Health Insurance and Consumption Risk

Lee M. Lockwood*

March 25, 2025

Abstract

The effect of health insurance on consumption risk depends in part on how it interacts with other risks beyond health care cost risk, such as income risk. Using a variety of approaches, I find that for U.S. households, the interaction with other risks transforms the risk protection from health insurance. Standard contracts amplify the impact of other risks, due to both subsidizing normal goods and undoing the protection against other risks from discounts, charity care, and bad debt. Alternative contracts that account for other risks, such as contracts that limit out-of-pocket spending relative to income, can provide better risk protection.

*University of Virginia and NBER, leelockwood@virginia.edu. I thank John Cochrane, Manasi Deshpande, Amy Finkelstein, Tal Gross, Nathaniel Hendren, John Bailey Jones, Daniel Kaliski, Anton Korinek, Amanda Kowalski, Kory Kroft, Kurt Lavetti, Tim Layton, Adam Leive, Gideon Magnus, Matthew Notowidigdo, Mark Shepard, Jonathan Skinner, Isaac Sorkin, referees, and many seminar participants 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 and $21,000 per household (U.S. Centers for Medicare & Medicaid Services, 2019; U.S. Census, 2020). A fundamental motivation for health insurance is risk protection: insuring living standards against financial risk by helping individuals more in states of the world in which consumption is lower and marginal utility is higher. Accordingly, important research has documented beneficial effects of health insurance on financial outcomes such as out-of-pocket spending and medical debt—to say nothing of the beneficial effects on health (see Finkelstein et al., 2018, for a review).

Yet while these important benefits of health insurance contribute to its overall risk protection, there is another contributor that has received less consideration: the interaction with other risks, such as from unemployment, wages, and asset prices. In contexts with other risks beyond health care costs, the overall risk protection from health insurance depends in part on its interaction with other risks, not just its protection from health care costs.

In this paper, I use a variety of approaches to investigate the risk protection from 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—accounting for other risks. I find that other risks transform the risk protection from health insurance. Standard contracts, which provide the same coverage regardless of circumstances, tend to amplify other risks: They help individuals more on average when the realization of other risks is better and less when it is worse, increasing the welfare cost of other risks. This amplification is strong enough, in fact, that for U.S. households, such contracts tend to increase consumption risk on net; their amplification of other risks outweighs their insurance of health care costs. Contracts that account for other risks, such as contracts that limit out-of-pocket spending relative to income, can provide better risk protection.[1]

Two factors cause standard contracts to amplify other risks. One is their coverage of less risky, more "discretionary" types of health care that are normal goods. Such coverage is less valuable when other circumstances are worse. For example, if in the absence of health insurance someone would postpone an elective surgery following a negative earnings shock, health insurance would be worth less in such states and thereby amplify earnings risk. The other factor is "implicit health insurance" from discounts, charity care, and bad debt. Implicit insurance provides significant protection against otherwise-uninsured health care costs, even for households in strong financial positions and especially when other circumstances are worse. For example, among uninsured households with over $20,000 of health care charges, out-of-pocket spending is about $5,000 on average and even lower when income or wealth is low. The greater protection when other circumstances are worse insures other risks. Receiving more charity care when unemployed, for example, partially offsets the income loss. Such protection against other risks from implicit insurance is undone by health insurance.

To illustrate, consider two households that typically earn $100k per year. Each experiences a heart attack that results in $20k of uninsured health care costs. In the "lucky" household, the affected member is able to continue working after a brief paid leave. They are billed $10k (after a $10k discount), and they pay the bill in full, leaving them net income of $90k. In the "unlucky" household, the affected member can no longer work, causing household income to fall to $50k. They are billed $5k (after $15k of discounts and charity care), but they never pay the bill, leaving their net income unchanged at $50k. Now introduce comprehensive coverage of all costs. For the lucky household, this increases net income by $10k. For the unlucky household, however, it brings no change in net income—it solely displaces discounts, charity care, and bad debt. This increases the gap in net income between the lucky and unlucky

[1]That standard contracts increase consumption risk in no way contradicts their financial and health benefits. Nor does it imply that they decrease welfare or are worse than other types of contracts; they have many other potential benefits. I discuss other benefits and the implications of the results in Section 5.

households from \$40k (\$90k – \$50k) to \$50k (\$100k – \$50k). Although the coverage itself does not depend on income, it helps the household less when income is lower and thereby amplifies the associated income risk.[2]

I start with a descriptive analysis of out-of-pocket 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 risk protection from health insurance. First, out-of-pocket spending risk is relatively small. 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, out-of-pocket spending buffers (i.e., partially offsets) other risks: It tends to be lower (i.e., more favorable) when the realization of other risks is less favorable, decreasing the welfare cost of other risks. For example, out-of-pocket spending drops when households experience negative income shocks such as from unemployment, which partially offsets the income losses. Third, out-of-pocket spending buffers consumption risk. Out-of-pocket spending is lower when consumption is lower. This suggests that if out-of-pocket spending were to disappear, the volatility of consumption would increase. Hence, the interaction with other risks causes out-of-pocket spending to reduce risk on net. The buffer against other risks outweighs the health care risk because other risks are larger. This makes it hard for standard health insurance contracts to reduce consumption risk; their protection against health care costs tends to be outweighed by their amplification of other risks.[3]

I then turn to estimating the value of the risk protection from different types of health

[2]While receiving discounts or charity care or failing to repay medical debt may have certain costs to the individual, the evidence suggests that such costs tend to be small. For example, in two large-scale randomized experiments, medical debt relief led to "no improvements in financial well-being or mental health" (Kluender et al., 2024, p. 7). Nevertheless, where relevant, I test the effects of such costs being large.

[3]While these patterns are key determinants of the risk protection from health insurance, additional considerations matter for welfare. For example, out-of-pocket spending buffers other risks in part because individuals forgo or postpone care when times are tight. Limiting such disruptions to care could be a major benefit of health insurance. See Section 5.

insurance: the extent to which the ex ante value exceeds the mean ex post value due to the budget-neutral reallocation of resources across states of the world, i.e., the "pure insurance" aspect. 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 under a wide range of assumptions, contracts that account for other risks would provide better risk protection than standard contracts. The estimates suggest that standard coverage, while valuable, is considerably less valuable to households ex ante than the same mean ex post value worth of cash—20–70% less valuable, depending on the perspective. By contrast, a contract that limits out-of-pocket spending to 10% of income can provide valuable risk protection.[4]

To better understand the underlying mechanisms, I construct a simple model guided by the empirical findings. It is an otherwise-standard model of health care cost risk except that it includes other risks. The model matches well the key empirical patterns, including those not targeted directly such as the sufficient statistic estimates. Here too, the conclusion that income-dependent health insurance would provide better risk protection than standard contracts is extremely robust. Counterfactual analyses highlight the crucial role of other risks, which reverse the risk protection ranking of standard versus income-dependent contracts. Were it not for other risks, standard contracts would provide slightly better risk protection. With realistic levels of other risks, however, income-dependent contracts provide considerably better risk protection. Whereas standard contracts amplify other risks at a welfare cost

[4]This is similar to the main contract proposed by Feldstein and Gruber (1995), though their aim was not to insure other risks, but to reduce moral hazard relative to more comprehensive contracts while ensuring that out-of-pocket spending is not too large relative to income.

of several hundred dollars per year, income-dependent contracts not only avoid amplifying other risks so much, they can even insure them. As a result, such contracts can provide valuable risk protection against health care costs and other risks alike.

That standard contracts increase consumption risk does not mean that they reduce welfare or that households are making mistakes by holding them. It just means that one component of their overall welfare effect is not the benefit previously thought but a cost. A full accounting must include the many important benefits of such contracts, including reduced reliance on implicit insurance and improved health (and perhaps reduced risk in health). Similarly, that alternative contracts could provide better risk protection does not imply that they would be better all things considered. I discuss these and other issues of interpretation in Section 5.

The main contribution of this paper is to analyze the risk protection from different types of health insurance, accounting for other risks. My findings build on and help reconcile two strands of related literature. The first seeks to understand the risk protection from health insurance. This strand is based on two types of evidence: structural analyses that seek to quantify risk protection value and empirical analyses of the effects of health insurance on financial outcomes such as out-of-pocket spending, medical debt, bankruptcy, and credit scores.[5] To the best of my knowledge, all previous studies have concluded that standard contracts provide valuable risk protection. In fact, a common view is that such contracts provide *too much* risk protection because of over-insurance due to subsidies.

These conclusions are largely based on the fact that health insurance decreases the volatility of out-of-pocket spending. My analysis highlights an important limitation of this evidence: It does not account for other risks. That causes the analysis to miss what turns out to be the

[5]On the former, see 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); Finkelstein et al. (2019a). On the latter, see Finkelstein and McKnight (2008); Engelhardt and Gruber (2011); Gross and Notowidigdo (2011); Finkelstein et al. (2012); Barcellos and Jacobson (2015); Mazumder and Miller (2016); Hu et al. (2018); Brevoort et al. (2020).

most consequential effect of standard contracts on consumption risk: their amplification of other risks. Though unexpected, that conclusion emerges clearly from diverse strands of evidence, including the "bottom-line" evidence that out-of-pocket spending covaries positively with consumption, evidence on the proximate mechanisms that out-of-pocket spending risk is relatively small and buffers other risks, and evidence on the ultimate mechanisms that adding other risks to an otherwise-standard model causes standard contracts to increase consumption risk under a wide range of parameter values.

The interaction with other risks also helps reconcile the literature on risk protection from health insurance with the second strand of related literature, which seeks to quantify the overall value of health insurance.[6] A key finding of this strand 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 typical contracts decrease consumption risk but accords well with my finding that they increase it.[7] My findings complement and extend earlier research by showing the crucial role of other risks in not only reducing the value of typical contracts but even potentially making it smaller than the mean net benefit. More generally, I find that the interaction with other risks reverses several conclusions about the risk protection from different types of coverage. Less comprehensive coverage not only has a lower moral hazard cost than more comprehensive coverage, it also provides better risk protection. Same for indemnity insurance that pays a fixed cash benefit based on a health diagnosis. These findings highlight an unappreciated cost of the type of comprehensive coverage that is encouraged by tax subsidies for health insurance. More broadly, my findings show how other risks beyond those directly targeted by an insurance policy can transform its risk protection. This has

[6]See 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.

been shown to be important for disability insurance (Deshpande and Lockwood, 2022) and could be important in other contexts as well.

2 Data, Institutions, and Empirical Approach

Data.— *PSID.—* The Panel Study of Income Dynamics (PSID) has many advantages for analyzing the risk protection from 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 11 waves from 1999–2019 inclusive. These waves occur every two years. 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 roles of implicit insurance and income effects of demand for health care 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 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 income measures include income from all sources, including social insurance and means-tested programs, to reflect the net risk after such insurance. My hospitalization measures 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 A contains details of variable construction, and 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 consider two main types of contracts. One is "standard contracts," which cover a fixed share of health care costs regardless of other circumstances. This 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 individual is billed for health care. While this can lead to over-consumption of health care, such contracts are thought to provide better risk protection than other types.[8] I also consider "income-dependent contracts" that limit out-of-pocket spending relative to realized income.[9]

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

[9]Feldstein and Gruber (1995) propose a contract with a stop-loss of 10% of income. Itemizing taxpayers can deduct qualified medical care expenses that exceed 7.5% of their adjusted gross income. Some health insurance exchange plans offer cost-sharing reductions that provide more coverage to low-income enrollees. Contracts that provide more coverage when income is lower provide some protection against income risk.

Implicit health insurance.— Discounts, charity care, and bad debt provide significant protection against otherwise-uninsured health care costs. Individuals without formal health insurance pay only about one-fifth to one-third of their health care costs out of pocket (Hadley et al., 2008; Coughlin et al., 2014; Finkelstein et al., 2019a), and, in two large-scale randomized experiments, medical debt relief led to "no improvements in financial well-being or mental health" (Kluender et al., 2024, p. 7).[10] 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 Appendix B). 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 *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 in others, it can transform health insurance targeting. This paper analyzes the risk protection from formal health insurance accounting for other risks and the displacement of implicit insurance.

Empirical approach.— *Conceptual experiment.—* Many of my analyses seek to characterize the effects and value of (hypothetical) health insurance coverage expansions: increases in coverage from status quo levels. I mostly follow the standard approach of focusing on the effects of health insurance on out-of-pocket spending. This is the main financial effect

[10] Charity care arises from both charitable motives and legal obligations. To qualify for certain tax exemptions, nonprofit hospitals (roughly 70% of 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 hospitals' legal obligation to provide emergency medical care on credit. Most hospitals provide non-emergency care on credit as well (see Mahoney, 2015, and references therein). Much of the care provided on credit is never paid for. Uninsured individuals pay only about 10–20% of what they are billed (LeCuyer and Singhal, 2007). Medical debt often is defaulted on implicitly rather than explicitly through bankruptcy. Whereas unpaid medical bills affect nearly one-fifth of consumers' credit reports and comprise a majority of collections lines (CFPB, 2014), each year less than 1% of Americans file for personal bankruptcy.

of health insurance and, under standard assumptions, is a first order approximation to its ex post value.[11] Hence, the risk protection from health insurance depends 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 risk protection from health insurance that I am aware of). Still, where relevant, I consider the effects on health and medical debt as well.

To fix ideas, consider the provision of full coverage to an uninsured household. The main financial effect would be to eliminate out-of-pocket spending. Other things equal, this would increase net income by status-quo out-of-pocket spending in each state of the world. This key effect of health insurance can be ascertained from knowledge of out-of-pocket spending in the status quo. Knowledge of other outcomes, including counterfactual outcomes away from the status quo (e.g., with the expanded coverage) and the causal effects of the contemplated change in coverage, is unnecessary. 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 states of the world they might experience.[12]

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 risk protection from

[11] 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 reduction in out-of-pocket spending it would cause if behavior were fixed, by the envelope theorem. Of course, optimization is a strong assumption. For example, recent evidence suggests that health care consumption might be excessively sensitive to liquidity (e.g., Gross et al., 2020). While my sufficient statistic analysis assumes optimization, my other analyses do not.

[12] This experiment varies the state of the world, not health insurance. Section 4 presents my approach to estimating risk protection value. Appendix C provides details, including on the close relationships to the widely used Baily-Chetty approach and Finkelstein et al.'s (2019a) "optimization approach." It also discusses the significant challenges facing methods based on exogenous variation in health insurance.

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 from the later perspective but also the additional risk of which "later perspective" one will experience.

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.[13] 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.[14] In a few instances, I consider the perspective of someone who knows their education level but nothing else. This

[13]The consumption version is the health insurance analogue of a common specification in the literature on unemployment insurance (e.g., Hendren, 2017).

[14]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 reduce the impact of any misspecification of price indices, household equivalence scales, and age-dependent utility. The consumption version is the same regression used by Finkelstein et al. (2019a) in their most closely related analysis of the value of Medicaid.

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.

3 Descriptive Analysis of Out-of-Pocket Spending

Finding 1: Out-of-pocket spending risk is relatively small.

Figure 1a shows a histogram and estimated kernel density of the distribution of annual out-of-pocket spending among non-elderly uninsured households, and Table A3 shows associated statistics. Out-of-pocket spending is small on average (average of $1,060) and only modestly variable (standard deviation of $2,720 and 99th percentile of $11,460). Income and consumption are much more variable. Among non-elderly households, the *within-household* standard deviations of income and consumption are $34,910 and $15,620, respectively.

Out-of-pocket spending risk is relatively small mainly because of implicit insurance. Among non-elderly households, total health care costs are large on average (average of $9,610) and highly variable (standard deviation of $23,030), albeit significantly less variable than income (see Table A3). 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 nonparametric 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.[15] The difference, presumably but in accordance with other evidence, arises from greater reliance on implicit insurance by uninsured households. The nominally uninsured, though lacking formal health insurance, have substantial implicit insurance from discounts, charity care, and bad debt. This implicit insurance resembles catastrophic coverage above a modest deductible (Mahoney, 2015).[16]

Implications.— Health insurance provides relatively limited protection against health care costs. The considerable protection from implicit insurance broadly resembles 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 risk protection would be outweighed by administrative and moral hazard costs. So while *total* protection against health care costs, including from implicit insurance, likely is highly valuable, additional, *marginal* protection on top of that provided by implicit insurance is unlikely to generate much risk protection value.

Finding 2: Out-of-pocket spending buffers other risks.

A major risk for many households is income risk. Figure 2a shows, for non-elderly households, a nonparametric estimate of the conditional mean of income as a function of out-of-pocket spending, controlling for household fixed effects, year dummies, and a cubic in age. Associ-

[15]While $5,000 is a lot to spend on health care, even such an extreme realization of out-of-pocket spending (around the 95th percentile) is small in comparison to many other risks, 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).

