)	(0.016)	(0.015)
Insurance value	-205	-107	-80	-55
(se)	(38)	(26)	(19)	(12)
Mean ex post value	1,016	284	454	249
Markup	-0.20	-0.38	-0.18	-0.22

Notes: Statistics related to the short run value of health insurance coverage of different types of health care for non-elderly uninsured households. Column (1) reproduces the main estimates of the short run value of coverage of all three types of health care to non-elderly uninsured households (see Table 1). Columns (2)–(4) show the short run value of coverage of each of the three sub-component types of health care: hospital care ("Hospital"), doctor/outpatient surgery/dental ("Doctor"), and prescriptions/in-home medical care/special facilities/other services ("Rx"). The ex post value V is out-of-pocket spending on the type of health care given by the column header. These are based on regressions of within-household changes in log consumption on within-household changes in $\log(1+V)$, plus year dummies and a cubic in age, where the changes are from one wave to the next. The aim is to capture the value of coverage from the perspective of immediately before the coverage begins. Insurance value is more negative from longer-run perspectives (see Table 1 and Appendix Figure 4). Corr(log(c), log(oop)) is the correlation between the change in log consumption and the change in $\log(1+oop)$, both residualized with the corresponding controls. "Insurance value," $Cov\left(\widehat{\lambda}, V\right)$, is $-\gamma\times\beta\times\frac{Var(V)}{E(V)}$, where $\gamma=3$ is the coefficient of relative risk aversion and β is the regression coefficient on the V term (see equation (8)). "Markup" is insurance value per dollar of mean ex post value, $Cov\left(\widehat{\lambda}, V\right)/E\left(V\right)$. Standard errors, which are clustered at the household level, reflect sampling uncertainty in β but treat $E(V)$ and $Var(V)$ as non-stochastic. Data are from the PSID. Monetary amounts are in real 2020 dollars per household per year. Non-elderly are households whose heads are 25–64.

Table A12: Sufficient statistic estimates: Heterogeneity across education groups and states

	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)
Corr(log(c),log(oop))	0.09	0.11	0.09	0.06	0.12	0.09	0.09	0.08	0.10	0.08	0.13
(se)	(0.017)	(0.035)	(0.034)	(0.029)	(0.050)	(0.020)	(0.028)	(0.025)	(0.023)	(0.020)	(0.035)
Insurance value	-205	-287	-205	-124	-250	-215	-177	-161	-234	-181	-307
(se)	(38)	(95)	(80)	(61)	(108)	(47)	(58)	(53)	(52)	(43)	(83)
Mean ex post value	1,016	899	934	929	1,267	772	1,220	790	1,192	983	1,277
Markup	-0.20	-0.32	-0.22	-0.13	-0.20	-0.28	-0.15	-0.20	-0.20	-0.18	-0.24

Notes: Statistics related to the short run value of health insurance coverage for different education groups and in different subsets of non-elderly uninsured states. Column (1) reproduces the main estimates of the short run value of coverage in all non-elderly uninsured states (see Table 1). Columns (2)–(5) show heterogeneity across education groups in the value of health insurance in non-elderly uninsured states. Columns (6)–(11) show heterogeneity across different subsets of non-elderly uninsured states in the value of health insurance in those states. These values are based on willingness to pay out of income in the relevant states for full coverage in those states. Columns (6) and (7) split non-elderly uninsured states into two sets: 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 $3,120 and about 30% of non-elderly households have less than or equal to $500 worth of this measure of liquidity. Lagged liquidity is liquidity in the preceding wave. Columns (8) and (9) split non-elderly uninsured states into two sets: those in which the household head's age is 25–39 or 40–64. Columns (10) and (11) split non-elderly uninsured states into two sets: those in which the household head's self-reported health status is (i) "excellent," "very good," or "good" ("Good") or (ii) "fair" or "poor" ("Bad"). These are based on regressions of within-household changes in log consumption on within-household changes in $\log(1+oop)$, plus year dummies and a cubic in age, where the changes are from one wave to the next. The aim is to capture the value of coverage from the perspective of immediately before the coverage begins. Insurance value is more negative from longer-run perspectives (see Table 1 and Appendix Figure 4). Corr(log(c), log(oop)) is the correlation between the change in log consumption and the change in $\log(1+oop)$, both residualized with the corresponding controls. "Insurance value," $Cov\left(\widehat{\lambda}, V\right)$, is $-\gamma \times \beta \times \frac{Var(V)}{E(V)}$, where $\gamma = 3$ is the coefficient of relative risk aversion and β is the regression coefficient on the V term (see equation (8)). "Markup" is insurance value per dollar of mean ex post value, $Cov\left(\widehat{\lambda}, V\right)/E\left(V\right)$. Standard errors, which are clustered at the household level, reflect sampling uncertainty in β but treat $E(V)$ and $Var(V)$ as non-stochastic. Data are from the PSID. Monetary amounts are in real 2020 dollars per household per year. Non-elderly are households whose heads are 25–64.

