

The Risk Protection Value of Moral Hazard^{*}

Angie Acquatella[§] Victoria Marone[†]

August, 2025

Abstract

Health insurance lowers the out-of-pocket price of healthcare, and it is well-established that this leads to higher utilization of care. This type of “moral hazard” is typically viewed as a social cost of insurance. Within a standard model, we show that there are two important ways in which the consumer’s ability to change her behavior in response to insurance can play a central role in the ability of insurance to protect her from risk. These are (i) by allowing optimal exploitation of real income gained in bad states, and (ii) by enabling more resources to be shifted to bad states than otherwise could be. We provide a theoretical characterization of these cases and quantify their importance empirically. Under standard parameterizations of demand for healthcare and health insurance, estimates in the literature imply that moral hazard accounts for as much as half of the total value of risk protection derived from insurance. Preventing consumers from changing their behavior in response to insurance would lower costs, but also result in a major loss of risk protection, on-net *reducing* consumer welfare in the population we study. Our results suggest that under-utilization of healthcare may thus be an equally important threat to welfare as over-utilization.

Keywords: moral hazard, risk protection, optimal insurance

JEL Codes: I13, D81, I18

^{*}We are grateful to Marika Cabral, Hector Chade, David Cutler, Liran Einav, Amy Finkelstein, Marty Gaynor, Francois Gerard, Christian Gollier, Igal Hendel, Lee Lockwood, Erzo Luttmer, Marius Ring, Amanda Starc, Nicholas Tilipman, Bob Town, Ben Vatter, and Andre Veiga for helpful comments and discussion. We also thank many participants in seminars at the 2024 ASHEcon conference, BI Norwegian Business School, and Toulouse School of Economics for helpful comments. Dario Gori and Junrui Lin provided excellent research assistance. [§]Department of Economics, Toulouse School of Economics (email: angelique.acquatella@tse-fr.eu) [†]Department of Economics, Yale University (email: victoria.marone@yale.edu)

“[M]oral hazard”...refer[s] to the responsiveness of healthcare spending to insurance coverage. ...[It] is of economic interest because it creates an obstacle to the consumption-smoothing purpose of insurance.

(Einav and Finkelstein, 2018)

Without the payout from my spousal life insurance policy when my husband died, my children and I could never have afforded to stay in our home.

– Unknown

I Introduction

The possibility that insurance distorts behavior after the realization of risk is typically viewed as a social cost. Across settings such as unemployment insurance, disability insurance, long-term care insurance, and health insurance, the manifestation of this type of “ex-post moral hazard” has been widely studied and broadly influential on policy debates around the world. In the realm of health insurance, in particular, controlling moral hazard has loomed large in the design of public health insurance schemes, since Pauly (1968) first described its potential to render some types of contracts non-optimal. The standard intuition of optimal insurance is that the value derived from risk protection must be weighed against the social cost from distorted incentives (Zeckhauser, 1970). If only the consumer could be prevented from changing her behavior, the logic goes, maximal risk protection could be achieved. This argument fails, however, in some important situations (as the widow quoted would be quick to point out).

This paper characterizes the situations in which the consumer’s ability to change her ex post behavior in response to insurance can play a central role in the ability of insurance to protect her from risk. Our aim is to bring this point—appreciated by many, but perhaps not by all—into a formal analytical and empirical framework familiar within the contemporary literature on health insurance. Though we view health insurance as among the most important markets plagued by ex-post moral hazard, most of our analysis is general enough to also apply to other forms of insurance. The framework developed provides a new lens through which to consider optimal insurance design, and in turn, a new way to interpret some aspects of insurance contracts observed in real-world public and private markets.

To begin, consider the simplest textbook model of insurance, featuring a consumer with concave utility over one good (let us say “money”), the quantity of which is subject to uncertainty. Risk aversion prompts her to trade some money up front (a premium) in exchange for a greater degree of certainty over the quantity of money she will have. Such a model is

predicated on an assumption that the insurance in question can be written to be contingent on the realized state of the world. But in most lines of real-world insurance, contracts are not truly state-contingent claims. Instead, they offer a promise of indemnification against financial losses resulting from covered perils.¹ And in many cases—perhaps health insurance most of all—the insured has some degree of control over the size of said losses. In other words, there is the potential for ex-post moral hazard.²

Now, any economic model with moral hazard must feature at least two goods. Only then can different tradeoffs be made in response to different incentives. With two goods, however, the goal of insurance can no longer be coherently stated as to “smooth the marginal utility of consumption,” since consumption takes two forms. Instead, one must recognize the goal as to smooth the *marginal indirect utility of income* across states of the world.³ Given our focus on health insurance, the consumer in our model does not face uncertainty over her money income, but rather over the intensity of her preferences for uses of her income. Just as in home or auto insurance, income will typically be more valuable when disaster strikes. Re-framing the *risk* from which the consumer demands protection to be over her marginal indirect utility of income, rather than over consumption directly, is a central feature of our approach.

The basic elements of the model are as follows. A consumer has a fixed initial income and a known distribution over health states. She has state-dependent preferences over two goods: healthcare utilization and non-health consumption. Absent insurance, she will typically have more valuable uses of income in sicker states. A health insurance contract lowers the consumer’s point-of-sale (“out-of-pocket”) cost of healthcare below its market price, but places otherwise no restrictions on her behavior.⁴ The realization of the health state is the con-

¹The principle of indemnity dates back to earliest origins of insurance. It holds that in the event of loss, the insured should be made whole, but not enriched (Wilson and Thomas, 2005). Indemnification against verifiable financial losses arose as a workable alternative to true state-contingent claims, which are often impossible. Enumerating, let alone verifying, the possible states of the world can be either technologically infeasible or prohibitively costly, in a way not wholly dissimilar to the ideas behind incomplete contracts (Hart and Moore, 1999).

²Ex-post versus ex-ante moral hazard can be distinguished by whether the agent acts after (ex-post) or before (ex-ante) the uncertain act of nature. The observation that the insured can take ex-post actions to influence the size of covered losses to a substantial degree in health insurance dates back to at least Arrow (1963) and was also discussed at length by Pauly (1969) and Zeckhauser (1970). It may be useful to consider the comparison to car insurance, where under the well-informed and watchful eye of the insurer’s claims-adjuster, the consumer and mechanic have relatively less ability to inflate the size of a covered loss.

³Of course, this is equivalent to smoothing the marginal utility of consumption when there is only one good.

⁴The health insurance literature has modeled “indemnification against financial losses from covered perils” as a contract that simply lowers the consumer’s price of healthcare since at least Pauly (1968), and we follow that convention. Though a useful simplification, it is not an accurate reflection of reality. Since the first instances of market-based health insurance, there has been “utilization management” of some kind, wherein

sumer’s private information, and she may act fully in accordance with her ex-post incentives under the contract. All consumer characteristics relevant to the insurer’s expected cost are commonly known. There is therefore no *ex-ante* asymmetric information (selection) in the model, only *ex-post* information asymmetry (moral hazard). Following the prevailing norm in the health insurance literature, we henceforth use the term “moral hazard” to refer to the full causal effect of insurance on consumer behavior (Pauly, 1968; Einav and Finkelstein, 2018).⁵ Premiums are set to be actuarially fair, such that the private and social value of insurance coincide. The singular source of value from insurance in the model is risk protection; the insurer has the technology to harness the Law of Large Numbers, while the individual does not.

Without the possibility to contract on realized health, the ability of health insurance to shift resources across states derives from the consumer’s differential tastes for healthcare across states. Making healthcare cheaper raises her real income in all states (gross of the premium), but crucially, it does so by *more* in states in which healthcare is more valuable to her. Though health insurance cannot transfer *nominal* income across states, it can thus still transfer *real* income across states. The matter of moral hazard—the consumer’s resulting change in behavior—is then simply a question of how this additional purchasing power is used. Any change in (uncompensated) demand for healthcare can be decomposed into the standard parts: an income effect and a substitution effect (Nyman, 1999a). Both effects represent “moral hazard” (as defined). The principal focus of this paper is to show that neither is necessarily welfare-decreasing, and that their contributions to the risk protection value of insurance can be economically large.

A central element of our approach is the notion of the first-best insurance contract. This contract arises in a world where asymmetric information is eliminated, rendering the health state contractible. Consistent with the intuition that efficient healthcare utilization accounts for its full social cost, the first-best contract does not distort prices. Instead, it allows the consumer to reallocate her initial income across states in an actuarially fair way—not unlike a life insurance policy. Under the first best, the consumer will equate her marginal indirect utility of income across states, achieving maximal risk protection (and welfare). But with income reallocated, there is clearly the possibility for new optimal consumption bundles. And,

the insurer verifies that treatment was in fact medically indicated before releasing funds. In other words, the health insurer plays the same role of “claims adjuster” as the home or auto insurer. Of course, there remains a great deal of consumer and doctor discretion—as the empirical literature on moral hazard in health insurance has repeatedly demonstrated—and it is on this utilization-side discretion that the model focuses.

⁵See footnote 28 for a discussion of this terminology.

since these changes in consumption behavior arise under first-best insurance, they can only have arisen in service of risk protection—prima facie evidence that moral hazard can provide risk protection. We define the *socially efficient level* of healthcare utilization in each state to be the privately optimal level of utilization under the first-best contract. This level represents what the consumer herself is *willing* to pay for at undistorted prices out of the income she chooses to allocate to that state in a first-best world.⁶ But critically, it may not correspond to what she is *able* to pay for at undistorted prices out of her initial income in the real world.

Returning to a traditional health insurance contract, moral hazard can now be decomposed into two parts. Absent insurance, utilization will typically fall *below* the efficient level, resulting in “under-utilization.” With insurance, it will typically *exceed* the efficient level, resulting in “over-utilization.” Moral hazard may thus involve both a reduction in under-utilization—what we will call an “efficient increase in utilization”—as well as some amount of over-utilization.⁷ In our model, these two components correspond exactly to the income and substitution effects of the healthcare price decrease that is health insurance.

As long as the consumer values insurance, any reduction in under-utilization will represent a generically positive contribution to social welfare.⁸ Because it represents an income effect, it will be more important when healthcare utilization is more responsive to income, and when the income allocated to some states under the first-best contract is far higher than initial income.⁹ Over-utilization—the substitution effect—will instead have two opposing effects

⁶The idea behind this level of healthcare utilization—the one the consumer would have chosen in a first-best world—is not new. The same (at least approximate) idea appears in [Zeckhauser \(1970\)](#), [de Meza \(1983\)](#), [Cutler and Zeckhauser \(2000\)](#), and [Nyman \(2002, 2004\)](#). Our analysis offers a formal characterization of the first best and its associated consumption bundles.

⁷To our knowledge, this decomposition of moral hazard first appeared in [Nyman et al. \(2018\)](#) in the context of a simpler but largely equivalent model. Those authors refer to the efficient increase in utilization as “efficient moral hazard” and over-utilization as “inefficient moral hazard.” We use slightly less judgmental terminology for the latter because, as we will show, its normative characterization of depends on consumer preferences.

⁸That the consumer will prefer actuarially fair insurance to no insurance is not a foregone conclusion for a few reasons. Most trivially, she could be risk neutral or face no uncertainty. More substantively, she could be happy with having less non-health consumption in sick states ([Finkelstein, Luttmer and Notowidigdo, 2013](#)). We call these types of consumers *Travelers*. Keeping careful track of the aspects of consumer preferences under which the welfare effects of moral hazard can be signed will be an important feature of our analysis.

⁹For an example of an efficient increase in utilization, suppose there is a 0.1 percent chance of becoming sick, meaning income could be fairly transferred from the healthy to the sick state at a rate of 1 to 999. Suppose further that in a first-best world, a consumer with initial income of \$100,000 would happily sacrifice \$1,000 when healthy in order to have \$1,000,000 when sick. Take the sickness to be cancer and a potential treatment (immunotherapy) to cost \$500,000. Suppose that with a million dollars, the sick consumer would buy immunotherapy. But without those additional resources, she clearly could not do so, and would instead settle with chemotherapy for just \$50,000. The increase in healthcare utilization from \$50,000 to \$500,000 represents the efficient increase in utilization.

on welfare. On one hand, it represents the classic welfare loss from insurance, inefficiently reallocating resources from non-health to healthcare consumption (Pauly, 1968). On the other hand, over-utilization may offer valuable risk protection for certain consumers—specifically, those who would wish to have higher non-health consumption when sick than when healthy. These consumers want not only to be made whole in the event of illness, but also enriched.¹⁰ Transferring resources to sick states may be so valuable to them—from a risk protection perspective—that getting extra healthcare when sick may be a worthwhile substitute for getting extra cash.¹¹ Even though over-utilization is not valued at cost ex post, it may thus still be valued at cost ex ante.

Our model offers a new perspective on a number of long-standing ambiguities in the health insurance literature. First, our analysis clarifies that gaining “access” (à la Nyman, 2002) to the efficient level of healthcare utilization is still very much a risk protection motive for insurance. Once the model is cast in the light of smoothing the marginal indirect utility of income, the risk-protection interpretation of income’s uses follows directly. Second, our analysis highlights that without strong assumptions on preferences, the (ex-post) demand curve for healthcare cannot alone be used to make ex ante welfare statements about the consumption of care. Since welfare from insurance must invariably be valued ex ante, the normative interpretation of its effects on consumer behavior must also be.¹² Third, our analysis emphasizes that the value of insurance—and the value of moral hazard in particular—relies critically on the nature of state-dependence in consumer preferences. Using the first-best contract as a benchmark allows us to offer a formal vocabulary for classifying preferences along

¹⁰Targeted cash transfers are impossible under the health insurance contract we study because, again, the state is not contractible. Indeed, cash transfers directly violate the principle of indemnity. Even so, many cash-transfer-like products exist in private health insurance markets around the world (for example, critical illness insurance), and these products are growing in popularity (Spherical Insights & Consulting, 2025), suggesting that there exist consumers with preferences of this type.

¹¹For example, a debilitated patient may be pleased to receive covered in-home nursing visits when recovering at home, even if she would have preferred to receive the cash equivalent so that her husband could take off work to care for her. Despite that she doesn’t value the nursing visits at-cost, even under first-best income—and thus, they are classifiable as “over-utilization”—they still make recovery less bad than it would be otherwise.

¹²The standard logic we challenge is well summarized here: “By the standard logic of moral hazard, if consumers optimally choose [a given level of medical spending absent insurance], they would value the ... [insurance]-induced medical spending at less than its cost, since they chose not to purchase that medical spending at an unsubsidized price.” (Finkelstein, Hendren and Luttmer, 2019, footnote 5). What this argument misses is that insurance-induced increases in medical spending also provide value ex ante, “from behind the veil of ignorance” (Hendren, 2021). This point has also been made repeatedly by Nyman (e.g., Nyman, 1999a) in discussing the importance of the income effect (the efficient increase in utilization) in driving welfare from insurance. Our analysis shows exactly why it is the case, and how the missing link is risk protection.

the dimensions relevant for normative evaluations of insurance, unifying discussions in prior work.¹³ Fourth, our results underscore that the elasticity of demand for healthcare with respect to insurance coverage is not a sufficient statistic for optimal insurance design. Considering the reason for changes in behavior is essential, and in some cases (as we demonstrate in our empirical analysis), doing so can reverse the standard intuition that greater demand response implies higher optimal cost-sharing. Finally, our approach makes clear that it is not moral hazard itself that makes first-best insurance infeasible, but rather the non-contractibility of the state. Though the issues are related, their implications are distinct. Indeed, it would be surprising if first-best insurance did not involve moral hazard.¹⁴

The second part of the paper takes the model to data. Our aim is to evaluate the quantitative importance of these forces in the context of empirical estimates of consumer preferences. To that end, we simulate a population of consumers using standard parameterizations of consumer utility and health risk, along with parameter estimates from [Einav et al. \(2013\)](#) and [Marone and Sabety \(2022\)](#). We find that the positive contribution of moral hazard to the welfare generated by insurance—the risk protection value of moral hazard—is economically large. Our main estimates imply that in a full insurance plan, it accounts for on average 10 percent of welfare. The corresponding negative contribution to welfare from the social cost of over-utilization amounts to on average only 5 percent of welfare. Eliminating moral hazard would thus have a net negative effect on welfare in this population. In other words, somehow preventing consumers from changing their behavior in response to insurance would *reduce* the social value that insurance provides.¹⁵ Fundamentally, our findings imply that, when given

¹³The critical feature is how the consumer’s marginal utility of non-health consumption varies with health status, *while fixing healthcare utilization at its socially efficient level in each state*. The three types of consumers we define are (i) *Travelers*, who would wish (under the first-best contract) to have greater non-health consumption in healthy rather than in sick states; (ii) *Homebodies*, who would wish to have equal non-health consumption in all states; and (iii) *Worriers*, who would wish to have greater non-health consumption in sick rather than in healthy states. *Homebodies* are the type most commonly featured in models in the literature, and are consistent with what is often referred to as “state-independent” preferences.

¹⁴Of course, the interpretation of this statement relies on the definition of moral hazard we have used (the entire change in behavior induced by insurance). It is probably not controversial to claim that absent insurance, some amount of under-utilization is likely to be ubiquitous. Few would be able to afford medically recommended care for brain cancer absent insurance, yet most would be willing to pay the fair premium to cover it. But the point still stands even if one wished to use the term “moral hazard” to mean only “over-utilization.” As discussed, consumers with certain preferences may prefer to keep over-utilization rather than commit to reduce it, even if such a thing were possible. The cause of their troubles would not be lack of commitment against over-utilization, but rather that sound insurance can offer no means by which to “enrich” oneself when loss events are non-verifiable.

¹⁵While there is no formal mechanism in our model to actually prevent moral hazard, there are many mechanisms that try to do so in the real world; for example, prior authorization, utilization review, referral requirements, and step therapy protocols. An important question is whether these mechanisms attempt

the opportunity to transfer resources towards sick states, consumers place a high value on the ability to spend those additional resources on additional healthcare, rather than entirely on additional non-health consumption.

There is substantial heterogeneity across consumers. Consumers with low income relative to their potential health outcomes are especially likely to wish to spend additional resources on healthcare in sick states—or in other words, are especially prone to under-utilization absent insurance. For these consumers, the risk protection value of moral hazard accounts for fully half of the welfare generated by insurance (while the social cost of over-utilization amounts to only 2 percent). Eliminating moral hazard in this group would thus be enormously detrimental. Among consumers who would value cash transfers in sick states, those who have the most elastic demand for healthcare derive substantial risk protection value from over-utilization. For 26 percent of consumers, this force is so strong that even solely eliminating over-utilization—while retaining the full efficient increase in utilization—would still lead to a net decrease in welfare (on average by 2 percent). Overall, we find that consumers who change their behavior most in response to insurance tend to have *higher* optimal coverage levels than those with less responsive demand for care. Our empirical analysis highlights the aspects of demographics and preferences that drive these results.

The study of moral hazard in health insurance has been an active area of economic research since at least the 1960’s (Cutler and Zeckhauser, 2000, p. 580). Quantifying its contribution to a social cost of insurance has been the focus of a now vast literature in health economics (see foundational work by Feldstein (1973) and Manning et al. (1987), and Einav and Finkelstein (2018) for a recent review).¹⁶ Our focus is on how moral hazard can also provide a social benefit, and on whether in some cases, this benefit can exceed its cost. Of course, many prior authors have suggested that there may be countervailing forces in consumers’ demand for healthcare, meaning increasing utilization with insurance could be beneficial. For example, there may be supra-competitive provider prices (Frick and Chernew, 2009), costly dynamic externalities (Cawley and Ruhm, 2011), liquidity constraints (Ericson, Jaspersen and Sydnor, 2025), or “behavioral hazard” under which misperceived benefits result in under-use (Baicker, Mullainathan and Schwartzstein, 2015; Chandra, Flack and Obermeyer, 2024). Our arguments

to target the level of utilization that would arise absent insurance, or whether they attempt to target the efficient level. Our analysis implies that in a competitive insurance market with no selection, they would target the efficient level, as that is where most value can be generated.