[16]Although implicit insurance provides more protection to households in worse financial positions, it provides considerable protection to households in strong financial positions as well (see Appendix B and Figures A1 and A2).

ated heterogeneity and robustness results are in Tables A4 and A5. For each type of state and perspective, out-of-pocket 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. Out-of-pocket spending tends to be lower when income is lower and higher when income is higher, partially offsetting income shocks. Out-of-pocket spending buffers income risk.[17]

That out-of-pocket spending buffers income risk is striking given that health risk is a force toward out-of-pocket spending amplifying income risk. Health shocks increase out-of-pocket spending and decrease income, which tends to make out-of-pocket spending higher when income is lower, which would amplify income risk. So the positive relationship between out-of-pocket spending and income must reflect stronger countervailing forces. One is that certain types of health care are normal goods (Acemoglu et al., 2013; Gross et al., 2020). This tends to make the ex post value of health insurance greater when the realization of other risks is more favorable (see Appendix D for details).

The other force is implicit insurance. Figure 2b shows a nonparametric estimate of the conditional mean of out-of-pocket 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. 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

[17]For example, out-of-pocket spending buffers unemployment risk. Out-of-pocket spending drops when households experience unemployment (see Table A6 and Figure A3), partially compensating for associated income losses, despite health insurance coverage dropping as well.

(seizable) net assets. Implicit insurance helps more when circumstances are worse. In this way, implicit insurance is not like standard catastrophic insurance that covers costs above a fixed deductible but rather like special catastrophic insurance with a state-contingent deductible that is lower (more coverage) when circumstances are worse. As a result, it implicitly insures risk in income, assets, and non-health care circumstances more generally.

Implications.— Standard health insurance coverage tends to amplify other risks. In this sense, holding standard coverage is like holding stock in one's employer: It tends to be worth less when income is lower and thereby increases the welfare cost of income risk. Standard coverage amplifies other risks by subsidizing normal goods and undoing the protection from implicit insurance. Alternative contracts that provide more coverage when other circumstances are worse would amplify other risks less or even insure them. For example, a contract that limits out-of-pocket spending to 10% of income would reduce out-of-pocket spending more on average when income is lower and thereby insure income risk (see Table A7).

Finding 3: Out-of-pocket spending buffers consumption risk.

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

How could out-of-pocket spending decrease consumption risk? Consider the variance of net income,

$$Var(y - oop) = Var(y) + \left[\underbrace{Var(oop)}_{\text{“Partial effect”}} - \underbrace{2Cov(y, oop)}_{\text{“Portfolio effect”}} \right]. \tag{1}$$

The bracketed term is the total effect of out-of-pocket spending. The “partial effect” reflects that, other things equal, greater out-of-pocket spending reduces net income, tending to increase 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 Finding 2, out-of-pocket spending buffers income risk: $Cov(y, oop) > 0$. Hence, out-of-pocket spending has opposing effects on the volatility of net income: As a variable expense, it tends to increase volatility, but as a buffer against income risk, it tends to decrease volatility. I find 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,” so out-of-pocket spending reduces the variance of net income.[18] Beyond net income, the results from Finding 3 suggest that for all combinations of states and perspectives, out-of-pocket spending decreases the volatility of consumption.[19] Due to the interaction with other risks, more exposure to out-of-pocket spending means less exposure to risk on net.

The buffer against other risks dominates the volatility of health care costs because other risks are larger, increasingly so from earlier perspectives where more risk remains. Among

[18]This occurs if $2Cov(y, oop) > Var(oop)$, equivalently if the slope of the regression of income on out-of-pocket spending exceeds one-half: $\beta_{y|oop} \equiv \frac{Cov(y,oop)}{Var(oop)} > 1/2$. In non-elderly states, \$1 higher out-of-pocket spending is associated with income being higher by \$1.03 (short run), \$3.04 (medium run), and \$8.38 (long run), and so with net income being higher by \$0.03, \$2.04, and \$7.38, respectively.

[19]If the marginal propensity to consume were $\alpha > 0$ in all states, then other things equal, eliminating out-of-pocket spending would increase the variance of consumption if $2\alpha Cov(c, oop) > \alpha^2 Var(oop)$, i.e., if $\beta_{c|oop} > -\alpha/2$ (since the variance of consumption 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). All estimates of $\beta_{c|oop}$ are positive and significant, and thus exceed $-\alpha/2$ for all $\alpha > 0$. In non-elderly states, \$1 higher out-of-pocket spending is associated with consumption being higher by 32 cents (short run), 95 cents (medium run), and \$2.36 (long run).

uninsured households, income alone has a standard deviation over 30 times that of out-of-pocket spending (see Table A3). This connects the three main findings of the analysis of out-of-pocket spending. Out-of-pocket spending decreases risk on net (Finding 3) because its buffer against other risks (Finding 2) dominates its partial effect of being risky itself, because out-of-pocket spending risk is relatively small (Finding 1). As a rough measure of magnitudes, simple calculations suggest that eliminating out-of-pocket spending would increase the standard deviation of consumption about 2–4 times what eliminating unemployment insurance would (see Figure A4).

Implications.— In terms of risk protection, catastrophic coverage likely is better than comprehensive coverage, and income-dependent coverage would likely be better still. The risk protection ranking is reversed by the interaction with other risks because standard coverage amplifies other risks, enough to outweigh its protection against health care costs. Alternative contracts that provide more coverage when other circumstances are worse can mitigate or even reverse the amplification of other risks and thereby provide valuable risk protection. For example, a contract that limits out-of-pocket spending to 10% of income would tend to reduce out-of-pocket spending more on average when consumption is lower and so likely provide valuable risk protection (see Table A10).

4 Risk Protection Value of Health Insurance

4.1 Risk protection 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

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

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 as well off as switching to the new constraint.[20]

Risk protection 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

$$E\left[u\left(c_0 + EAV, a_0; \theta\right)\right] = E\left[u\left(c_1, a_1; \theta\right)\right], \tag{3}$$

where expectations are over possible states of the world, $\theta \sim F(\theta)$. EAV is the increase in consumption in all states that makes the individual as well off ex ante as they would be under the new constraints. "Risk protection value," what Finkelstein et al. (2019a) call "pure-insurance value," is the excess of ex ante value over mean ex post value:

$$EAV = E(V) + \text{Risk protection value}. \tag{4}$$

Risk protection value is the "pure-insurance" 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

$$\underbrace{EAV}_{\text{Ex ante value}} \approx \frac{E\left(\lambda \times V\right)}{E\left(\lambda\right)} = \underbrace{E\left(V\right)}_{\text{Mean ex post value}} + \underbrace{Cov\left(\widehat{\lambda},\ V\right)}_{\text{Risk protection value}}, \tag{5}$$

[20]The theory applies to dynamic settings, as in Chetty (2006). I present a static setting for simplicity. Value is measured in terms of consumption rather than income to avoid the measure itself having insurance effects, since income is implicitly taxed by implicit insurance more in some states of the world than in others.

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

$$\underbrace{Cov\left(\widehat{\lambda}, V\right)}_{\text{Risk protection value}} = E\left[\underbrace{\left(\widehat{\lambda} - E\left(\widehat{\lambda}\right)\right)}_{\text{Marginal utility gap}} \times \underbrace{(V - E(V))}_{\text{Value gap}}\right], \tag{6}$$

is a first order approximation to the risk protection value of the change. Risk protection value is increasing in the extent to which V is an *indicator* of marginal utility, that is, in the extent to which the change in constraints benefits the individual more when marginal utility is higher. The change in constraints has positive risk protection value, i.e., is worth more ex ante than its mean ex post value, 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 risk protection value; it is worth less ex ante than its mean ex post value.[22]

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 such factors or others might have on risk protection value manifest themselves through this covariance. For instance, if persistent health shocks not only increase out-of-pocket spending but also

[21]This follows from plugging equation (2) into equation (3) and taking first order approximations around the original allocation (see Appendix C.1). The risk protection value covariance—closely related to that in Finkelstein et al. (2019a) and analogous to a redistribution value covariance in optimal taxation—generalizes the risk protection part of the Baily-Chetty social insurance analysis (Baily, 1978; Chetty, 2006) to cases where the ex post value of the change in constraints can take more than two values (see Appendix C.2).

[22]As is clear from equations (4) and (5), negative risk protection value does not imply negative ex ante value, just ex ante value that is smaller than the mean ex post value.

decrease current and future income, that would tend to increase the risk protection value of health insurance. The key advantage of aiming to recover a first order approximation rather than the exact value is the 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 households optimize.[23]

To be clear, this notion of risk protection—the traditional one in economics and a key determinant of welfare—is distinct from certain intuitive notions of risk protection from health insurance. One is reduced risk in health or well-being, in the sense of reduced volatility of health or well-being across states of the world. For example, individuals who are risk averse in health would value reductions in health risk (Lakdawalla and Phelps, 2020), not just financial risk. Another is protection from the risk (here meaning "safeguard against the possibility") that if one experiences a health shock, one might face the dilemma of forgoing desired care or risking a huge bill, or even be unable to secure desired care at any price (e.g., if a special treatment is available only to individuals with health insurance). This "access motive," which could be extremely valuable (Nyman, 1999), contributes to the risk protection value of health insurance to the extent that the ex post value of the access motive is correlated with marginal utility (as can be seen from equation (5)).[24]

[23]With optimization, a change in constraints can be valued to first order using only the constraints and the status quo allocation; behavioral responses have no first-order effect on utility by the envelope theorem. Economic logic and results of the structural model both suggest that the approximation error of the sufficient statistic tends to work against the main conclusions. Intuitively, it tends to overstate the benefit of insuring health care costs and understate the cost of amplifying other risks, since it neglects that the marginal benefit of reducing a distortion declines as the distortion shrinks (and vice versa).

[24]Access is presumably extremely valuable ex post in certain states of the world and may have considerable ex ante value as well (Nyman, 1999). However, the sign of its contribution to risk protection value is theoretically ambiguous due to opposing effects. On one hand, worse health increases both marginal utility and access value—a force toward a positive correlation. On the other hand, worse non-health circumstances increase marginal utility but *decrease* access value (by increasing implicit insurance support and reducing the

4.2 Sufficient statistic estimates

Implementation.— I estimate the risk protection value of hypothetical increases in health insurance coverage from the status quo. So the sufficient statistic, a generalization of the risk protection part of the Baily-Chetty approach, depends only on marginal utility and the ex post value in the status quo. As discussed in Section 2, I estimate risk protection value from three main perspectives. The short run perspective of immediately before the coverage begins is based on the following regression:

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

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 C.5).[25] 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

value of health in terms of the household's scarce resources ex post)—a force toward a negative correlation.

[25]My regressions based on equation (7) are the health insurance analogue of a common specification in the literature on unemployment insurance (e.g., Hendren, 2017). As discussed in Section 2, the goal is to estimate a covariance across states of the world, not a causal effect of health insurance.

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 cubic in age, and a quadratic in household size.[26]

Consumption and marginal utility, c_{it}, λ_{it}.— My main specifications follow the common practice of modeling marginal utility as a decreasing function of consumption. 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 a baseline, I assume state-independent, constant relative risk aversion utility, $\lambda_{it} = c_{it}^{-\gamma}$, with $\gamma = 3$. I test robustness to many alternative assumptions about marginal utility, including different measures and models of consumption and different assumptions about state-dependent utility. Using measured rather than modeled consumption ensures that the key relationship, between consumption and out-of-pocket 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_{it} is $V_{it} = \max\{0,\ oop_{it} - d_{it}\}$. Standard contracts provide the same coverage regardless of circumstances: $d_{it} = \overline{d}$. Full coverage is the special case with $\overline{d} = 0$: $V_{it} = oop_{it}$. I also consider contracts with a stop-loss that is increasing in realized income. For example, a contract that limits out-of-pocket spending to 10% of income has $d_{it} = 0.10 \times y_{it}$.[27] As discussed in Section 2, this

[26] 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.

[27] Feldstein and Gruber (1995) propose similar contracts, except with coverage below the stop-loss as

follows the standard approach of focusing on out-of-pocket 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.— *Comprehensive coverage.—* Table 1 presents estimates of the risk protection 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 risk protection 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 a risk protection value 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, though valuable, is worth less ex ante than the mean ex post value worth of cash in the same states by 20%, 43%, and 71%, respectively. Filling the gaps in coverage in insured states has a risk protection cost that is about half that of providing full coverage in uninsured states in the short and medium runs but similar in the long run. Figure A5 shows that the estimated risk protection value decreases roughly linearly in the time until coverage begins, as more and more risk remains to be revealed.

Other types of coverage.— Each of a wide variety of types of standard coverage I have investigated has negative risk protection value.[28] Risk protection value is negative for coverage of different types of health care, including hospital care (see Table A11); for different levels of coverage, from minimal catastrophic to full (Figure A6); and for each education group and each of several subsets of the state space, including high or low liquidity and good or bad health (Table A12). That standard coverage has negative risk protection value is robust

well. I focus on full coverage above a threshold in part because its effect on out-of-pocket spending is straightforward to infer even with unobserved, nonlinear implicit taxation by implicit insurance.

[28]Because the estimated risk protection value of standard coverage becomes more negative as more risk is included, to be conservative the heterogeneity and robustness analyses focus on the short run perspective.

to a wide range of assumptions about marginal utility, including health-dependent utility (Table A13); to a wide range of changes in the regression specification (Table A14); and to large private benefits from improved health and reduced medical debt (Appendix C.6 and Table A15). The key empirical relationship underlying this conclusion is that depicted in Figure 3: On average, higher out-of-pocket spending is associated with higher consumption.[29]

Contracts that account for other risks could provide better risk protection. Table A16 considers one contract with a stop-loss of 10% of income ($d_{it} = 0.10 \times y_{it}$) and one with a stop-loss \$500 below that (to provide more coverage to improve statistical precision: $d_{it} = \max\{0, 0.10 \times y_{it} - 500\}$). The estimates, though somewhat imprecise given how rarely out-of-pocket spending exceeds these thresholds, are suggestive that such contracts could provide valuable risk protection. For example, providing full coverage above a stop-loss of \$500 less than 10% of income in uninsured states has a risk protection value of \$10 from the perspective of immediately before the coverage begins, \$280 from ten years before the coverage begins, and \$110 from behind the veil, despite a mean ex post value of just \$250.[30]

4.3 Structural analysis of mechanisms

To better understand the underlying mechanisms and assess generalizability, this section develops and analyzes a simple model guided by the key empirical regularities. The model is based on standard models of health care cost risk but adds other risks.

[29]Finkelstein et al. (2019a) also estimate that the risk protection value of comprehensive health insurance is robustly negative, using a variety of specifications in both the PSID and the Consumer Expenditure Survey. 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 risk protection value, Appendix E 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.

[30]To better understand the results, Appendix F uses a simple model to derive connections between the sufficient statistic results and the three key findings of the descriptive analysis of out-of-pocket spending.

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

$$c(h, y; HI) = \max\{\underline{c},\ y - [tot(h) - hi(h, y; HI) - ihi(h, y; HI)]\}, \tag{9}$$

where $\underline{c}$ is the consumption floor, $tot(h)$ is the total cost of the household's health care consumption, $hi(h, y; HI)$ is the health insurance benefit, if any, and $ihi(h, y; HI)$ is "implicit health insurance" support. Ex ante expected utility is the expected value of a state-independent constant relative risk aversion function of consumption,

$$v(h, y; HI) = \frac{c(h, y; HI)^{1-\gamma}}{1-\gamma}, \tag{10}$$

where γ is the coefficient of relative risk aversion.

Empirical inputs.— The key ingredients are the joint distribution across states of the world 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. Income is the maximum of residualized total annual income and an income floor of \$15,000 (about the tenth percentile). Total health care costs are as before (see Section 2), 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 likely the two most important elements of the budget constraint in this context: out-of-pocket spending and income, including income losses from bad health and income support from unemployment insurance and other sources.[31]

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

$$ihi(h, y; HI) = \max\{0,\ tot(h) - hi(h, y; HI) - d_{ihi}(y)\}. \tag{11}$$

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, shown in Figure A7). 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 consider three main health insurance contracts: full coverage, catastrophic coverage above a \$5,000 deductible, and catastrophic coverage above 10% of income.

Remarks.— As in the simplest standard approach, everything is driven by the budget constraint and consumption equals income minus out-of-pocket spending. The key difference is that here, out-of-pocket spending is not exogenous with respect to other non-consumption elements of the constraint. Out-of-pocket 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

[31] The generous income floor tends to understate income risk, and moral hazard-inclusive health care consumption tends to overstate health care cost risk. Both tend to overstate risk protection from standard contracts. The moral hazard factor of 25% is the change in utilization from health insurance in the Oregon Health Insurance Experiment (Finkelstein et al., 2012).