Table A13: Sufficient statistic estimates: Robustness to assumptions about marginal utility

	Baseline	Log utility	Food consumption	Consumption proxy $c = y - oop$	State-dependent utility			
					50% lower if health bad	50% higher if health bad	50% lower if hosp=1	50% higher if hosp=1
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Corr(log(c),log(oop))	0.09	0.09	0.04	0.03	0.09	0.09	0.11	0.09
(se)	(0.017)	(0.017)	(0.018)	(0.017)	(0.017)	(0.017)	(0.017)	(0.017)
Insurance value	-205	-68	-122	-96	-211	-202	-263	-201
(se)	(38)	(13)	(51)	(58)	(39)	(38)	(40)	(40)
Mean ex post value	1,016	1,016	1,016	1,016	1,015	1,015	1,042	1,042
Markup	-0.20	-0.07	-0.12	-0.09	-0.21	-0.20	-0.25	-0.19

Notes: Statistics related to the short run value of full health insurance coverage for non-elderly uninsured households under different assumptions about marginal utility. Column (1) reproduces the main short run results for non-elderly uninsured households (see Table 1). Column (2) uses a coefficient of relative risk aversion of one (log utility) rather than the baseline value of three. Column (3) assumes that marginal utility is a function of food consumption rather than total non-health consumption. Column (4) assumes that marginal utility is a function of the "consumption proxy" of income less out-of-pocket spending with a floor, $c = \max\{\$5,000,\ y - oop\}$. Columns (5)–(8) make different assumptions about state-dependent utility. Column (5) assumes that the marginal utility of a given level of consumption is 50% lower if the household head's self-reported health status is "fair" or "poor" (rather than "excellent," "very good," or "good"), whereas column (6) assumes that marginal utility is 50% higher in those states. Column (7) assumes that the marginal utility of a given level of consumption is 50% lower if the household head or spouse experiences a hospitalization and there is no child under two years old (to exclude hospitalizations related to childbirth), whereas column (8) assumes that marginal utility is 50% higher in those states. These are meant to be relatively extreme assumptions about the extent of state-dependent utility. As a benchmark, Finkelstein et al. (2013) estimate that a one-standard deviation increase in the number of chronic diseases is associated with a 10%–25% decrease in marginal utility. State-dependent utility makes relatively little difference because bad health is only weakly related to out-of-pocket spending (correlation of 0.02) and hospitalization is only weakly related to consumption (correlation of -0.02). For additional evidence on hospitalization, see the analysis at the end of Section 5.2. These are based on regressions of within-household changes in log consumption on within-household changes in $\log(1 + oop)$, plus year dummies and a cubic in age, where the changes are from one wave to the next. The aim is to capture the value of coverage from the perspective of immediately before the coverage begins. Insurance value is more negative from longer-run perspectives (see Table 1 and Appendix Figure 4). Corr(log(c), log(oop)) is the correlation between the change in log consumption and the change in $\log(1 + oop)$, both residualized with the corresponding controls. "Insurance value," $Cov\left(\widehat{\lambda}, V\right)$, is $-\gamma \times \beta \times \frac{Var(V)}{E(V)}$, where γ is the coefficient of relative risk aversion and β is the regression coefficient on the V term (see equation (8)). The coefficient of relative risk aversion is three except in column (2), in which it is one. "Markup" is insurance value per dollar of mean ex post value, $Cov\left(\widehat{\lambda}, V\right)/E\left(V\right)$. Standard errors, which are clustered at the household level, reflect sampling uncertainty in β but treat $E(V)$ and $Var(V)$ as non-stochastic. Data are from the PSID. Monetary amounts are in real 2020 dollars per household per year. Non-elderly are households whose heads are 25–64.