¹⁶It has also been an important object of study in disability insurance (Deshpande and Lockwood, 2022), unemployment insurance (Chetty and Finkelstein, 2013), and long-term care insurance (Lieber and Lockwood, 2019; Kesternich et al., 2025).

for social benefits from moral hazard are distinct. In our setting, the benefit of moral hazard arises from fully rational behavior in a perfectly competitive and static environment. The key necessary condition is *risk*, and the mechanism of welfare benefit is risk protection.¹⁷

Beyond healthcare, the risk protection value of moral hazard arises also in a number of other settings. Most immediately, the model we study is formally equivalent to one of optimal income taxation for redistribution (Mirrlees, 1971). As in our setting, the realized state (ability) is unobserved by the insurer (government), leaving price distortion (taxes) as the second-best instrument for risk-spreading (redistribution). Our analysis suggests that if the value of leisure is accounted for in social welfare, and if leisure is valued differently by different people, labor supply distortions need not represent a pure social cost.¹⁸ These same ideas appear also in the context of social insurance programs that use in-kind transfers, which reduce the eligible consumer’s price of some goods below their social cost (Lieber and Lockwood, 2019). Our model parallels this structure of these programs closely. Insurance-induced increases in the use of covered benefits—for example, an increase in the use of formal care in response to long-term care insurance—may thus not necessarily represent a social welfare loss. Finally, our work also connects naturally to studies of unemployment insurance (UI), where it is now understood that providing the liquidity to allow workers to better smooth utility within unemployment spells is socially beneficial (Chetty, 2008; Shimer and Werning, 2008; Landais, 2015). Whereas the UI problem is formally one of ex-ante moral hazard, we provide parallel results in the context of ex-post moral hazard.

The paper proceeds as follows. Section 2 presents our model and introduces the key concepts for welfare analysis. Section 3 presents the empirical analysis of the model. Section 4 discusses the implications of our results. Section 5 concludes.

¹⁷In contrast, a social benefit from moral hazard could arise in all the papers cited above even if the consumer faced no uncertainty about her health. In our setting, that would be impossible. This characterization of the link between the existence of risk and the existence of an “efficient increase in healthcare utilization” is also central to how we build on the body of work by J. Nyman. As discussed above, the “access motive” (Nyman, 1999b) can only provide social value in the presence of uncertainty. Otherwise, there is no scope for a transfer of resources across states.

¹⁸Similar parallels would exist in the context of long-term care insurance that covers formal care but not informal care (Mommaerts, 2025), or disability insurance that conditions benefits on earned income (Maestas, Mullen and Strand, 2013).

II Model

An expected-utility-maximizing consumer faces uncertainty about her health state $l \in \mathbb{R}$, which is distributed according to cumulative distribution function F . The consumer values two goods: healthcare utilization $m \in \mathbb{R}_+$ and non-health consumption $y \in \mathbb{R}_+$. The price (and social cost) of both goods is fixed at one throughout the analysis, such that m represents the money-metric resource cost of healthcare, and y represents a composite non-health consumption good. The consumer is initially endowed with income $w_0 \in \mathbb{R}_+$. Her preferences are represented by utility function $u(y, m; l)$, where u is strictly increasing and concave in y , strictly concave in m , and continuously differentiable in both y and m .¹⁹ Utility features a (state-dependent) bliss point with respect to m , reflecting the intuition that the returns to healthcare utilization eventually become negative. Incremental healthcare is more valuable when the consumer is in worse health, and higher l indicates worse health, so u_m is strictly increasing in l .²⁰

A health insurance contract x is characterized by an out-of-pocket cost function $c(m; x)$, where $c(m; x) \in [0, m]$ represents the consumer’s point-of-service cost of healthcare utilization level m .²¹ We restrict attention to the set of contracts \mathcal{X} for which $c(\cdot; x)$ is increasing and concave. An uninsured consumer is effectively enrolled in the null contract x_0 , under which $c(m; x_0) = m$. Insurance thus lowers the consumer’s point-of-service price of healthcare, according to the price schedule $c(m; x)$. The out-of-pocket cost function cannot depend directly on the health state l due to ex-post asymmetric information between the consumer and the insurer about which state has realized.

¹⁹Throughout the paper we use increasing and decreasing in the weak sense of non-decreasing and non-increasing, adding “strictly” when needed, and similarly with positive and negative, and concave and convex.

²⁰This model has a number of antecedents in the health economics literature. In more recent times, the literature has built directly from [Cardon and Hendel \(2001\)](#), but that model is in turn closely related to the [Grossman \(1972\)](#) framework. The notation has varied across papers, despite the model remaining essentially unchanged.

²¹A central feature of our model (and all those similar) is the classification of goods in the economy into category m (“healthcare utilization”) or category y (“non-health consumption”). Notice that what formally distinguishes them is whether a particular good is *covered* under health insurance plan x , and thus whether its point-of-sale cost is lowered under the contract. In reality, there could easily be goods that are *not* covered under contract x , but which may still be reasonably classified as “healthcare utilization” (for example, services rendered at a hospital excluded from the network of an HMO plan). In our model, any non-covered goods (e.g., out-of-network providers, cosmetic procedures, housing costs, groceries) are classified as y . This important interpretational point is explored further in [Marone and Huitfeldt \(2025\)](#). A good way to interpret the framing of the model in the present paper, where we will continue to refer to y as “non-health consumption,” is that the contract x covers *all* technologically feasible healthcare goods in the economy. Traditional Medicare in the U.S., for example, would approximately fit this description.

When the consumer is enrolled in contract x and realizes health state l , she makes an optimal choice of healthcare utilization m by trading off its benefit against the cost of foregone non-health consumption y . The optimal choice of healthcare utilization is given by:

$$m^*(l, x, p) = \arg \max_{m \geq 0} u(y, m; l) \quad s.t. \quad 0 \leq y \leq w_0 - p - c(m; x). \quad (1)$$

where $p \in \mathbb{R}_+$ represents the contract premium.²² Optimal healthcare utilization implies an optimal amount of non-health consumption $y^*(l, x, p) = w_0 - p - c(m^*(l, x, p); x)$. Prior to the realization of the health state, expected utility from the lottery induced by insurance is

$$U(x, p) = \mathbb{E}_l [u(y^*(l, x, p), m^*(l, x, p); l)].$$

For simplicity, we will restrict attention to actuarially fair pricing, which has the useful implication that the private and social value of insurance coincide in the model. In other words, there is no *ex-ante* asymmetric information between the consumer and the insurer, only *ex-post* information asymmetry. The fair premium for contract x is defined (implicitly) by

$$\bar{p}(x) = \mathbb{E}_l [m^*(l, x, \bar{p}(x)) - c(m^*(l, x, \bar{p}(x)); x)].$$

To simplify notation going forward, let $m^*(l, x) = m^*(l, x, \bar{p}(x))$ denote the consumer's optimal healthcare spending under the actuarially fair premium.

II.A Welfare

As is usual, the welfare generated from an insurance contract is equal to the difference between the consumer's willingness to pay for the contract and the insurer's expected cost of providing it. The consumer's willingness to pay corresponds to the premium that makes her indifferent between the lottery induced by insurance and the lottery faced absent insurance. That is, her willingness to pay for contract x relative to the null contract x_0 is equal to the \tilde{p} that solves:

$$U(x, \tilde{p}) = U(x_0, 0).$$

Going forward, our analysis will be more transparent if the welfare generated by the provision of an insurance contract is invariant to price (premium) at which it is provided. To that end,

²²We suppress dependence on w_0 as it will remain fixed throughout the analysis. Optimal healthcare utilization m^* will be unique in a given health state so long as $u_{ym} \geq 0$ for all m less than the bliss point.

we proceed under the following assumption.²³

ASSUMPTION 1 (Multiplicative Separability and Constant Absolute Risk Aversion in y). *The consumer’s utility function takes the form $u(y, m; l) = \tilde{u}(y) \cdot \tilde{b}(m; l)$ where $\tilde{u}(y) = -\exp(-\psi y)$ with $\psi \in \mathbb{R}_{++}$.*

The force of Assumption 1 is that the contract premium does not affect the curvature of the consumer’s utility function, and thus the value of risk protection she derives from insurance.²⁴ Constant curvature in y also implies that once uncertainty has been resolved, she acts as if she has quasilinear utility in y . This can be seen easily by rewriting utility as $u(y, m; l) = \tilde{u}(y + b(m; l))$, where the function $b(m; l) = -\psi^{-1} \log(\tilde{b}(m; l))$ can be interpreted as the state-dependent, money-metric valuation of healthcare utilization.²⁵ Since $\tilde{u}(\cdot)$ is strictly increasing, $y + b(m; l)$ is all that matters ex post. Taken together, Assumption 1 rules out local income effects of the premium on the consumer’s degree of risk aversion as well as on her demand for healthcare. Naturally, this means it is most reasonable when the fair premium is low relative to initial income.²⁶

Given the parameterization in Assumption 1, it will further simplify notation to introduce a notion of the consumer’s money-metric equilibrium “payoff,” $y + b(m; l)$, under a given

²³That social welfare depends only on allocations, and not on prices, is among the most commonly maintained assumptions in economic models. It follows directly from the standard assumption that indirect utility is quasilinear in price, which will be guaranteed in our model under Assumption 1, as long as the consumer’s initial income is not too low relative to her expected healthcare needs (see Appendix A.1). While this assumption is weak enough to incorporate most existing models in the health insurance literature, it is not without loss of generality. We return to this point in footnote 26.

²⁴Note that this assumption does not rule out dependencies between risk aversion and initial income (w_0) or health status ($\mathbb{E}[l]$); ψ can vary arbitrarily with ex-ante objects.

²⁵Given our original assumptions on utility, $b(m; l)$ inherits the properties that it is strictly concave in m , features a state-dependent bliss point in m , and has $b_m(m; l)$ increasing in l . Note that some authors (e.g. Grossman, 1972; Cardon and Hendel, 2001; Finkelstein, Hendren and Luttmer, 2019) would call b the “health production function” and say that consumers directly values “health” itself, rather than healthcare utilization.

²⁶While this parameterization of utility is standard in the health insurance literature and clearly advantageous for the study of insurance, it is somewhat atypical among economic models more broadly. In particular, note that the marginal utility of y does not go to infinity as its consumption goes to zero. A corner solution could thus arise at which zero non-health consumption is consumed, despite the absurdity of such an allocation. A common way around this issue is to simply assume initial income is sufficiently high that it never arises (e.g., Einav et al., 2013; Marone and Sabety, 2022). This solution is somewhat unsatisfying, however, when one reflects on the reality that absent insurance—and in the U.S., even sometimes with insurance—the cost of many healthcare services could be far above most incomes. A defining feature of our present model is that it does not rule out this possibility. As we will soon see, this will imply that the uninsured or under-insured consumer may in some states choose to spend all of her money on healthcare (m). Our preferred interpretation of this possibility is that in our model, y represents non-health consumption *beyond* some initial (unmodeled) consumption floor, and that w_0 represents the consumer’s initial income *in excess* of that required to fund the consumption floor.

insurance contract in a given health state:

$$z^*(l, x) = w_0 - \bar{p}(x) - c(m^*(l, x); x) + b(m^*(l, x); l). \quad (2)$$

Prior to the realization of health state l , an insurance contract then simply represents a lottery over equilibrium payoffs, with expected value, certainty equivalent, and risk premium given by:

$$\begin{aligned} EV(x) &= \mathbb{E}_l [z^*(l, x)] \\ CE(x) &= \tilde{u}^{-1}(\mathbb{E}_l [\tilde{u}(z^*(l, x))]) \\ RP(x) &= EV(x) - CE(x). \end{aligned} \quad (3)$$

Appendix A.1 shows that under Assumption 1, the welfare generated by a focal contract x relative to the null contract x_0 can be expressed as:

$$V(x) = CE(x) - CE(x_0) = \underbrace{RP(x_0) - RP(x)}_{\substack{\text{Value of risk} \\ \text{protection,} \\ \Psi(x)}} + \underbrace{EV(x) - EV(x_0)}_{\substack{\text{Net social cost of} \\ \text{moral hazard,} \\ \Delta EV(x)}}. \quad (4)$$

The first term on the right-hand side, Ψ , represents the value of risk protection provided by insurance. If there were no uncertainty about the health state, or if $\tilde{u}(\cdot)$ were linear, it would equal zero. Otherwise, if the null contract provides a riskier distribution of payoffs than contract x , Ψ will be positive. The second term, ΔEV , is the incremental *expected* payoff derived from insurance. It can therefore be nonzero even if the health state is degenerate. If the two contracts had equal premiums, ΔEV would also be positive, since higher coverage delivers higher payoffs before accounting for premiums. However, since higher coverage comes with higher insurer costs, it also comes with a commensurately higher fair premium. If healthcare utilization were held fixed across contracts, these two forces would exactly offset, and ΔEV would equal zero.²⁷ Otherwise, ΔEV will reflect the extent to which insurance (in expectation) leads to ex-post moral hazard. It will be negative if the consumer does not value her expected change in utilization at cost. As is widely appreciated in health economics, moral hazard can thus create a social welfare cost of insurance (Pauly, 1968). The purpose of our analysis going forward is to explore the extent to which eliminating moral hazard may also detract from the social benefit of insurance.

²⁷To see this, note that under actuarially fair premiums, $EV(x) = w_0 + \mathbb{E}_l[b(m^*(l, x); l)] - \mathbb{E}_l[m^*(l, x)]$.

II.B Decomposition of Healthcare Utilization

Our goal is to understand the role of moral hazard in driving social benefit from insurance, in addition to its traditional role in driving social cost. To that end, we decompose healthcare utilization under a given insurance contract into two parts: *moral hazard spending* (i.e., causally resulting from insurance) and *non-moral hazard spending* (i.e., would have arisen even absent insurance).²⁸

Non-moral hazard spending is defined as the amount of healthcare a consumer would find optimal absent insurance:

$$m^N(l) \equiv m^*(l, x_0) = \arg \max_{m \geq 0} u(y, m; l) \quad s.t. \quad 0 \leq y \leq w_0 - m. \quad (5)$$

It is invariant to the insurance contract the consumer ultimately enrolls in. Moral hazard spending—under a given insurance contract x —is then simply all remaining spending that the consumer finds optimal: $m^*(l, x) - m^N(l)$. Because healthcare spending is generically higher under insurance than without insurance, moral hazard spending is positive in every health state (See Appendix A.3 for a proof). As we explore next, however, this amount is not entirely wasteful spending.

Decomposition of Moral Hazard. Moral hazard is typically characterized as a phenomenon that would be best prevented. Absent insurance, the argument goes, consumers are fully exposed to the resource cost of healthcare, and will thus trade off its social benefit with its social cost. There is a small literature in health economics, however, that has argued that some amount of moral hazard spending may actually be socially efficient (Zeckhauser, 1970; de Meza, 1983; Nyman, 1999a,b; Cutler and Zeckhauser, 2000; Nyman, 2002, 2007, 2008; Nyman et al., 2018). This idea has a formal manifestation in our model.

Despite the fact that the uninsured consumer is fully rational, fully informed, and faces prices that reflect true social costs, her level of non-moral hazard spending $m^N(l)$ need not coincide with the socially efficient level. In fact, in some health states, it may fall far below the efficient level. The intuition for this possibility follows exactly the arguments laid out by the

²⁸It is worth emphasizing that the definition of “moral hazard spending” has been a matter of some debate in the health economics literature. We follow what has become the most common interpretation of its meaning, which is that ex-post moral hazard refers to *all incremental spending induced by insurance* (Einav and Finkelstein, 2018, p. 959). An alternative definition has been used by de Meza (1983); Nyman (1999a, 2002) and Cutler and Zeckhauser (2000), where moral hazard is said to refer only to the “likely malfeasance of an individual making purchases that are partly or fully paid for by others” (Ibid.). This aspect of moral hazard is what we will refer to below as “over-utilization” (or the “substitution effect”).

above-mentioned authors.²⁹ The amount of healthcare the consumer chooses to utilize absent insurance represents her *ability* to pay for care given initial resources. But this amount does not necessarily reflect her *willingness* to pay for care in that health state, had she had the ex-ante opportunity to transfer resources to that state at actuarially fair rates. Her choices under this latter scenario would have revealed the socially efficient level of healthcare utilization.

To formalize this intuition, we define the *socially efficient level* of healthcare utilization in a given health state to be the amount of utilization that would arise in that state under the first-best insurance contract (Cutler and Zeckhauser, 2000). In the first best world, the health state is contractible (i.e., there is no longer ex-post asymmetric information between the consumer and the insurer). The consumer can thus choose state-contingent income transfers to maximize her expected utility, subject to the constraint that she can afford these transfers in expectation. In other words, the first-best insurance contract allows the consumer to reallocate her initial income across states in order to maximize her expected utility. Her first-best vector of state-contingent incomes is thus given by:

$$(w^{FB}(l)) = \arg \max_{(w(l))} \mathbb{E}_l [v(w(l); l)] \quad s.t. \quad \mathbb{E}_l [w(l)] \leq w_0, \quad (6)$$

where $v(w; l) = u(y^*(w; l), m^*(w; l); l)$ is the consumer's baseline indirect utility function, with optimal bundles given by $m^*(w; l) = \arg \max_{0 \leq m \leq w} u(w - m, m; l)$ and $y^*(w; l) = w - m^*(w; l)$. Going forward, we refer to the socially efficient levels of health and non-health consumption by $m^{FB}(l)$ and $y^{FB}(l)$, which correspond to the first-best bundles $m^*(w^{FB}(l); l)$ and $y^*(w^{FB}(l); l)$. At these bundles, the consumer's marginal indirect utility of income v_w will be equated across states.³⁰ In this precise sense, the sole reason the consumer demands actuarially fair insurance is to smooth her marginal indirect utility of income across health states, or in other words, to protect herself from risk.

This formulation of the first-best contract highlights that in health insurance models with ex-post moral hazard, the consumer fundamentally faces *preference* risk as opposed to *income*

²⁹For example, from de Meza (1983), "To illustrate, in the absence of insurance most people would be unable to afford major open heart surgery. Given the opportunity to take out insurance to cover the expense they would jump at the chance. In the woeful event that an operation is needed, if the insurance company were to offer the choice between surgery or a payment equal to its cost (thereby allowing some cheaper treatment to be substituted), the cash option would normally be rejected. If so, the expansion in demand has no efficiency cost."

³⁰Unless she finds it preferable to allocate zero income to some state(s), in which case v_w would only be equated across states with non-zero income.

risk.³¹ Absent insurance, her marginal indirect utility of income $v_w(w_0; l)$ varies across states—just like in a model with income risk and a single consumption good—but does so because of differential “needs” (or “wants”), rather than because of differential resources. She may therefore benefit from reallocating her income across states. But naturally, changing the consumer’s income in a given state may also change her optimal bundle in that state. It is in this sense that the consumer’s actions absent insurance represent her ability to pay, while her actions under the first-best contract represent her willingness to pay.

As long as initial income is not excessively low relative to expected healthcare utilization, non-moral hazard spending $m^N(l)$ will fall weakly below socially efficient healthcare spending $m^{FB}(l)$ in every health state (see Appendix A.2 for a proof). The reason is that, absent insurance, the consumer must respect her resource constraint w_0 on a *state-by-state basis*. In contrast, under the first-best insurance contract, she need only respect her resource constraint *in expectation*.³² This flexibility allows the consumer to transfer potentially huge amounts of resources to very unlikely health states. In turn, she then has the possibility to purchase far more expensive treatments with insurance than she could have afforded absent insurance. The discrepancy between $m^N(l)$ and $m^{FB}(l)$ is exactly the focus of John Nyman’s body of work on the interpretation of moral hazard, and represents one important channel through which some moral hazard spending may represent a social benefit.

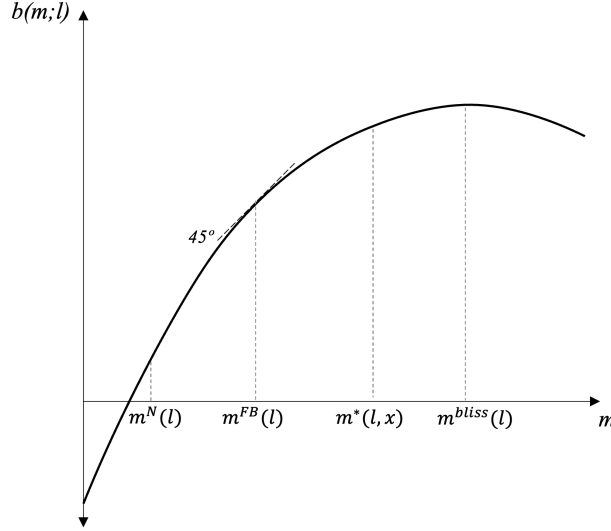
Figure 1 provides a depiction of the relationship between these levels of utilization and the consumer’s money-metric valuation of healthcare $b(m; l)$ in a given health state. Recall that $b(m; l)$ is strictly concave in m and features a bliss point beyond which additional care only reduces utility. The socially efficient level of healthcare utilization achieves where $b_m(m; l) = 1$, or in other words, where the marginal benefit of healthcare (b_m) equals its marginal social cost

³¹This distinction clarifies much of the intuition from Nyman’s work, who (rightly) took issue with arguments that the only motivation for purchasing health insurance was reduce risk of income loss: “[T]here is a fundamental difference between insurance against a loss of property ... and insurance against a ‘loss’ represented by the quid pro quo purchase of medical care ... [W]hile the loss of income or wealth as medical expenditures is captured by conventional theory, the other side of the transaction is not, and the access to medical care can represent an extremely valuable gain. This suggests that conventional theory may be inappropriately applied to health insurance and that an alternative theory is needed to explain better the demand for this important commodity.” (Nyman et al., 2018). Our analysis shows that all that is needed to explain the “access motive” for purchasing insurance is a model with moral hazard (two goods), and clarifies that this motive is still predicated directly on risk.