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: It insures health care risk but may amplify income risk.[32]

Results.— *Risk protection.—* The model matches well the key empirical patterns, including those not targeted directly such as the sufficient statistic estimates. Here too, I find that catastrophic coverage provides better risk protection than comprehensive coverage and income-dependent coverage would provide better risk protection still (see Table A17). Here too, standard coverage increases consumption risk. With the baseline parameters, the risk protection from comprehensive coverage is as costly as reducing consumption in all states of the world by $490 per year. This is about 10% of gross benefits and 19% of net benefits (mean ex post value). Even standard catastrophic coverage has (modestly) negative risk protection value. Hence, the small change to the standard approach of accounting for other risks can explain why standard contracts would increase consumption risk, to an extent broadly similar to that implied by the sufficient statistic estimates.[33] Income-dependent coverage, by contrast, can provide valuable risk protection. In the baseline model, catastrophic coverage above 10% of income provides risk protection worth $730 despite having a mean ex post value of just $100. The key qualitative conclusions are highly robust. They hold for households with less income risk and less implicit insurance than typical households (see Figure 4 and Table A18) and for many other changes to the model (Table A19).

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

[33] 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 of the risk protection value of comprehensive coverage are −$440 and −$720, respectively (versus −$490 in the structural model).

Other risks.— The interaction with income risk reverses the risk protection ranking of the three contracts (see Figure 4a and Table A17). Without income risk, comprehensive coverage provides the best risk protection ($70 more valuable than that of income-dependent coverage). But with non-negligible income risk, income-dependent coverage provides the best risk protection ($1,220 more valuable than that of comprehensive coverage with the baseline income risk). This reversal comes from the opposite-signed effects of income risk on the risk protection from standard versus income-dependent contracts. As income risk grows, the risk protection from standard contracts becomes less valuable, more so for more comprehensive contracts. With non-negligible income risk, the amplification of such risk by standard contracts outweighs the protection against health care costs. In the baseline model, the amplification of income risk by comprehensive coverage has a welfare cost of about $560 per year (risk protection value of −$490 in the baseline versus $70 without income risk), eight times the value of its protection against health care costs. For income-dependent coverage, by contrast, greater income risk increases its risk protection value by increasing the value of its buffer against income risk. This buffer is why such coverage provides such valuable risk protection despite providing only catastrophic protection against health care costs.

Implicit insurance.— Implicit insurance significantly reduces the risk protection value of each type of health insurance without affecting the ranking (see Figure 4b and Table A17). The ranking is preserved due to the interaction with income risk, since income-dependent coverage insures it whereas standard contracts amplify it.[34] Without implicit insurance, all three contracts would provide highly valuable risk protection, with an annual risk protection surplus of $3,150 for income-dependent coverage, $1,500 for catastrophic coverage, and $1,310 for comprehensive coverage (versus $730, −$50, and −$490, respectively, with implicit insurance). The large impact of implicit insurance on risk protection value—though not on the ranking of the contracts—reflects two reinforcing effects. First, implicit insurance de-

[34]Standard contracts amplify income risk even without implicit insurance due to subsidizing normal goods and displacing the implicit income insurance from the consumption floor.

creases the extent to which income-dependent coverage insures income risk and increases the extent to which standard contracts amplify it. Second, implicit insurance transforms health insurance protection against health care costs from highly valuable to a matter of near indifference. Without income risk, risk protection from comprehensive coverage is worth $3,060 per year without implicit insurance but just $70 per year with it. So health insurance protection against health care costs is of little value not because protection against 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 relatively small.

How generalizable are the results? Other risks and implicit insurance vary significantly across households and economies. Still, the results suggest that they would have to be quite different from those of typical U.S. households to reverse the conclusions that less comprehensive coverage provides better risk protection than more comprehensive coverage and that income-dependent coverage would provide better risk protection still (see Figure 4 and Table A18). Standard contracts would amplify other risks in a variety of settings due to subsidizing normal goods and undoing the protection from implicit insurance. Of course, the smaller the other risks, the less important any interaction with them. But other risks would have to be much smaller than those facing typical U.S. households to reverse the risk protection ranking of different types of coverage. 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. More generally, many societies have uneven safety nets that provide more protection against health care costs than other risks. Where other risks exceed health care cost risk, even a small interaction with other risks can outweigh the effect on health care costs.

5 Implications

What type of health insurance would be best for 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 less comprehensive coverage would be better than more comprehensive coverage, and income-dependent coverage would be better still. Rough, back-of-the-envelope calculations suggest that the annual net increase in risk protection value less moral hazard cost from switching from comprehensive coverage to catastrophic or income-dependent coverage would exceed $2,000 per household or $275 billion in the U.S.[35]

But risk protection and moral hazard are *not* the only considerations: Coverage of certain types of health care can have a variety of benefits beyond risk protection. It reduces reliance on implicit insurance and the associated costs (e.g., Garthwaite et al., 2018). It can help correct under-consumption of valuable care (e.g., Baicker et al., 2015). It can help individuals secure care that otherwise would be unaffordable (Nyman, 1999). It could reduce risk in health itself, which would confer a different kind of risk protection benefit if individuals are risk averse over health (Lakdawalla and Phelps, 2020). It can help internalize externalities from infectious disease and health care innovation. It can confer positive fiscal externalities on the government through lower health care costs and greater tax revenue due to improved health and increased productivity.[36] It can help satisfy altruistic feelings about the health of others. It reduces the role of ability to pay in determining the consumption of goods and

[35]This assumes a reduction in moral hazard costs of $1,500 per household and an increase in risk protection value of over $500. Finkelstein et al. (2019a) estimate an average per-person moral hazard cost of comprehensive coverage of around $750, which I double to get a rough estimate for households. My sufficient statistic estimates of long-run risk protection value imply gains of over $600 per household.

[36]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 recipients at negative cost to the government, are desirable under a wide range of social welfare functions, regardless of the risk protection value.

services that many view as a right or moral imperative. It might facilitate more redistribution than would be politically feasible in cash. Although much remains unknown about many such benefits, a variety of evidence suggests that the overall magnitude could be quite large. Hence, risk protection is just one consideration among many, and contracts that provide better risk protection may not be better all things considered.

Which income-dependent contracts best balance risk protection and feasibility? While income-dependent coverage can provide valuable risk protection, implementation requires aligning coverage with a verifiable income measure. The choice of income measure involves a tradeoff between risk protection and feasibility. Typically, finer or more-current income measures offer better risk protection but incur greater implementation costs than coarser or less-current measures. For example, cost-sharing reductions (CSRs) for low-income enrollees in health insurance exchange plans use a relatively coarse income measure that reflects both predicted and actual current income (see Appendix G). I find that the coarse income dependence of CSRs does not significantly reduce risk protection, but basing coverage on predicted or past income does (see Appendix G and Tables A21 and A22). This highlights a key tradeoff: Basing coverage on current income is best for risk protection but necessitates an ex post adjustment to align coverage with actual income.[37]

Is it possible to better insure health risk? Individuals with and without health insurance alike are exposed to considerable risk from health shocks (e.g., French and Jones, 2004; Dobkin et al., 2018; Meyer and Mok, 2019). Much of this risk comes from a source beyond the reach of standard health insurance coverage: income losses from bad health. Supplementing coverage of health care costs with indemnity insurance that pays a fixed cash benefit based on one's health diagnosis, or proxies thereof, could help.[38] I estimate that a

[37]This adjustment could follow the general approach of ACA premium tax credit reconciliation to retroactively adjust cost-sharing based on actual income. See Appendix G for further discussion.

[38]For example, a critical illness insurance policy might pay $10,000 in the event of a heart attack. Indemnity insurance is widely available, typically as a supplement to standard health insurance.

hospital day indemnity would generate considerable risk protection value: 59 cents per dollar of expected value in non-elderly uninsured states, 81 cents in non-elderly insured states, and 28 cents in elderly states (see Table A23).

What are the implications for policy? The analysis indicates the importance of accounting for other risks when designing health insurance contracts and related policies. The risk protection from health insurance depends on not only health care risk but individuals' broader economic vulnerability. Because of the interaction with other risks, comprehensive coverage tends to increase the dispersion in consumption both across states of the world and across the income distribution. Less comprehensive coverage would mitigate these effects, and income-dependent coverage could reverse them. Concentrating coverage on care that is less income elastic could mitigate these effects as well.

The analysis provides further evidence of the importance of the Samaritan's dilemma for health policy. Policies to address the externality from implicit insurance could potentially aim to encourage contracts that provide greater coverage when circumstances are worse, rather than comprehensive coverage regardless of circumstances. In addition to potentially targeting the externality more precisely, this could improve risk protection and moral hazard. That implicit insurance provides valuable protection against other risks reflects the unevenness of the safety net; the safety net provides more protection against health care costs than other risks. More uniform protection might provide better risk protection. With the current safety net, standard coverage is complementary with reducing other risk exposures.

6 Conclusion

The risk protection from health insurance is transformed by the interaction with other risks beyond health care costs. Standard contracts amplify other risks, due to subsidizing normal goods and undoing the protection against other risks from implicit insurance. Alternative contracts that provide more coverage when other circumstances are worse, such as contracts that limit out-of-pocket spending relative to income, would amplify other risks less and potentially even insure them. Such contracts can provide valuable protection against health care costs and other risks alike. Because of the interaction with other risks, catastrophic coverage tends to provide better risk protection than comprehensive coverage, and income-dependent coverage would tend to provide better risk protection still.

An important priority for future research is to quantify other major components of the overall welfare effect of different health insurance contracts, especially the non-insurance benefits of standard contracts and the costs and benefits of alternative contracts that account for other risks. That the interaction with other risks reverses widely-held views about the risk protection from different types of contracts raises the possibility of identifying changes that would improve individual well-being and social welfare.

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.

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.

Barcellos, Silvia Helena and Mireille Jacobson, "The Effects of Medicare on Medical Expenditure Risk and Financial Strain," *American Economic Journal: Economic Policy*, November 2015, *7* (4), 41–70.

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.

Brevoort, Kenneth, Daniel Grodzicki, and Martin B. Hackmann, "The credit consequences of unpaid medical bills," *Journal of Public Economics*, 2020, *187* (C).

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.

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: Nonhealth risk and the value of disability insurance," *Econometrica*, 2022, *90* (4), 1781–1810.

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.

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.

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.

Gabel, Jon, Heidi Whitmore, Matthew Green, Adrienne Call, Sam Stromberg, and Rebecca Oran, "The ACA's Cost-Sharing Reduction Plans: A Key to Affordable Health Coverage for Millions of US Workers," *Issue Brief (Commonwealth Fund)*, 2016, *35*, 1–12.

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 and Matthew Notowidigdo, "Health insurance and the consumer bankruptcy decision: Evidence from expansions of Medicaid," *Journal of Public Economics*, 2011, *95* (7), 767–778.

_ **, 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.

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.

Hu, Luojia, Robert Kaestner, Bhashkar Mazumder, Sarah Miller, and Ashley Wong, "The effect of the Affordable Care Act Medicaid expansions on financial wellbeing," *Journal of Public Economics*, 2018, *163*, 99–112.

Kluender, Raymond, Neale Mahoney, Francis Wong, and Wesley Yin, "The Effects of Medical Debt Relief: Evidence from Two Randomized Experiments," NBER Working Papers 32315, National Bureau of Economic Research, Inc April 2024.

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.

Lakdawalla, Darius N and Charles E Phelps, "Health technology assessment with risk aversion in health," *Journal of Health Economics*, 2020, *72* (C).

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

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.

Mazumder, Bhashkar and Sarah Miller, "The Effects of the Massachusetts Health Reform on Household Financial Distress," *American Economic Journal: Economic Policy*, August 2016, *8* (3), 284–313.

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.

Nyman, John A., "The value of health insurance: the access motive," *Journal of Health Economics*, April 1999, *18* (2), 141–152.

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.

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

U.S. 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.

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.

Figures and Tables

Figure 1: Out-of-pocket spending risk is relatively small

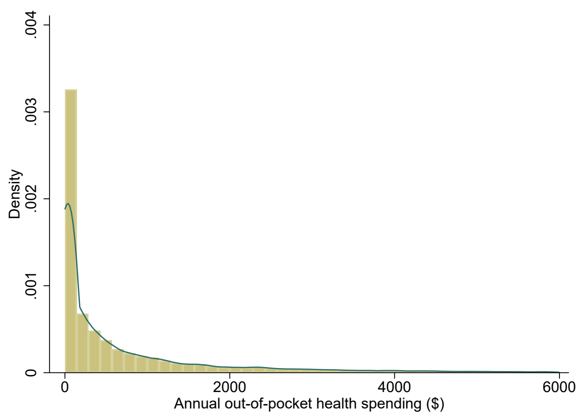

(a) OOP spending among uninsured HHs

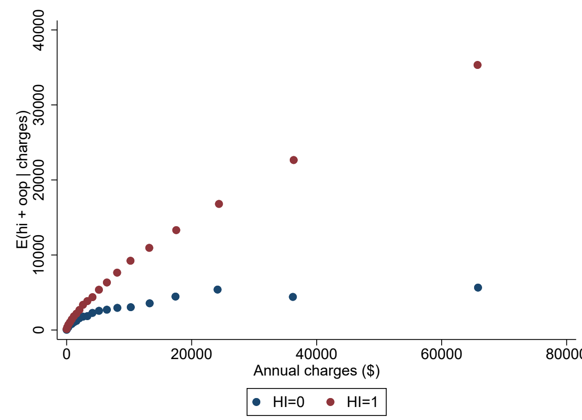

(b) Implicit health insurance

Notes: Left panel: Histogram and estimated kernel density function of annual out-of-pocket 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 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.

Figure 2: Out-of-pocket spending buffers other risks

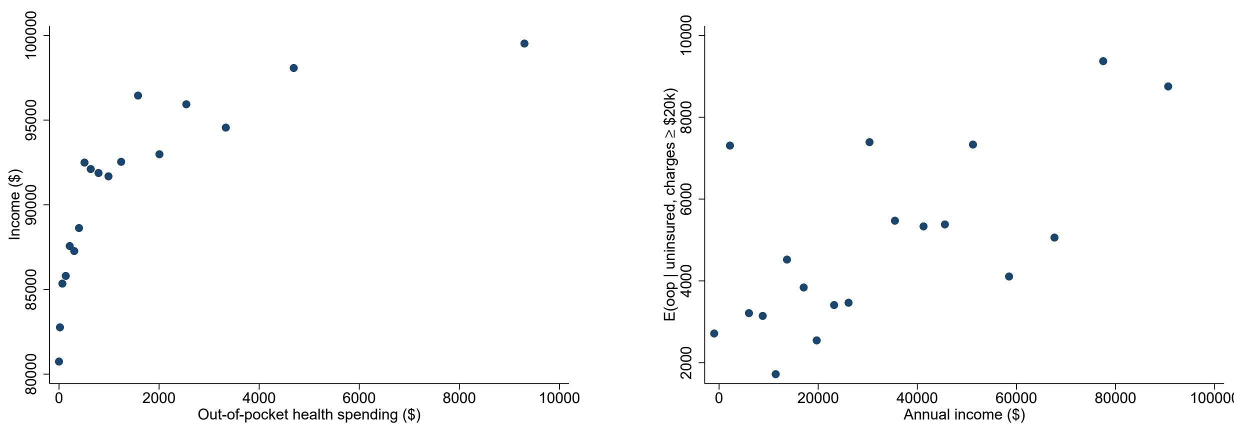

(a) Out-of-pocket spending buffers income risk (b) Implicit HI support decreases in income

Notes: Left panel: Conditional mean of income as a function of out-of-pocket 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). The average height is mean income. The level of a dot may differ from mean income in that bin because of the controls. Right panel: Conditional mean of out-of-pocket 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).