Table A14: Sufficient statistic estimates: Robustness to regression specification

	Baseline	Control for quintic in income	Regress $\Delta\widehat{\lambda}$ on ΔV	Regress $\Delta \log \widehat{\lambda}$ on ΔV	Fixed effects	Fixed effects longer run
	(1)	(2)	(3)	(4)	(5)	(6)
Corr(x, y)	0.09	0.09	0.02	0.06	0.08	0.12
(se)	(0.017)	(0.017)	(0.014)	(0.016)	(0.018)	(0.028)
Insurance value	-205	-203	-172	-202	-237	-368
(se)	(38)	(38)	(103)	(52)	(50)	(86)
Mean ex post value	1,016	1,044	1,016	1,016	1,020	1,149
Markup	-0.20	-0.19	-0.17	-0.20	-0.23	-0.32

Notes: Statistics related to the short run (columns (1)–(4)) and medium run (columns (5)–(6)) value of full health insurance coverage for non-elderly uninsured households under different assumptions. Column (1) reproduces the main short run results for non-elderly uninsured households (see Table 1). These are based on regressions of within-household changes in log consumption on within-household changes in $\log(1 + oop)$, plus year dummies and a cubic in age, where the changes are from one wave to the next. The aim is to capture the value of coverage from the perspective of immediately before the coverage begins. Insurance value is more negative from longer-run perspectives (see Table 1 and Appendix Figure 4). Column (2) adds a quintic in income to the controls. Column (3) is based on regressions of within-household first differences in normalized marginal utility on within-household first differences in out-of-pocket spending and year dummies and a cubic in age. Column (4) is based on regressions of within-household first differences in the log of normalized marginal utility on within-household first differences in out-of-pocket spending and year dummies and a cubic in age. Column (5) is based on regressions of the log of consumption on out-of-pocket spending and household fixed effects, year dummies, and a cubic in age. Given the coverage of the panel, this should capture risk between the short run (one year) and medium run (ten year) perspectives discussed in Section 3 and so somewhat longer-term risk than is captured by columns (1)–(4). Column (6) is based on the same regression specification as in column (5) but limits the sample to the subset of households who are tracked continuously throughout the entire sample period from 1999–2019 inclusive. This specification therefore captures longer-term risk than is captured by the other columns in this table (though not as long as that captured by the long-run columns in Table 1). Corr(x, y) is the correlation between the dependent variable and the key independent variable (the ex post value variable), both residualized with that column's controls. "Insurance value," $Cov\left(\widehat{\lambda}, V\right)$, is $-\gamma \times \beta \times \frac{Var(V)}{E(V)}$, where $\gamma = 3$ is the coefficient of relative risk aversion, β is the regression coefficient on the out-of-pocket spending term, and $V = oop$ (see equation (8)). "Markup" is insurance value per dollar of mean ex post value, $Cov\left(\widehat{\lambda}, V\right)/E\left(V\right)$. Standard errors, which are clustered at the household level, reflect sampling uncertainty in β but treat $E(V)$ and $Var(V)$ as non-stochastic. Data are from the PSID. Monetary amounts are in real 2020 dollars per household per year. Non-elderly are households whose heads are 25–64.