³²This logic also helps explain why $m^N(l)$ and $m^{FB}(l)$ will only differ if the consumer faces *uncertainty* about her health state. Absent uncertainty, $m^N(l)$ and $m^{FB}(l)$ will coincide. This fact represents a key distinction between the “under-utilization” that can arise in our model (for which *risk* is a necessary condition), and the “under-utilization” that can arise in models with information frictions or externalities—as in, for example, Baicker, Mullainathan and Schwartzstein (2015) or Frick and Chernew (2009)—for which risk is not a necessary condition.

(1).³³ In this sense, the traditional argument that it is efficient to expose consumers to the full social cost of care is correct. What is missing from this argument is that in the absence of insurance, the consumer may not be able to afford to reach this interior solution. She may instead find herself at a corner, at which she has spent all her money on healthcare, but its shadow value still exceeds its marginal cost.³⁴

Figure 1. Benefit of Healthcare Utilization: $b(m; l)$



Notes: This figure shows the money-metric valuation of healthcare $b(m; l)$ as a function of healthcare utilization m , in a given health state l . The socially efficient level of healthcare is $m^{FB}(l)$, which corresponds to the level at which the marginal benefit of healthcare equals its marginal social cost: $b_m(m; l) = 1$. The equilibrium level of healthcare utilization under contract x is $m^*(l, x)$, which corresponds to the level at which the marginal benefit of healthcare equals its marginal out-of-pocket cost: $b_m(m; l) = c_m(m; x)$. The level of healthcare utilization absent insurance is $m^N(l)$. If this level falls below $m^{FB}(l)$, it must be because the consumer was constrained in that state absent insurance, i.e., $m^N(l) = w_0$.

Returning back to the reality of traditional health insurance in which the state is not contractible, Appendix A.3 then shows that as long as coverage is not excessively low, equilibrium spending $m^*(l, x)$ will weakly exceed first-best spending $m^{FB}(l)$ in all health states. Beyond a potential efficient increase in utilization, insurance will therefore still also lead to some amount of over-utilization of care. The equilibrium level of healthcare utilization will equate the marginal benefit of healthcare (b_m) to its marginal out-of-pocket cost (c_m): $b_m(m; l) = c_m(m; x)$. As marginal out-of-pocket cost lies between zero and one, equilibrium utilization $m^*(l, x)$ lies between $m^{FB}(l)$ and $m^{bliss}(l)$ (the satiation level of m). Figure 1 illustrates this relationship as well.

³³See Lemma 1 in Appendix A.2.

³⁴See footnote 26 for a discussion of how zero non-health consumption can be interpreted.

Taken together, for a given health state realization, the moral hazard spending amount $m^*(l, x) - m^N(l)$ represents two distinct components: an *efficient increase* in utilization induced by insurance $m^{FB}(l) - m^N(l)$, plus some amount of *over-utilization* $m^*(l, x) - m^{FB}(l)$. This decomposition will be central to interpreting how moral hazard contributes to the social value of insurance.

II.C Decomposition of Welfare

Quantifying moral hazard spending—and its contribution to a social cost of insurance—has been the focus of a now vast literature in health economics (see [Einav and Finkelstein \(2018\)](#) for a recent review). Our focus is on how it can also provide social benefit, and on whether in some cases, this benefit can exceed its cost. We approach this question by evaluating how the social welfare derived from insurance would change if moral hazard spending could be prevented.

The one remaining building block is to state the payoffs that would arise under each of the two alternative levels of utilization described above, $m^N(l)$ and $m^{FB}(l)$. If the consumer were enrolled in focal contract x and chose to utilize the same level of healthcare as she would have done absent insurance, the payoff she would derive in health state l is given by

$$z^N(l, x) = w_0 - \bar{p}^N(x) - c(m^N(l); x) + b(m^N(l); l), \quad (7)$$

where $\bar{p}^N(x) = \mathbb{E}_l[m^N(l) - c(m^N(l); x)]$ is the corresponding actuarially fair premium. Similarly, the payoff she would derive from choosing the efficient level of spending is given by

$$z^{FB}(l, x) = w_0 - \bar{p}^{FB}(x) - c(m^{FB}(l); x) + b(m^{FB}(l); l), \quad (8)$$

where $\bar{p}^{FB}(x) = \mathbb{E}_l[m^{FB}(l) - c(m^{FB}(l); x)]$. The associated components of the value of insurance can thus be written:

$$\begin{aligned} EV_N(x) &= \mathbb{E}_l z^N(l, x) & EV_{FB}(x) &= \mathbb{E}_l z^{FB}(l, x) \\ CE_N(x) &= \tilde{u}^{-1}(\mathbb{E}_l \tilde{u}(z^N(l, x))) & CE_{FB}(x) &= \tilde{u}^{-1}(\mathbb{E}_l \tilde{u}(z^{FB}(l, x))) \\ RP_N(x) &= EV_N(x) - CE_N(x) & RP_{FB}(x) &= EV_{FB}(x) - CE_{FB}(x). \end{aligned}$$

Using these definitions, we can now further decompose each of the two terms comprising the welfare generated by insurance, as stated in Equation 4: $V(x) = \Psi(x) + \Delta EV(x)$.

Net Social Cost of Moral Hazard. Because it is more straightforward, we begin with the net social cost of moral hazard, $\Delta EV(x)$. This term is in its entirety related to the presence of moral hazard spending, but its sign is ambiguous. Mirroring the decomposition of moral hazard discussed above, we can split this term into two: one component that is generically negative, and one that is generically positive:

$$\Delta EV(x) = \underbrace{EV(x) - EV_{FB}(x)}_{\substack{\text{Social cost} \\ \text{of over-utilization} \\ \Delta EV_{over}(x) \\ (-)}} + \underbrace{EV_{FB}(x) - EV(x_0)}_{\substack{\text{Social benefit of efficient} \\ \text{increase in utilization} \\ \Delta EV_{eff}(x) \\ (+)}} . \quad (9)$$

Recall that under actuarially fair premiums, the consumer pays the full *expected* social cost of all healthcare she uses. Under a given vector of state-contingent healthcare utilization levels ($m(l)$), the consumer's expected payoff under a fair premium can thus be expressed as $EV = w_0 + \mathbb{E}_l[b(m(l); l)] - \mathbb{E}_l[m(l)]$. As a result, if healthcare utilization is held fixed, the *difference* in EV 's derived from any two fairly-priced contracts would be zero.³⁵ But, if she *does* change her utilization, then the difference in EV 's will reflect the extent to which she does or does not value that change in utilization at cost.

The first term in Equation 9, ΔEV_{over} , represents the consumer's expected valuation of moving from her socially-efficient level of healthcare utilization to the equilibrium utilization level induced by insurance contract x . It is clear from Figure 1 that if these levels do not coincide, then the consumer does *not* value the increase in utilization at cost. The social cost of over-utilization represents the expected amount by which equilibrium utilization is valued below its cost. ΔEV_{over} is therefore generically negative. The second term, ΔEV_{eff} , represents the consumer's expected valuation of moving from the utilization she would choose absent insurance to her socially-efficient level of utilization. Figure 1 also makes clear that in any state in which these levels do not coincide, then the consumer *does* value the increase in utilization at cost. ΔEV_{eff} is therefore generically positive.³⁶

Taken together, if the magnitude of ΔEV_{eff} exceeds that of ΔEV_{over} , the consumer's expected payoff from actuarially fair insurance will be positive. In other words, even for a consumer who is risk neutral with respect to gambles over non-health consumption (i.e., for whom $\tilde{u}(\cdot)$ is linear), the fact that she can change her behavior in response to insurance can still represent a net social benefit. This possibility has everything to do with the extent of

³⁵Thus, $EV_N(x) = EV(x_0)$ for all x .

³⁶Appendix A.5 provides a proof of both statements in this paragraph.

curvature in the consumer’s utility from healthcare utilization (represented by the function $b(m; l)$). Naturally, curvature in b can translate into curvature in the consumer’s indirect utility from income, even if her direct utility from non-health consumption is linear.³⁷ We return to this point in Section II.D below.

Value of Risk Protection. Consider now the value of risk protection, $\Psi(x)$. It can be decomposed into three parts: (i) the component that would arise absent moral hazard, (ii) the incremental amount associated with the efficient increase in healthcare utilization, and (iii) the incremental amount associated with over-utilization:

$$\Psi(x) = \underbrace{RP(x_0) - RP_N(x)}_{\substack{\text{Value of risk protection} \\ \text{absent moral hazard,} \\ \Psi_N(x)}} + \underbrace{RP_N(x) - RP_{FB}(x)}_{\substack{\text{Value of risk protection} \\ \text{from efficient increase} \\ \text{in spending,} \\ \Psi_{eff}(x)}} + \underbrace{RP_{FB}(x) - RP(x)}_{\substack{\text{Value of risk protection} \\ \text{from over-utilization,} \\ \Psi_{over}(x)}}. \quad (10)$$

Perhaps surprisingly, it is not possible to determinatively sign any of these terms, nor $\Psi(x)$ as a whole, given only the assumptions established so far. The reason is that we have placed no restriction on whether the consumer would prefer to have more non-health consumption in healthy or in sick states (or to have y equal in all states). While it is natural to assume (as we have) that the consumer will want more healthcare when she is more sick, it is not at all clear what to expect about non-health consumption. If sickness level (l) and non-health consumption (y) are *complementary*, then one would prefer to have more y when one is more sick (for example, to hire an in-home aide or purchase greater material comforts). If sickness level and non-health consumption are *substitutable*, then one might want to have more y when one is less sick (for example, to travel). This relationship is central to whether the value of risk protection provided by health insurance is in fact positive.³⁸

³⁷The related literature (for example [Marone and Sabety, 2022](#)) nearly universally discusses the consumer’s risk aversion parameter ψ as if it fully represented her risk aversion, in the sense of the claim: if ψ were equal to zero, the consumer should optimally have no insurance. What our analysis here makes clear is that this claim is not correct. When the consumer values multiple goods and faces risk over her tastes for them, quantifying her “risk aversion” becomes more complex. In our model, curvature in b also represents risk aversion, in the sense that greater curvature in b increases the risk premium associated with a given gamble, holding all else equal.

³⁸There are a limited number of empirical studies investigating the relationship between sickness and the marginal utility of non-health consumption. Recent studies have found evidence suggesting the relationship is negative (that is, that sickness level and non-health consumption are substitutable), albeit heterogeneous across types of illnesses ([Finkelstein, Luttmer and Notowidigdo, 2013](#); [Blundell et al., 2024](#)). Earlier evidence yielded a wide range of estimates ([Viscusi and Evans, 1990](#); [Evans and Viscusi, 1991](#); [Sloan et al., 1998](#)). As discussed in the introduction, the increasing uptake of insurance products that offer cash payouts when sick (for example, critical illness insurance) suggests also that there are at least some consumers for whom

If sickness level and non-health consumption are substitutable, then the consumer is fairly happy not to have insurance. When she gets sick, she allocates a greater fraction of resources to healthcare. Consistent with her preferences, this leaves less non-health consumption in sick states, but—since she is not paying an insurance premium—greater non-health consumption in healthy states. Conversely, if sickness level and non-health consumption are complementary, insurance is highly valuable, as it helps transfer non-health consumption toward sick states. We formalize these distinctions with the following definitions:

DEFINITION 1 (Consumer Types). *Consider the consumer's socially-optimal allocation of goods across health states: $(y^{FB}(l), m^{FB}(l))$:*

- (i) *The consumer is a **Traveler** if $y^{FB}(l)$ is decreasing in l .*
- (ii) *The consumer is a **Homebody** if $y^{FB}(l)$ is constant in l .*
- (iii) *The consumer is a **Worrier** if $y^{FB}(l)$ is increasing in l .*

That is, *Travelers* prefer to have greater non-health consumption when they are healthier, while *Worriers* prefer to have greater non-health consumption when they are sicker, and *Homebodies* prefer to have equal non-health consumption in all states. Note that in all cases, healthcare utilization is fixed at its efficient level. The types thus do not define how the consumer wants to allocate her *income* across states, but rather how she wants to allocate her *residual* income net of optimal spending on healthcare. Whether the value of risk protection provided by insurance is determinatively positive will depend on the consumer's type.

Returning to Equation 10, the first term, Ψ_N , is the value of risk protection that would arise even absent moral hazard. Its magnitude reflects the value to the consumer of smoothing her non-health consumption across states, holding her healthcare utilization behavior fixed. As long as the consumer is a *Homebody* or a *Worrier* (i.e., as long as she in fact wishes to have equal non-health consumption across states, or a higher amount in sicker states), this term is positive.³⁹

The second term, Ψ_{eff} , is the incremental value of risk protection associated with any efficient increase in healthcare utilization induced by insurance. If there were no under-utilization absent insurance (i.e., if $m^N = m^{FB}$ in all health states), this term would equal zero. But if healthcare utilization absent insurance falls below the efficient level in any state,

sickness level and non-health consumption are complementary.

³⁹Proofs of statements in this and the following paragraphs are provided in Appendix A.5.

this term will be strictly positive, reflecting mechanically the fact that the marginal indirect utility of income is fully equated only at the efficient levels of healthcare spending.

Finally, the third term, Ψ_{over} , reflects the incremental value of risk protection associated with any over-utilization induced by insurance. Even conditional on being a *Homebody* or a *Worrier*, the sign of Ψ_{over} hinges on whether the consumer values her over-utilization more when she is sick or when she is healthy. If equilibrium over-utilization is more valuable when she is healthy, its presence will erode the beneficial risk smoothing being done by the insurance contract. If it is more valuable when she is sick, its presence can enhance the beneficial risk smoothing being done by the contract. If it were equally valuable in all health states, Ψ_{over} will just equal zero. Our model permits any of these possibilities. Most existing models in the structural health insurance literature have embedded assumptions that Ψ_{over} is either zero or positive.⁴⁰

II.D The Risk Protection Value of Moral Hazard

We now have all the building blocks to characterize the mechanisms by which the consumer’s ability to change her behavior in response to insurance directly contributes to the ability of insurance to protect her from risk. We refer to the direct contribution of moral hazard to risk protection as the “risk protection value of moral hazard.” Echoing the discussion above, it has two main components, representing two distinct reasons why moral hazard can be socially valuable.

To build intuition, it is useful to first take stock of how traditional health insurance—through lowering the price of healthcare—operationally provides risk protection. Consider the framing of the first-best problem, under which the consumer chooses a set of state-contingent income transfers. Canonically, she will wish to transfer income towards sick states, so that she can afford the healthcare she needs when sick, without depriving herself of non-health consumption. Though traditional health insurance cannot achieve the same *nominal* income transfer, it can (through price effects) achieve a *real* income transfer, which may be nearly as good. In particular, insurance will decrease real income in healthy states (when the premium was paid but cheaper healthcare was of little value) and increase it in sick states (when cheaper

⁴⁰To our knowledge, however, there has not been formal analysis of this point in the literature, most likely due to the considerable complications in setting up a valid empirical test. One would need to learn about both the elasticity and level of demand for healthcare in (verifiably) different health states. Shedding light on these relationships within a well-identified empirical setting seems to us like a productive direction for future research.

healthcare was highly valuable).

Moral hazard is then simply a question of how these changes purchasing power are used. If healthcare utilization were forced to be held fixed, any changes in real income would be commensurate changes in non-health consumption. This would achieve a transfer of non-health consumption towards sick states, which could be highly valuable. Indeed, we discussed above that if the consumer is a *Homebody/Worrier*, risk protection absent moral hazard, Ψ_N , will be generically positive (and Section III below will show that it may quantitatively represent the vast majority of the value of health insurance). But in some sick states, spending all additional income on non-health consumption may not be what the consumer would have *optimally* done, even at non-distorted prices. She may have preferred to purchase some additional healthcare.

As any economics textbook will describe, lowering the price of a good will have two effects on demand: an income effect and a substitution effect. If transferring income towards sick states aligns with what the consumer would want to do under the first-best contract, then the portion of the change in healthcare utilization attributable to the income effect will contribute positively to risk protection. It is exactly this part of moral hazard that we call the “efficient increase in utilization.” Notice that its existence relies on *risk*. If the health state were degenerate, there could be no (either real or nominal) income transfers across states.⁴¹ The two components of welfare that are attributable to this “income effect” component of moral hazard, Ψ_{eff} and ΔEV_{eff} , are thus both classifiable as part of the *risk* protection value of moral hazard.

The substitution effect, on the other hand, can play two roles. First, and most notoriously, it represents an inefficient reallocation of consumption from y to m . It is for this reason we refer to this component of moral hazard as “over-utilization.” In expectation, it produces (for any type of consumer) a social welfare loss, ΔEV_{over} . This loss would arise even if the consumer faced no risk, in which case a welfare evaluation of insurance would be no different than a study of the effect of a tax (Friedman, 1953; Harberger, 1964; Pauly, 1969). The second role of over-utilization is more subtle. Consider again the transfer of real income achieved by health insurance. Even under the maximum possible coverage—a full insurance plan in which healthcare were completely free—there is a limit on the amount that can be transferred towards sick states.⁴² Since the state is not contractible, additional cash transfers

⁴¹Indeed, as discussed in footnote 32, if the health state distribution were degenerate, then the socially efficient level of healthcare utilization would equal the consumers’ optimal utilization absent insurance ($m^{FB} = m^N$), and there would be no efficient increase in utilization.

⁴²In particular, the maximum gain in real income a consumer could receive in a given health state would be equal to the amount she was initially spending on healthcare absent insurance (less the fair premium).

over and above the cost of utilized care are simply not possible. But if the consumer is a *Worrier*, that is exactly what she would want to arrange. Nonetheless, though she cannot transfer additional resources to sick states in the form of cash, she can still do so in the form of healthcare utilization. As long as m and y are at all substitutable, over-utilization can thus provide an imperfect substitute for cash transfers, allowing the consumer to raise her utility in bad states by more than she would be able to do absent the opportunity for over-utilization.⁴³

Taking stock, there are two reasons that the consumer’s ability to change her behavior in response to insurance may provide valuable risk protection. First, for consumers who do not want to have lower non-health consumption in sick states—*Homebodies* or *Worriers*—any efficient increase in utilization represents valuable protection from risk.⁴⁴ This piece is represented by Ψ_{eff} and ΔEV_{eff} together. Second, for consumers who want to have greater non-health consumption in sick states—*Worriers*—over-utilization can serve as an imperfect substitute for a cash transfer, allowing the consumer to better smooth her marginal indirect utility from income, and thus also represent valuable risk protection. This piece is represented by Ψ_{over} . Re-writing the social value generated by insurance using these interpretations, we have:

$$V(x) = \underbrace{\Psi_N(x)}_{\text{Value of risk protection absent moral hazard}} + \underbrace{\Psi_{eff}(x) + \Delta EV_{eff}(x) + \Psi_{over}(x)}_{\text{Risk protection value of moral hazard}} + \underbrace{\Delta EV_{over}(x)}_{\text{Social cost of over-utilization}}. \quad (11)$$

Our goals in the remainder of the paper are twofold. First, we want to evaluate the quantitative importance of each of these forces in standard parameterizations of consumer preferences that have been estimated and used in the health insurance literature. Our second goal is to ask whether these estimates imply that the social benefit of moral hazard (its risk protection value) can ever be worth its social cost.

⁴³For example, even if the consumer would have preferred a cash-equivalent payment when sick so that her spouse could stay home from work to take care of her, receiving free in-home nursing visits may be a reasonably good substitute for her.

⁴⁴Note that this component could be positive even for *Travelers*, it will just not be generically so. It will depend on “how big of travelers” they are.