Figure 3: Out-of-pocket spending buffers consumption risk

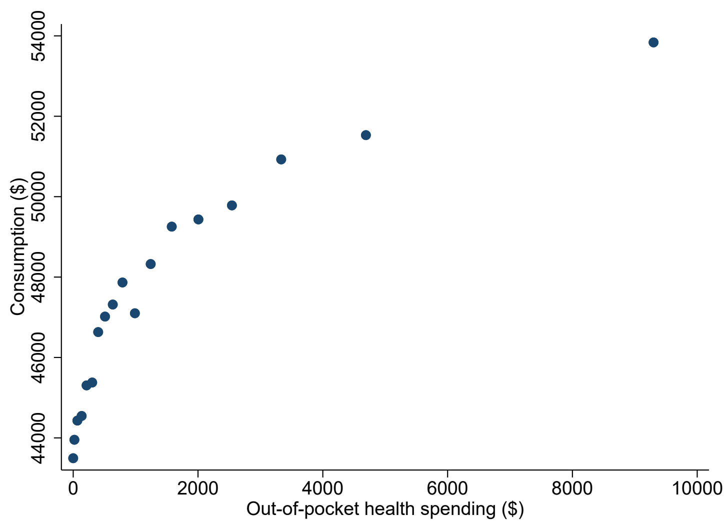

Notes: Conditional mean of non-health consumption spending as a function of out-of-pocket spending among non-elderly households in the PSID, controlling for household fixed effects, year dummies, and a cubic in age. This is a binned scatter plot using the methods of Cattaneo et al. (2019). The average height is mean consumption. The level of a dot may differ from mean consumption in that bin because of the controls.

Figure 4: Structural analysis of mechanisms

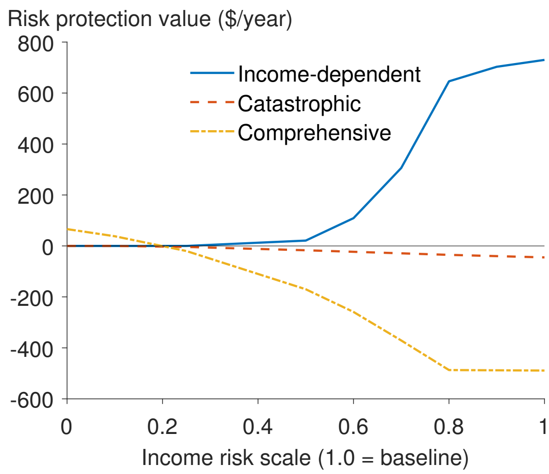

(a) Other risks

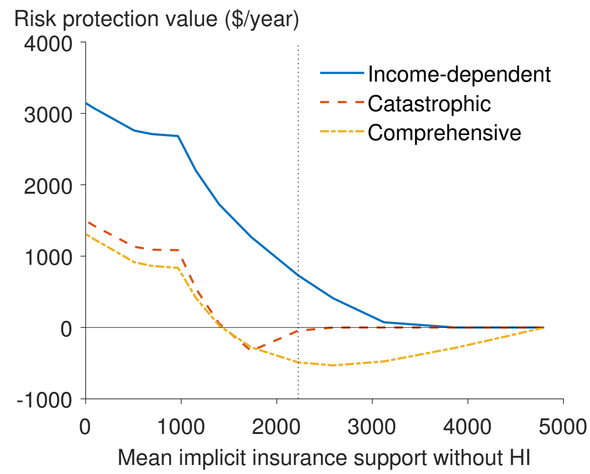

(b) Implicit insurance

Notes: Risk protection value of different types of health insurance as a function of income risk (left panel) and implicit insurance (right panel). In the left panel, adjusted income in state of the world ω is $\widetilde{y}_\omega = \alpha y_\omega + (1-\alpha)y_{med}$, where y_{med} is median income and α is the income risk scale. So when the income risk scale is zero, there is no income risk; income in every state equals median income. An income risk scale of one is the baseline risk process, 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. In the right panel, implicit insurance is varied by multiplying its deductible function, $d_{ihi}(y)$, by a scaling factor. The measure of implicit insurance is the mean of implicit insurance support across states of the world if the individual does not have formal health insurance. This mean is $2,220 in the baseline calibration, shown by the vertical dotted line. Zero corresponds to no implicit insurance. The rightmost value of $4,800 corresponds to complete coverage of all costs. Risk protection value is the amount by which the ex ante equivalent variation of health insurance exceeds its mean ex post value (see equation (4)), using consumption-based equivalent variation. See Table A17 for related statistics.

Table 1: Risk protection 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))	.09	.17	.25	.05	.13	.29	.03	.07	.23
(*se*)	(.017)	(.027)	(.014)	(.007)	(.011)	(.008)	(.015)	(.018)	(.014)
Risk protection 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	-.20	-.43	-.71	-.06	-.19	-.50	-.04	-.10	-.38

Notes: Statistics related to the value of completing health insurance in 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. "Risk protection 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 risk protection 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.

Online Appendix

Health Insurance and Consumption Risk

Lee M. Lockwood

A Data Appendix

A.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 were collected, such as many of the measures of consumption and out-of-pocket spending. 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 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 in the following, aligns better with the timing of the consumption and income variables, the results are similar but stronger, implying larger risk protection costs of standard contracts.

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 greater risk protection costs of standard contracts. 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., Table A13). Appendix E presents 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 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 insured or uninsured using the main measures of household-level health insurance coverage in the PSID. What could be a complicating factor that turns out to be useful is that this main measure of health insurance coverage at the household level changes 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, because of implicit insurance, the insured and uninsured are fairly similar in terms of their protection against health care costs. 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 instead I use raw rather than winsorized measures of consumption and out-of-pocket spending, my main sufficient statistic estimate of the short run risk protection value of comprehensive health insurance for non-elderly uninsured households decreases 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 Table A1.

A.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 less than 25 years old. The resulting sample has 268,235 family-year observations.

Out-of-pocket 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 A.1, there is no single ideal way to classify households as 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 A.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 time 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 Table A2.

B Implicit Insurance Protection for Households in Strong Financial Positions

Although implicit insurance provides more protection to households in worse financial positions, it provides considerable protection to households in strong financial positions as well. 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). 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 or charity care or would not pay a medical bill in full.

C Sufficient Statistic Approximation to Risk Protection Value

C.1 Derivation of equation (5)

Recall equation (5),

$$\underbrace{EAV}_{\text{Ex ante value}} \approx \frac{E(\lambda \times V)}{E(\lambda)} = \underbrace{E(V)}_{\text{Mean ex post value}} + \underbrace{Cov\left(\hat{\lambda},\ V\right)}_{\text{Risk protection value}}. \tag{5}$$

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 (3),

$$E\left[u\left(c_0 + EAV, a_0; \theta\right)\right] = E\left[u\left(c_1, a_1; \theta\right)\right]. \tag{3}$$

Use equation (2),

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

to write the right-hand-side of equation (3) in terms of V:

$$E\left[u\left(c_0 + EAV, a_0; \theta\right)\right] = E\left[u\left(c_0 + V, a_0; \theta\right)\right]. \tag{12}$$

Take first-order approximations to the utility levels inside the expectations on both sides of equation (12) around the allocation under the original constraint, (c_0, a_0):

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

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 $E[u(c_0, a_0; \theta)]$ from both sides and moving the constant EAV outside the expectation yields

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

Solving for EAV yields

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

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

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

Passing $E[u_c(c_0, a_0; \theta)]$ 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

$$EAV \approx E(V) + Cov\left(\widehat{\lambda}, V\right), \tag{17}$$

which is equation (5), as was to be shown.

C.2 Relationship to Baily-Chetty

The "risk protection value" covariance, $Cov\left(\widehat{\lambda}, V\right)$, generalizes the risk protection 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 risk protection value covariance can be written,

$$\begin{aligned} Cov\left(\widehat{\lambda}, V\right) &= E\left[\left(\widehat{\lambda} - E\left(\widehat{\lambda}\right)\right)(V - E(V))\right] \\ &= p\left(V_H - E(V)\right)\left[E\left(\widehat{\lambda}|V = V_H\right) - E\left(\widehat{\lambda}\right)\right] \\ &\quad + (1-p)\left(V_L - E(V)\right)\left[E\left(\widehat{\lambda}|V = V_L\right) - E\left(\widehat{\lambda}\right)\right]. \end{aligned} \tag{18}$$

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

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

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 risk protection value covariance is about the value of the change in constraints, not the cost, so causal effects of the change are not necessary.

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

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

In both cases, the key statistic is the "risk protection value covariance" of equation (5) (which Finkelstein et al. (2019a) call "pure-insurance value"): 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.[39]

The difference is that Finkelstein et al. (2019a) make two assumptions to go beyond a first

[39] Finkelstein et al. (2019a) use exogenous variation in health insurance, generated by the Oregon Health Insurance Experiment, for several purposes, just not for estimating risk protection value in their optimization approaches. For example, they use it to estimate the cost of providing the coverage, the value of the coverage in their "complete-information approach" (described in the next section), and the private value of the moral hazard response in their optimization approaches.

order approximation. They assume that (i) the marginal risk protection 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 risk protection 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 risk protection value of going from no coverage to full coverage—which is the integral of the marginal risk protection value over that range of coverage—is one-half the marginal risk protection 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 risk protection value at an unobserved counterfactual coverage level with an assumed functional form of the marginal risk protection 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 risk protection value at any given coverage level is ambiguous.[40] Another option would be to combine my estimates of the marginal risk protection value of increasing coverage from its status quo level among households with versus

[40]This is true even at full coverage. Although the marginal risk protection 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 risk protection value. This marginal risk protection 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 risk protection 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).

without health insurance, plus an assumed functional form of the marginal risk protection 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 risk protection 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 the key conclusions. Economic logic and quantitative results of the structural model of Section 4.3 both suggest that the approximation error tends to make the sufficient statistic overstate the risk protection value of standard health insurance coverage. 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 amplifying other risks by ignoring that the marginal cost of increasing a distortion increases as the size of the distortion increases. In the structural model, this bias is substantial. The sufficient statistic approximation in the simulated data is −$280, significantly less negative than the exact value of −$490. This suggests that the sufficient statistic estimates are upper bounds on the risk protection value of standard health insurance coverage (i.e., they are less negative than the true value). Analogous reasoning and results of the structural model are also suggestive that the approximation error tends to make the sufficient statistic understate the risk protection value of income-dependent health insurance coverage, again working against the key conclusions.

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.[41] 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,

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

where $\bar{c}$ 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 out-of-pocket spending (an instance of a common class of models in the literature on health care cost 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 risk protection value.

C.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 previous 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 based most closely on the causal effects of health insurance.

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

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 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 increase in 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 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 is 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).[42] These considerations are why I focus on an alternative optimization approach instead.

[42]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 of health insurance 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 of health insurance 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.

C.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 log-linearization of marginal utility:

$$\log\left(\widehat{\lambda}\right) \approx \widehat{\lambda} - 1, \tag{20}$$

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

$$\widehat{\lambda} \approx 1 + \log\left(\widehat{\lambda}\right) = (1 - \log[E(\lambda)]) + \log(\lambda) = (1 - \log[E(\lambda)]) - \gamma\log(c), \tag{21}$$

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:

$$\log(V) \approx \log(E(V)) + \frac{1}{E(V)}(V - E(V)), \tag{22}$$

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

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

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:

$$Cov\left(\widehat{\lambda}, V\right) \approx -\gamma Cov\left(\log(c), \log(V)\right) E(V) = -\gamma\beta Var(\log(V))E(V) \approx -\gamma\beta \frac{Var(V)}{E(V)}, \quad (24)$$

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.

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

The main effects of health insurance on individuals are reduced out-of-pocket spending, improved health, and reduced medical debt (Finkelstein et al., 2018). My baseline specifications focus on out-of-pocket spending. 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 (see footnote 11). Perhaps in part from such considerations, the vast majority of analyses of the risk protection value of health insurance focus on out-of-pocket spending.[43]

That health insurance also improves health and reduces medical debt may increase its ex ante value and mean ex post value a great deal. Although the ex post value of improved health, from moral hazard effects on health care consumption, is second order for optimizing individuals, individuals might fail to optimize and 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

[43] Important exceptions include Gross and Notowidigdo's (2011) analysis of bankruptcy, Finkelstein et al.'s (2019a) "complete-information approach" to estimating the value of Medicaid (which I discuss in Appendix C.4), and Brevoort et al.'s (2020) analysis of medical debt.

even absent any reduction in out-of-pocket spending or medical debt. And although reduced medical debt has clearer benefits to creditors than to individuals and evidence of benefits to individuals "remains limited" (Finkelstein et al., 2018, p. 270), that does not imply that the benefits to individuals are always small.[44]

Though the effects of health insurance on health and medical debt may be of considerable value, their effect on risk protection is more subtle. They do not directly affect risk in net income or consumption, and the sign of their contribution to the risk protection value of health insurance is ambiguous in theory. Their contribution to risk protection value is the covariance between marginal utility and the ex post value of the health improvements and medical debt reductions (as can be seen from equation (5)). So large ex post values do not necessarily translate into a large or even positive contribution to risk protection value. Rather, what matters is how the *differential* value in some states of the world relative to others covaries with marginal utility. The sign of this covariance is ambiguous in theory due to opposing pro- and anti-insurance effects: a pro-insurance effect from variation in health and an anti-insurance effect from variation in non-health circumstances. On one hand, health improvements and medical debt reductions likely are concentrated in states of the world in which health is worse and marginal utility is high. This is a force toward a positive covariance. On the other hand, the ex post value to the household (in terms of resources in that state of the world) of a given health improvement and medical debt reduction is lower, other things equal, when other circumstances are worse and the marginal utility of consumption is higher. This is a force toward a negative covariance.[45]

[44]While Kluender et al. (2024) find no impact of medical debt relief on credit access, credit utilization, financial distress, or mental health on average in two large-scale randomized experiments, Brevoort et al. (2020) find evidence that the Medicaid expansion from the Affordable Care Act led to better terms of credit.

[45]Recall that risk protection value is the excess of the ex ante value over the mean ex post value (equation (4)). As such, it is distinct from reduced risk in health or well-being, in the sense of reduced volatility of health or well-being across states of the world. Risk protection value includes the individual's valuation of changes in health and well-being but weights these (and other sources of ex post value) by marginal utility in

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 risk protection value of comprehensive health insurance coverage. Columns (2)–(7) increase the ex post value of health insurance by $20,000 in the states of the world in which the private benefit of improved health is likely to be largest relative to that in other states, 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 the 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) in these states relative to other states.[46] The estimated risk protection 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.

The main reason 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 increases

each state. Although standard health insurance contracts increase financial risk, they might decrease health risk (though the mechanisms that cause them to increase financial risk tend to cause them to increase health risk as well) and might even decrease well-being risk (if the possible decrease in health risk outweighs the increase in financial risk).

[46]Not only is $20,000 a large value of the differential ex post surplus to the household from moral hazard in these states relative to other states, 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 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 only account for the pro-insurance effect.

marginal utility by greatly limiting a key channel by which it otherwise would: increased out-of-pocket spending. Regardless of the underlying mechanisms, bad health is a much weaker indicator of marginal utility than, for example, unemployment. That, in turn, is another manifestation of the key proximate reason that standard contracts increase consumption risk: Other risks are much less well-insured than health care costs, even among households without health insurance. As a result, for standard contracts the amplification of other risks outweighs the protection against health care costs.[47]

Columns (8) and (9) test the potential effects of large private benefits of reduced medical debt on the risk protection value of comprehensive 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. 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 reduced medical debt in these states of the world relative to other states.[48] In both cases, the estimated risk protection value remains significantly negative.

[47]The estimated risk protection 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 risk protection value increasing in absolute value (as can be seen from equation (8)). However, the more important result is that even large values of improved health do not change the sign of the key covariance: the ex post value of standard contracts is positively related to consumption (and so negatively related to marginal utility).

[48]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. In practice, the value to households of reducing medical debt appears to be considerably lower than this upper bound (Kluender et al., 2024). Another sense in which these robustness tests are conservative is that they ignore 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 stigma or a given improvement in future credit access).

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 standard contracts have negative risk protection value, i.e., that they are worth less ex ante than their mean ex post value.

D Income Effects of Demand for Health Care: A Force Toward Health Insurance Amplifying Other Risks

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 amplifying 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 during unemployment, health insurance would be worth less in unemployed states of the world and thereby amplify that risk.

To illustrate, suppose two households are considering an elective surgery that costs \$10k and their health insurance covers 50% of the cost. The "lucky" household receives a raise at work and chooses to get the surgery, spending \$5k out of pocket. The "unlucky" household does not receive a raise and chooses to postpone the surgery, spending \$0 out of pocket. Now consider a supplemental insurance policy that covers the remaining 50% of the cost. For the lucky household, this policy increases net income by \$5k. For the unlucky household, however, this policy has no effect on net income. Although the policy provides the same coverage to both households, it increases the net income of the lucky household more, increasing the gap in net income between them.[49]

[49]Of course, the unlucky household may still benefit from the policy if it gets the surgery. If its choice

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 Table A24). 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, like 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 a durable rather than the investment flow, so temporarily postponing investment in a durable can be a lower-cost way of making ends meet when times are tight than cutting non-durable consumption.

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 in keeping 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).