Table A15: Sufficient statistic estimates: Robustness to large private values of improved health and reduced medical debt

	Baseline (1)	New cancer diagnosis (2)	Ever cancer diagnosis (3)	Health much worse (4)	Health newly bad (5)	Health bad (6)	Hospitalization (7)	Medical bills (8)	Medical bills up to $10k (9)
Corr(log(c),log(V))	0.09	0.09	0.09	0.07	0.06	0.04	0.08	0.04	0.05
(se)	(0.017)	(0.017)	(0.017)	(0.017)	(0.018)	(0.018)	(0.018)	(0.024)	(0.023)
Insurance value	-205	-330	-535	-372	-295	-243	-496	-4,270	-143
(se)	(38)	(64)	(99)	(85)	(89)	(98)	(110)	(2,583)	(67)
Mean ex post value	1,016	1,269	1,907	1,685	2,586	5,225	2,646	4,119	1,889
Markup	-0.20	-0.26	-0.28	-0.22	-0.11	-0.05	-0.19	-1.04	-0.08

Notes: Statistics related to the short run value of full health insurance coverage for non-elderly uninsured households under different assumptions about the value of improved health and reduced medical debt. Column (1) reproduces the main short run results for non-elderly uninsured households (see Table 1). These are based on regressions of within-household changes in log consumption on within-household changes in $\log(1 + oop)$, plus year dummies and a cubic in age, where the changes are from one wave to the next. The aim is to capture the value of coverage from the perspective of immediately before the coverage begins. Insurance value is more negative from longer-run perspectives (see Table 1 and Appendix Figure 4). Columns (2)–(7) increase the ex post value of health insurance V by \$20,000 in the states given by the column header. The aim is to overstate any additional ex post value of health insurance to the household, over and above that from reduced out-of-pocket spending, from improved health (from moral hazard). The estimated insurance values remain significantly negative even when V is increased by \$100,000 in these states. "Health much worse" is a dummy that equals one if either the household head or spouse reports that their health is "much worse" than it was two years ago (as opposed to "better," "about the same," or "somewhat worse"). "Health newly bad" is a dummy that equals one if the household head reports that their health is "fair" or "poor" (as opposed to "excellent," "very good," or "good") after reporting that it was "excellent," "very good," or "good" in the previous wave. "Health bad" is a dummy that equals one if the household head reports that their health is "fair" or "poor." Column (8) adds the amount of the household's outstanding medical bills to the ex post value of health insurance. Column (9) adds the lesser of this amount and \$10,000. Both aim to overstate any additional ex post value of health insurance to the household, over and above that from reduced out-of-pocket spending, from reduced medical debt. In theory, reducing debt by \$X should be worth at most \$X to the household, since it could simply repay \$X to achieve that. Other options include not repaying—the most common choice—or discharging through bankruptcy. Corr(log(c), log(V)) is the correlation between the first differences of log consumption and the log of one plus V, both residualized with the corresponding controls. "Insurance value," $Cov\left(\hat{\lambda}, V\right)$, is $-\gamma \times \beta \times \frac{Var(V)}{E(V)}$, where $\gamma = 3$ is the coefficient of relative risk aversion and β is the regression coefficient on the V term (see equation (8)). "Markup" is insurance value per dollar of mean ex post value, $Cov\left(\hat{\lambda}, V\right)/E\left(V\right)$. Standard errors, which are clustered at the household level, reflect sampling uncertainty in β but treat $E(V)$ and $Var(V)$ as non-stochastic. Data are from the PSID. Monetary amounts are in real 2020 dollars per household per year. Non-elderly are households whose heads are 25–64. See Appendix A.6 for a discussion of these results.