III Mapping to Empirical Evidence

Existing evidence on the welfare impacts of insurance in the health economics literature has taken roughly one of three forms. First, there is a long tradition of estimating just the welfare *loss* from health insurance, dating back to the seminal paper by [Pauly \(1969\)](#). This paper first introduced the idea that because health insurance lowers the price of healthcare below its cost, consumers will naturally over-utilize insured care, and the welfare loss can be quantified using the same “Harberger triangle” approach that one might use to evaluate the effects of a tax. Subsequent authors have followed the same approach, while bringing to bear better measurements of the demand curve for healthcare using either experimental (e.g., the RAND HIE, [Manning et al. \(1987\)](#); the Oregon HIE, [Finkelstein, Hendren and Luttmer \(2019\)](#)) or quasi-experimental methods (see [McGuire \(2011\)](#) for a review).

In the language of our framework, the object being estimated in these studies is ΔEV_{over} . Recalling from above that this object is generically negative, all authors thus rightly conclude that there is a social welfare loss from moral hazard. A key characteristic of these analyses is that they rely only on measurements of consumer’s demand for *healthcare*, and not of her demand for health *insurance*. They therefore do not take a stand on the consumer’s attitudes towards risk nor the amount of risk she is facing, and thus do not measure the benefit side of the equation. Indeed, ΔEV_{over} is the *only* object within our framework that can be measured using only data on ex-post actions (demand for healthcare alone). As is clear from our analysis, however, this term alone is not sufficient for understanding the *entire* effect of moral hazard on welfare.

The second and third forms of evidence in the literature study both the benefit and cost of insurance, but do so using distinct approaches: a demand function approach or a utility function approach ([Berry and Haile, 2021](#)). With a demand function approach, the researcher uses quasi-experimental variation to directly measure the demand curve for a particular insurance contract, as well as the insurer’s marginal (expected) cost of producing that contract. The classic reference is [Einav, Finkelstein and Cullen \(2010\)](#). While this approach is advantageous for many research questions, it cannot speak to primitives underlying the demand and cost curves, and thus it cannot tell us about what these objects would look like absent moral hazard.

With a utility function approach, the researcher specifies a consumer utility function and probability distribution over health outcomes as the primitive objects, and treats consumer demand for healthcare/insurance and insurer costs as derived objects. Here, [Zeckhauser \(1970\)](#)

and Cardon and Hendel (2001) are foundational papers, the latter of which spawned a now long string of studies using approximately the same model (e.g., Einav et al., 2013; Bajari et al., 2014; Finkelstein, Hendren and Luttmer, 2019; Marone and Sabety, 2022; Ho and Lee, 2023). In exchange for stronger assumptions, the utility approach allows a full dissection of the social value of insurance into the component parts we describe. Typically, however, the breakdown is presented as between the value of risk protection and the (net) social cost of moral hazard. But as discussed, this breakdown does not provide an answer to how the value of insurance would change if moral hazard could be prevented.

Our goal in this section is to provide what is, to our knowledge, the first breakdown of welfare into the components described in Equation 11: (a) the value that insurance can provide even without allowing consumers to change their behavior (Ψ_N), (b) the additional risk protection it can provide when consumers are allowed to change their behavior ($\Psi_{eff} + \Psi_{over} + \Delta EV_{eff}$), and (c) the welfare loss of expected over-utilization (ΔEV_{over}). If moral hazard were eliminated, only Ψ_N would remain.

III.A Parameterization of Utility

Given Assumption 1—that utility takes the form $\tilde{u}(y + b(m; l))$ —the only parameterization one must make within our model is of the function $b(\cdot)$, representing the consumer’s money-metric value of healthcare. Following the literature, we impose that it is quadratic in healthcare utilization (Cardon and Hendel, 2001; Einav et al., 2013), implying that consumers have linear demand for healthcare (consistent with Pauly’s original analysis).

We focus on two parameterizations of utility. The first represents a *Homebody* for whom the level of over-utilization is invariant to health status (and thus the elasticity of demand for healthcare is lower when sicker). This specification is equivalent to the “additive” model used in the main text of Einav et al. (2013). The second represents a *Worrier* for whom elasticity of demand for healthcare is constant in health status (and thus the level of over-utilization is higher when sicker). This specification corresponds to the “multiplicative” model used in Appendix E of Einav et al. (2013).⁴⁵

⁴⁵The close reader will notice that the original “multiplicative” specification in Einav et al. (2013) would have actually also corresponded to a *Homebody*, since payoffs were constant across health states at the efficient level of utilization. We adjust the parameterization because the original results in the undesirable property that the value of risk protection from insurance is not everywhere increasing in coverage level (see Appendix A.2 of Marone and Sabety (2022) for a formal definition of coverage level ordering). Since the consumer derives greater value from over-utilization in sicker states under the “multiplicative” model of moral hazard, equilibrium payoffs under the original parameterization begin becoming *higher* when she is

Table 1 presents the two parameterizations. In each case, there is one additional parameter introduced, which governs the magnitude of the consumer’s elasticity of demand for healthcare. In specification (1), $\omega \in \mathbb{R}_{++}$ represents the dollar amount of over-utilization that would result under a full insurance contract (and this amount is invariant to health status). In specification (2), $\tilde{\omega} \in \mathbb{R}_{++}$ represents the proportional amount of over-utilization that would result under a full insurance contract (and this proportion is invariant to health status). In other words, if the consumer is twice as sick, she engage in twice as much over-utilization. The final column provides the implied demand for healthcare as a function of a (linear) out-of-pocket price $\bar{c} \in [0, 1]$, where $\bar{c} = 0$ represents full insurance, and $\bar{c} = 1$ represents the null contract.

Table 1. Specifications of Utility

Spec.	Description	Parametrization of $b(m; l)$	Implied demand for healthcare
(1)	<i>Homebody</i> with elasticity of demand for healthcare lower when sicker	$(m - l) - \frac{1}{2\omega}(m - l)^2 - \frac{\omega}{2}$	$l + \omega(1 - \bar{c})$
(2)	<i>Worrier</i> with elasticity of demand for healthcare constant in health status	$(m - l) - \frac{1}{2\tilde{\omega}l}(m - l)^2 - \frac{\tilde{\omega}l}{2}$	$l + \tilde{\omega}l(1 - \bar{c})$

Notes: The table shows the parameterizations of the function $b(m; l)$ we will use in our numerical analysis. The final column provides the implied demand functions for healthcare, $m(\bar{c}; l)$, where $\bar{c} \in [0, 1]$ represents a scalar price of healthcare. In specification (1), the implied elasticity of demand for healthcare is given by $\varepsilon = -\bar{c}\omega/(l + (1 - \bar{c})\omega)$. In specification (2), the implied elasticity of demand for healthcare is given by $\varepsilon = -\bar{c}\omega/(1 + (1 - \bar{c})\omega)$.

In both specifications, the efficient level of healthcare utilization is normalized to equal the health state: $m^{FB}(l) = l$. In specification (1), the consumer thus derives a constant benefit from her efficient healthcare utilization in every health state: $b(m^{FB}(l); l) = -\frac{\omega}{2} \forall l$. As a result, her marginal utility from non-health consumption is constant across health states at the efficient level of health utilization. In other words, she is a *Homebody*. In specification (2), the key difference is that holding y fixed, the consumer’s utility level is lower when she is sicker. At the efficient level of utilization, she thus has state-dependent marginal utility of non-health consumption: $b(m^{FB}(l); l) = -\frac{\tilde{\omega}l}{2} \forall l$. In particular, her marginal utility from non-health consumption is higher when she is sicker, meaning she is a *Worrier*.

sick than when she is healthy once coverage is high enough. At full insurance, she may have completely over-done it, exposing herself to more payoff risk than she would have done at a lower level of coverage. Adding the $-\tilde{\omega}l/2$ term exactly offsets this effect, ensuring that higher coverage always yields less payoff risk (and thus greater risk protection). While we don’t formally impose this restriction on utility in Section II, risk protection increasing in coverage level is an intuitive property for an insurance model to have, and its absence complicates the interpretation of our numerical results. In both the specifications we use, maximal risk protection—that is, fully smoothed marginal indirect utility of income—is achieved under full insurance, consistent with the standard intuition.

We focus our primary analysis on the parameter estimates of [Marone and Sabety \(2022\)](#). We construct a distribution of underlying population demographics to match the non-elderly US population based on data from the American Community Survey. Each consumer faces a log-normal distribution of health states. Consumers are fully described by their initial income w_0 , their health state distribution F , their risk aversion parameter ψ , and their “over-utilization parameter” ω .⁴⁶ Appendix B provides additional details about the construction of the population. Appendix Table B.1 provides summary statistics. A parallel analysis using parameter estimates from [Einav et al. \(2013\)](#) is presented in Appendix C.

III.B Results

Table 2 presents the decomposition of welfare V generated by a full insurance contract into the five terms described in Equation 11. Panel A presents the results for utility specification (1), while Panel B presents the results for specification (2). In both specifications, all consumers prefer full insurance to no insurance, so V is positive for every consumer in both cases. The population’s underlying distribution of consumer types is the same in each panel; all that has changed is the parameterization of utility (through b).

Each of the five columns under “Decomposition of welfare” reports the proportion of total welfare V represented by that term, on average across consumers. For example, an entry under Ψ_N represents the average value of $\frac{\Psi_N}{V}$ across consumers. The five columns thus sum to one in each row. In both specifications, the average proportion of welfare attributable to non-moral hazard spending is 94 percent. For the average consumer, the risk protection value of insurance *absent* moral hazard thus accounts for the overwhelming majority of the welfare generated by insurance. The contribution of moral hazard is captured by the remaining four terms. The risk protection value of moral hazard is represented by $\Psi_{eff} + \Psi_{over} + \Delta EV_{eff}$ together. Under both specifications, it represents a non-trivial proportion of welfare. Finally, the negative contribution to welfare from the social cost of over-utilization is reported under ΔEV_{over} . It is on average equal to 5 percent of welfare under specification (1), and 3 percent under specification (2). These amounts represent the social cost of over-utilization induced by full insurance.

Within each panel, results are reported for the population as a whole, as well as in subgroups.

⁴⁶To maintain comparability across the specifications, we set $\tilde{\omega} = \omega/\mathbb{E}_l[l]$, such that a consumer’s expected amount of over-utilization under full insurance is the same across the two specifications. The only difference across specifications is thus whether over-utilization is spread evenly across health states or concentrated on sicker states.

Table 2. Decomposition of Welfare

Consumer group	Decomposition of welfare (Under full insurance)					Relative change in welfare under		
	Ψ_N	Ψ_{eff}	Ψ_{over}	ΔEV_{eff}	ΔEV_{over}	Efficient level of coverage	Eliminate over-utiliz.	Eliminate all moral hazard
Panel A. Spec (1): Homebodies								
All	0.94	0.10	–	<0.01	-0.05	0.02	0.05	-0.06
<i>Prob. constrained</i> [†]								
Low	1.05	<0.01	–	<0.01	-0.05	0.02	0.05	0.05
Medium	0.69	0.33	–	<0.01	-0.02	<0.01	0.02	-0.31
High	0.49	0.49	–	0.05	-0.02	<0.01	0.02	-0.51
<i>Tertile of ω</i>								
Lowest	0.89	0.12	–	<0.01	-0.02	<0.01	0.02	-0.11
Middle	0.92	0.11	–	<0.01	-0.03	<0.01	0.03	-0.08
Highest	1.02	0.06	–	<0.01	-0.09	0.04	0.09	0.02
Panel B. Spec (2): Worriers								
All	0.94	0.06	0.02	<0.01	-0.03	<0.01	<0.01	-0.06
<i>Prob. constrained</i> [†]								
Low	1.01	<0.01	0.02	<0.01	-0.03	<0.01	<0.01	<0.01
Medium	0.82	0.17	0.03	<0.01	-0.02	<0.01	-0.01	-0.18
High	0.62	0.36	<0.01	0.03	-0.03	<0.01	0.02	-0.38
<i>Tertile of ω</i>								
Lowest	0.92	0.07	0.02	<0.01	-0.02	<0.01	-0.00	-0.08
Middle	0.93	0.07	0.03	<0.01	-0.03	<0.01	<0.01	-0.07
Highest	0.98	0.04	0.03	<0.01	-0.05	<0.01	0.02	-0.02

Notes: This table shows the decomposition of welfare generated by full insurance into the five terms described in Equation 11, for each of the parameterizations of consumer utility described in Table 1. Each of the five columns under “Decomposition of welfare” represents the proportion of total welfare V represented by that term, on average across consumers. The columns thus sum to one within each row. The final three columns report the relative changes in welfare that would occur in three scenarios: (a) if consumers were enrolled in their efficient level of coverage, rather than full insurance; (b) if over-utilization were prevented, and (c) if all moral hazard were prevented. A relative change of 0.02 means a 2% increase in welfare. [†]Low probability of being constrained means a <1% change (true for 78 percent of consumers); medium means a 1%–5% chance (11 percent of consumers); high means $\geq 5\%$ chance (11 percent of consumers).

The first grouping is based on the likelihood a consumer experiences a health outcome in which their efficient level of utilization exceeds their initial income. In other words, it is based on the probability that the consumer would be *constrained* in their consumption of healthcare if they did not have health insurance. Even though all consumers can by definition afford their efficient level of utilization *in expectation*, not all could do so in every state of the world without insurance. Those with a low probability of being constrained (<1 percent chance) are those who have a very high income relative to their potential health outcomes. Those with a medium (1–5 percent) or high (≥ 5 percent) chance are commensurately more likely to

experience these “tail events” in which they are constrained, and thus more likely to experience inefficient under-utilization absent insurance. Consumers with a higher probability of being constrained are exactly those for whom insurance will thus provide greater risk protection through the terms associated with the “efficient increase in utilization”: Ψ_{eff} and ΔEV_{eff} . Indeed, among those with a high probability of being constrained, $\Psi_{eff} + \Delta EV_{eff}$ account for on average 54 percent of welfare in specification (1) and 39 percent in specification (2).⁴⁷

The second grouping of consumers is into tertiles of their propensity for over-utilization, measured by the parameter ω . Consumers with higher ω have a more elastic demand for healthcare and thus higher over-utilization with insurance. In specification (1), the social cost of over-utilization (ΔEV_{over}) represents 2 percent of welfare among the third of consumers with the lowest ω , while it represents 9 percent of welfare for those with the highest. Note, however, that consumers with a high propensity for over-utilization are also those for whom over-utilization is most valuable.⁴⁸ And for *Worriers*, if over-utilization is more prevalent in sick states—as in specification (2)—this over-utilization can provide valuable additional risk protection.⁴⁹ In specification (2), Ψ_{over} is thus naturally higher for higher- ω consumers. Even so, over-utilization contributes less to the value of risk protection than the efficient increase in utilization.

The final three columns present the relative change in welfare that would arise in three scenarios: (a) if consumers were enrolled in their efficient coverage level (second-best contract), rather than in full insurance;⁵⁰ (b) if over-utilization were prevented; and (c) if all moral hazard were prevented. First, enrolling consumers in their efficient (second-best) contract rather than full insurance will, of course, weakly increase welfare. Our estimates suggest that it will not,

⁴⁷Among consumers with a high (≥ 5 percent) probability of being constrained, the average probability of being constrained is 13 percent, and the median is 6 percent. For these consumers, the efficient increase in utilization accounts for 31 percent of moral hazard spending on average, while over-utilization accounts for the remaining 69 percent. Across all consumers, the efficient increase in utilization accounts for 4 percent of moral hazard spending on average. For comparison, [Nyman et al. \(2018\)](#) find that the efficient increase in utilization accounts for 13 percent of moral hazard spending.

⁴⁸The fact that over-utilization is more valuable to consumers with a higher propensity for over-utilization is a direct consequence of our parameterization of $b(m; l)$. The quadratic specification mechanically links the amount of over-utilization to its value. In principle, the value of over-utilization may be large for small amounts of over-utilization ($b(m)$ narrow and steep), or vice-versa ($b(m)$ wide and flat). To our knowledge, there has been no direct empirical investigation of these patterns.

⁴⁹The “additive” model of over-utilization in specification (1) has the feature that over-utilization is equally valuable in all states of the world. It thus results in exactly zero risk protection value from over-utilization regardless of the consumer’s type.

⁵⁰A consumer’s efficient level of coverage is defined as the welfare-maximizing contract from among a pre-specified set ([Zeckhauser, 1970](#); [Marone and Sabety, 2022](#)). The set of contracts we use is shown in Appendix Figure B.1.

however, have a large impact quantitatively, echoing findings reported in [Marone and Sabety \(2022\)](#) and [Ho and Lee \(2023\)](#). The (small) magnitude of welfare effects in this first scenario provides a benchmark for thinking about the relative importance of others. In terms of the coverage levels themselves, the average efficient coverage level in the population provides an actuarial value of 79 percent. As is usual, consumers who have higher efficient coverage levels tend to face more uncertainty about their health status (variance of F) and have higher risk aversion parameters (ψ). They do not, however, necessarily have a smaller increase in healthcare utilization in response to insurance coverage, in particular among consumers with lower incomes. Appendix Figure [A.1](#) shows that among consumers in the bottom third of the income distribution, the relationship between efficient coverage level and expected moral hazard spending is U-shaped. At relatively low levels of moral hazard spending, the relationship with efficient coverage level is negative (as expected). But for consumers with high levels of moral hazard spending, the relationship becomes positive. A primary indication for why moral hazard spending may be quite high is if the consumer were under-consuming healthcare absent insurance, or if over-utilization were highly valuable to her. In both cases, moral hazard spending will tend to be worth its cost, increasing the efficient coverage level.⁵¹

The second scenario eliminates over-utilization of healthcare while still providing full insurance. This scenario coincides with the first-best outcome for *Homebodies*.⁵² In effect, this counterfactual simply eliminates the contribution to welfare represented by Ψ_{over} and ΔEV_{over} . In specification (1), welfare increases by an average of 5 percent. In specification (2), eliminating over-utilization is not unequivocally good, as it was also providing some valuable risk protection. On balance, the net effect on welfare is near zero (0.8 percent average increase). That said, there is heterogeneity across consumers. For 26 percent of consumers, eliminating over-utilization actually *reduces* welfare. Transferring resources towards sick states of the world is so valuable for these consumers that the inefficient transfer mechanism of over-utilization is still better than no transfer. For the other 74 percent of consumers, the transfer motive is not so strong to overcome the standard expected cost of over-utilization, and welfare

⁵¹These patterns have been absent in prior papers for two reasons. First, prior formulations have not enforced that non-health consumption be positive. In effect, this allows the consumer to circumvent her budget constraint and achieve her efficient level of healthcare utilization regardless of whether she can afford it, leaving no possibility for inefficient underutilization (see, e.g., [Einav et al. \(2013\)](#), Equation 2, page 184; [Azevedo and Gottlieb \(2017\)](#), page 73; [Marone and Sabety \(2022\)](#), page 309). The second is that most prior papers assume the consumer is a *Homebody*, meaning that over-utilization is unambiguously bad.

⁵²It does not coincide with the first-best outcome for *Worriers* because recall that *Worriers* would transfer additional non-health consumption towards sicker states under the first-best contract, something that a “full insurance” health insurance contract—which simply lowers the cost of healthcare to zero—cannot provide.

increases when over-utilization is prevented.⁵³

Finally, the third scenario eliminates moral hazard entirely, including both over-utilization as well as the efficient increase in spending. Full insurance is still provided, but healthcare utilization is held fixed at the same level that prevailed absent insurance. Here, substantial risk protection value is lost, and in both specifications, welfare is decreased. Naturally, preventing consumers from changing their behavior in response to insurance is especially costly for those with a substantial degree of under-utilization absent insurance (i.e., those with a high probability of being constrained). In specification (1), welfare for the average consumer would be higher by 8 percent under their efficient contract rather than under a full insurance contract that prevented them from changing their behavior in response to insurance. For consumers with a high probability of being constrained, welfare would be a striking 51 percent higher. At the same time as the efficient contract delivers higher average welfare, the costs to the insurer for providing it are on average 15 percent lower. A lower fair premium thus yields higher risk protection, implying that allowing moral hazard represents a more efficient way to produce risk protection.