E Measurement Error

Measurement error is a concern for any analysis. The concern is amplified when certain results are contrary to priors. Although classical measurement error would tend to attenuate

to postpone the surgery was privately optimal, the policy increases not only the gap in net income but the gap in well-being. But if its choices are not privately optimal, the policy could potentially reduce the gap in well-being despite increasing the gap in net income. Where the effect on well-being, not just net income, is relevant, I test the robustness of the conclusions to large private benefits from improved health and reduced medical debt (see Appendix C.6).

the results rather than bias them toward health insurance increasing consumption risk, this section considers the possibility that non-classical measurement error in certain key variables might bias the results toward health insurance increasing consumption risk.

A key result in the descriptive analysis of out-of-pocket spending, and a key driver of the sufficient statistic estimates of the risk protection value of health insurance, is that the correlation between out-of-pocket 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.[50] If measurement error in out-of-pocket spending and consumption were positively correlated—i.e., if positive (negative) errors in out-of-pocket 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 different in that they allow respondents to choose

[50]Specifically, Finkelstein et al. (2019a) estimate the correlation between out-of-pocket spending and marginal utility, which they model as a decreasing function of measured consumption spending, or the logs thereof. In all cases across a wide variety of specifications, the estimated correlation is negative.

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 (risk protection value of standard contracts 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 Tables A4 and A5 and the discussion on page 16). Because of this, even setting aside the consumption measures in the PSID and Consumer Expenditure Survey, I estimate negative risk protection value of standard coverage based on simple consumption-proxy measures, for example using a consumption proxy of income minus out-of-pocket spending (see Table A13).[51] Third, I find

[51]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

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 not only in terms of means and standard deviations (see Table A3) but also in terms of correlations with income (see Tables A4 and A24). Fourth, making two small changes to the workhorse model of health care cost 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 (see 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 seems unlikely to explain why such a wide range of evidence based on a variety of approaches—from descriptive evidence about out-of-pocket spending (including not only its marginal distribution but how it relates to consumption, in-

to affect the income-based results.

come, 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—points to the same, robust conclusions.

F Connecting the Sufficient Statistic Estimates to the Descriptive Analysis of Out-of-Pocket Spending

F.1 Analysis

To better understand the results so far and the connections between them, consider a simple model in which consumption equals income minus out-of-pocket spending, $c = y - oop$, and marginal utility is linear in consumption, $u'(c) = u'(\bar{c}) + u''(\bar{c})(c - \bar{c})$. In this case, to first order the risk protection value of health insurance is

$$Cov\left(\widehat{\lambda}, V\right) = -\frac{\gamma(\bar{c})}{\bar{c}} Cov\left(c, V\right) = \frac{\gamma(\bar{c})}{\bar{c}} \left[\underbrace{Cov(oop, V)}_{\text{Out-of-pocket spending}} - \underbrace{Cov\left(y, V\right)}_{\text{Other risks}} \right], \tag{25}$$

where V is the ex post value of health insurance, $\bar{c} \equiv E(c)$ 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 F.2 for derivations of all equations in this section). Health insurance provides better risk protection the more strongly positively its ex post value covaries with out-of-pocket spending (which protects against out-of-pocket spending risk) and the more strongly negatively its ex post value covaries with income (which protects against income risk).

Standard contracts.— The risk protection value of full coverage ($V = oop$) is approximately

$$Cov\left(\widehat{\lambda}, V\right) = -\frac{\gamma(\overline{c})}{\overline{c}} Cov\left(c, oop\right) = \frac{\gamma(\overline{c})}{\overline{c}} \left[\underbrace{Var(oop)}_{\text{"Partial effect"}} - \underbrace{Cov\left(y, oop\right)}_{\text{"Portfolio effect"}} \right]. \quad (26)$$

The first equality connects the sufficient statistic to the relationship between out-of-pocket spending and consumption. The second equality, which is related to equation (1), connects the sufficient statistic to the extent of out-of-pocket spending risk and the relationship between out-of-pocket spending and other risks. Interpreting the empirical results through the lens of this model suggests the following. Standard contracts increase consumption risk (Finding 3 and the sufficient statistic estimates) because the valuable partial effect of insuring health care costs is outweighed by a costly portfolio effect of amplifying other risks. The partial effect is small because out-of-pocket spending risk is relatively small (Finding 1), which in turn is because implicit insurance provides considerable protection against health care costs. The portfolio effect is costly because out-of-pocket spending buffers other risks (Finding 2), which in turn is because certain types of health care are normal goods and because implicit insurance provides greater protection when other circumstances are worse. The portfolio effect cost is significant because households face substantial risk in income, assets, and living expenses, so even a modest amplification of these risks can have a large welfare cost.[52]

[52]In terms of the economics of the second best and equation (6), 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 risk protection cost is increasing in the amount of risk that remains to be revealed because the costly portfolio effect grows more rapidly than the valuable partial effect, since other risks are more persistent than health care costs.

Income-dependent coverage.— The risk protection value of full coverage above an income-dependent deductible ($V = \max\{0, oop - \beta y\}$) is approximately

$$Cov\left(\widehat{\lambda}, V\right) = \frac{\gamma(\overline{c})}{\overline{c}} \left[\underbrace{Var(oop)}_{\text{"Partial effect"}} - \underbrace{Cov\left(y, oop\right) + \beta \left[Var(y) - Cov(y, oop)\right]}_{\text{"Portfolio effect"}} \right], \quad (27)$$

if out-of-pocket spending would otherwise, without the coverage, exceed the deductible in all states of the world, $oop > \beta y$. Comparison of equations (27) and (26) reveals that in this case, a sufficient condition for income-dependent coverage with $\beta > 0$ to provide better risk protection than comprehensive coverage is that income is more variable than out-of-pocket spending.[53] Interpreting the empirical results through the lens of this model suggests that income-dependent coverage provides better risk protection than standard contracts (Finding 3 and the sufficient statistic estimates) by amplifying other risks less or even insuring them (Finding 2). By providing more protection against health care costs when circumstances are worse, such coverage can insure not only income losses from bad health but other risks more generally. The better protection against other risks than comprehensive coverage outweighs the worse protection against health care costs because other risks are much larger than below-deductible health care cost risk (Finding 1).

[53] The difference in risk protection value between income-dependent coverage and comprehensive coverage in this case is $\beta \left[Var(y) - Cov(y, oop)\right] = \beta Var(y) \left(1 - \frac{Sd(oop)}{Sd(y)} Corr(y, oop)\right)$, where $Sd(X)$ is the standard deviation of X. This is positive if $\frac{Sd(oop)}{Sd(y)} Corr(y, oop) < 1$, which is guaranteed if $Sd(y) > Sd(oop)$. In the more general case with below-deductible risk, income-dependent coverage also has a disadvantage relative to full coverage: less protection against below-deductible risk in health care costs.

F.2 Derivation of equations in Section F.1

Equation (25).— Recall equation (25),

$$Cov\left(\widehat{\lambda},V\right)=-\frac{\gamma(\overline{c})}{\overline{c}}Cov\left(c,V\right)=\frac{\gamma(\overline{c})}{\overline{c}}\left[\underbrace{Cov(oop,V)}_{\text{Out-of-pocket spending}}-\underbrace{Cov\left(y,V\right)}_{\text{Other risks}}\right]. \tag{25}$$

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

$$\widehat{\lambda}\equiv\frac{u'(c)}{E[u'(c)]}=\frac{u'(\overline{c})+u''(\overline{c})(c-\overline{c})}{u'(\overline{c})}=\alpha+\frac{u''(\overline{c})}{u'(\overline{c})}c=\alpha-\frac{\gamma(\overline{c})}{\overline{c}}c, \tag{28}$$

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 (28) into $Cov\left(\widehat{\lambda},V\right)$ yields

$$Cov\left(\widehat{\lambda},V\right)=Cov\left(\alpha-\frac{\gamma(\overline{c})}{\overline{c}}c,V\right)=Cov\left(-\frac{\gamma(\overline{c})}{\overline{c}}c,V\right)=-\frac{\gamma(\overline{c})}{\overline{c}}Cov\left(c,V\right), \tag{29}$$

which is the first equality of equation (25). For the second equality, plug $c=y-oop$ into the right-hand-side of equation (29) and rearrange to find

$$-\frac{\gamma(\overline{c})}{\overline{c}}Cov\left(c,V\right)=-\frac{\gamma(\overline{c})}{\overline{c}}Cov\left(y-oop,V\right)=\frac{\gamma(\overline{c})}{\overline{c}}\left[Cov\left(oop,V\right)-Cov\left(y,V\right)\right], \tag{30}$$

which is the second equality of equation (25), as was to be shown.

Equation (26).— Equation (26) follows immediately from plugging $V = oop$ into equation (25):

$$Cov\left(\widehat{\lambda}, V\right) = -\frac{\gamma(\bar{c})}{\bar{c}} Cov\left(c, oop\right) = \frac{\gamma(\bar{c})}{\bar{c}} \left[\underbrace{Var(oop)}_{\text{"Partial effect"}} - \underbrace{Cov\left(y, oop\right)}_{\text{"Portfolio effect"}} \right]. \tag{26}$$

Equation (27).— Recall equation (27):

$$Cov\left(\widehat{\lambda}, V\right) = \frac{\gamma(\bar{c})}{\bar{c}} \left[\underbrace{Var(oop)}_{\text{"Partial effect"}} - \underbrace{Cov\left(y, oop\right) + \beta\left[Var(y) - Cov(y, oop)\right]}_{\text{"Portfolio effect"}} \right]. \tag{27}$$

This is a first order approximation to the risk protection value of full coverage above an income-dependent deductible ($V = \max\{0, oop - \beta y\}$) in the simple model described above if out-of-pocket spending would otherwise, without the coverage, exceed the deductible in all states of the world, $oop > \beta y$.

To see this, first consider the "out-of-pocket spending" term of equation (25), $Cov(oop, V)$. Plugging in $V = \max\{0, oop - \beta y\}$ yields

$$\begin{aligned} Cov\left(oop, V\right) &= Cov\left(oop, \max\{0, oop - \beta y\}\right) \\ &= Pr(oop > \beta y)\Big\{Cov(oop, oop - \beta y | oop > \beta y) \\ &\quad + E(oop - \beta y | oop > \beta y)\left[E(oop | oop > \beta y) - E(oop)\right]\Big\} \\ &= Pr(oop > \beta y)\Big\{Var(oop | oop > \beta y) - \beta Cov(oop, y | oop > \beta y) \\ &\quad + E(oop - \beta y | oop > \beta y)\left[E(oop | oop > \beta y) - E(oop)\right]\Big\}, \end{aligned} \tag{31}$$

where the second equation used that

$$
\begin{aligned}
Cov\left(X, \max\{0, Y\}\right) &= E\left[(X - E(X))(\max\{0, Y\} - E(\max\{0, Y\}))\right] \\
&= E\left[(X - E(X)) \max\{0, Y\}\right] \\
&= Pr(Y > 0) E\left[(X - E(X)) Y | Y > 0\right] \\
&= Pr(Y > 0) \Bigg\{ \underbrace{E(XY|Y > 0) - E(X)E(Y|Y > 0)}_{E[(X - E(X))Y|Y>0]} \\
&\quad + \underbrace{E(X|Y > 0)E(Y|Y > 0) - E(X|Y > 0)E(Y|Y > 0)}_{0} \Bigg\} \\
&= Pr(Y > 0) \Bigg\{ \underbrace{E(XY|Y > 0) - E(X|Y > 0)E(Y|Y > 0)}_{Cov(X,Y|Y>0)} \\
&\quad + E(Y|Y > 0)\left[E(X|Y > 0) - E(X)\right] \Bigg\} \\
&= Pr(Y > 0) \Big\{ E(Y|Y > 0)\left[E(X|Y > 0) - E(X)\right] + Cov(X, Y|Y > 0) \Big\}.
\end{aligned} \tag{32}
$$

Now consider the "other risks" term of equation (25), $-Cov(y, V)$. Plugging in $V = \max\{0, oop - \beta y\}$ yields

$$
\begin{aligned}
-Cov\left(y, V\right) &= -Cov\left(y, \max\{0, oop - \beta y\}\right) \\
&= -Pr(oop > \beta y) \Big\{ Cov(y, oop - \beta y | oop > \beta y) \\
&\quad + E(oop - \beta y | oop > \beta y)\left[E(y|oop > \beta y) - E(y)\right] \Big\} \\
&= Pr(oop > \beta y) \Big\{ \beta Var(y|oop > \beta y) - Cov(oop, y|oop > \beta y) \\
&\quad + E(oop - \beta y | oop > \beta y)\left[E(y) - E(y|oop > \beta y)\right] \Big\},
\end{aligned} \tag{33}
$$

where the second equation used equation (32).

Plugging equations (31) and (33) into equation (25) yields

$$\begin{aligned}
Cov\left(\widehat{\lambda}, V\right) &= \frac{\gamma(\bar{c})}{\bar{c}}\left[Cov(oop, V) - Cov\left(y, V\right)\right] \\
&= \frac{\gamma(\bar{c})}{\bar{c}} Pr(oop > \beta y) \Big\{ Var(oop|oop > \beta y) \\
&\quad + E(oop - \beta y|oop > \beta y)\left[E(oop|oop > \beta y) - E(oop)\right] \\
&\quad - (1+\beta) Cov(oop, y|oop > \beta y) + \beta Var(y|oop > \beta y) \\
&\quad + E(oop - \beta y|oop > \beta y)\left[E(y) - E(y|oop > \beta y)\right] \Big\}.
\end{aligned} \tag{34}$$

First note that in the special case in which out-of-pocket spending would otherwise, without the coverage, exceed the deductible in all states of the world, $oop > \beta y$, equation (34) becomes

$$\begin{aligned}
Cov\left(\widehat{\lambda}, V\right) &= \frac{\gamma(\bar{c})}{\bar{c}}\left[Var(oop) - (1+\beta) Cov(oop, y) + \beta Var(y)\right] \\
&= \frac{\gamma(\bar{c})}{\bar{c}}\left\{Var(oop) - Cov(y, oop) + \beta[Var(y) - Cov(y, oop)]\right\},
\end{aligned} \tag{35}$$

which is equation (27), as was to be shown.

In the more general case with below-deductible risk, i.e., in which out-of-pocket spending is below the deductible in at least some states of the world, $oop < \beta y$, equation (34) is the relevant one. The first two terms inside the curly brackets, $Var(oop|oop > \beta y) + E(oop - \beta y|oop > \beta y)\left[E(oop|oop > \beta y) - E(oop)\right]$, are the valuable insurance of out-of-pocket spending risk. This is positive regardless of income risk. The other terms are from the interaction with income risk. In addition to the effect from any covariance between income and out-of-pocket spending, the greater coverage when income is lower tends to insure not only the income cost of bad health (the $E(oop-\beta y|oop > \beta y)\left[E(y) - E(y|oop > \beta y)\right]$ term) but income risk more generally (the $\beta Var(y|oop > \beta y)$ term).

G Income-Dependent Coverage: Risk Protection and Feasibility

This section analyzes the risk protection from alternative income-dependent health insurance contracts and explores implementation approaches, building on the main text's finding that such coverage can provide valuable risk protection. I examine the tradeoff that, typically, coverage based on finer or more-current income measures offers better risk protection but incurs greater implementation costs than coverage based on coarser or less-current income measures. I examine this tradeoff in the context of two different types of income-dependent coverage: contracts that limit out-of-pocket spending to 10% of income and contracts based on the cost-sharing reductions (CSRs) for low-income enrollees in health insurance exchange plans.

G.1 Risk protection

Approaches.— I analyze the risk protection from alternative income-dependent contracts using both the sufficient statistic and structural approaches. Although the sufficient statistic is my preferred approach for estimating risk protection value when feasible, it cannot consider as wide a range of contracts as the structural analysis. The sufficient statistic approach requires estimating the "mechanical effect" of the proposed change in coverage: the reduction in out-of-pocket spending in the absence of behavioral responses. This is straightforward to estimate for contracts with full coverage above a threshold, but is complicated for more complex contracts due to the nonlinear effects of implicit insurance.

Contracts.— *Finer versus coarser income measures.—* I consider two broad types of income-dependent contracts. One limits out-of-pocket spending to 10% of income—full coverage above a stop-loss of 10% of income. The other is based on what is perhaps the leading example of income-dependent coverage in the U.S.: health insurance exchange plans with cost-sharing reductions (CSRs) for low-income enrollees. CSRs adjust enrollees' costs based on their income during the coverage period. Lower-income enrollees receive more comprehensive coverage. Specifically, while the standard "Silver" plan covers 70% of costs on average, households with income between 200–250% of the federal poverty level (FPL) receive enhanced coverage that covers 73% of costs on average, households with income between 150–200% of the FPL receive enhanced coverage that covers 87% of costs on average, and households with income between 100–150% of the FPL receive enhanced coverage that covers 94% of costs on average.[54] So CSR coverage is based on a coarse income measure of five income bins, whereas the contract that limits out-of-pocket spending to 10% of income is based on a fine income measure—income itself.