Table A16: Health insurance and hospitalization risk: Decomposition estimates

	Baseline				No implicit health insurance			
	Insurance value	Contrib to markup	Transfer value	Benefit	Insurance value	Contrib to markup	Transfer value	Benefit
Non-elderly								
Total	-100	-0.87	116		1,629	2.30	708	
Non-hosp to hosp	27	0.24	116	1,171	2,063	2.91	708	7,171
Low- to high-y within hosp	-24	-0.21	0	1,108	0	0	0	0
Within hosp-by-y	-103	-0.89	0		-434	-0.61	0	
Non-elderly uninsured								
Total	-80	-0.62	130		1,590	2.38	669	
Non-hosp to hosp	43	0.33	130	1,446	1,912	2.86	669	7,446
Low- to high-y within hosp	-21	-0.16	0	1,578	0	0	0	0
Within hosp-by-y	-102	-0.78	0		-323	-0.48	0	

Notes: Statistics related to the value of completing health insurance coverage of hospitalization risk, from the perspective of an individual who knows their education but nothing else about the realization of risk. These correspond to the decomposition depicted in Figure A6, which decomposes the transfers into hospitalization states from completing health insurance into the sum of three components, as described in Appendix Section E. The first, "Non-hosp to hosp," is the 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 the 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 year dummies, a cubic in age, a quadratic in household size, and education category dummies. This equals the difference in average out-of-pocket spending between hospitalization-and-high-income and hospitalization-and-low-income states. The third, "Within hosp-by-y," is the transfers within the two hospitalization-and-income-is-high/low categories, from states in which completing health insurance is less to more valuable. In each hospitalization state, this transfer equals the difference between out-of-pocket spending in that state and average out-of-pocket 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 spending and consumption (for marginal utility). The "No implicit health insurance" columns use a simple model of counterfactual consumption and out-of-pocket spending were there no implicit health insurance. Counterfactual out-of-pocket spending is $oop_\omega^{no} = oop_\omega + ihi(hosp_\omega, y_\omega)$, where oop_ω is observed out-of-pocket spending and $ihi(hosp_\omega, y_\omega)$ is such that (i) average counterfactual out-of-pocket spending is the same in both sets of hospitalization states and (ii) average implicit health insurance 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 consumption is $c_\omega^{no} = \max\{\underline{c},\ c - ihi(hosp_\omega, y_\omega)\}$. This coarse model of implicit health insurance, which assumes that implicit health insurance support is uniform within each hospitalization-by-low/high-income groups, likely understates implicit health insurance support in lower- relative to higher-income states within these groups and so likely leaves much of implicit health insurance's transformation of health spending risk intact. This could explain why the "No implicit health insurance" results still have health insurance increasing within-group risk and is a force toward these results understating the value of health insurance in the "No implicit health insurance" counterfactual. Insurance value is the excess of the ex ante value over the mean ex post value (see equation (4)). The markup is the insurance value per dollar of mean ex post value. The mean ex post 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 child under two years old present in the household (to exclude hospitalizations related to childbirth). The underlying regressions are of normalized marginal utility on V (not log consumption on $\log(1+V)$), plus year dummies, a cubic in age, a quadratic in household size, and education category dummies. This "exact" specification, which aims to approximate relatively long run risk (all risk within education groups but not the risk of being in one education group as opposed to another), guarantees that the sum of the values of each component sum to the overall value, whereas specifications based on approximations to normalized marginal utility or V do not. Data are from the PSID.

Table A17: Structural analysis of mechanisms

	Baseline (1)	No income risk (2)	No income risk and no implicit HI (3)
Insurance value	-489	66	3,060
Mean ex post value	2,573	2,587	4,797
Markup	-0.19	0.03	0.64
$Corr\left(\widehat{\lambda}, V\right)$	-0.06	0.999	0.46
$Corr\left(\widehat{\lambda}, hi\right)$	0.08	0.59	0.46
$Corr(hi, y)$	-0.002	0	0
$Corr(V, y)$	0.17	0	0