We draw two main conclusions from this analysis. First, it is clear that the most important driver of welfare from health insurance in this population is the value of risk protection that arises even absent moral hazard (Ψ_N). The ability to transfer non-health consumption across states, holding behavior fixed, is almost all that matters for most consumers. The welfare effects of moral hazard—both positive and negative—are an order of magnitude smaller than the welfare effect of gaining insurance at all. Second, the welfare effects of moral hazard are nonetheless economically significant when measured relative to other counterfactual scenarios of interest. In specification (1), eliminating all ex-ante asymmetric information (enabling the assignment of all consumers to their efficient coverage level) would increase welfare by less than half as much as eliminating all ex-post asymmetric information (enabling all over-utilization to be prevented). Moreover, eliminating *both* ex ante and ex post asymmetric information (enabling the first-best) would produce a welfare benefit about equivalent in magnitude to the welfare harm produced by a policy that eliminated moral hazard entirely. The welfare effects of limiting moral hazard, either in part or in whole, are thus very much on par with those of other policy-relevant objects.

An important caveat to both of these conclusions is that there is substantial heterogeneity across consumers. Consumers with lower incomes, for example, derive a substantially greater

⁵³In specification (2), welfare would be 3 percent higher under the first best (when income—and thus both health *and* non-health consumption—can be transferred arbitrarily across states).

value from the ability to use a greater amount of healthcare with insurance, rather than just to increase their non-health consumption in sick states. Among those in the lowest third of the income distribution, moral hazard accounts for 30 percent of the value of risk protection provided by full insurance (27 percent in specification 2). In contrast, among those in the highest third of income, moral accounts for none of the value of risk protection (2 percent in specification 2).

These quantitative results are of course specific to the parameterizations of utility and the estimates of structural primitives we have used. The distribution of consumer types (w_0, ψ, F, ω) as well as the functional form for b (specifically, how its level and curvature change across health states) all play an important role. Our empirical analysis can thus be viewed as highlighting that existing estimates of the primitives governing demand for healthcare and health insurance fall within the space of theoretical possibilities where moral hazard plays a central role in the risk protection value of insurance. We view further investigation of these primitives—especially the functional form of b , which so far has been simply assumed—to be an important direction for future work.

IV Discussion

Our analysis offers a new perspective on a number of healthcare policy issues. First, it challenges the conventional view that in designing public health insurance, greater cost-sharing should be imposed on services for which demand is more responsive to coverage. While this logic is rooted in the goal of discouraging over-utilization, it overlooks the fact that part of the increase in utilization caused by insurance is efficient. If individuals are resource constrained in some health states absent insurance, efforts to contain over-utilization may unwittingly exacerbate under-utilization, ultimately doing more harm than good. Our analysis suggests that healthcare services which are expensive and rarely needed are more likely to be under-utilized than over-utilized, suggesting caution in placing any substantial cost-sharing on their use.

A related implication is the importance of the consumer’s income level in relation to their efficient levels of healthcare utilization. As discussed above, if the consumer is likely to not be able to afford her efficient level of utilization in many health states, insurance sufficiently generous to allow her to do so will avoid welfare losses from under-utilization. As long as bad health outcomes are sufficiently unlikely, lower-income individuals are therefore more

likely to have higher efficient levels of insurance coverage. Using income as a tag on which to differentiate the terms of public insurance coverage is consistent with this idea. Providing more generous insurance to lower-income individuals thus need not be predicated on distributional concerns alone, but can also be grounded in an efficient argument.

Our model also provides a way to interpret some features of market-based health insurance, for example prior authorization, case management, referral requirements, utilization review, and even vertical integration. All of these tools can be considered efforts by the insurer to verify the realized state of the world, and thus to enact the efficient level of utilization while preventing over-utilization. Managed care plans (HMOs) can be viewed as contracts under which consumers willingly grants the insurer even more tools with which to monitor and verify the realized health state, thus further reducing the extent of over-utilization (and the premium along with it). Different consumers will perceive the costs of this health-state verification exercise differently. One consumer may prefer to avoid it by paying the premium increase necessary to fund her own over-utilization (for example, under a PPO plan). Another may happily submit to hassle costs in exchange for a lower premium.

Finally, our model makes plain that unlike in the case of adverse selection, the market failure that leads to moral hazard cannot be remedied by any powers held by a benevolent regulator. In the face of ex-ante asymmetric information, the government can combat adverse selection by using the power of taxation to compel otherwise-unwilling consumers to take-up insurance. But this power is not of use in the face of ex-post asymmetric information. Though the government can imitate tools used in private market, we see no strong reason to think the government would have a technological advantage over the market in eliciting private information from consumers about realized health status. This observation might be placed in the “pro” column of an evaluation of privately-administered (versus publicly-administered) tax-funded health insurance.

V Conclusion

This paper provides a novel treatment of a standard health insurance model with ex-post moral hazard. The model is substantially general, and likely applies to a number of other insurance (or redistribution) settings. Given the presence of multiple consumption goods (a necessary condition for ex-post moral hazard), we recast the consumer’s fundamental *risk* to be over her marginal indirect utility of income, rather than over income or consumption directly. This

approach allows a number of novel transparencies, including (i) a formal characterization of the first-best contract, (ii) a new vocabulary for describing the features of consumer preferences relevant for the value of insurance, and (iii) a formal decomposition of the welfare generated by insurance into terms that isolate the impact of moral hazard.

We find that there are two important channels through which a consumer’s ability to change her behavior in response to insurance can contribute positively to social welfare. The first is through the ability to optimally exploit changes in real income brought about by insurance. Cheaper healthcare increases real income in sick states, and the ability to re-optimize behavior in response to this increase can be highly valuable. The second is the ability to shift more resources towards sick states than would otherwise be possible. When the state is non-contractible, health insurance provides no way for consumers to shift more than a certain amount of non-health consumption towards sick states. Even so, if healthcare is at all substitutable with non-health consumption, the ability to over-utilize healthcare in sick states can serve as an imperfect substitute for additional non-health consumption. In both instances, accounting for the implications of moral hazard from an *ex ante* perspective is central to our results. Quantitatively, our empirical analysis reveals that standard estimates in the literature imply a non-trivial role for these forces.

We conclude from our analysis that one should rather not think of insurance as offering a tradeoff between risk protection and moral hazard, but rather between risk protection ($\Psi_N + \Psi_{eff} + \Psi_{over} + \Delta EV_{eff}$) and the social cost of over-utilization (ΔEV_{over}) (see Equation 11). As we discuss, signing these terms relies critically on assumptions about the consumer’s preferences for consumption (of both health and non-health goods) in different health states. The value of risk protection provided by insurance will only be generically positive so long as the consumer is a *Homebody* or a *Worrier*. Over-utilization will only provide risk protection so long as the consumer is a *Worrier*.

There are a number of important issues with which this paper does not engage. First, we limit attention to parameterizations of utility in which the consumer has constant absolute risk aversion in non-health consumption (Assumption 1). This choice enables the decomposition of welfare provided in Equation 11, but means that income affects the demand for healthcare only through the budget constraint. At the expense of transparency, a more general analysis would allow the consumer’s interior optimal amount of healthcare utilization to also vary with income. Second, our model features only two goods—healthcare goods and non-healthcare goods—despite the fact that the degree of substitutability between them is of central importance to one mechanism through which moral hazard may be valuable. A richer analysis

could split non-healthcare goods into two: a set of goods that are not at all substitutable with healthcare (e.g., housing), and a set of goods that are substitutable (e.g., non-covered healthcare services). Lastly, healthcare utilization decisions in our model are assumed to be a single, unidimensional choice over a continuous support. Reality is lumpier. These decisions are also rarely made by patients alone, but rather arise from a joint decision-making process between patients and physicians, shaped by norms, incentives, and asymmetric information. Our analysis abstracts from these issues. As such, it is only an imperfect approximation of a fundamentally discrete and institutionally mediated process. While these simplifications facilitate tractability, they also highlight the need for further research.

References

- Arrow, K. J. (1963). Uncertainty and the welfare economics of medical care. *The American Economic Review* 53(5), 941–973.
- Azevedo, E. M. and D. Gottlieb (2017). Perfect Competition in Markets With Adverse Selection. *Econometrica* 85(1), 67–105.
- Baicker, K., S. Mullainathan, and J. Schwartzstein (2015). Behavioral hazard in health insurance. *The Quarterly Journal of Economics* 130(4), 1623–1667.
- Bajari, P., H. Hong, A. Khwaja, and C. Dalton (2014). Moral hazard, adverse selection, and health expenditures: A semiparametric analysis. *The RAND Journal of Economics* 45(4), 747–763.
- Berry, S. T. and P. A. Haile (2021, September). Foundations of demand estimation. Technical Report 29305, National Bureau of Economic Research.
- Blundell, R., M. Borella, J. Commault, and M. De Nardi (2024, February). Old Age Risks, Consumption, and Insurance. *American Economic Review* 114(2), 575–613.
- Cardon, J. H. and I. Hendel (2001). Asymmetric information in health insurance: Evidence from the national medical expenditure survey. *The RAND Journal of Economics* 32(3), 408–427.
- Cawley, J. and C. J. Ruhm (2011). The economics of risky health behaviors. In M. V. Pauly, T. G. McGuire, and P. P. Barros (Eds.), *Handbook of Health Economics*, Volume 2, Chapter 3, pp. 95–199. Elsevier.
- Chade, H., V. Marone, A. Starc, and J. Swinkels (2023, December). Multidimensional screening and menu design in health insurance markets. Working Paper.
- Chandra, A., E. Flack, and Z. Obermeyer (2024, November). The health costs of cost sharing. *The Quarterly Journal of Economics* 139(4), 2037–2082.
- Chetty, R. (2008). Moral hazard versus liquidity and optimal unemployment insurance. *Journal of Political Economy* 116(2), 173–234.
- Chetty, R. and A. Finkelstein (2013). Social insurance: Connecting theory to data. In A. J. Auerbach and M. Feldstein (Eds.), *Handbook of Public Economics*, Volume 5 of *Handbook of Public Economics*, pp. 111–193. Elsevier.
- Cutler, D. M. and R. J. Zeckhauser (2000, January). The Anatomy of Health Insurance*. In A. J. Culyer and J. P. Newhouse (Eds.), *Handbook of Health Economics*, Volume 1 of *Handbook of Health Economics*, pp. 563–643. Elsevier.

- de Meza, D. (1983). Health insurance and the demand for medical care. *Journal of Health Economics* 2(1), 47–54.
- Deshpande, M. and L. M. Lockwood (2022, July). Beyond health: Nonhealth risk and the value of disability insurance. *Econometrica* 90(4), 1781–1810.
- Einav, L. and A. Finkelstein (2018, 05). Moral Hazard in Health Insurance: What We Know and How We Know It. *Journal of the European Economic Association* 16(4), 957–982.
- Einav, L., A. Finkelstein, S. P. Ryan, P. Schrimpf, and M. R. Cullen (2013). Selection on moral hazard in health insurance. *American Economic Review* 103(1), 178–219.
- Einav, L., A. N. Finkelstein, and M. Cullen (2010). Estimating Welfare in Insurance Markets Using Variation in Prices. *The Quarterly Journal of Economics* CXV(3).
- Ericson, K. M., J. G. Jaspersen, and J. R. Sydnor (2025, April). Moral hazard under liquidity constraints. Working Paper 33648, National Bureau of Economic Research.
- Evans, W. N. and W. K. Viscusi (1991). Estimation of state-dependent utility functions using survey data. *The Review of Economics and Statistics* 73(1), 94–104.
- Feldstein, M. (1973). The welfare loss of excess health insurance. *Journal of Political Economy* 81(2), 251–80.
- Finkelstein, A., N. Hendren, and E. F. P. Luttmer (2019). The value of medicaid: Interpreting results from the oregon health insurance experiment. *Journal of Political Economy* 127(6), 2836–2874.
- Finkelstein, A., E. F. P. Luttmer, and M. J. Notowidigdo (2013). What good is wealth without health? the effect of health on the marginal utility of consumption. *Journal of the European Economic Association* 11, 221–258.
- Frick, K. D. and M. E. Chernew (2009). Beneficial moral hazard and the theory of the second best. *Inquiry: A Journal of Medical Care Organization, Provision and Financing* 46(2), 229–240.
- Friedman, M. (1953). The 'welfare' effects of an income tax and an excise tax. In *Essays in Positive Economics*, pp. 100–113. Chicago: University of Chicago Press.
- Grossman, M. (1972). On the concept of health capital and the demand for health. *Journal of Political Economy* 80(2), 223–255.
- Harberger, A. C. (1964). Taxation, resource allocation, and welfare. In *The Role of Direct and Indirect Taxes in the Federal Revenue System*. Princeton: Princeton University Press.
- Hart, O. and J. Moore (1999). Foundations of incomplete contracts. *The Review of Economic Studies* 66(1), 115–138.
- Hendren, N. (2021). Measuring ex ante welfare in insurance markets. *Review of Economic Studies* 88(3), 1193–1223.
- Ho, K. and R. S. Lee (2023). Health insurance menu design for large employers. *RAND Journal of Economics* 54(4), 598–637.
- Kesternich, I., A. Romahn, J. Van Biesebroeck, and M. Van Damme (2025, May). Cash or care? insights from the german long-term care system. CESifo Working Paper 11875, CESifo, Munich. Available at SSRN: <https://ssrn.com/abstract=5254204>.
- Landais, C. (2015). Assessing the welfare effects of unemployment benefits using the regression kink design. *American Economic Journal: Economic Policy* 7(4), 243–278.
- Lieber, E. M. J. and L. M. Lockwood (2019). Targeting with in-kind transfers: Evidence from medicaid home care. *American Economic Review* 109(4), 1461–1485.
- Maestas, N., K. J. Mullen, and A. Strand (2013, August). Does disability insurance receipt discourage work? using examiner assignment to estimate causal effects of ssdi receipt. *American Economic Review* 103(5), 1797–1829.

- Manning, W. G., J. P. Newhouse, N. Duan, E. B. Keeler, and A. Leibowitz (1987). Health insurance and the demand for medical care: Evidence from a randomized experiment. *The American Economic Review* 77(3), 251–277.
- Marone, V. R. and I. Huitfelt (2025). Health insurance for healthcare amenities.
- Marone, V. R. and A. Sabety (2022, January). When should there be vertical choice in health insurance markets? *American Economic Review* 112(1), 304–42.
- McGuire, T. (2011, 12). Demand for health insurance. *Handbook of Health Economics* 2, 317–396.
- Mirrlees, J. A. (1971). An exploration in the theory of optimum income taxation. *Review of Economic Studies* 38(2), 175–208.
- Mommaerts, C. (2025). Long-term care insurance and the family. *Journal of Political Economy* 133(1), 1–52.
- Nyman, J. A. (1999a, December). The economics of moral hazard revisited. *Journal of Health Economics* 18(6), 811–824.
- Nyman, J. A. (1999b). The value of health insurance: the access motive. *Journal of Health Economics* 18(2), 141–152.
- Nyman, J. A. (2002). *The Theory of Demand for Health Insurance*. Stanford, CA: Stanford University Press.
- Nyman, J. A. (2004). Is ‘moral hazard’ inefficient? the policy implications of a new theory. *Health Affairs* 23(5), 194–199.
- Nyman, J. A. (2007). American health policy: Cracks in the foundation. *Journal of Health Politics, Policy and Law* 32(5), 759–783.
- Nyman, J. A. (2008). Health Insurance Theory: The Case of the Missing Welfare Gain. *The European Journal of Health Economics* 9(4), 369–380. Publisher: Springer.
- Nyman, J. A., C. Koc, B. E. Dowd, E. McCreedy, and H. M. Trenz (2018). Decomposition of moral hazard. *Journal of Health Economics* 57, 168–178.
- Oregon Public Employees Retirement System (2024, May). PERS by the Numbers. <https://www.oregon.gov/pers/Documents/General-Information/PERS-by-the-Numbers.pdf>. Accessed: 2025-07-17.
- Pauly, M. V. (1968). The economics of moral hazard: Comment. *American Economic Review* 58(3), 531–537.
- Pauly, M. V. (1969, Winter). A measure of the welfare cost of health insurance. *Health Serv Res* 4(4), 281–292.
- Shimer, R. and I. Werning (2008). Liquidity and insurance for the unemployed. *American Economic Review* 98(5), 1922–1942.
- Sloan, F. A., W. K. Viscusi, H. W. Chesson, C. J. Conover, and K. Whetten-Goldstein (1998). Alternative approaches to valuing intangible health losses: The evidence for multiple sclerosis. *Journal of Health Economics* 17(4), 475–497.
- Spherical Insights & Consulting (2025, March). Global critical illness insurance market size to worth usd 431.9 billion by 2033: Market insight report. Accessed: 2025-07-10.
- U.S. Census Bureau (2023). American community survey 1-year estimates. <https://www.census.gov/programs-surveys/acs>. Accessed on [insert date of access].
- Viscusi, W. K. and W. N. Evans (1990). Utility functions that depend on health status: Estimates and economic implications. *American Economic Review* 80(3), 353–374.
- Wilson, B. and J. E. Thomas (2005). The indemnity principle: From a financial to a function paradigm. *The Journal of Risk Management and Insurance* 10, 30.

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

Appendix A Theoretical Appendix

A.1 Derivation of Welfare

Consider the model presented in Section II. The welfare generated from an insurance contract $x \in \mathcal{X}$ is equal to the difference between the consumer's willingness to pay for the contract and the insurer's expected cost of providing the contract. The consumer's willingness to pay $WTP(x)$ is given by the premium \tilde{p} at which her expected utility with insurance equals her expected utility absent insurance:

$$U(x, \tilde{p}) = U(x_0, 0).$$

Obtaining a closed-form solution for \tilde{p} will be possible if the value generated by insurance is quasilinear in the premium at which it is supplied. The premium will be additively separable in this way so long as it does not directly affect the consumer's ex-ante value of risk protection derived from insurance nor her ex-post optimal choice of healthcare utilization. Assumption 1 in Section II guarantees that changes in the premium will not affect the curvature of the consumer's utility function, nor her *interior* optimal choice of healthcare utilization. It does not, however, guarantee that her *optimal* choice of healthcare utilization will be invariant to the contract premium. It is easy to see a trivial example of why this is the case by considering a small premium versus an arbitrarily high premium. At a small premium, the consumer will have some remaining income to purchase some y and m in her realized health state. At a premium that exceeds her initial income, she will have no money left to purchase any y or m , and thus, the premium has affected her demand for healthcare utilization. The purpose of the following assumption is to rule out this type of possibility, where the premium is so high relative to initial income that it moves the consumer from her interior optimal choice of healthcare utilization to a corner solution. The assumption essentially enforces that the distance between $WTP(x)$ and $\bar{p}(x)$ is not too large relative to initial income, and thus, that varying the premium within this range will not affect the consumer's demand for healthcare utilization.

ASSUMPTION 2 (No Income Effects Between Relevant Premia). *The consumer's optimal choice of healthcare utilization under the actuarially fair premium is equal to that under a premium equal to her full willingness to pay: $m^*(l, x) \equiv m^*(l, x, \bar{p}(x)) = m^*(l, x, WTP(x)) \forall l \in \mathcal{L}, \forall x \in \mathcal{X}$.*

Now, using the parameterization of utility implied by Assumptions 1 and the condition defined in Assumption 2, we can rewrite the equation above as:

$$\mathbb{E}_l [-\exp(-\psi(z^*(l, x) + \bar{p}(x) - \tilde{p}))] = \mathbb{E}_l [-\exp(-\psi z^*(l, x_0))]$$

where $z^*(l, x) = w_0 - \bar{p}(x) - c(m^*(l, x); x) + b(m^*(l, x); l)$ as defined by Equation 2, and where $m^*(l, x) = m^*(l, x, \bar{p}(x)) = m^*(l, x, \tilde{p})$ by Assumption 2. Solving for \tilde{p} then yields a closed form expression for willingness to pay for contract x relative to the null contract x_0 :

$$\begin{aligned} WTP(x) &\equiv \tilde{p} = \bar{p}(x) + \psi^{-1} \cdot [\log(\mathbb{E}_l [\exp(-\psi z^*(l, x_0))]) - \log(\mathbb{E}_l [\exp(-\psi z^*(l, x))])] \\ &= \bar{p}(x) + CE(x) - CE(x_0), \end{aligned}$$

where $CE(x) = \tilde{u}^{-1}(\mathbb{E}_l [\tilde{u}(z^*(l, x))])$ as defined in Equation 3. The insurer's expected cost of providing contract x is by definition equal to the actuarially fair premium $\bar{p}(x)$. As a result, the welfare derived from insurance is equal to $V(x) \equiv WTP(x) - \bar{p}(x) = CE(x) - CE(x_0)$.

