American Economic Journal: Microeconomics 2026, 18(3): 426–453
https://doi.org/10.1257/mic.20240042

Liquidity Constraints and the Value of Insurance†

By Keith Marzilli Ericson and Justin Sydnor*

Insurance moves resources across both time and states. We study the consumption-smoothing benefits of insurance under liquidity constraints in a model where contracts span multiple consumption periods. The normative benchmarks for insurance demand under liquidity constraints differ qualitatively and quantitatively from the standard model: Individuals may only partially insure at actuarially fair prices, may benefit from insurance when premiums are very high and even sometimes when dominated, and may value insurance against events that will surely happen. Using simulations for health insurance, we highlight how these findings generate insights about how cost-sharing should be designed differently for liquidity-constrained populations. (*JEL* D86, G21, G51, G52)

Foundational insurance theories use a static expected utility model in which insurance contracts equate to different lotteries over terminal wealth (e.g., Arrow 1963; Rothschild and Stiglitz 1976). This standard model provides the framework for the majority of empirical studies of adverse selection (e.g., Einav, Finkelstein, and Cullen 2010; Handel 2013; Handel, Hendel, and Whinston 2015) and evaluations of the quality of insurance choices (e.g., Abaluck and Gruber 2011; Handel and Kolstad 2015; Bhargava, Loewenstein, and Sydnor 2017; Abaluck and Gruber 2023). The literature on the evaluation of insurance choices has highlighted that empirical patters systematically deviate from basic predictions of the standard model. Those findings help to motivate a desire to better understand and model the considerations that affect how people choose insurance and the value they get from insurance products.

In this paper, we address an important limitation of the standard model of insurance—namely, that it cannot address how insurance decisions translate to consumption shocks (Gollier 2003). It is empirically well established that many people face liquidity constraints of some sort. Aguiar, Bils, and Boar (2025) use data from the Panel Study of Income Dynamics and classify around 40 percent of American adults as living "hand-to-mouth" based on holding limited liquid assets. These individuals have high marginal propensities to consume, and there is substantial persistence

*Ericson: Boston University and NBER (email: kericson@bu.edu); Sydnor: University of Wisconsin and NBER (email: justin.sydnor@wisc.edu). The author order is alphabetical. Kareen Rozen was coeditor for this article. We thank Katsuhiro Komatsu for invaluable research assistance. We are also grateful to Jason Abaluck, Nikhil Agarwal, Doug Bernheim, Jacob Bor, Marika Cabral, Randy Ellis, Amy Finkelstein, Neale Mahoney, Corina Mommaerts, Michaela Pagel, Jim Rebitzer, and Jeremy Tobacman for helpful comments, as well as to workshop participants at various universities and the manuscript reviewers. This study was approved by the University of Wisconsin, Madison Education and Social/Behavioral Science IRB under protocol number 2015-1223.

†Go to https://doi.org/10.1257/mic.20240042 to visit the article page for additional materials and author disclosure statement(s).

in hand-to-mouth situations over time. A similar share of around 40 percent of Americans report that they would not be able to pay for an emergency $400 expense using only cash (Federal Reserve 2018). Of those living hand-to-mouth, about one-third appear to have near-zero levels of absolute assets, while the remainder carry more substantial illiquid assets, such as housing equity (Kaplan, Violante, and Weidner 2014; Aguiar, Bils, and Boar 2025). Moreover, many people with apparent access to assets and liquidity nevertheless act as if they consume out of whatever cash is available on hand, a phenomenon that has been dubbed the "liquid hand-to-mouth" (Olafsson and Pagel 2018; Hundtofte, Olafsson, and Pagel 2019). We show that the standard approach of modeling insurance as lotteries over final wealth and ignoring financial flows within the contract term can abstract from important considerations when people are unable to effectively smooth consumption.

We investigate the value of insurance in a dynamic consumption-utility framework that incorporates liquidity constraints.[1] We allow for insurance contracts that span multiple consumption periods—for example, a yearlong insurance contract for an individual who has monthly consumption periods. This feature of our framework distinguishes it from prior literature that has examined risk aversion and insurance demand in life cycle models with liquidity constraints when insurance contracts coincide with individual consumption periods (e.g., Gollier 1994, 2003). We take as given the empirical evidence that many people are liquidity constrained (or act *as if* they were liquidity constrained) and describe how such individuals should in theory value insurance.

We provide both formal theoretical results and numerical calibrations. Our formal theoretical results are developed for a cash-on-hand individual who can neither borrow nor accumulate savings. This represents the starkest case where the consumption-utility model will diverge from the standard static utility-of-wealth framework. The cash-on-hand case also simplifies the analysis by eliminating dynamic considerations about consumption and borrowing. While an extreme case, cash-on-hand behavior is arguably a reasonable representation for individuals who persistently fail to accumulate assets. We complement these theoretical results with numerical calibrations that allow us to examine more general cases involving different costs of borrowing and the ability to save.