Notes: Statistics related to the insurance value of health insurance in three versions of the structural model: the baseline model, no income risk, and neither income risk nor implicit health insurance. The "No income risk" counterfactual has health risk as in the data, $h \sim F(h)$, but income equal to median income in all states, $y \equiv y_{med}$. The "No income risk and no implicit health insurance" counterfactual has no income risk and no implicit health insurance: $y \equiv y_{med}$, $ihi(h, y; HI) \equiv 0$. Insurance value is the amount by which the ex ante equivalent variation of health insurance exceeds its mean ex post value (see equation (3)), using consumption-based equivalent variation (the amount by which the consumption of a household without health insurance must be increased in all states to be as well off as it would be with health insurance). The markup is the ratio of insurance value to mean ex post value. All results are for non-elderly households. The baseline risk process aims to approximate relatively long run risk: all risk within education groups but not the risk of being in one education group as opposed to another.

Table A18: Structural analysis heterogeneity by age and education

	Non-elderly			Elderly		
	$<$ High school (1)	High school (2)	$\geqslant$ College (3)	$<$ High school (4)	High school (5)	$\geqslant$ College (6)
Insurance value	-437	-520	-544	-578	-502	-375
Mean ex post value	2,379	2,561	2,688	3,204	3,239	3,160
Markup	-0.18	-0.20	-0.20	-0.18	-0.15	-0.12
$Corr\left(\widehat{\lambda}, V\right)$	-0.04	-0.06	-0.08	-0.11	-0.12	-0.12
$Corr\left(\widehat{\lambda}, hi\right)$	0.04	0.07	0.11	0.07	0.08	0.09
$Corr(hi, y)$	-0.01	-0.02	0.01	-0.03	-0.002	-0.0001
$Corr(V, y)$	0.15	0.15	0.18	0.17	0.21	0.26

Notes: Statistics related to the insurance value of health insurance in the structural model for different "risk types." Insurance value is the amount by which the ex ante equivalent variation of health insurance exceeds its mean ex post value (see equation (3)), using consumption-based equivalent variation (the amount by which the consumption of a household without health insurance must be increased in all states to be as well off as it would be with health insurance). The markup is the ratio of insurance value to mean ex post value. The risk process in a given column aims to approximate relatively long run risk: all risk within that age-by-education group but not the risk of being in that group as opposed to another.

Table A19: Structural analysis robustness to implicit HI coverage and income risk

	Baseline implicit HI		High cost of bankruptcy	
	Baseline y risk (1)	Half y risk (2)	Baseline y risk (3)	Half y risk (4)
Insurance value	-489	-170	-235	-56
Mean ex post value	2,573	2,592	3,127	3,150
Markup	-0.19	-0.07	-0.08	-0.02
$Corr\left(\widehat{\lambda}, V\right)$	-0.06	-0.04	0.04	0.05
$Corr\left(\widehat{\lambda}, hi\right)$	0.08	0.08	0.14	0.13
$Corr(hi, y)$	-0.002	-0.003	-0.002	-0.003
$Corr(V, y)$	0.17	0.12	0.12	0.09

Notes: Statistics related to the insurance value of health insurance in the structural model for different levels of implicit health insurance coverage and income risk. Columns (1) and (2) use the baseline implicit health insurance calibration. This baseline calibration tends to understate the amount of support from implicit health insurance received by the typical uninsured household. For example, among uninsured households $E(oop)$ is \$2,060 in this calibration versus \$990 in the data and $E(oop|tot > 10k)$ is about \$4,440 in this calibration versus \$3,940 in the data. Columns (3) and (4) scale up the implicit health insurance deductibles—reducing implicit health insurance support—to match mean out-of-pocket health spending in the top ventile of charges among uninsured households with high (>\$50,000) financial costs of bankruptcy as estimated by Mahoney (2015), which is about \$7,000. This calibration aims to approximate the implicit health insurance protection against otherwise-uninsured health care costs available to households with significant assets or low willingness to rely on implicit health insurance. (Granted, the model is ill-suited to quantify insurance value for households with significant assets since it assumes that consumption equals net income in each state of the world, with no consumption smoothing over time.) Columns (2) and (4) halve income risk by setting income in each state of the world to a weighted average of its observed value and median income, with half of the weight on each: $\widetilde{y} = 0.5 \times y + 0.5 \times y_{median}$. All of the non-baseline columns result in correlations between out-of-pocket spending and income ($Corr(V, y)$) below the corresponding empirical estimates (see, e.g., Appendix Table A4). Insurance value is the amount by which the ex ante equivalent variation of health insurance exceeds its mean ex post value (see equation (3)), using consumption-based equivalent variation (the amount by which the consumption of a household without health insurance must be increased in all states to be as well off as it would be with health insurance). The markup is the ratio of insurance value to mean ex post value. The baseline risk process aims to approximate relatively long run risk: all risk within education groups but not the risk of being in one education group as opposed to another.