A.2 First-Best Allocation

Consider the first-best allocation of income to health states, as defined in Equation 6 and repeated here:

$$(w^{FB}(l)) = \arg \max_{(w(l))} \mathbb{E}_l [v(w(l); l)] \quad s.t. \quad \mathbb{E}_l [w(l)] \leq w_0,$$

where $v(w; l) = u(y^*(w; l), m^*(w; l); l)$ denotes the consumer's baseline indirect utility function, with optimal bundles given by $m^*(w; l) = \arg \max_{0 \leq m \leq w} u(w - m, m; l)$ and $y^*(w; l) = w - m^*(w; l)$. As in the main text, we refer to the first-best levels of health and non-health consumption using $m^{FB}(l) = m^*(w^{FB}(l); l)$ and $y^{FB}(l) = y^*(w^{FB}(l); l)$. Our goal in this section is to provide conditions under which we can characterize the first-best level of healthcare utilization, $m^{FB}(l)$.

We take as given that utility is parameterized according to Assumption 1, such that

$$u(y, m; l) = \tilde{u}(y + b(m; l)),$$

where $\tilde{u}(\cdot)$ is a CARA utility function, and where $b(m; l)$ is strictly concave and features a (state-dependent) bliss point in m . Plugging into Equation 5, the consumer's optimal

healthcare utilization absent insurance is given by:

$$m^N(l) = m^*(l, x_0) = \arg \max_{m \geq 0} \tilde{u}(w_0 - m + b(m; l)) \quad s.t. \quad m \leq w_0, \quad (12)$$

where the budget constraint will bind with equality ($y = w_0 - m$) since utility is strictly increasing in y .

Three ancillary definitions will now be useful. The first defines the solution to a relaxed version of the consumer’s healthcare utilization problem, under which she may set $m > w_0$, or in other words, she may have negative non-health consumption.⁵⁴

DEFINITION 2 (Unconstrained Optimal Healthcare Spending). *Let $m^U(l)$ denote the solution to the consumer’s relaxed healthcare utilization problem absent insurance: $m^U(l) \equiv \arg \max_{m \geq 0} \tilde{u}(w_0 - m + b(m; l)) = \arg \max_{m \geq 0} (b(m; l) - m)$.*

Notice that since $\tilde{u}(\cdot)$ is strictly increasing, $m^U(l)$ simply solves the first order condition $b_m(m^U(l); l) = 1$.⁵⁵ In other words, in a given health state, $m^U(l)$ represents the level of healthcare utilization at which its marginal social benefit (b_m) equals its marginal social cost (1).

The second definition is of the “least needy” health state: Suppose that in each state, the consumer were provided her unconstrained optimal healthcare spending $m^U(l)$, such that her healthcare needs were sufficiently provided for, but she were given zero non-health consumption. In which health state (or set of health states) would the first dollar of non-health consumption be the *least* valuable?⁵⁶

DEFINITION 3 (“Least Needy” Health State). *Let l_0 denote the health state (or one of the health states) with the smallest marginal utility of y when evaluated at $y = 0$ and $m = m^U(l)$: $l_0 \in \arg \min_{\mathcal{L}} \tilde{u}'(0 + b(m^U(l); l)) = \arg \max_{\mathcal{L}} b(m^U(l); l)$*

⁵⁴This relaxed version of the problem corresponds to the version used in [Einav et al. \(2013\)](#) (c.f. Equation 2, page 184), [Azevedo and Gottlieb \(2017\)](#) (c.f. page 73), [Marone and Sabety \(2022\)](#) (c.f. page 309), and [Ho and Lee \(2023\)](#) (c.f. Equation 5, page 614). In contrast, [Cardon and Hendel \(2001\)](#) (c.f. Equation 1, page 414) and [Bajari et al. \(2014\)](#) (c.f. Equation 2.2, page 750) specify the “non-relaxed” version of the problem (as in Equation 1 above).

⁵⁵Or, if the marginal benefit of healthcare utilization falls below one at $m = 0$, then $m^U(l) = 0$.

⁵⁶Whether the health state in which incremental non-health consumption is “least-needed” (least valuable) corresponds to the healthiest or the sickest state is at the heart of whether health insurance provides positive risk protection and is socially valuable. This issue is discussed in greater detail in Section II.C in the main text. Indeed, whether the “least needy” health state corresponds to the healthiest or sickest state (or whether all health states are “equally needy”) maps exactly to the three types of consumers defined in Definition 1.

Since \tilde{u} is strictly increasing and concave, the least needy health state corresponds to the state(s) in which $b(m^U(l); l)$ is highest.

The final definition specifies the minimum amount of income the consumer would need in each health state in order to equate her marginal utility of income across states *while still consuming* her unconstrained optimal healthcare utilization. In other words, what amount of income would be required in each health state in order to afford both her unconstrained optimal healthcare spending $m^U(l)$ *and* just enough non-health consumption y that she could equate her marginal utility of income in that state with that in the least-needy health state when evaluated at $y = 0$? Or said formally, what state-contingent income level $w(l)$ solves: $\tilde{u}'(w(l) - m^U(l) + b(m^U(l); l)) = \tilde{u}'(b(m^U(l_0); l_0))$?

DEFINITION 4 (Minimum Income to Equate Marginal Utility). Let $w^{\min}(l) \equiv m^U(l) + b(m^U(l_0); l_0) - b(m^U(l); l)$.

Note that by definition of the least-needy health state, the minimum income amount $w^{\min}(l)$ weakly exceeds the amount required to purchase the unconstrained optimal healthcare spending amount in all health states, i.e., $b(m^U(l_0); l_0) - b(m^U(l); l) \geq 0 \forall l$.

Intuitively, if the consumer's level of initial income is too low, she may not only be unable to afford her unconstrained optimal healthcare utilization in some health states, she may also be unable to afford it *in expectation*. Moreover, if the marginal utility of income (evaluated at $m^U(l)$) is extremely different across states, she may wish to allocate her scarce income in such a way that she puts herself at a corner solution: allocating zero income to some state in which v_w was very low, rather than equating her marginal utilities across all states. In order to analytically characterize socially efficient utilization, we need to rule out both of these possibilities. To that end, we introduce an assumption on the consumer's initial income with respect to her preferences.

ASSUMPTION 3 (Sufficiently High Initial Income). Initial income is sufficiently high that the vector of minimum income amounts is affordable in expectation: $w_0 \geq \mathbb{E}_l[w^{\min}(l)]$.

Assumption 3 provides a sufficient condition under which we can formally characterize the first best allocation. In particular, it implies that in every health state, the socially efficient level of healthcare utilization coincides with $m^U(l)$, the level at which its marginal social benefit equals its marginal social cost.

LEMMA 1 (First-Best Allocation). Under Assumptions 1 and 3, the first-best income

vector is $w^{FB}(l) = w^{min}(l) + w_0 - \mathbb{E}_l[w^{min}(l)]$, first best healthcare spending is $m^{FB}(l) = m^U(l)$, and first best non-health consumption is $y^{FB}(l) = w^{FB}(l) - m^{FB}(l)$.

Proof. We begin by characterizing the consumer's indirect utility function $v(w; l)$ and marginal indirect utility of income $v_w(w; l)$, and then proceed to characterizing the first-best income vector $w^{FB}(l)$. Given arbitrary income level $w \geq 0$ in arbitrary health state l , the consumer's demand for healthcare utilization is given by

$$m^*(w; l) = \begin{cases} w & \text{if } w \leq m^U(l) \\ m^U(l) & \text{otherwise.} \end{cases} \quad (13)$$

If w falls below her unconstrained optimal healthcare spending level $m^U(l)$, the consumer spends all of her money on healthcare. If she can afford $m^U(l)$, she sets $m = m^U(l)$ and spends all her remaining money on non-health consumption.⁵⁷ The consumer's marginal indirect utility of income is thus given by:

$$v_w(w; l) = \begin{cases} \tilde{u}'(b(w; l))b_m(w; l) & \text{if } w \leq m^U(l) \\ \tilde{u}'(b(m^U(l); l) + w - m^U(l)) & \text{otherwise.} \end{cases}$$

Since $\tilde{u}(\cdot)$ is strictly increasing and differentiable, and $b_m(w; l)$ is strictly positive for $w \leq m^U(l)$, it follows that marginal indirect utility of income always exists and is strictly positive: $v_w(w; l) > 0 \forall l, \forall w$. Moreover, the indirect utility function is strictly concave: in the constrained region ($w \leq m^U(l)$), $v_{ww}(w; l) = \tilde{u}''(b(w; l))b_m(w; l) + \tilde{u}'(b(w; l))b_{mm}(w; l) < 0$; and in the unconstrained region $v_{ww}(w; l) = \tilde{u}''(\cdot) < 0$. As a result, the resource constraint will bind with equality at the solution to the first best problem: $\mathbb{E}_l[w^{FB}(l)] = w_0$; and the solution will be fully characterized by the first-order condition of the problem.

The first-best problem stated in Equation 6 has a corresponding Lagrangian formulation given by: $\mathcal{L} = \mathbb{E}_l[v(w(l); l)] + \mu(w_0 - \mathbb{E}_l[w(l)])$, where $\mu \geq 0$ is the Lagrange multiplier. The first-order condition for the first-best income transfer in state l is: $f(l)v_w(w^{FB}(l); l) - \mu f(l) = 0$; and it implies that the marginal indirect utility of income is equalized across all states: $v_w(w(l); l) = \mu$ for all l .

We next show that if the vector of minimum income amounts is affordable in expectation ($w_0 \geq \mathbb{E}_l[w^{min}(l)]$), then first best income will exceed the minimum income amount on a state-

⁵⁷This can be seen clearly by examination of Figure 1.

by-state basis: $w^{FB}(l) \geq w^{min}(l) \forall l$. For the sake of contradiction, suppose there exists a state l' in which first best income does not exceed the minimum income amount: $w^{FB}(l') < w^{min}(l')$. In order to exhaust the resource constraint, there must therefore also be a state l'' in which first best income *does* exceed the minimum income amount: $w^{FB}(l'') > w^{min}(l'')$. Recall that by definition, the consumer's marginal utility of income is equated across states at the minimum amount: $v_w(w^{min}(l'), l') = v_w(w^{min}(l''), l'')$. Since $v_{ww}(w; l) \leq 0$, it must be the case that marginal utility in state l' exceeds marginal utility in state l'' : $v_w(w^{FB}(l'), l') > v_w(w^{FB}(l''), l'')$. But then, the consumer could strictly benefit from reallocating an amount $\epsilon > 0$ of income from state l'' to state l' . This implies a contradiction, as the conjectured income $w^{FB}(l') < w^{min}(l')$ could not have been optimal. Therefore, since the state considered was arbitrary, it follows that $w^{FB}(l) \geq w^{min}(l) \forall l$.

Finally, it follows from $w^{FB}(l) \geq w^{min}(l)$ that the first best healthcare spending coincides with unconstrained optimal healthcare spending in state l . Note also that, by definition, the minimum income amount exceeds unconstrained optimal healthcare spending in all states: $w^{min}(l) \geq m^U(l) \forall l$. Thus, referring to Equation 13, $m^*(w^{FB}(l); l) = m^U(l)$. \square

It then follows directly that first-best payoff is constant.

COROLLARY 1 (Constant First Best Payoff). *Under Assumptions 1 and 3, $z^{FB}(l)$ is constant across health states.*

Proof. By definition, the first best payoff is $z^{FB}(l) = y^{FB}(l) + b(m^{FB}(l); l)$, where $y^{FB}(l) = w^{FB}(l) - m^{FB}(l)$. By Lemma 1, Definition 3, and Definition 4, we can rewrite:

$$\begin{aligned} y^{FB}(l) &= w_0 + \underbrace{m^{FB}(l) + b(m^{FB}(l_0); l_0) - b(m^{FB}(l); l)}_{w^{min}(l)} - \mathbb{E}_l[w^{min}(l)] - m^{FB}(l) \\ &= w_0 + b(m^{FB}(l_0); l_0) - b(m^{FB}(l); l) - \mathbb{E}_l[m^{FB}(l) + b(m^{FB}(l_0); l_0) - b(m^{FB}(l); l)] \\ &= w_0 - \mathbb{E}_l[m^{FB}(l)] + \mathbb{E}_l[b(m^{FB}(l); l)] - b(m^{FB}(l); l). \end{aligned}$$

Therefore, $z^{FB}(l) = w_0 - \mathbb{E}_l[m^{FB}(l)] + \mathbb{E}_l[b(m^{FB}(l); l)] - b(m^{FB}(l); l) + b(m^{FB}(l); l) = w_0 - \mathbb{E}_l[m^{FB}(l) + b(m^{FB}(l); l)]$, which does not depend on l . \square

Further, it follows that healthcare spending weakly exceeds non-moral hazard healthcare spending.

COROLLARY 2 (First Best Spending Weakly Exceeds Non-Moral Hazard Spending). $m^N(l) \leq m^{FB}(l)$

Proof. As defined in Equation 5 and expressed in Equations 12 and 13, $m^N(l) = m^*(l, x_0) = m^*(w_0; l)$, where the consumer is subject to the resource constraint of her initial income w_0 . Per Equation 13, $m^*(w; l)$ will fall weakly below $m^U(l)$ in all health states and for all income levels, so $m^*(w_0; l) \leq m^U(l) \forall l$. So, $m^{FB}(l) = m^U(l) \geq m^*(w_0; l) = m^N(l)$. \square

A.3 Equilibrium Healthcare Spending

Consider the consumer's demand for healthcare utilization under contract $x \in \mathcal{X}$ obtained at premium $p \in \mathbb{R}_+$, as defined in Equation 1 and repeated here:

$$m^*(l, x, p) = \arg \max_{m \geq 0} u(y, m; l) \quad s.t. \quad 0 \leq y \leq w_0 - p - c(m; x).$$

Since utility is strictly increasing in y , the consumer's budget constraint will bind with equality: $y = w_0 - p - c(m; x)$. Applying the parameterization of utility given in Assumption 1, we can therefore rewrite this problem as:

$$m^*(l, x, p) = \arg \max_{m \geq 0} \tilde{u}(w_0 - p - c(m; x) + b(m; l)) \quad s.t. \quad 0 \leq w_0 - p - c(m; x), \quad (14)$$

where recall that $\tilde{u}(\cdot)$ is a CARA utility function, and where $b(m; l)$ is strictly concave and features a (state-dependent) bliss point in m .

Our goal in this section is to provide conditions under which we can characterize the equilibrium level of healthcare utilization $m^*(l, x) = m^*(l, x, \bar{p}(x))$ under a given insurance contract x obtained at the actuarially fair premium. Similar to Definition 2 in Section A.2, it will be useful here also to define the solution to a relaxed version of the consumer's healthcare utilization problem, under which she may circumvent her budget constraint by having negative non-health consumption (i.e., by allowing $y < 0$).

DEFINITION 5 (Unconstrained Optimal Healthcare Spending under Contract x). Let $m^U(l, x)$ denote the solution to the consumer's relaxed healthcare utilization problem under contract x when obtained at premium p : $m^U(l, x) \equiv \arg \max_{m \geq 0} \tilde{u}(w_0 - p - c(m; x) + b(m; l)) = \arg \max_{m \geq 0} (b(m; l) - c(m; x))$.

Notice that since $\tilde{u}(\cdot)$ is monotonically increasing, $m^U(l, x)$ simply solves the first order condition $b_m(m^U(l, x); l) = c_m(m^U(l, x); x)$. In other words, $m^U(l, x)$ represents the level of healthcare utilization at which the marginal benefit of care equals its marginal out-of-pocket cost to

the consumer, and thus represents the consumer's (interior) privately optimal level of health-care utilization. It therefore does not depend on the consumer's initial income or the contract premium. When evaluated at the null contract, $m^U(l, x_0) = m^U(l)$, as defined in Definition 2.

We can now provide a sufficient condition under which the consumer's equilibrium optimal healthcare spending $m^*(l, x)$ coincides with her unconstrained optimal spending $m^U(l, x)$. The basic idea is simply that coverage must be sufficiently high that the consumer can afford the out-of-pocket costs that would arise in any realization of the health state. We denote the set of such contracts by \mathcal{X}' .

LEMMA 2 (Equilibrium Healthcare Spending). *Under Assumption 1, for any contract $x \in \mathcal{X}' \subseteq \mathcal{X}$ for which $c(m^U(l; x); x) \leq w_0 - \bar{p}^U(x) \forall l$, where $\bar{p}^U(x) = \mathbb{E}_l [m^U(l; x) - c(m^U(l; x); x)]$, the consumer's equilibrium optimal healthcare spending is equal to her unconstrained optimal healthcare spending in every health state: $m^*(l, x) = m^U(l, x) \forall l$.*

Proof. The consumer's unconstrained optimal healthcare utilization is given by $m^U(l, x) \equiv \arg \max_{m \geq 0} \tilde{u}(w_0 - p - c(m; x) + b(m; l))$. Supposing that the consumer chose this amount of healthcare utilization in each state and paid an actuarially fair premium for that amount of utilization, the corresponding amount of non-health consumption she would have is $y^U(l, x) = w_0 - \bar{p}^U(x) - c(m^U(l, x); x)$. By assumption of the Lemma, this amount is positive in every state. As a result, the budget constraint was non-binding in the original problem (Equation 14), and the solution to the relaxed problem must therefore coincide with the solution to the original problem. \square

It then follows directly that equilibrium healthcare spending weakly exceeds first best healthcare spending.

COROLLARY 3 (Equilibrium Healthcare Spending Weakly Exceeds First Best Spending). $m^{FB}(l) \leq m^*(l, x)$

Proof. It follows from Definition 2 that unconstrained optimal healthcare spending absent insurance is characterized implicitly by the condition $b_m(m^U(l); l) = 1$. Similarly, from Definition 5, unconstrained optimal healthcare spending under contract x is characterized implicitly by the condition $b_m(m^U(l, x); l) = c_m(m^U(l, x); x)$. Our original assumptions on contracts imply that for all $x \in \mathcal{X}$, $c(m; x) \leq m$ and $c_{mm}(m; x) \leq 0$, and thus, that $c_m(m; x) \leq$

1. Therefore:

$$\begin{aligned} b_m(m^U(l, x_0); l) &= 1 \\ &\geq c_m(m^U(l, x); x) = b_m(m^U(l, x); l). \end{aligned}$$

Moreover, the benefit of healthcare utilization is concave. Therefore, it follows that $m^U(l) = m^U(l, x_0) \leq m^U(l, x)$ for any $x \in \mathcal{X}'$. By Lemma 1, $m^{FB}(l) = m^U(l)$, and by Lemma 2, $m^*(l, x) = m^U(l, x)$. Therefore, $m^{FB}(l) \leq m^*(l, x)$ for any $x \in \mathcal{X}'$. \square

It now follows from Corollaries 2 and 3 that $m^N(l) \leq m^{FB}(l) \leq m^*(l, x)$.

A.4 Higher Healthcare Utilization When Sicker

It is natural to expect that healthcare utilization will be higher when the consumer is sicker. The notion of “sicker” in our model corresponds to a higher health state l . The key assumption disciplining consumer preferences with respect to the level of l is that the marginal utility of healthcare utilization is higher when the consumer is in worse health, i.e., $u_m(y, m; l)$ strictly increasing in l . Together with Assumption 1, this condition allows us to derive the following results:

1. Unconstrained healthcare spending $m^U(l, x)$ is increasing in l for all $x \in \mathcal{X}'$.
2. Non-moral hazard spending $m^N(l)$ is increasing in l .

Proofs of each of these statements are provided below. It then follows directly that for any contract $x \in \mathcal{X}'$, equilibrium healthcare spending $m^*(l, x)$ is also increasing in l , since in such case $m^*(l, x) = m^U(l, x)$. Moreover, if Assumption 3 holds, then first-best spending $m^{FB}(l)$ is also increasing in l , since in such case $m^{FB}(l) = m^U(l, x_0)$.

Proof of Result 1. Unconstrained healthcare spending $m^U(l, x)$ is increasing in l for all $x \in \mathcal{X}'$.

Following Definition 5, unconstrained optimal healthcare spending is given by:

$$m^U(l, x) = \begin{cases} 0 & \text{if } b_m(0; l) \leq c_m(0; x) \\ \arg \max_m (b(m; l) - c(m; x)) & \text{otherwise} \end{cases} \quad (15)$$

Given that $c_m(0; x)$ is invariant to l and $b_m(m; l)$ is strictly increasing in l by assumption, $m^U(l, x)$ will be equal to zero for low l . Once l is sufficiently high, $m^U(l, x)$ will be defined

implicitly as the solution to $b_m(m; l) = c_m(m; x)$. In this case, $m^U(l, x)$ will be strictly positive and unique, since $b_m(0; l) > c_m(0; x)$, with $b_m(m; l)$ and $c_m(m; x)$ continuous and decreasing in m (with $b_m(m; l)$ strictly so).