Our first set of results demonstrates that the timing of when expenditures arrive within the policy period affects willingness to pay (WTP) for insurance (even when individuals are perfectly patient). For instance, losses can occur either smoothly over time (for instance, prescription drugs refilled each month) or concentrated in time (for instance, a single visit to an urgent care center). Liquidity-constrained individuals should value insurance against losses concentrated in time more than they value insurance against losses that occur smoothly across time. Turning to premiums, in some cases, premiums must be paid up front, but in others, premiums are paid smoothly throughout the policy period. We revisit Mossin's classic theorem (Mossin 1968), developed in the standard static framework, which states that

[1] While liquidity constraints are important for how insured people use insured services and thus make claims, this paper aims to contribute new insights about the baseline value of insurance in the absence of moral hazard concerns. For extensions to moral hazard, see Ericson, Jaspersen, and Sydnor (2025).

individuals will demand full insurance if and only if premiums are actuarially fair. We show that while the theorem extends to cash-on-hand individuals when premiums are paid smoothly over the policy period, it does not hold when premiums are paid up front. Up-front premiums generate a large consumption shock in the initial period and lead a cash-on-hand individual to desire less than full insurance even when premiums are actuarially fair. The result that individuals facing a liquidity constraint will demand less than full insurance when premiums are paid up front has previously been established by Liu and Myers (2016) (see Proposition 2), and there is a closely related result in Casaburi and Willis (2018) as well.[2] An implication of this result that has not been explored in prior literature, which we demonstrate in numerical simulations, is that individuals with liquidity constraints can have low demand and appear risk loving when offered insurance options that generate large one-time premium shocks.

Our second set of results highlights that insurance not only has the classic benefit of spreading risk across states of the world but also helps the liquidity-constrained smooth consumption across periods. As a result, the value of additional insurance can be much higher for those with liquidity constraints than for unconstrained individuals.[3] However, the effects of liquidity constraints are different than individuals simply becoming more risk averse. In fact, we provide a novel result that because of the consumption-smoothing benefit of insurance, there are conditions under which liquidity-constrained individuals should be willing to purchase insurance even when the premiums are so costly that insurance appears *dominated* from the perspective of the standard static model of insurance demand. Additionally, the consumption smoothing benefits of insurance lead liquidity-constrained individuals to desire to purchase insurance even against events that are sure to happen.

Our third set of results highlights that liquidity constraints affect the value of insurance contract designs. While the simplest insurance design is a full-insurance contract that removes all risk, this is rarely observed in practice because various costs and moral hazard concerns often make full coverage suboptimal. In the standard analysis of partial-insurance design, the key consideration for risk protection is how large could uninsured losses be in the worst case (Arrow 1963; Gollier and Schlesinger 1996). Individuals with liquidity constraints, however, are concerned not only with the total size of uninsured losses under the policy but also with how these uninsured shocks are borne over time. We investigate the significance of these differences for the design of health insurance policies, which has been the subject of considerable academic and policy interest. We use a large dataset on the health insurance claims of adults from Truven Marketscan (Merative 2018) to extract individual-level profiles of monthly medical spending realizations over

[2] Both of these papers discuss insurance demand in a developing-country context where the full cost of insurance is paid all at once. They show that the demand for insurance is higher if the premium can be timed to occur during a period of higher liquidity. Casaburi and Willis (2018) empirically find that the take-up of crop insurance is dramatically higher when premiums are paid at harvest time, and there are some related findings in Giné, Townsend, and Vickery (2008) and Cole et al. (2013). Our results show that premiums paid smoothly over the course of the policy period can provide a similar liquidity-matching role for the cash-on-hand agent and that the classic result by Mossin is obtained in this case.

[3] See also Dionne and Eeckhoudt (1984); Gollier (1994, 2003); and Casaburi and Willis (2018) for related models making this point.

the course of a year. We then simulate the value of health insurance plans with different cost-sharing designs involving deductibles, coinsurance, and maximum out-of-pocket limits for individuals who are "behind the veil of ignorance" and will have their medical spending drawn from one of these realized profiles. We compare these insurance values between perfectly liquid and cash-on-hand individuals.