Table A20: Structural analysis additional robustness tests

	Baseline (1)	Income floor \$50k (2)	One-fourth income risk (3)	HI deductible \$3k (4)	Log utility (5)	Hospitalization targeting only (6)
Insurance value	-489	-141	-20	-351	-236	-117
Mean ex post value	2,573	2,617	2,594	583	2,573	2,601
Markup	-0.19	-0.05	-0.01	-0.60	-0.09	-0.04
$Corr\left(\widehat{\lambda}, V\right)$	-0.06	-0.04	0.06	-0.15	-0.13	-0.07
$Corr\left(\widehat{\lambda}, hi\right)$	0.08	0.09	0.12	0.09	0.06	-0.05
$Corr(hi, y)$	-0.002	0.001	-0.003	-0.01	-0.002	-0.15
$Corr(V, y)$	0.17	0.12	0.09	0.27	0.17	-0.14

Notes: Statistics related to the insurance value of health insurance in the structural model under different assumptions. Column (1) is the baseline specification. Columns (2) and (3) reduce income risk. Column (2) increases the income floor from \$15,000 to \$50,000. Column (3) 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, with three-fourths of the weight on median income: $\widetilde{y} = 0.25 \times y + 0.75 \times y_{median}$. Column (4) 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 (5) uses log utility (a coefficient of relative risk aversion of one). Column (6) reports statistics related to the hospitalization-related targeting of health insurance, based on Dobkin et al.'s (2018) ("DFKN") estimates of the health care costs and earnings losses associated with hospitalization. Start from the joint distribution of residualized total health care costs and residualized income among non-elderly households, both residualized with year dummies, a cubic in age, a quadratic in household size, and education category dummies. If the household experienced a hospitalization, (i) its total health care costs h are increased by \$18,750 (DFKN's estimate of the increase in total annual health care costs in the first three years following a hospitalization), and (ii) its (before-income-floor) income is probabilistically decreased as follows. Conditional on hospitalization, with probability 10% income is decreased by \$45,000 (based on DFKN's estimate of a 10pp reduction in employment from hospitalization, and average pre-hospitalization earnings of \$45,000 [inferred from the fact that DFKN's estimate of a \$9,000 decrease in earnings represents a decrease of about 20%]) and otherwise income is decreased by \$5,000 (so that the average income loss is \$9,000, DFKN's estimate of the decrease in annual earnings in the first few years following a hospitalization). The results are almost identical with any other feasible attribution of the income losses beyond those from reduced employment. If the household does not experience a hospitalization, its total health care costs are set to the lesser of median total health care costs and the minimum implicit health insurance deductible in order to "shut down" health insurance targeting within non-hospitalization states. This means there is only targeting from non-hospitalization states to hospitalization states and within hospitalization states. Insurance value is the amount by which the ex ante equivalent variation of health insurance exceeds its mean ex post value (see equation (3)), using consumption-based equivalent variation (the amount by which the consumption of a household without health insurance must be increased in all states to be as well off as it would be with health insurance). The markup is the ratio of insurance value to mean ex post value. All results are for non-elderly households. The baseline risk process aims to approximate relatively long run risk: all risk within education groups but not the risk of being in one education group as opposed to another.