Fixing a contract $x \in \mathcal{X}'$, consider two arbitrary states $l', l'' \in \mathcal{L}$, where $l' < l''$. We must consider three cases, which—recalling that $b_m(m; l') < b_m(m; l'')$ for all m —are mutually exclusive and exhaustive:

- (i) Suppose $c_m(0; x) < b_m(0; l') < b_m(0; l'')$. Because $b_m(m; l') < b_m(m; l'')$ for all m , the unique solution to $b_m(m; l'') = c_m(m; x)$ will exceed the unique solution to $b_m(m; l') = c_m(m; x)$, and $m^U(l', x) < m^U(l'', x)$.
- (ii) Suppose $b_m(0; l') < c_m(0; x) < b_m(0; l'')$. Then, $m^U(l', x) = 0 < m^U(l'', x)$, since $m^U(l'', x)$ is strictly positive by the fact that $b_m(0; l'') > c_m(0; x)$.
- (iii) Suppose $b_m(0; l') < b_m(0; l'') < c_m(0; x)$. Then, $m^U(l', x) = m^U(l'', x) = 0$.

In any case, $m^U(l', x) \leq m^U(l'', x)$. □

Proof of Result 2. Non-moral hazard spending $m^N(l)$ is increasing in l .

In a given state $l \in \mathcal{L}$, non-moral hazard spending is given by:

$$m^N(l) = m^*(w_0, l) = \begin{cases} w_0 & \text{if } w_0 \leq m^U(l) \\ m^U(l) & \text{otherwise,} \end{cases} \quad (16)$$

as described in Equation 13. Recall that $m^U(l) = m^U(l, x_0)$ is the consumers unconstrained optimal healthcare utilization under the null contract (cf. Definition 2).

Consider two arbitrary states $l', l'' \in \mathcal{L}$, where $l' < l''$. Since $x_0 \in \mathcal{X}$, it follows from Result 1 that $m^U(l)$ is increasing in l : $m^U(l') \leq m^U(l'')$. To order the corresponding $m^N(l)$, we must consider three mutually exclusive and exhaustive cases:

- (i) Suppose $w_0 \leq m^U(l') \leq m^U(l'')$. Then, $m^N(l') = m^N(l'') = w_0$.
- (ii) Suppose $m^U(l') \leq w_0 \leq m^U(l'')$. Then, $m^N(l') = m^U(l') \leq w_0 = m^N(l'')$.
- (iii) Suppose $m^U(l') \leq m^U(l'') \leq w_0$. Then, $m^N(l') = m^U(l') \leq m^U(l'') = m^N(l'')$.

In any case, $m^N(l') \leq m^N(l'')$. □

A.5 Signing Terms in Welfare Decomposition

Our goal in this section is to establish conditions under which we can sign each of the five components of social welfare discussed in Section II.C.

Decomposition of Expected Payoff. Consider the decomposition of the expected payoff derived from a contract $x \in \mathcal{X}'$ relative to the null contract x_0 , as defined in Equation 9 and repeated here:

$$\Delta EV(x) = \underbrace{EV(x) - EV_{FB}}_{\substack{\text{Social cost} \\ \text{of over-utilization} \\ \Delta EV_{over}(x)}} + \underbrace{EV_{FB} - EV(x_0)}_{\substack{\text{Social benefit of efficient} \\ \text{increase in utilization} \\ \Delta EV_{eff}}},$$

noting that due to actuarially fair pricing, EV_{FB} does not depend on the contract x .

As discussed in the main text, $\Delta EV(x)$ captures the extent to which the consumer's ex ante expected value of changes in utilization exceed their expected cost. Its components correspond to the components of utilization changes. The first term, $\Delta EV_{over}(x)$, reflects the net expected value of moving from non-moral hazard spending to socially efficient spending in each health state (accounting for the commensurate change in the fair premium). The second term, ΔEV_{eff} , reflects the net expected value of moving from socially efficient spending to the equilibrium level of spending in each health state (again accounting for the change in premium). Under Assumptions 1, 2, and 3, $\Delta EV_{over}(x)$ is always negative, while ΔEV_{eff} is always positive.

Proof. $\Delta EV_{over}(x) \leq 0$ and $\Delta EV_{eff} \geq 0$. Simplifying the expression for $\Delta EV_{over}(x)$ yields

$$\Delta EV_{over}(x) = \mathbb{E}_l [b(m^*(l, x); l) - m^*(l, x) - b(m^{FB}(l); l) + m^{FB}(l)].$$

Since m^{FB} maximizes $b(m; l) - m$ by construction, $b(m^{FB}(l); l) - m^{FB}(l) \geq b(m^*(l, x); l) - m^*(l, x)$ in every state, and thus $\Delta EV_{over}(x) \leq 0$. Similarly, simplifying the expression for ΔEV_{eff} yields

$$\Delta EV_{eff} = \mathbb{E}_l [b(m^{FB}(l); l) - m^{FB}(l) - b(m^N(l); l) + m^N(l)].$$

Since m^{FB} maximizes $b(m; l) - m$ by construction, $b(m^{FB}(l); l) - m^{FB}(l) \geq b(m^N(l); l) - m^N(l)$ in every state, and thus $\Delta EV_{eff} \geq 0$. \square

Decomposition of Risk Protection. Consider the decomposition of the value of risk protection derived from a contract $x \in \mathcal{X}'$ relative to the null contract x_0 , as defined in Equation 10 and repeated here:

$$\Psi(x) = \underbrace{RP(x_0) - RP_N(x)}_{\substack{\text{Value of risk protection} \\ \text{from non-MH spending,} \\ \Psi_N(x)}} + \underbrace{RP_N(x) - RP_{FB}(x)}_{\substack{\text{Value of risk protection} \\ \text{from efficient increase} \\ \text{in spending,} \\ \Psi_{eff}(x)}} + \underbrace{RP_{FB}(x) - RP(x)}_{\substack{\text{Value of risk protection} \\ \text{from over-utilization,} \\ \Psi_{over}(x)}} .$$

To build intuition about the sign of these objects, consider the general expression for the risk premium RP associated with a lottery induced by a given vector of state-dependent payoffs $(z(l))$, in light of Assumption 1:

$$\begin{aligned} RP &= EV - CE \\ &= \mathbb{E}_l z(l) - \tilde{u}^{-1}[\mathbb{E}_l \tilde{u}(z(l))] \\ &= -\tilde{u}^{-1}[\mathbb{E}_l \tilde{u}(z(l) - \bar{z})]. \end{aligned}$$

The last equality follows from the fact that $\tilde{u}(\cdot)$ is the exponential function, and where $\bar{z} \equiv \mathbb{E}_l z(l)$ to save on notation. This simplified expression for the risk premium makes clear that the value the consumer would place on certain outcome \bar{z} relative to uncertain outcome $z(l)$ is equal to the absolute value of the certainty equivalent of the mean-zero lottery consisting of deviations from mean payoffs. The (incremental) value of risk protection derived between any two lotteries is then simply the *difference* in risk premia induced by each lottery. The three terms in $\Psi(x)$ are three such differences.

Given Assumptions 1, 2, and 3, along with the definitions of consumer types in Section II.C (Definition 1), we are able to derive the following results:

1. If the consumer is a *Homebody* or a *Worrier*—that is, if $y^{FB}(l)$ is non-decreasing in l —then $\Psi_N(x) \geq 0$.
2. If the consumer is a *Homebody* or a *Worrier*—that is, if $y^{FB}(l)$ is non-decreasing in l —then $\Psi_{eff}(x) \geq 0$.
3. Suppose that contract x induces an equilibrium payoff $z^*(l, x)$ which is decreasing in l at a slower rate than $z^{FB}(l, x)$. Then: $\Psi_{over}(x)$ is positive.

As discussed in Section II.C, whether the value of risk protection is positive depends on

how the consumer's preferences over m and y vary with her health status. Definition 1 provides a classification into three “types,” where a *Traveler* prefers to have greater non-health consumption when she is healthy, a *Homebody* prefers to have equal non-health consumption in all states, and a *Worrier* prefers to have greater non-health consumption when she is sicker. The first and second results are that, as long as the consumer is a *Homebody* or a *Worrier*, the value of risk protection from allowing her to increase her healthcare utilization up to the efficient level is unambiguously positive.

The third result pertains to the sign of Ψ_{over} . While the three consumer “types” contain conditions on payoffs at *efficient* health and non-health consumption, signing Ψ_{over} requires conditions on payoffs at *equilibrium* consumption bundles. The statement of the third results provides a general formulation of such conditions. These conditions would hold, for example, if the consumer were a *Worrier* for whom $b(m^{bliss}; x)$ (her value of bliss-point healthcare utilization) were constant across health states, for any linear insurance contract. This is precisely the example considered in Section III (Specification 2).

Proofs of these results are provided below. As it will be useful in all three proofs, we first establish one ancillary lemma.

LEMMA 3. *Suppose the consumer is a Homebody or a Worrier. Then, under Assumptions 1, 2, and 3, the payoffs $z^*(l, x_0)$, $z^N(l, x)$, and $z^{FB}(l, x)$ are decreasing in l .*

Proof. Recall from Equations 2, 7, and 8 that:

$$\begin{aligned} z^*(l, x_0) &= w_0 - m^N(l) + b(m^N(l); l) \\ z^N(l, x) &= w_0 - \bar{p}^N(x) - c(m^N(l); x) + b(m^N(l); l) \\ z^{FB}(l, x) &= w_0 - \bar{p}^{FB}(x) - c(m^{FB}(l); x) + b(m^{FB}(l); l). \end{aligned}$$

In order to sign the derivative of each with respect to l , we must characterize the derivatives of $m^N(l)$, $m^{FB}(l)$, $b(m^{FB}(l); l)$, and $b(m^N(l); l)$ with respect to l . It has already been established that $m^N(l)$ is increasing in l (see Result 2 in Section A.4); and that $m^{FB}(l)$ is increasing in l (see Lemma 1 along with Result 1 in Appendix A.4). Since the out-of-pocket cost function is increasing in m , $c(m^N(l); x)$ and $c(m^{FB}(l); x)$ are likewise increasing with respect to l .

Turning to $b(m^{FB}(l); l)$, notice that if $y^{FB}(l)$ is non-decreasing in l (as for a *Homebody* and *Worrier*), then $b(m^{FB}(l); l)$ is non-increasing in l . Since first-best payoffs $y^{FB}(l) + b(m^{FB}(l); l)$ are constant across states (see Corollary 1), the two terms must perfectly counteract one another.

Finally, we can then show that $b(m^N(l); l)$ must be decreasing in l if $b(m^{FB}(l); l)$ is decreasing in l . Note that from Equation 16, $m^N(l)$ is only different from $m^{FB}(l)$ if $m^N(l) = w_0$. Given two arbitrary states $l' < l''$, we can therefore consider three cases, which are mutually exclusive and exhaustive:

- (i) Suppose $m^{FB}(l'') \leq w_0$. Then, $m^N(l') = m^{FB}(l')$ and $m^N(l'') = m^{FB}(l'')$. It then follows that $b(m^N(l'); l') - b(m^N(l''); l'') = b(m^{FB}(l'); l') - b(m^{FB}(l''); l'') \geq 0$. That is, $b(m^{FB}(l); l)$ decreasing in l directly implies that $b(m^N(l); l)$ is decreasing in l in this case.
- (ii) Suppose $m^{FB}(l') \leq w_0 \leq m^{FB}(l'')$. Then, $b(m^N(l'); l') - b(m^N(l''); l'') = b(m^{FB}(l'); l') - b(w_0; l'') \geq b(m^{FB}(l'); l') - b(m^{FB}(l''); l'') \geq 0$, where the penultimate inequality follows from the fact that $b(m; l'')$ is increasing for $m \leq m^{FB}(l'')$ by construction, and the last inequality holds because $b(m^{FB}(l); l)$ is decreasing in l . Thus, $b(m^N(l); l)$ is also decreasing in l in this case.
- (iii) Suppose $w_0 \leq m^{FB}(l')$. Then $b(m^N(l'); l') - b(m^N(l''); l'') = b(w_0; l') - b(w_0; l'')$. To establish the sign of this difference, we will leverage the relationship across $b(m^{FB}(l); l)$, along with the facts that $m^{FB}(l)$ is increasing in l for all consumers (Lemma 1 and Result 1 in Appendix A.4), and that $b_m(m; l)$ is increasing in l (see Footnote 25). Notice that: $b(m^{FB}(l'); l') = b(w_0; l') + \int_{w_0}^{m^{FB}(l')} b_m(m; l') dm$, and similarly $b(m^{FB}(l''); l'') = b(w_0; l'') + \int_{w_0}^{m^{FB}(l'')} b_m(m; l'') dm$. Using the fact that $m^{FB}(l'') \geq m^{FB}(l')$, we can write:

$$\begin{aligned} b(m^{FB}(l'); l') - b(m^{FB}(l''); l'') &= - \int_{w_0}^{m^{FB}(l')} (b_m(m; l'') - b_m(m; l')) dm - \int_{m^{FB}(l')}^{m^{FB}(l'')} b_m(m; l'') dm \\ &\quad + b(w_0; l') - b(w_0; l''). \end{aligned}$$

For Homebodies and Worriers, $b(m^{FB}(l'); l') - b(m^{FB}(l''); l'') \geq 0$, which implies that $b(w_0; l') - b(w_0; l'') \geq \int_{w_0}^{m^{FB}(l')} (b_m(m; l'') - b_m(m; l')) dm + \int_{m^{FB}(l')}^{m^{FB}(l'')} b_m(m; l'') dm \geq 0$, where the last inequality follows from the fact that from the fact that $b_m(m^{FB}(l); l)$ and $m^{FB}(l)$ are both increasing in l . Therefore, $b(m^N(l); l)$ is also decreasing in l in this final case.

Taking these conclusions together, it is clear that all the terms in $z^*(l, x_0)$, $z^N(l, x)$, and $z^{FB}(l, x)$ are either constant or decreasing in l . \square

Proof of Result 1. If the consumer is a Homebody or a Worrier, then $\Psi_N(x) \geq 0$.

By definition, $\Psi_N(x) = RP(x_0) - RP_N(x)$, where:

$$\begin{aligned} RP(x_0) &= -\tilde{u}^{-1}[\mathbb{E}_l \tilde{u}(z^*(l, x_0) - \bar{z}^*(x_0))] \\ RP_N(x) &= -\tilde{u}^{-1}[\mathbb{E}_l \tilde{u}(z^N(l, x) - \bar{z}^N(x))], \end{aligned}$$

Now re-writing:

$$\Psi_N(x) = \tilde{u}^{-1}[\mathbb{E}_l \underbrace{\tilde{u}(z^N(l, x) - \bar{z}^N(x))}_{\Delta z^N(l, x)}] - \tilde{u}^{-1}[\mathbb{E}_l \underbrace{\tilde{u}(z^*(l, x_0) - \bar{z}^*(x_0))}_{\Delta z^*(l, x_0)}].$$

Let $\Delta z^N(l, x) \equiv z^N(l, x) - \bar{z}^N(x)$ and $\Delta z^*(l, x_0) \equiv z^*(l, x_0) - \bar{z}^*(x_0)$ denote the mean-zero lotteries induced by the consumer's distribution over health states $l \sim F$. Given that $\tilde{u}^{-1}(\cdot)$ is monotonically increasing and $\tilde{u}(\cdot)$ concave, $\Psi_N(x)$ will be positive if the lottery $\Delta z^*(l, x_0)$ (induced by the null contract x_0) is a mean-preserving spread of the lottery $\Delta z^N(l, x)$ (induced by the insurance contract x). Lemma 3 has already established that $\Delta z^*(l, x_0)$ and $\Delta z^N(l, x)$ are both decreasing in l . $\Psi_N(x)$ will thus be positive so long as $\Delta z^*(l, x_0)$ is decreasing in l faster than is $\Delta z^N(l, x)$ (since their means are equal).

Consider two arbitrary states $l' < l''$. It is immediate to see that, plugging in for the definitions, rearranging, and simplifying, yields:

$$\begin{aligned} &\Delta z^*(l', x_0) - \Delta z^*(l'', x_0) - (\Delta z^N(l', x) - \Delta z^N(l'', x)) \\ &= m^N(l'') - m^N(l') - (c(m^N(l''); x) - c(m^N(l'); x)) \\ &\geq 0, \end{aligned}$$

because $c(m; x) \leq m$ and is weakly concave in m by construction. $\Delta z^*(l, x_0)$ is thus decreasing in l everywhere faster than is $\Delta z^N(l, x)$, and $\Delta z^*(l, x_0)$ is mean-preserving spread of $\Delta z^N(l, x)$. \square

Proof of Result 2. If the consumer is a Homebody or a Worrier, then $\Psi_{eff}(x) \geq 0$.

By definition, $\Psi_{eff}(x) = RP_N(x) - RP_{FB}(x)$, where:

$$\begin{aligned} RP_N(x) &= -\tilde{u}^{-1}[\mathbb{E}_l \tilde{u}(z^N(l, x) - \bar{z}^N(x))] \\ RP_{FB}(x) &= -\tilde{u}^{-1}[\mathbb{E}_l \tilde{u}(z^{FB}(l, x) - \bar{z}^{FB}(x))], \end{aligned}$$

Now re-writing:

$$\Psi_{eff}(x) = \tilde{u}^{-1}[\mathbb{E}_l \tilde{u}(\underbrace{z^{FB}(l, x) - \bar{z}^{FB}(x)}_{\Delta z^{FB}(l, x)})] - \tilde{u}^{-1}[\mathbb{E}_l \tilde{u}(\underbrace{z^N(l, x) - \bar{z}^N(x)}_{\Delta z^N(l, x)})].$$

Let $\Delta z^{FB}(l, x) = z^{FB}(l, x) - \bar{z}^{FB}(x)$ and $\Delta z^N(l, x) = z^N(l, x) - \bar{z}^N(x)$ denote the mean-zero lotteries induced by the consumer's distribution over health states $l \sim F$. Given that $\tilde{u}^{-1}(\cdot)$ is monotonically increasing and $\tilde{u}(\cdot)$ concave, $\Psi_{eff}(x)$ will be positive if the lottery $\Delta z^N(l, x)$ (induced by payoffs derived from non-moral hazard spending) is a mean-preserving spread of the lottery $\Delta z^{FB}(l, x)$ (induced by payoffs derived from first-best spending). Lemma 3 has already established that $\Delta z^N(l, x)$ and $\Delta z^{FB}(l, x)$ are both decreasing in l . $\Psi_{eff}(x)$ will thus be positive so long as $\Delta z^N(l, x)$ is decreasing in l *faster* than $\Delta z^{FB}(l, x)$ (since their means are equal).

Note that from Equation 16, $m^N(l)$ is only different from $m^{FB}(l)$ if $m^N(l) = w_0$. Given two arbitrary states $l' < l''$, we can therefore consider three cases, which are mutually exclusive and exhaustive:

- (i) Suppose $m^{FB}(l'') \leq w_0$. Then, $m^N(l') = m^{FB}(l')$ and $m^N(l'') = m^{FB}(l'')$ and so, $z^N(l', x) - z^N(l'', x) - (z^{FB}(l', x) - z^{FB}(l'', x)) = 0$. $\Delta z^N(l, x)$ and $\Delta z^{FB}(l, x)$ are decreasing at the same speed in this case.
- (ii) Suppose $m^{FB}(l') \leq w_0 \leq m^{FB}(l'')$. Then, $m^{FB}(l') = m^N(l')$, which implies that $z^N(l', x) - z^N(l'', x) - (z^{FB}(l', x) - z^{FB}(l'', x)) = z^{FB}(l'', x) - z^N(l'', x) = b(m^{FB}(l''); l'') - c(m^{FB}(l''); x) - (b(m^N(l''); l'') - c(m^N(l''); x))$. This expression is clearly positive because $b_m(m; l)$ is greater than or equal to 1 anywhere to the left of $m^{FB}(l)$ and $c_m(m; x)$ is always less than 1.
- (iii) Suppose $w_0 \leq m^{FB}(l')$. Then, $m^N(l') = m^N(l'') = w_0$, and $z^N(l', x) - z^N(l'', x) - (z^{FB}(l', x) - z^{FB}(l'', x)) = -(b(w_0; l'') - b(w_0; l')) + b(m^{FB}(l''); l'') - b(m^{FB}(l'); l') - (c(m^{FB}(l''); x) - c(m^{FB}(l'); x))$. Notice that, as in the proof of Lemma 3 case (iii), we can rewrite this expression as: $z^N(l', x) - z^N(l'', x) - (z^{FB}(l', x) - z^{FB}(l'', x)) =$

$\int_{w_0}^{m^{FB}(l'')} (b_m(m; l'') - b_m(m; l')) dm + \int_{m^{FB}(l'')}^{m^{FB}(l'')} (b_m(m; l'') - c_m(m; x)) dm$. The first term is positive by the fact that $b_m(m; l)$ is increasing in l (Footnote 25), and the second term is positive by the fact that $b_m(m; l) \geq 1$ for $m \leq m^{FB}(l)$ and that $c_m(m; x) \leq 1$ for all m , so the expression is positive.