Table A20 shows the average cost-sharing parameters for individual coverage through ACA Silver Plans with CSRs in 2016.[55] Lower-income individuals enjoy greater coverage through lower deductibles, lower coinsurance rates, and lower stop-losses. For example, the average deductible decreases from $3,304 for individuals with household income above 250% of the FPL to $265 for individuals with household income between 100–150% of the FPL. By providing greater coverage when income is lower, CSRs reduce the amplification of income risk relative to standard contracts and may even insure it.

In the structural analysis, I consider typical CSR (plus Medicaid) coverage, where the CSR deductibles, coinsurance rates, and stop-losses in each income bin take typical values of actual CSR coverage. I assume that households with income below 100% of the FPL receive full

[54] Households with income below 100% of the FPL are not eligible for CSRs. In states that expanded Medicaid, they are eligible for Medicaid. In states that did not expand Medicaid, they may fall into a "coverage gap" in which they earn too little to qualify for CSRs but too much to qualify for Medicaid. In that case, they have only 70% of their costs covered—the standard actuarial value of "Silver" plans.

[55] I selected 2016 because it is toward the middle of the sample period without being too soon after ACA implementation, to avoid initial transition issues.

coverage of their health care costs, an approximation to the case in which they are eligible for and enroll in Medicaid.[56] I base the CSR contracts on the typical CSR contract parameters shown in Table A20, adjusted as follows. I multiply the deductibles and out-of-pocket maxima by two to reflect typical family coverage. For coverage between the deductible and out-of-pocket maximum, I take the midpoint of the range of coinsurance rates (for different types of health care) in the table.[57]

In the structural and sufficient statistic analyses, I also consider a rough approximation to actual CSR (plus Medicaid) coverage that provides full coverage above a stop-loss, where the stop-loss depends on the CSR income bin in which the household falls. I set the stop-loss in each income bin to be halfway between the average deductible and the average out-of-pocket maximum for family coverage in that bin. This approximates the partial coverage that the actual contracts provide between the deductible and out-of-pocket maximum. As before, I assume that the deductibles and out-of-pocket maxima for family coverage are twice those for individual coverage reported in Table A20. The main reason to consider such "approximate" CSR coverage is that it is straightforward to estimate the mechanical effect, so risk protection value can be estimated with the sufficient statistic approach. In the structural analysis, I find that the risk protection value of the approximate CSR coverage is very similar to that of typical CSR coverage. This suggests that the sufficient statistic estimates of the risk protection value of the approximate CSR coverage are informative about the risk protection value of actual CSR coverage.

More- versus less-current income measures.— For both types of contracts, I consider two

[56]The case in which such households are not eligible for Medicaid involves less risk protection, since coverage is not monotonically decreasing in income.

[57]To isolate the effects of the income-dependence of CSRs from changes in federal poverty levels and other factors over time, I classify household poverty status based on the (inflation-adjusted) federal poverty level thresholds in 2009, the center of my PSID data period of 1999–2019. I use the thresholds for the 48 contiguous states plus Washington DC (`https://aspe.hhs.gov/2009-poverty-guidelines-federal-register-notice`).

different income measures on which to base coverage. The first is current or contemporaneous income: annual income in the year in which the coverage is in force. The second is a measure of *past* income. Specifically, I consider contracts that condition coverage of health care costs in year t on income in year $(t-2)$. For example, the contract that limits out-of-pocket spending to 10% of income provides full coverage of any health care costs in year t above a stop-loss of 10% of income in year $(t-2)$.[58] The "mechanical" reduction in out-of-pocket spending from such a contract is $V_{i,t} = \max\{0, oop_{i,t} - 0.10y_{i,t-2}\}$. The main advantage of contracts based on past income is that they are easier to implement than contracts based on current income, an issue I address in Appendix Section G.2. The main disadvantage is that they may not provide as much risk protection. For example, the contract based on income in year $(t-2)$ provides no protection against income risk that is realized in year $(t-1)$ or t (conditional on income in year $(t-2)$).

In practice, actual CSRs in the health insurance exchange are based partly on ex ante predicted income and partly on ex post realized income. During initial enrollment, individuals must report their predicted income in the coverage year to determine their initial coverage. During the coverage year, individuals must report significant income changes to determine their coverage for the following months. If the individual reports a significant change, the exchange may adjust the individual's plan and cost-sharing level for the following months. But such changes are not retroactive, benefits already received are not recalculated, and

[58] I use annual income in $(t-2)$ rather than $(t-1)$—i.e., not the year immediately before the coverage but the year before that—mainly due to data quality considerations. The PSID, whose waves occur in every second year during my sample period, cautions that the "off" year variables (e.g., income in year $(t-1)$) "may be subject to significantly higher measurement error (e.g. through recall errors) than the corresponding core PSID variables" (see `https://psidonline.isr.umich.edu/data/sl/T-2Income-PSID.pdf`, noting that the PSID refers to the "off" year as $(T-2)$, not $(t-1)$). Another reason to use income two years before coverage is feasibility. Various legal measures of annual income are available only later in the following year. For example, households learn their adjusted gross income (AGI) when they calculate and file their income taxes, which typically happens in the first few months of the following year.

cost-sharing reductions, unlike premium tax credits, are not reconciled at tax time. Hence, the final cost-sharing arrangement reflects a mixture of ex ante predicted income and ex post realized income, which likely reflects a mixture of past and current income. Although data limitations prevent me from analyzing this type of income dependence directly, its risk protection value is likely bounded by those of CSR coverage based on past income and CSR coverage based on current income. I discuss the implementation of alternative income measures in Appendix Section G.2 below.

Results.— Table A21 shows the results of the sufficient statistic estimations, and Table A22 shows the results of the structural analysis. There are two key findings. First, CSR coverage provides risk protection broadly comparable to that of the contract that limits out-of-pocket spending to 10% of income, despite its coarse income dependence. For example, while the contract that limits out-of-pocket spending to 10% of income has a markup in uninsured states of 5% from the perspective of immediately before the coverage begins, 89% from ten years before the coverage begins, and −20% from behind the veil, the corresponding markups for the approximate CSR coverage are −4%, 88%, and 158%, respectively.[59] This suggests that even the coarse income dependence of CSRs is sufficient to avoid much of the costly amplification of other risks by standard income-independent coverage and thereby provide valuable risk protection.

The second key finding is that both types of income-dependent coverage provide better risk protection when their coverage depends on current income rather than past income. For both types of contracts, the risk protection value and markups are greater from all three perspectives when coverage is based on current income rather than past income. The

[59]In the structural analysis, I find that the risk protection value of the approximate CSR coverage is very similar to that of typical CSR coverage: $1,446 and $1,430, respectively. This suggests that the sufficient statistic estimates of the risk protection value of the approximate CSR coverage are informative about the risk protection value of actual CSR coverage, though these estimates are somewhat imprecise given how rarely out-of-pocket spending exceeds the relevant stop-losses.

differences are especially stark from the short-run perspective. From that perspective, past income has already been realized and is no longer risky, so the risk protection from income-dependent coverage based on past income is similar to that from otherwise-similar income-*in*dependent coverage. Even from longer-run perspectives, however, contracts based on past income provide less risk protection than contracts based on current income. This is because contracts based on past income do not provide protection against income risk realized *after* the measure of past income, so there is nothing to offset the amplification of more recent income risk from their coverage of health care costs. Despite that, these contracts provide better risk protection than standard coverage from the medium- and long-run perspectives because their greater coverage when past income is lower provides some protection against income risk up to that point.

G.2 Implementation

More- versus less-current income measures.— The results highlight a key tradeoff: Basing coverage on current income (annual income in the year of the coverage itself) is best for risk protection but necessitates an ex post adjustment to align coverage with actual income. Such an ex post adjustment is akin to that of the tax deduction for out-of-pocket spending above 7.5% of income or the premium tax credit for health insurance exchange plans, both of which depend on current income and are reconciled when the household files its tax return. In principle, the implementation could be similar. At the end of the year or at tax time, households could report their final annual income, out-of-pocket spending, and the total cost of health care they received, and insurers (or government administrators) could credit or bill any difference between the actual cost sharing due and what was paid throughout the year. Households whose income turned out to be lower than expected would receive a rebate for any out-of-pocket spending that should have been covered (given their low income) but was not (because of their higher expected income). Households whose income turned out to be higher than expected would have to repay any coverage they received (because of their lower expected income) but should not have (given their high income).

As mentioned above, the CSRs of actual health insurance exchange plans do not require an ex post reconciliation, since they are not based purely on current income but rather on a mixture of predicted and current income. The lack of ex post reconciliation avoids the associated costs, though the reporting of initial predicted income and of significant income changes throughout the year imposes costs of a similar kind. The coarseness of the income measure mitigates the costs of reporting income changes, as only "significant" changes (that change the individual's income bin) need be reported.

The details of the reconciliation process affect risk protection. For example, a reconciliation based solely on total health costs and income, without accounting for out-of-pocket spending (and thus implicit insurance support), would likely provide greater risk protection than estimated in this analysis.[60]

In summary, while contracts based on current income provide better risk protection than contracts based on past or predicted income, they can be more complex to administer. The CSRs of health insurance exchange plans would likely provide better risk protection if they were based fully on current income, at the cost of requiring an ex post reconciliation (though, depending on the implementation, this might eliminate the need for notifications of changes during the coverage period).

[60]While implicit insurance is a secondary payer to the original coverage based on predicted income, the ex post reconciliation is a secondary payer to implicit insurance if it accounts for out-of-pocket spending (and so implicit insurance support) but not if it accounts for only total health care costs (and income). Administratively, the simplest reconciliation is based solely on total health care costs and income. In that case, reconciliation payments are not implicitly taxed by implicit insurance and so provide better risk protection. My analysis assumes that reconciliation is based on out-of-pocket spending (necessarily due to data constraints), so the results understate the value of coverage that is reconciled based solely on total health care costs and income.

Finer versus coarser income measures.— That CSR coverage provides risk protection broadly comparable to that of the contract that limits out-of-pocket spending to 10% of income suggests that coarse income dependence is not too great an impediment to providing valuable risk protection. Coverage based on a coarser income measure has the advantage of being less likely to require later adjustments. Hence, basing coverage on coarse income measures could be an effective way to trade off risk protection against administrative costs.

Income verification.— To condition coverage on income, the insurer needs verifiable information on income. Presumably, this could be achieved with the same mechanisms used in other contexts to incentivize truthful information disclosure from individuals to governments or insurers. For example, the ACA's premium tax credits are reconciled based on information in the individual's federal income tax return. As another example, many life insurance contracts require a medical underwriting process in which the individual submits extensive information about their health. Misrepresentation or omission of material facts can invalidate an insurance contract and lead to denial of claims or policy cancellation. In cases of deliberate fraud or misrepresentation, insurers can pursue legal action against the insured to recover losses and potentially press criminal charges.

Labor distortion.— Although income-dependent contracts implicitly tax income, the extent of such taxation by the contracts I study is relatively small. For example, the average implicit marginal tax rate on income from the contract that limits out-of-pocket spending to 10% of income is 0.8% for otherwise-uninsured households.[61] This implicit taxation is small both in absolute terms and relative to typical marginal income tax rates. This suggests that the associated labor supply distortion is likely to be relatively small. Of course, income-dependent coverage that was more sensitive to income or that provided more coverage would tend to implicitly tax income to a greater extent.

[61] The implicit marginal tax rate is 10% in the roughly 8% of uninsured states in which the stop-loss is reached and 0% otherwise.

Summary.— Overall, this analysis suggests that there could be significant risk-protection gains from switching from standard to income-dependent health insurance coverage. Whether these and other gains outweigh the administrative and other costs is an important topic for future research.

Appendix Figures and Tables

Figure A1: Implicit health insurance support by education

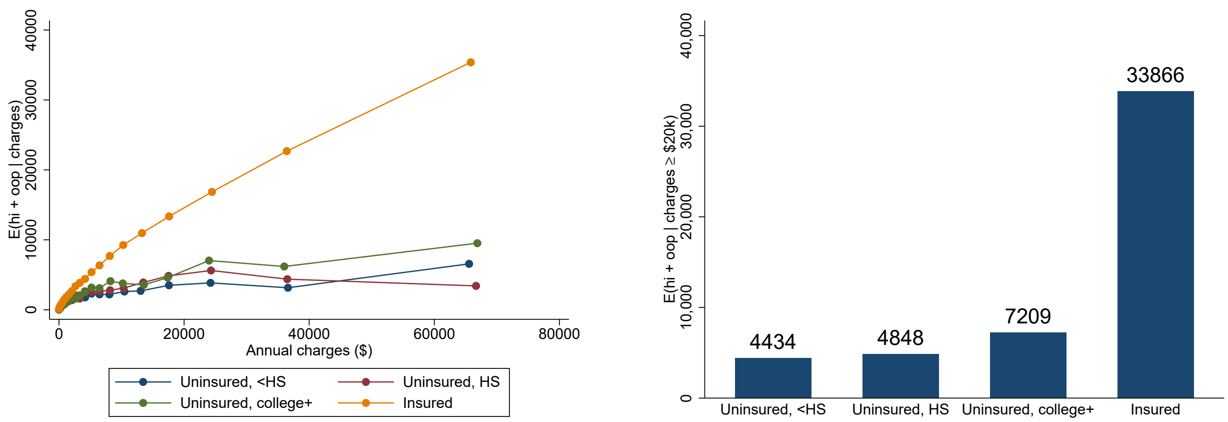

(a) Payments as a function of charges

(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 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 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

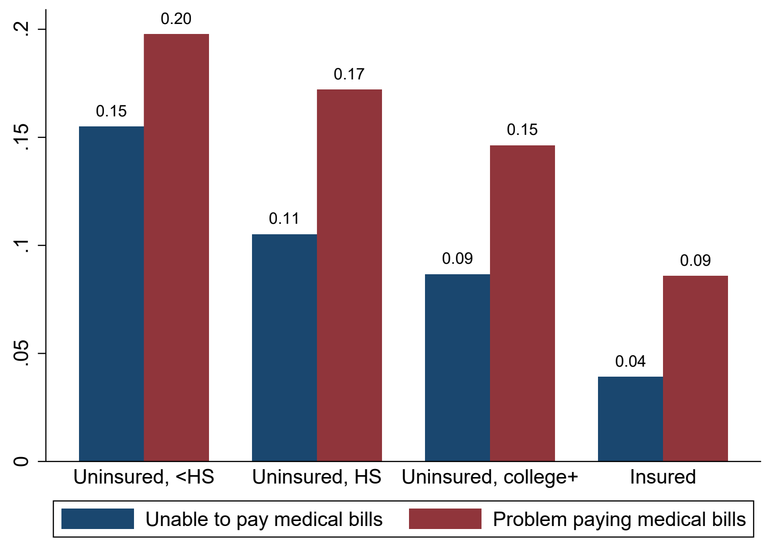

Notes: Figure shows the share of each group of non-elderly households who respond "Yes" to (i) "Does anyone in your family currently have any medical bills that you are unable to pay at all?" (which I label "Unable to pay medical bills") and, separately, (ii) "In the past 12 months did anyone in the 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.)