The results differ by liquidity constraints in three important ways. First, cash-on-hand individuals would benefit from different plan designs than those with perfect liquidity. Consistent with classic theoretical results (Arrow 1963), holding fixed the expected cost of the contract, those with perfect liquidity prefer to limit the maximum uninsured spending risk with a "straight-deductible" design. Cash-on-hand individuals, however, get significant value from lowering the deductible and taking on a higher maximum-out-of-pocket limit. Even though they are exposed to a higher total spending risk, they benefit from limiting the size of their out-of-pocket shock at any moment in time. Second, because cash-on-hand individuals are constrained, the precise design of the contract matters much more to them. In the contract regions we explore, the value of different plan designs with the same expected cost varies by hundreds of dollars for the cash on hand but only by a few dollars for those with perfect liquidity. As such, the costs of implementing the "wrong" plan design are much larger for the cash on hand. Third, how losses are spread out over the policy year is important for the cash-on-hand individuals. The value of the designs we study changes dramatically for cash-on-hand individuals if losses arrived in a more lumpy fashion all at once in the year or if they were more smoothly spread out throughout the year. Complex nonlinear contracts have primarily been rationalized in prior literature as a compromise between risk protection and incentives to combat moral hazard (e.g., Zeckhauser 1970). Our analysis shows that even in the absence of moral hazard concerns, liquidity constraints offer another justification for these insurance designs. Our results also show, however, that the details of these nonlinear plan designs can have significant impacts on the value of insurance for those with liquidity constraints.

Our paper is motivated by recent empirical work showing that liquidity constraints affect take up of insurance (Giné, Townsend, and Vickery 2008; Cole et al. 2013; Casaburi and Willis 2018), as well as utilization and lapse in insurance contracts (Gottlieb and Smetters 2021; Gross, Layton, and Prinz 2022). Our study contributes to a broader literature discussing how liquidity constraints affect the value of insurance. A number of papers have highlighted that insurance is most valuable when people have limited assets and savings and that precautionary savings should substitute for costly market insurance (Dionne and Eeckhoudt 1984; Gollier 1994, 2003).[4] A related stream of literature on optimal unemployment insurance examines how social insurance affects consumption flows (Hansen and Imrohoroğlu 1992; Gruber 1997; Chetty 2006, 2008). In particular, Chetty (2008) emphasizes that unemployment insurance provides a liquidity benefit that raises the value of this insurance. Other work has also emphasized that the ability to borrow

[4] See also Hofmann and Peter (2016); Peter (2017); and Huber (2022) for related work on the link between savings and investments in reducing the likelihood or potential severity of losses, sometimes referred to as "self-insurance" and "self-protection,"

and save is a substitute for insurance (Handel, Hendel, and Whinston 2015; Malani and Jaffe 2018; Jaffe, Malani, and Reif 2026). Chetty and Szeidl (2007) further show that consumption commitments can raise the effective level of risk aversion because individuals cannot costlessly reallocate their consumption profiles in the face of shocks.

Our work complements these results and shares a general theme that strong liquidity constraints can heighten the value of insurance. Our approach differs from much of the prior literature, though, in considering insurance contracts that span across multiple underlying consumption periods and considering the timing of costs and shocks within those insurance contracts. This leads to new insights on the interaction of liquidity constraints and insurance: the importance of the way in which expenses are paid over time, the possibility for violations of classic notions of dominance, and new insights on the optimal design of insurance contracts.

Our findings also have implications for the normative benchmarks used in recent studies to evaluate the quality of insurance choices. Extreme WTP for insurance and violations of dominance have been documented and interpreted as evidence of mistakes in insurance choice (Sydnor 2010; Handel 2013; Bhargava, Loewenstein, and Sydnor 2017). In a series of papers, Abaluck and Gruber (2011, 2016, 2023) find that, on average, people violate two normative principles from the standard model of demand—namely, that (i) individuals should value premium reductions similarly to out-of-pocket cost reductions and that (ii) they should not value particular contract features, such as deductible, beyond their impact on the expected level and variance of total spending. We show that for those with strong liquidity constraints, high WTP—and even, in some cases, dominated choice—and a concern for contract features like deductibles that affect how loss arrives over time is not per se irrational.[5] This does not mean that liquidity constraints fully explain documented insurance "anomalies." Indeed, prior research clearly shows that people are confused about insurance, and some choice patterns relate to confusion (e.g., Johnson et al. 2013; Loewenstein et al. 2013; Handel and Kolstad 2015; Bhargava, Loewenstein, and Sydnor 2017; Samek and Sydnor forthcoming). The key point is that some behaviors that have been seen as clear signs of "mistakes" in prior literature have some underlying rational basis for those with liquidity constraints.[6] At the same time, confusion about insurance also means that people with liquidity constraints may not always recognize the potential consumption-smoothing impacts of insurance contracts and might also choose suboptimally for their situation.