In any case, $\Delta z^N(l, x)$ is decreasing in l weakly faster than $\Delta z^{FB}(l, x)$, and therefore $\Delta z^N(l, x)$ is mean-preserving spread of $\Delta z^{FB}(l, x)$. \square

Proof of Result 3. By definition, $\Psi_{over}(x) = RP_{FB}(x) - RP(x)$, where:

$$\begin{aligned} RP_{FB}(x) &= -\tilde{u}^{-1}[\mathbb{E}_l \tilde{u}(z^{FB}(l, x) - \bar{z}^{FB}(x))] \\ RP(x) &= -\tilde{u}^{-1}[\mathbb{E}_l \tilde{u}(z^*(l, x) - \bar{z}^*(x))], \end{aligned}$$

Now re-writing:

$$\Psi_{over}(x) = \tilde{u}^{-1}[\mathbb{E}_l \underbrace{\tilde{u}(z^*(l, x) - \bar{z}^*(x))}_{\Delta z^*(l, x)}] - \tilde{u}^{-1}[\mathbb{E}_l \underbrace{\tilde{u}(z^{FB}(l, x) - \bar{z}^{FB}(x))}_{\Delta z^{FB}(l, x)}].$$

Let $\Delta z^*(l, x) = z^*(l, x) - \bar{z}^*(x)$ and $\Delta z^{FB}(l, x) = z^{FB}(l, x) - \bar{z}^{FB}(x)$ denote the mean-zero lotteries induced by the consumer's distribution over health states $l \sim F$. Given that $\tilde{u}^{-1}(\cdot)$ is monotonically increasing and $\tilde{u}(\cdot)$ concave, $\Psi_{over}(x)$ will be positive if the lottery $\Delta z^{FB}(l, x)$ (induced by payoffs derived from first-best spending) is a mean-preserving spread of the lottery $\Delta z^*(l, x)$ (induced by payoffs derived from equilibrium spending). Given that we have assumed that $z^*(l, x)$ is decreasing slower than $z^{FB}(l, x)$ for contract x , it follows that $\Delta z^{FB}(l, x)$ is a mean-preserving spread of $\Delta z^*(l, x)$. \square

Appendix B Computational Appendix

Population of Consumers. We model consumers as households, where a household is a group of individuals. Each individual is characterized by an age, a gender, and a health risk score. We construct a population of households to match characteristics of the non-elderly US population reported in the American Community Survey (ACS) for 2023 (U.S. Census Bureau, 2023). The construction of each household begins with a “head of household,” who is female

with 50 percent probability and has a uniform distribution of age between 22 and 65. We assume that households have a spouse present with a 90 percent chance, and that the spouse is of the opposite gender to the head of household when present. The age of the spouses is drawn from a normal distribution with mean equal to the age of the head of household and a standard deviation of 4, subject to bounds between 22 and 65. Each household can have between 1 and 4 children, where each child independently exists with 15 percent probability. Conditional on existing, a child is female with 50 percent probability and their age is drawn from a uniform distribution between 0 and 18. All individuals draw a risk score from a log-normal distribution with mean positively related to age, such that for individual i : $\log(riskscore_i) \sim \mathcal{N}(\frac{age_i}{20}, 1)$. The right tail of the risk score distribution is censored such that no individual has a risk score greater than five standard deviations above the uncensored mean. Households draw an income from a log-normal distribution with mean \$55,000 and standard deviation \$22,000.⁵⁸ Our baseline population contains 10,000 households. Increasing the number of households does not change our results.

We then construct the key parameters that describe a consumer in the model: the health state distribution F , the risk aversion parameter ψ , the “over-utilization parameter” ω , and initial income w_0 . We use the parameter estimates reported in Column 3 of Table 3 and Appendix Table A.8 of [Marone and Sabety \(2022\)](#). We make one adjustment, which is to cap the risk aversion parameter at a value of 5.⁵⁹

Summary statistics on the population distribution of demographics and resulting household types are reported in Table B.1. Facing an equal odds gamble between \$0 and \$100, the average household would have a certainty equivalent of \$48.8, reflecting risk aversion. For the average household, there is a 2 percent probability of realizing a health in which their efficient level of healthcare utilization exceeds their initial income. This measure is highly skewed: the median household is never constrained, and the 90th percentile household has a 4 percent probability of being constrained. The average household would have total healthcare spending equal

⁵⁸We calibrate the mean based on the ACS mean income for employees in Oregon in 2010; the rationale is that [Marone and Sabety \(2022\)](#) estimates of the health state distribution come from a sample of Oregon public school employees between 2008 and 2013. The mean household income for Oregon employees in 2010 was \$61,095 (ACS Table S1902). For our income draws, we adjust the mean down to \$55,000 to capture the fact that public school employees have below-median wages ([Oregon Public Employees Retirement System, 2024](#)). Given this construction, it is possible that some households end up with an income draw which is so low relative to the health state distribution, that the household cannot in expectation afford their first best level of healthcare utilization. For these cases, income is subsequently adjusted upward numerically. In other words, income is adjusted to satisfy Assumption 3.

⁵⁹We express monetary amounts in thousands of dollars, so dividing our coefficients of absolute risk aversion by 1,000 makes them comparable to other settings where monetary amounts are measured in dollars.

to \$10,372 under a full insurance contract, but only \$8,926 under a null contract, reflecting moral hazard. Under first-best contracts, the average household would have total healthcare spending of \$9,024.

Table B.1. Population Demographics

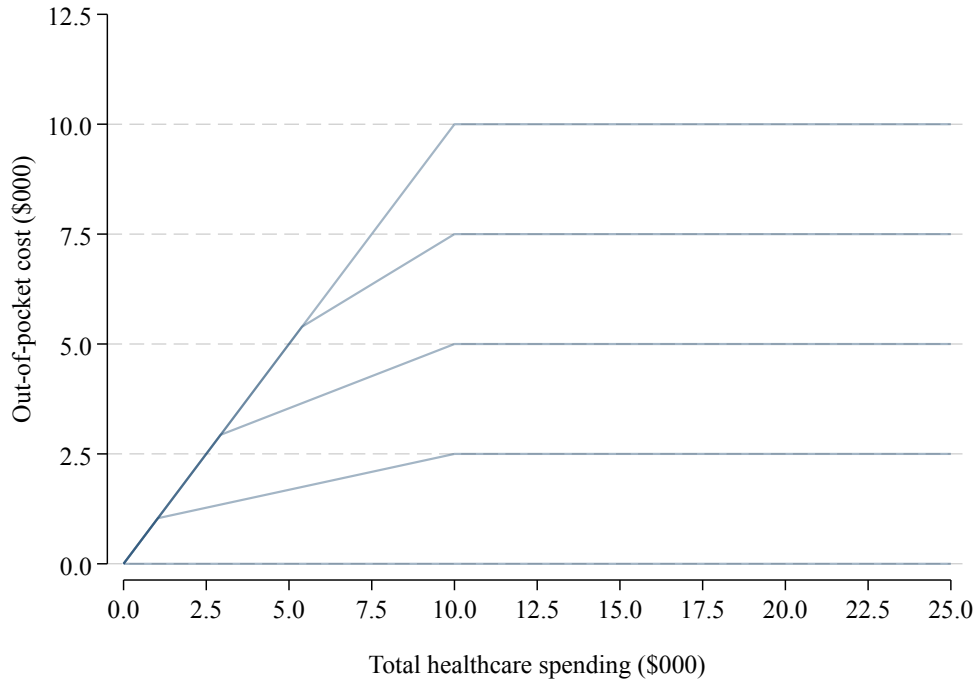
Sample demographic	Mean	Percentile				
		10	25	Median	75	90
<i>Demographics</i>						
Number of adults	1.9	2.0	2.0	2.0	2.0	2.0
Number of children	0.6	–	–	–	1.0	2.0
Average age of household adults	43.4	26.5	32.5	43.4	54.1	61.1
Initial income, w_0 (\$000)	55.5	31.9	40.2	51.5	66.7	86.9
Risk aversion parameter, ψ	0.9	0.2	0.4	0.7	1.2	2.0
Spec (1) moral hazard parameter, ω (\$000)	1.3	0.8	1.0	1.3	1.7	2.0
Spec (2) moral hazard parameter, $\tilde{\omega}$	0.2	0.1	0.1	0.2	0.3	0.5
<i>Resulting characteristics</i>						
CE of equal-odds gamble between \$0 and \$100 (\$)	48.8	47.6	48.5	49.2	49.5	49.7
Prob. of realizing health state in which $w_0 < m^{FB}(l)$	0.02	–	–	–	–	0.06
Utilization under null contract, $\mathbb{E}[m^*(l, x_0)]$ (\$000)	8.9	2.9	4.3	7.4	12.3	17.6
full insurance, $\mathbb{E}[m^*(l, x_{full})]$ (\$000)	10.4	4.0	5.6	8.9	13.8	19.3
FB contract, $\mathbb{E}[m^{FB}(l)]$ (\$000)	9.0	2.9	4.3	7.5	12.5	17.9

Notes: This table shows the distribution of consumer demographics and types used in our numerical analysis. The population consists of 10,000 consumers (households). “CE” stands for certainty equivalent. “FB” stands for first best. Spec (1) and Spec (2) are in reference to the parameterizations of utility presented in Table 1 in the main text.

Insurance Contracts. We consider a set of five contracts that are piecewise linear, with a deductible, coinsurance region, and out-of-pocket maximum design. The lowest level of coverage considered is a “Catastrophic” contract with a deductible and out-of-pocket maximum of \$10,000. The highest level of coverage is full insurance. The out-of-pocket cost functions for the five contracts are depicted in Figure B.1. Analysis in Chade et al. (2023) suggests that adding additional contracts would not have material effects on achievable second-best welfare.

Table 2 reports the average impact on welfare of moving a consumer from full insurance to their welfare-maximizing contract within this set. In specification (1), moving from lowest to highest coverage, the percent of consumers with each contract as their efficient contract is .1%, 2%, 21%, 54%, and 24%. Moving consumers from full insurance to their efficient contracts increases welfare by on average 2 percent. In specification (1), moving from lowest to highest coverage, the percent of consumers with each contract as their efficient contract is .1%, .1%, 1%, 21%, and 77%. Here, moving consumers to their efficient contracts increases welfare by on average 2 percent.

Figure B.1. Potential Contracts



Notes: The figure shows the set of potential contracts from which the consumers' efficient level of coverage is determined.

Appendix C Alternative Estimates: EFRSC

We repeat the analysis and welfare decomposition using alternative estimates of moral hazard, the health state distribution, and risk aversion from [Einav et al. \(2013\)](#) (hereby EFRSC). To make the populations comparable, we maintain the household size and gender distributions described in Appendix B. We construct the income distribution to match the mean household income of the EFRSC population, preserving the same coefficient of variation as in our main specification. Given that [Einav et al. \(2013\)](#) report only the mean *individual* income (Table 1), we scale this number by the mean number of adults per household.⁶⁰ Table C.1 reports summary statistics on the population distribution of demographics and household types in this alternative population. By construction, the first three demographic characteristics are identical to those in reported in Table B.1. The mean income and risk aversion parameter

⁶⁰The household size draws are held fixed across the different simulations, and thus the average number of adults per household is the same across the two populations. The average annual income reported in [Einav et al. \(2013\)](#) is \$31,292, and we scale this number by 1.9 to arrive at an mean household income of \$59,454 to construct our income draws. As in our main specification, income draws that violate Assumption 3 are subsequently adjusted to satisfy so that each household can afford their interior optimum level of healthcare utilization in expectation.

are both slightly larger, and the moral hazard parameters are slightly smaller, relative to our baseline population.

Facing an equal odds gamble between \$0 and \$100, the average household would have a certainty equivalent of \$48.7, reflecting the slightly larger degree of risk aversion than in our baseline population. For the average household, there is a 2 percent probability of realizing a health state in which their efficient level of healthcare utilization exceeds their initial income. This measure is highly skewed. The median household is never constrained, and the 95th percentile household has a 1.3 percent probability of being constrained. Among the top 5 percent of households, the average probability of being constrained is 3.6 percent. The average household would have total healthcare spending equal to \$5,279 under a full insurance contract, but only \$4,229 under a null contract, reflecting moral hazard. Under first-best contracts, the average household would have total healthcare spending of \$4,241.

Table C.1. Population Demographics: [Einav et al. \(2013\)](#) Estimates

Sample demographic	Mean	Percentile				
		10	25	Median	75	90
<i>Demographics</i>						
Number of adults	1.9	2.0	2.0	2.0	2.0	2.0
Number of children	0.6	–	–	–	1.0	2.0
Average age of household adults	43.4	26.5	32.5	43.4	54.1	61.1
Initial income, w_0 (\$000)	59.6	34.3	43.2	55.4	71.7	93.4
Risk aversion parameter, ψ	1.1	0.1	0.3	0.7	1.5	2.4
Spec (1) moral hazard parameter, ω (\$000)	1.1	0.2	0.2	0.3	1.0	2.9
Spec (2) moral hazard parameter, $\tilde{\omega}$	0.3	0.0	0.1	0.1	0.3	0.8
<i>Resulting characteristics</i>						
CE of equal-odds gamble between \$0 and \$100 (\$)	48.7	47.0	48.2	49.1	49.6	49.9
Prob. of realizing health state in which $w_0 < m^{FB}(l)$	<0.01	–	–	–	–	<0.01
Utilization under null contract, $\mathbb{E}[m^*(l, x_0)]$ (\$000)	4.5	0.7	1.7	3.3	6.0	10.0
full insurance, $\mathbb{E}[m^*(l, x_{full})]$ (\$000)	5.7	0.9	2.1	3.9	7.6	12.6
FB contract, $\mathbb{E}[m^{FB}(l)]$ (\$000)	4.5	0.7	1.7	3.3	6.1	10.0

Notes: This table shows the distribution of consumer demographics and types resulting from the [Einav et al. \(2013\)](#) estimates. The population consists of 10,000 consumers (households). “CE” stands for certainty equivalent. “FB” stands for first best. Spec (1) and Spec (2) are in reference to the parameterizations of utility presented in Table 1 in the main text.

Table C.2 presents the corresponding decompositions of welfare, parallel to Table 2 in the main text. The results are qualitatively similar. The vast majority of welfare is derived from Ψ_N , the risk protection value of insurance absent moral hazard. There are a small number of consumers who have substantially lower incomes and thus a far higher probability of being constrained absent insurance. For the 2 percent of consumers with a “High” (\geq

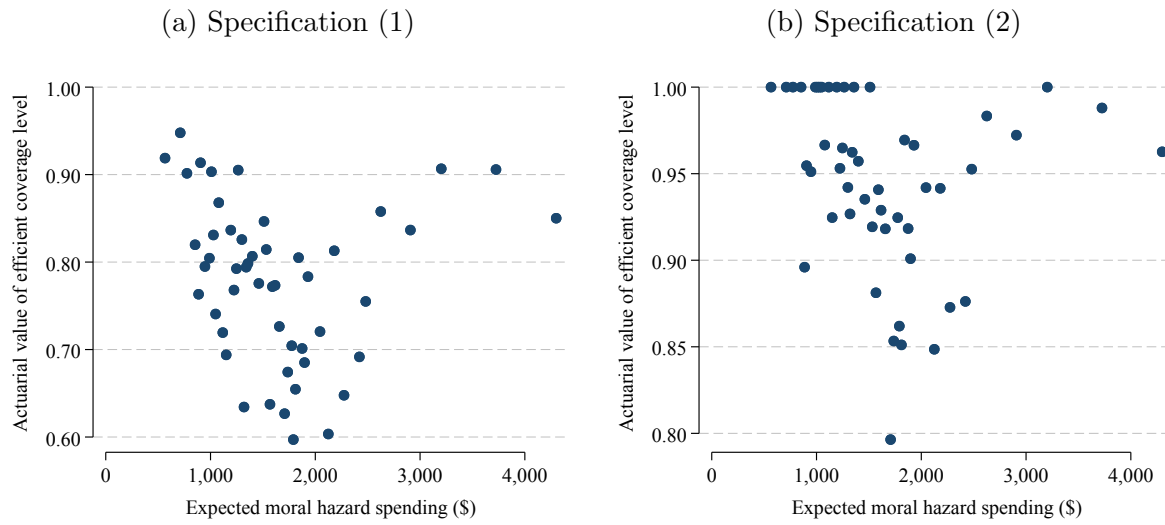
5 percent) probability of being constrained, the average probability is 7.4 percent, 37 times more likely than the average consumer. For these consumers, the risk protection value of moral hazard represents a huge fraction of the welfare generated by insurance. These consumers are responsible for driving the overall average of 6 percent of welfare attributable to ΔEV_{eff} .

Table C.2. Decomposition of Welfare: [Einav et al. \(2013\)](#) Estimates

Consumer group	Decomposition of welfare (Under full insurance)					Relative change in welfare under		
	Ψ_N	Ψ_{eff}	Ψ_{over}	ΔEV_{eff}	ΔEV_{over}	Efficient level of coverage	Eliminate over-utiliz.	Eliminate all moral hazard
Panel A. Spec (1): <i>Homebodies</i>								
All	0.93	0.12	–	0.06	-0.11	0.06	0.11	-0.07
<i>Prob. constrained</i> [†]								
Low	0.96	0.10	–	<0.01	-0.06	0.02	0.06	-0.04
Medium	0.59	0.37	–	0.13	-0.09	0.06	0.09	-0.41
High	0.45	0.52	–	1.63	-1.60	0.97	1.60	-0.55
<i>Tertile of ω</i>								
Lowest	0.85	0.17	–	<0.01	-0.02	0.01	0.02	-0.15
Middle	0.93	0.11	–	0.29	-0.33	0.20	0.33	-0.07
Highest	1.03	0.07	–	0.02	-0.12	0.04	0.12	0.03
Panel B. Spec (2): <i>Worriers</i>								
All	0.94	0.05	0.04	<0.01	-0.04	<0.01	-0.00	-0.06
<i>Prob. constrained</i> [†]								
Low	0.96	0.03	0.04	<0.01	-0.03	<0.01	-0.01	-0.04
Medium	0.81	0.22	0.03	0.08	-0.14	0.05	0.11	-0.19
High	0.62	0.40	0.01	0.02	-0.05	<0.01	0.04	-0.38
<i>Tertile of ω</i>								
Lowest	0.93	0.06	0.01	<0.01	-0.01	<0.01	<0.01	-0.07
Middle	0.98	0.05	0.02	<0.01	-0.06	0.02	0.04	-0.02
Highest	0.93	0.04	0.09	<0.01	-0.06	<0.01	-0.03	-0.07

Notes: This table shows the decomposition of welfare generated by full insurance into the five terms described in Equation 11, for each of the parameterizations of consumer utility described in Table 1. Each of the five columns under “Decomposition of welfare” represents the proportion of total welfare V represented by that term, on average across consumers. The columns thus sum to one within each row. The final three columns report the relative changes in welfare that would occur in three scenarios: (a) if consumers were enrolled in their efficient level of coverage, rather than full insurance; (b) if over-utilization were prevented, and (c) if all moral hazard were prevented. A relative change of 0.06 means a 6% increase in welfare. [†]Low probability of being constrained means a <1% change (true for 94.5 percent of consumers); medium means a 1%–5% chance (3.5 percent of consumers); high means $\geq 5\%$ chance (2 percent of consumers).

Figure A.1. Efficient Coverage Level versus Expected Moral Hazard Spending



Notes: The figure shows the relationship between a consumer's efficient coverage level (measured in actuarial value) and their expected increase in healthcare utilization when moving from no insurance to full insurance (expected moral hazard spending). Both panels are binned scatter plots using 50 points, across the 333 households in the bottom third of the income distribution. Panel (a) corresponds to utility specification (1), while panel (b) corresponds to specification (2). This figure is referenced in Section [III.B](#)