Figure A3: Association between unemployment and the mechanical reduction in out-of-pocket spending from different types of health insurance coverage (in dollars)

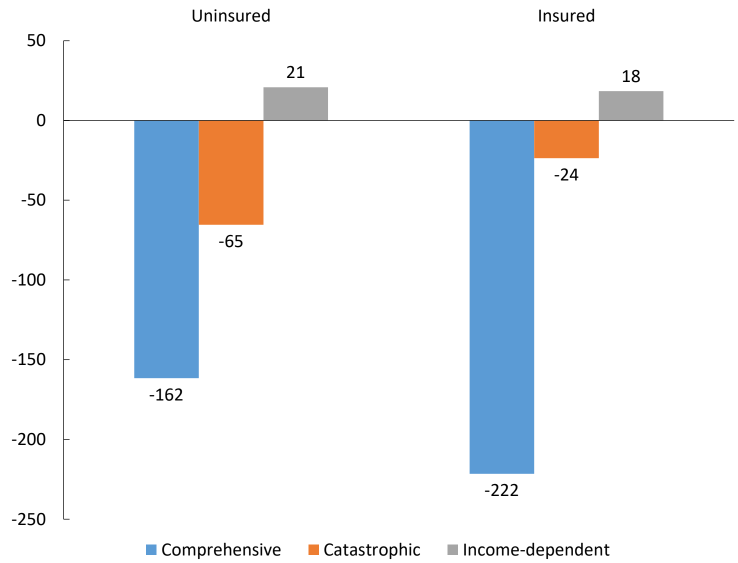

Notes: Figure shows the average association between unemployment and the mechanical reduction in out-of-pocket spending from each of three types of health insurance coverage in two sets of states: uninsured and insured non-elderly states. The mechanical reduction in out-of-pocket spending from comprehensive coverage is status quo out-of-pocket spending: $V = oop$. The mechanical reduction in out-of-pocket spending from catastrophic coverage of costs above \$5,000 per year is the excess of status quo out-of-pocket spending over \$5,000: $V = \max\{0,\ oop - 5,000\}$. The mechanical reduction in out-of-pocket spending from catastrophic coverage of costs above 10% of income ("Income-dependent") is the excess of status quo out-of-pocket spending over 10% of income: $V = \max\{0,\ oop - 0.10 \times y\}$. Each bar is the simple average of three regression coefficient estimates, from regressions of the mechanical reduction in out-of-pocket spending (V) on an indicator of unemployment and controls. The three regression specifications are those of the short run, medium run, and long run perspectives, described, for example, in Table A6 (which shows the underlying regression results for the case of comprehensive coverage [the $\widehat{\beta}_{oop|ue}$ estimates]). 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.

The results show that on average across the short run, medium run, and long run perspectives, the mechanical reduction in out-of-pocket spending from comprehensive coverage is about \$200 lower when the household is unemployed than not, that from catastrophic coverage is about \$50 lower when the household is unemployed than not, and that from income-dependent coverage is about \$20 higher when the household is unemployed than not. Hence, comprehensive coverage amplifies unemployment risk more than catastrophic coverage does, and income-dependent coverage provides a small buffer against it.

Figure A4: Simulated effects of eliminating out-of-pocket spending versus eliminating unemployment insurance on the volatility of consumption and net income

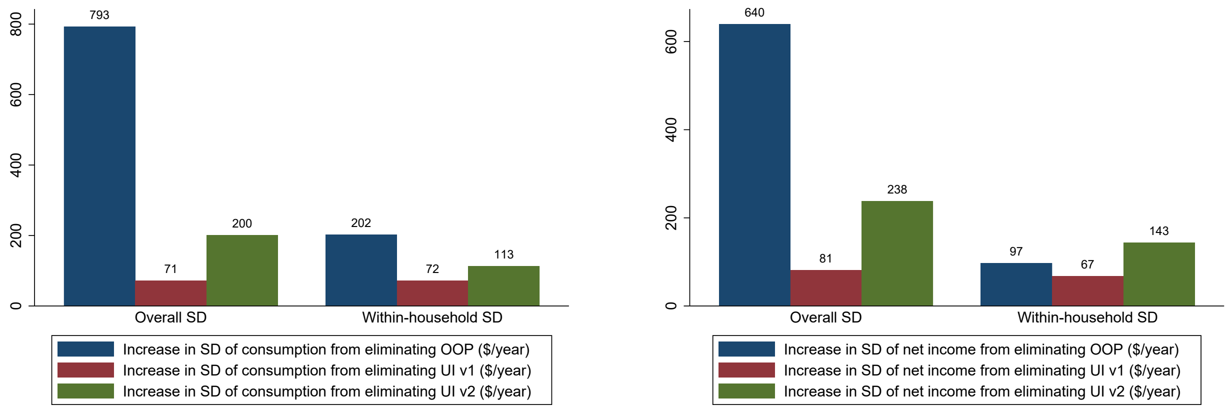

(a) Consumption ($/year)

(b) Net income ($/year)

Notes: Simulated effects on the overall standard deviation and the within-household standard deviation of annual consumption (panel (a)) and annual 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 A5: Risk protection value of completing health insurance coverage in future years

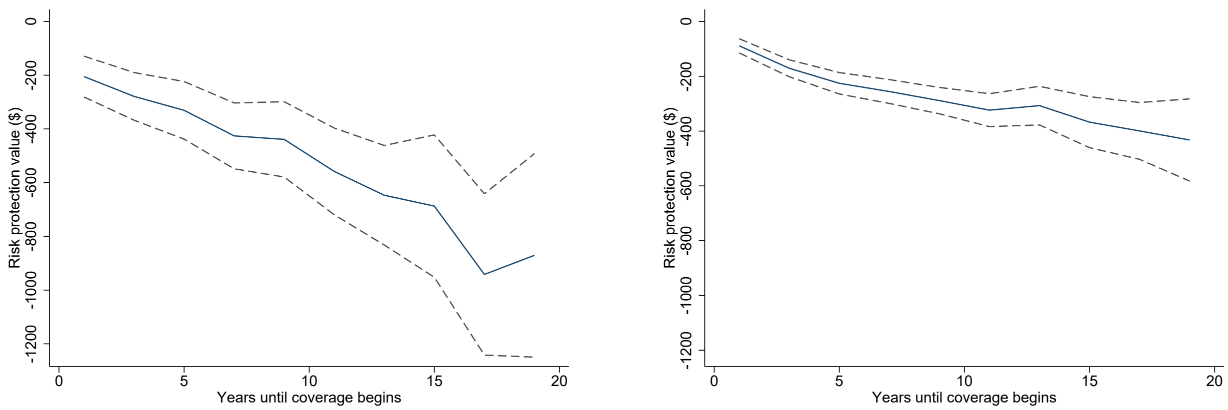

(a) Non-elderly uninsured

(b) Non-elderly insured

Notes: Risk protection value of completing health insurance coverage 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. "Risk protection 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)). Dashed lines are two standard errors above and below the estimates. 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" risk protection values to someone behind the veil are $-\$720$ and $-\$760$ for the uninsured and insured, respectively (see Table 1). Data are from the PSID.

Figure A6: Markup on standard health insurance coverage as a function of the level of coverage

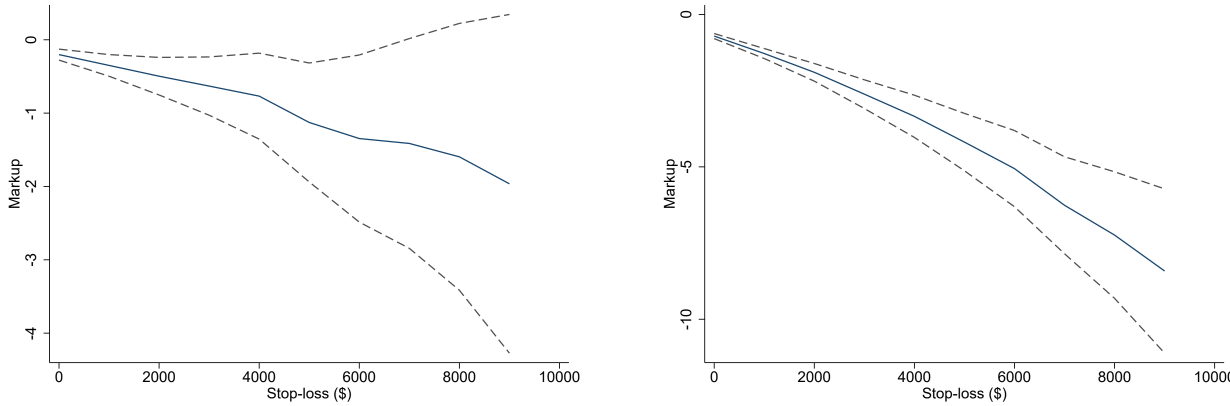

(a) Short run

(b) Long run

Notes: Sufficient statistic estimates of the markup (risk protection value per dollar of mean ex post value, $M = Cov\left(\widehat{\lambda}, V\right)/E(V)$) on different levels of standard 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. "Risk protection 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 estimates. 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 A7: Deductible as a function of income for implicit insurance versus income-dependent HI

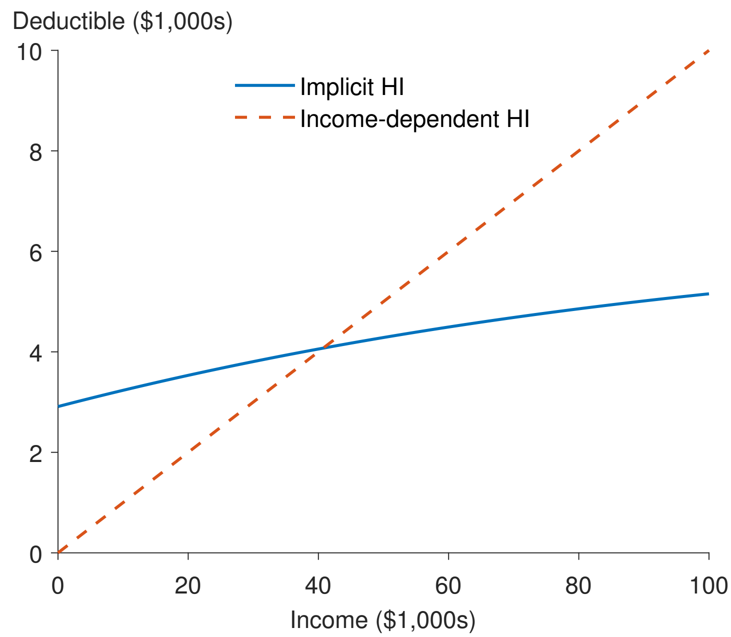

Notes: Baseline implicit insurance deducible function in the structural model, $d_{ihi}(y)$, and baseline income-dependent health insurance deductible function, $d(y) = 0.10 \times y$. Both functions relate the annual "deductible" above which there is full coverage of health care costs (and below which there is no coverage) to realized annual income. The implicit insurance deducible function in the structural model, $d_{ihi}(y)$, 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). The baseline income-dependent health insurance deductible function, $d(y) = 0.10 \times y$, is the stop-loss part of the main contract proposed by Feldstein and Gruber (1995).

The baseline implicit insurance deductible function, though a cubic in income, is well-approximated by a line with a slope of 0.02, considerably smaller than the slope of the income-dependent health insurance contract of 0.10.

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	90,908	46,538	98,123	64,137
Consumption	47,563	32,968	49,934	31,833
Out-of-pocket 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 A).

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 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 A).

Table A3: Out-of-pocket 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 spending (Oop), total health care costs (Tot), income, and consumption. Monetary variables are 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 A. 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).

Table A4: Out-of-pocket spending buffers 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 2 and reported in 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 16). If $\widehat{\beta}_{y\mid oop} > 1$, out-of-pocket spending covaries positively not only with income but even with net income. The elderly face significant risk in earnings (among those still working) and asset income (see, e.g., Blundell et al., 2020).

Table A5: Out-of-pocket spending buffers 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)
$\widehat{\beta}_{\log(y)\|\log(oop)}$	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 $\widehat{\beta}_{y\|oop}$	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 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 16). If $\widehat{\beta}_{y|oop} > 1$, out-of-pocket spending covaries positively not only with income but even with net income. The elderly face significant risk in earnings (among those still working) and asset income (see, e.g., Blundell et al., 2020).

Table A6: Out-of-pocket spending buffers 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 A7: Income-dependent health insurance would buffer income risk

	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)
$\widehat{\beta}_{\log(y)\mid\log(V)}$	-0.088	-0.112	-0.136	-0.098	-0.121	-0.102	-0.086	-0.108	-0.134	-0.052	-0.064	-0.068
(se)	(0.006)	(0.009)	(0.006)	(0.012)	(0.016)	(0.011)	(0.006)	(0.010)	(0.006)	(0.005)	(0.007)	(0.005)
Corr(log(y), log(V))	-0.22	-0.23	-0.20	-0.21	-0.25	-0.17	-0.22	-0.23	-0.20	-0.23	-0.23	-0.19
Implied $\widehat{\beta}_{y\mid V}$	-69.5	-88.4	-107.0	-21.9	-27.1	-22.8	-83.7	-105.1	-130.6	-9.7	-11.8	-12.7
(se)	(4.4)	(6.8)	(4.4)	(2.7)	(3.6)	(2.5)	(5.7)	(9.5)	(5.9)	(1.0)	(1.3)	(1.0)

Notes: Results from regressions of income variables on variables related to the ex post value of an income-dependent health insurance contract for each of four sets of states: non-elderly, non-elderly uninsured, non-elderly insured, and elderly insured. The contract provides full coverage above a stop-loss of 10% of income. The "mechanical effect" reduction in out-of-pocket spending from such a contract is $V = \max\{0, oop - 0.10y\}$. Short run and medium run columns are based on regressions of within-household changes in log income 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 (short run) or from one wave to five waves later (medium run). Long run is based on regressions of log income on $\log(1+V)$, 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. The first row shows the coefficient estimate on $\log(1+V)$. 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 $\log(1+V)$, both residualized with the relevant controls. The fourth row is the implied slope of income with respect to V, evaluated at the means of income and V: $\widehat{\beta}_{y\mid V} \equiv \widehat{\beta}_{\log(y)\mid\log(V)} \times \frac{E(y)}{E(V)}$. The fifth row is its standard error, which is the product of the standard error in the second row and $\frac{E(y)}{E(V)}$. Data are from the PSID.

While this contract provides greater coverage of health care costs when income is lower, it is not automatic that its ex post value to the household would be higher when income is lower on average across states of the world, since that also depends on the relationships between income and health care consumption and between income and implicit insurance support.

Table A8: Out-of-pocket spending buffers consumption risk

	Non-elderly (1)	Non-elderly uninsured (2)	Non-elderly insured (3)	Elderly insured (4)
$\widehat{\beta}_{\log(c)\mid\log(oop)}$	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 $\widehat{\beta}_{c\mid oop}$	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 2 and reported in 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 spending buffers consumption 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)
$\widehat{\beta}_{\log(c)\mid\log(oop)}$	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 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: Income-dependent health insurance would tend to buffer consumption risk

	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)
$\widehat{\beta}_{\log(c)\mid\log(V)}$	-0.002	-0.007	-0.020	-0.001	-0.010	0.002	-0.003	-0.005	-0.020	-0.001	-0.002	-0.004
(se)	(0.002)	(0.003)	(0.002)	(0.004)	(0.007)	(0.005)	(0.002)	(0.003)	(0.003)	(0.003)	(0.004)	(0.003)
Corr(log(c), log(V))	-0.01	-0.02	-0.05	-0.003	-0.04	0.01	-0.01	-0.02	-0.05	-0.01	-0.01	-0.02
Implied $\widehat{\beta}_{c\mid V}$	-0.89	-2.87	-8.07	-0.09	-1.53	0.34	-1.28	-2.56	-9.94	-0.13	-0.17	-0.41
(se)	(0.69)	(1.13)	(1.01)	(0.62)	(1.07)	(0.76)	(0.89)	(1.48)	(1.35)	(0.25)	(0.33)	(0.32)
Implied Corr(c, V)	-0.02	-0.08	-0.24	-0.005	-0.08	0.02	-0.03	-0.06	-0.26	-0.01	-0.01	-0.03

Notes: Results from regressions of consumption variables on variables related to the ex post value of an income-dependent health insurance contract for each of four sets of states: non-elderly, non-elderly uninsured, non-elderly insured, and elderly insured. The income-dependent health insurance contract provides full coverage above a stop-loss of 10% of income. The "mechanical effect" reduction in out-of-pocket spending from such a contract is $V = \max\{0, oop - 0.10y\}$. Short run and medium run columns 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 (short run) or from one wave to five waves later (medium run). 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. 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. The first row shows the coefficient estimate on $\log(1+V)$. 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 $\log(1+V)$, both residualized with the relevant controls. The fourth row is the implied slope of consumption with respect to V, evaluated at the means of consumption and V: $\widehat{\beta}_{c\mid V} \equiv \widehat{\beta}_{\log(c)\mid\log(V)} \times \frac{E(c)}{E(V)}$. The fifth row is its standard error, which is the product of the standard error in the second row and $\frac{E(c)}{E(V)}$. The sixth row is the implied correlation between consumption and V, defined as the product of the implied slope of consumption with respect to V, $\widehat{\beta}_{c\mid V}$, and the ratio of the standard deviation of V to the standard deviation of consumption, each residualized with the relevant controls. Data are from the PSID.