Table A21: Sufficient statistic estimates: Indemnity insurance

	Hospitalization indemnity									Hospital days indemnity								
	Non-elderly uninsured			Non-elderly insured			Elderly insured			Non-elderly uninsured			Non-elderly insured			Elderly insured		
	Short run (1)	Medium run (2)	Long run (3)	Short run (4)	Medium run (5)	Long run (6)	Short run (7)	Medium run (8)	Long run (9)	Short run (1)	Medium run (2)	Long run (3)	Short run (4)	Medium run (5)	Long run (6)	Short run (7)	Medium run (8)	Long run (9)
Corr(log(c),V)	-0.01	-0.02	-0.04	0.01	-0.02	-0.07	0.00	-0.03	-0.02	-0.01	-0.03	-0.04	-0.01	-0.02	-0.06	-0.01	-0.03	-0.03
(se)	(0.02)	(0.03)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.02)	(0.01)	(0.02)	(0.03)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.02)	(0.01)
Mean ex post value	0.07	0.11	0.09	0.08	0.10	0.10	0.20	0.21	0.24	0.71	1.07	0.69	0.76	0.84	0.78	2.01	2.02	2.08
Markup	0.07	0.09	0.24	-0.02	0.11	0.38	-0.01	0.09	0.07	0.07	0.36	0.59	0.03	0.16	0.81	0.05	0.19	0.28
(se)	(0.09)	(0.15)	(0.09)	(0.03)	(0.06)	(0.04)	(0.04)	(0.06)	(0.04)	(0.15)	(0.35)	(0.19)	(0.05)	(0.10)	(0.11)	(0.08)	(0.16)	(0.12)

Notes: Statistics related to the markup on hypothetical indemnity insurance contracts that pay a fixed cash benefit based on hospitalization or hospital days. The hospitalization indemnity pays out \$1 if the household head or spouse experienced a hospitalization in the past year *and* there is no child under two years old present in the household (to exclude hospitalizations related to childbirth) and zero otherwise. The hospital days indemnity pays out \$1 for each day the household head or spouse is hospitalized. This table assumes that ex post the individual benefits one-for-one from the indemnity benefit, i.e., that implicit health insurance does not implicitly tax these benefits. This extreme assumption, made for illustrative purposes, is based on the idea that indemnity insurance would displace less support from implicit health insurance than standard "service benefit" policies (see footnote 54 on page 33). Short run and medium run columns are based on regressions of within-household changes in log consumption on within-household changes in V, plus year dummies and a cubic in age, where the changes are from one wave to the next (short run) or from one wave to five waves later (medium run). Long run is based on regressions of log consumption on V, plus year dummies, a cubic in age, and a quadratic in household size. Short run aims to capture the value of coverage from the perspective of immediately before the coverage begins, medium run from ten years before the coverage begins, and long run from behind the veil of ignorance. Short and medium run specifications limit the sample to households who did not experience a hospitalization in the lagged period. These are analogous to a common specification in the unemployment insurance literature (e.g., Hendren, 2017). Corr(log(c), V) is the correlation between the relevant changes in (short and medium run) or levels of (long run) log consumption and V, both residualized with the corresponding controls. "Markup" is insurance value per dollar of mean ex post value, $Cov\left(\widehat{\lambda}, V\right)/E\left(V\right)$. "Insurance value," $Cov\left(\widehat{\lambda}, V\right)$, is $-\gamma \times \beta \times \frac{Var(V)}{E(V)}$, where $\gamma = 3$ is the coefficient of relative risk aversion and β is the regression coefficient on the V term (see equation (8)). Standard errors, which are clustered at the household level, reflect sampling uncertainty in β but treat $E(V)$ and $Var(V)$ as non-stochastic. Data are from the PSID. Monetary amounts are in real 2020 dollars per household per year. Non-elderly (elderly) are households whose heads are 25–64 (65+).