Table A11: Sufficient statistic estimates: Coverage of different types of health care

	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)
Risk protection 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 comprehensive coverage of different types of health care for non-elderly uninsured households. Column (1) reproduces the main estimates of the short run value of comprehensive 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 comprehensive coverage of each of the three subcomponent 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. Risk protection value is more negative from longer-run perspectives (see Table 1 and Figure A5). 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. "Risk protection 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 risk protection 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: Comprehensive coverage for different education groups and in different 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)
Risk protection 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 comprehensive 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 comprehensive coverage in all non-elderly uninsured states (see Table 1). Columns (2)–(5) show heterogeneity across education groups in the value of comprehensive coverage 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 among non-elderly households is about \$3,120 and about 30% of such 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. Risk protection value is more negative from longer-run perspectives (see Table 1 and Figure A5). 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. "Risk protection 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 risk protection 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)
Risk protection 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 comprehensive 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). 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. Risk protection value is more negative from longer-run perspectives (see Table 1 and Figure A5). 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. "Risk protection 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 risk protection 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 (1)	Control for quintic in income (2)	Regress $\Delta\widehat{\lambda}$ on ΔV (3)	Regress $\Delta \log \widehat{\lambda}$ on ΔV (4)	Fixed effects (5)	Fixed effects longer run (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)
Risk protection 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 comprehensive 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. Risk protection value is more negative from longer-run perspectives (see Table 1 and Figure A5). 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 2 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. "Risk protection 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 risk protection value per dollar of mean ex post value, $Cov\left(\widehat{\lambda}, V\right)/E(V)$. 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 benefits from improved health and reduced medical debt

	Baseline	New cancer diagnosis	Ever cancer diagnosis	Health much worse	Health newly bad	Health bad	Hospitalization	Medical bills	Medical bills up to \$10k
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(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)
Risk protection 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 comprehensive health insurance coverage for non-elderly uninsured households under different assumptions about the benefits from 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. Risk protection value is more negative from longer-run perspectives (see Table 1 and Figure A5). 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 risk protection 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. 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 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. "Risk protection 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 risk protection 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. See Appendix C.6 for a discussion of these results.

Table A16: Sufficient statistic estimates: Income-dependent health insurance

	Stop-loss is 10% of income									Stop-loss is $500 less than 10% of income								
	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 (10)	Medium run (11)	Long run (12)	Short run (13)	Medium run (14)	Long run (15)	Short run (16)	Medium run (17)	Long run (18)
Corr(log(c),log(V))	-0.003	-0.04	0.01	-0.01	-0.02	-0.04	-0.01	-0.01	-0.02	-0.003	-0.06	-0.02	-0.01	-0.03	-0.07	-0.01	-0.01	-0.04
(se)	(0.017)	(0.026)	(0.014)	(0.007)	(0.009)	(0.006)	(0.016)	(0.017)	(0.012)	(0.018)	(0.025)	(0.014)	(0.008)	(0.010)	(0.006)	(0.016)	(0.017)	(0.012)
Risk protection value	11	185	-41	44	88	341	25	33	83	10	277	114	49	140	468	36	47	191
(se)	(75)	(129)	(92)	(30)	(51)	(46)	(51)	(66)	(64)	(65)	(110)	(82)	(28)	(48)	(43)	(45)	(59)	(59)
Mean ex post value	208	208	208	100	100	100	346	346	346	246	246	246	120	120	120	409	409	409
Markup	0.05	0.89	-0.20	0.44	0.88	3.40	0.07	0.10	0.24	0.04	1.13	0.46	0.41	1.17	3.89	0.09	0.11	0.47

Notes: Statistics related to the value of income-dependent health insurance that provides full coverage above a stop-loss that depends on the realization of income and no coverage below that. Columns (1)–(9) show results based on a stop-loss of 10% of income: $V = \max\{0, oop - 0.10y\}$. Columns (10)–(18) show results based on a stop-loss $500 below that: $V = \max\{0, oop - \max\{0, 0.10y - 500\}\}$. This provides somewhat more coverage to improve statistical precision given how rarely out-of-pocket spending exceeds 10% of income (8% of non-elderly uninsured household-waves, 4% of non-elderly insured, and 12% of elderly). Short run and medium run columns 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 (short run) or from one wave to five waves later (medium run). 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. 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(V)) is the correlation between the relevant changes in (short and medium run) or levels of (long run) log consumption and $\log(1+V)$, both residualized with the corresponding controls. "Risk protection 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 risk protection value per dollar of mean ex post value, $Cov\left(\widehat{\lambda}, V\right)/E(V)$. 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+). An important part of why the markup can be so high is the interaction with income risk. Whereas out-of-pocket spending risk is relatively limited because of implicit insurance, income risk is much greater, so even a small buffer against income risk (and correlated risks) can provide highly valuable insurance.

Table A17: Structural analysis of mechanisms

	Baseline			No income risk			No implicit insurance		
	Full (1)	Cat (2)	Y-dep (3)	Full (4)	Cat (5)	Y-dep (6)	Full (7)	Cat (8)	Y-dep (9)
Risk protection value	-489	-45	730	66	0	0	1,311	1,498	3,145
Mean ex post value	2,573	46	100	2,587	0	0	4,558	1,900	1,566
Markup	-0.19	-0.99	7.28	0.03	N/A	N/A	0.29	0.79	2.01
$Corr\left(\widehat{\lambda}, V\right)$	-0.06	-0.08	0.80	0.999	0	0	0.28	0.26	0.40
$Corr\left(\widehat{\lambda}, hi\right)$	0.08	0.08	0.19	0.59	0.39	0.34	0.47	0.49	0.60
$Corr(hi, y)$	-0.002	-0.02	-0.14	N/A	N/A	N/A	-0.002	-0.02	-0.14
$Corr(V, y)$	0.17	0.43	-0.32	N/A	N/A	N/A	0.04	0.03	-0.11

Notes: Statistics related to the risk protection value of three health insurance contracts in three versions of the structural model: the baseline model, no income risk, and no implicit health insurance. The contracts are full coverage of all costs ("Full"), catastrophic coverage of all costs above \$5,000 with no coverage below that ("Cat"), and catastrophic coverage of all costs above 10% of income with no coverage below that ("Y-dep"). 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 of the world, $y \equiv y_{med}$. The "No implicit insurance" counterfactual has no implicit health insurance: $ihi(h, y; HI) \equiv 0$. Risk protection value is the amount by which the ex ante equivalent variation of health insurance exceeds its mean ex post value (see equation (4)), using consumption-based equivalent variation (the amount by which the consumption of a household without health insurance must be increased in all states of the world to be as well off ex ante as it would be with health insurance). The markup is the ratio of risk protection 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. In the "No income risk" counterfactual, the catastrophic and income-dependent contracts provide strictly less coverage than implicit insurance (their deductibles exceed the implicit insurance deductible at median income) and so have zero ex post value in all states of the world.

Table A18: Structural analysis robustness to income risk and implicit health insurance

	Baseline implicit health insurance						Less implicit health insurance					
	Baseline income risk			Half income risk			Baseline income risk			Half income risk		
	Full (1)	Cat (2)	Y-dep (3)	Full (4)	Cat (5)	Y-dep (6)	Full (7)	Cat (8)	Y-dep (9)	Full (10)	Cat (11)	Y-dep (12)
Risk protection value	-489	-45	730	-170	-17	21	-235	-274	1,344	-56	-10	285
Mean ex post value	2,573	46	100	2,592	20	4	3,127	469	297	3,150	491	137
Markup	-0.19	-0.99	7.28	-0.07	-0.88	5.04	-0.08	-0.58	4.52	-0.02	-0.02	2.07
$Corr\left(\widehat{\lambda}, V\right)$	-0.06	-0.08	0.80	-0.04	-0.15	0.52	0.04	-0.08	0.67	0.05	0.01	0.59
$Corr\left(\widehat{\lambda}, hi\right)$	0.08	0.08	0.19	0.08	0.08	0.14	0.14	0.13	0.25	0.13	0.13	0.19
$Corr(hi, y)$	-0.002	-0.02	-0.14	-0.003	-0.02	-0.09	-0.002	-0.02	-0.14	-0.003	-0.02	-0.09
$Corr(V, y)$	0.17	0.43	-0.32	0.12	0.43	-0.17	0.12	0.20	-0.35	0.09	0.12	-0.30

Notes: Statistics related to the risk protection value of three health insurance contracts in the structural model for different levels of implicit health insurance coverage and income risk. The contracts are full coverage of all costs ("Full"), catastrophic coverage of all costs above \$5,000 with no coverage below that ("Cat"), and catastrophic coverage of all costs above 10% of income with no coverage below that ("Y-dep"). Columns (1)–(6) 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 (7)–(12) scale up the implicit health insurance deductibles—reducing implicit health insurance support—to match mean out-of-pocket 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 insurance available to households with significant assets or low willingness to rely on implicit insurance. (Granted, the model is ill-suited to quantify risk protection 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 (4)–(6) and (10)–(12) 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: $\widetilde{y} = 0.5 \times y + 0.5 \times y_{median}$. Risk protection value is the amount by which the ex ante equivalent variation of health insurance exceeds its mean ex post value (see equation (4)), using consumption-based equivalent variation (the amount by which the consumption of a household without health insurance must be increased in all states of the world to be as well off ex ante as it would be with health insurance). The markup is the ratio of risk protection 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 A19: Structural analysis additional robustness tests

	Baseline			Log utility ($\gamma = 1$)			Hosp targeting only			Independent risks		
	Full	Cat	Y-dep	Full	Cat	Y-dep	Full	Cat	Y-dep	Full	Cat	Y-dep
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Risk protection value	-489	-45	730	-236	-32	219	-117	0	635	-205	-42	787
Mean ex post value	2,573	46	100	2,573	46	100	2,601	0	80	2,434	43	109
Markup	-0.19	-0.99	7.28	-0.09	-0.70	2.18	-0.04	N/A	7.98	-0.08	-0.99	7.22
$Corr\left(\widehat{\lambda}, V\right)$	-0.06	-0.08	0.80	-0.13	-0.18	0.71	-0.07	0	0.90	0.00	-0.08	0.82
$Corr\left(\widehat{\lambda}, hi\right)$	0.08	0.08	0.19	0.06	0.07	0.19	-0.05	-0.05	0.01	0.04	0.02	0.13
$Corr(hi, y)$	-0.002	-0.02	-0.14	-0.002	-0.02	-0.14	-0.15	-0.15	-0.18	0.004	0.005	-0.11
$Corr(V, y)$	0.17	0.43	-0.32	0.17	0.43	-0.32	-0.14	0	-0.38	0.10	0.40	-0.34

Notes: Statistics related to the risk protection value of three health insurance contracts in the structural model under different assumptions. The contracts are full coverage of all costs ("Full"), catastrophic coverage of all costs above \$5,000 with no coverage below that ("Cat"), and catastrophic coverage of all costs above 10% of income with no coverage below that ("Y-dep"). Columns (1)–(3) are the baseline specification. Columns (4)–(6) use log utility (a coefficient of relative risk aversion of one). Columns (7)–(9) isolate 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. Columns (10)–(12) use a risk process, $F(h, y)$, in which health care consumption and income are independent but the marginal distributions, $F(h)$ and $F(y)$, are as in the baseline risk process. Risk protection value is the amount by which the ex ante equivalent variation of health insurance exceeds its mean ex post value (see equation (4)), using consumption-based equivalent variation (the amount by which the consumption of a household without health insurance must be increased in all states of the world to be as well off ex ante as it would be with health insurance). The markup is the ratio of risk protection 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. In the analysis of hospitalization-related targeting only, the catastrophic contract provides strictly less coverage than implicit insurance in hospitalization states (its deductible exceeds the implicit insurance deductible in all hospitalization states) and so has zero ex post value in all hospitalization states.

Table A20: Average Cost-Sharing Parameters for Individual Coverage through ACA Silver Plans with CSRs in 2016 (2020 dollars)

Income Range	Actuarial Value	Deductible	Coinsurance Rate	Out-of-Pocket Max
>250% FPL	70%	$3,304	10.6–15.3%	$6,714
200–250% FPL	73%	$2,650	11.3–14.7%	$5,299
150–200% FPL	87%	$772	10.5–13.8%	$2,208
100–150% FPL	94%	$265	9.6–12.8%	$1,189

Notes: Income ranges are expressed in terms of the Federal Poverty Level (FPL). Actuarial value is the share of total health care costs covered by the contract on average. The actuarial value of a standard Silver Plan is 70%. Coinsurance rate ranges correspond to the average coinsurance rate on primary care (minimum of each range) and the average coinsurance rate on specialty care visits (maximum of each range). The coinsurance rates for other services, such as prescription drugs, do not necessarily fall within these ranges. Households with income below 100% of the FPL are not eligible for CSRs and may or may not be eligible for Medicaid, depending on their circumstances (see Appendix G). Source: Gabel et al. (2016), converted to 2020 dollars

Table A21: Sufficient statistic estimates: Income-dependent coverage based on current versus past income

	Full coverage			Stop-loss 10% of income			Stop-loss 10% of past income			CSRs based on income			CSRs based on past income		
	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)	Short run (13)	Medium run (14)	Long run (15)
Corr(log(c),log(V))	0.09	0.17	0.25	-0.003	-0.04	0.01	0.07	-0.01	0.04	0.003	-0.06	-0.08	0.08	0.02	-0.03
(se)	(0.017)	(0.027)	(0.014)	(0.017)	(0.026)	(0.014)	(0.019)	(0.029)	(0.015)	(0.017)	(0.022)	(0.014)	(0.017)	(0.023)	(0.014)
Risk protection value	-205	-439	-721	11	185	-41	-270	67	-246	-10	212	378	-248	-67	139
(se)	(38)	(70)	(42)	(75)	(129)	(92)	(79)	(147)	(103)	(53)	(80)	(67)	50	85	72
Mean ex post value	1,016	1,016	1,016	208	208	208	212	212	212	239	239	239	209	209	209
Markup	-0.20	-0.43	-0.71	0.05	0.89	-0.20	-1.27	0.32	-1.16	-0.04	0.88	1.58	-1.19	-0.32	0.67

Notes: Statistics related to the value of different types of health insurance for non-elderly uninsured households. Columns (1)–(3) reproduce the main results for full (income-independent) coverage (see Table 1). Columns (4)–(6) show results for a stop-loss of 10% of current income (that is, annual income in the year of the coverage itself). Columns (7)–(9) show results for a stop-loss of 10% of past income two years before the coverage period. Columns (10)–(12) show results for an approximation to ACA CSR contracts based on current income. Columns (13)–(15) show results for an approximation to ACA CSR contracts based on past income two years before the coverage period. The approximation to ACA CSR contracts and the timing of the income measures are described in detail in Appendix G. Short run and medium run columns 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 (short run) or from one wave to five waves later (medium run). 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. 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(V)) is the correlation between the relevant changes in (short and medium run) or levels of (long run) log consumption and $\log(1+V)$, both residualized with the corresponding controls. "Risk protection 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 risk protection 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 (elderly) are households whose heads are 25–64 (65+).

Table A22: Structural analysis of income-dependent coverage

	Full coverage (1)	Stop-loss of 10% of income (2)	Typical CSR contracts (3)	Approximate CSR contracts (4)
Risk protection value	-489	730	1,430	1,446
Mean ex post value	2,573	100	209	265
Markup	-0.19	7.28	6.83	5.45
$Corr\left(\widehat{\lambda}, V\right)$	-0.06	0.80	0.86	0.79
$Corr\left(\widehat{\lambda}, hi\right)$	0.08	0.19	0.27	0.24
$Corr(hi, y)$	-0.002	-0.14	-0.13	-0.12
$Corr(V, y)$	0.17	-0.32	-0.35	-0.39

Notes: Statistics related to the risk protection value of four health insurance contracts: full coverage of all costs (which is not income-dependent coverage but is shown for comparison purposes), catastrophic coverage of all costs above 10% of income with no coverage below that ("Stop-loss of 10% of income"), typical ACA CSR contracts (which provide greater coverage for lower-income households; see Appendix G), and an approximation to typical ACA CSR contracts that provides full-coverage above income-dependent stop-losses but no coverage below that (see Appendix G). The approximate CSR contracts are included to assess how closely the approximate contracts mimic the risk protection properties of the actual contracts since the corresponding sufficient statistic analysis (results of which are reported in Table A21) can handle only the approximate contracts. Risk protection value is the amount by which the ex ante equivalent variation of health insurance exceeds its mean ex post value (see equation (4)), using consumption-based equivalent variation (the amount by which the consumption of a household without health insurance must be increased in all states of the world to be as well off ex ante as it would be with health insurance). The markup is the ratio of risk protection 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 A23: 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 household benefits one-for-one from the indemnity benefit, i.e., that such benefits are not implicitly taxed by implicit insurance. The aim is to understand the likely effects of indemnity insurance that supplements direct coverage of health care costs (though even without such coverage, indemnity insurance likely 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). 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 risk protection value per dollar of mean ex post value, $Cov\left(\widehat{\lambda}, V\right)/E\left(V\right)$. "Risk protection 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+).

Table A24: 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.