It would be valuable for empirical work to distinguish between confusion and liquidity-related motivations for insurance choices. The approach to distinguishing the role of confusion versus liquidity constraints will differ with the setting but could

[5] Our approach shares some similarities to the consideration set literature, in which individuals choose optimally from a set, but do not consider all options in the set (e.g., Abaluck and Adams-Prassl 2021) and thus may appear to be making a mistake. Here, liquidity-constrained individuals choose optimally but do not have access to all the consumption paths that an economist would assume they have access to under perfect liquidity. In contrast to the consideration set approach, our approach makes a series of predictions about what choices will emerge, how people will respond to changes in credit, and how to design insurance plans in light of individuals' constraints.

[6] Variation in liquidity constraints may also explain part of the substantial heterogeneity insurance demand across individuals with similar expected losses (Finkelstein and McGarry 2006; Geruso 2017) and across domains (Barseghyan, Prince, and Teitelbaum 2011; Einav, Finkelstein, Pascu, and Cullen 2012).

include directly measuring liquidity situations, exploiting exogenous shocks to liquidity constraints, and interventions to clarify insurance trade-offs. Recognizing these different forces can be crucial for delivering effective policies. For example, interventions aimed at "nudging" individuals to make choices in line with standard normative benchmarks might harm individuals with liquidity constraints by inducing them to take on more risk than is optimal given the constraints they face on smoothing consumption. Our results provide a framework for accounting for liquidity constraints in establishing the normative benchmarks used to guide policy making in insurance.

I. Model

A. *Consumption Utility Model Setup*

We consider an individual who has standard separable discounted utility. Lifetime utility is given by

$$U = \sum_{t=0}^{T} \delta^t u(c_t), \tag{1}$$

where $0 < \delta \leq 1$ is the constant exponential discount factor, c_t is the level of consumption in period t, and u is the continuous instantaneous utility of consumption (with $u' > 0$ and $u'' < 0$). We assume $\delta = 1$ throughout to highlight the role of liquidity constraints apart from time preference; moreover, with the monthly time periods that we think are the most natural way to apply our model, δ should be near 1 (Ericson and Laibson 2019).

The individual lives for T periods and can purchase an insurance contract Z with a duration that spans N periods, with $T \geq N > 1$. For example, N might represent 12 months, which is common in health and property insurance markets, and the total lifetime T would be much larger than N.

In our baseline model, there is a single monetary loss L that may occur during the insurance contract duration.[7] Let the ex ante probability of the loss in each period t be given by π_t, which sums up to the probability π that the loss occurs at some point during the N periods covered by the insurance policy. We will typically assume that the ex ante probability of the loss is equal in any period ($\pi_t = \pi/N$). We assume that the loss size and probability are fixed, and so we set issues of moral hazard aside.

The insurance contract pays an indemnity $I(L)$ when a loss happens and zero otherwise. The insurance contract has a total cost of P, and the amount paid in each period is given by p_t. The timing of how that premium is paid can matter. We consider two cases. One possibility is for "smooth premiums" that are paid in equal

[7] For instance, consider insuring against the risk of a surgery. Surgery, if needed to happen, is equally likely to happen in any period. But once surgery happens, it won't need to happen again. The model's conclusions would not be substantially changed if the loss was i.i.d. and equally likely to happen in any period. However, the analytic expressions become more complicated and will depend on the form insurance takes. With multiple losses, insurance can take the form of a deductible for the first loss and full insurance thereafter. In this case, the likelihood of facing the first loss declines over time within the insurance contract period.

installments across the N contract periods, such that $p \equiv p_t = \frac{P}{N}$. The other possibility is that the total premium is due in a single period—typically, the first period $(p_1 = P)$—which we will refer to as an "up-front premium."

The individual makes consumption decisions subject to budget constraints. The individual earns constant income y in each period. We denote financial assets at time t as a_t. The individual earns gross interest on positive financial assets of R_s each period and incurs a gross borrowing interest rate of R_b on negative assets. We let $l_t \in \{0,1\}$ be an indicator function for whether the loss occurs in period t. Assets then evolve according to the following equation:

$$(2)\quad a_{t+1} = \begin{cases} R_s(a_t + y - p_t - c_t) - l_{t+1}(L - I(L)) & \text{if } a_t + y - p_t - c_t \geq 0 \\ R_b(a_t + y - p_t - c_t) - l_{t+1}(L - I(L)) & \text{if } a_t + y - p_t - c_t < 0 \end{cases}.$$

The individual chooses consumption each period to maximize the expected discounted utility flow, subject to the law of motion for assets. Consumption is chosen after observing whether the loss occurred in that period. This dynamic programming problem involves two state variables: assets and whether the loss has occurred yet.

It can be written as

$$(3)\qquad W_t\left(a_t \middle| \max_{\tau \in 0,\ldots,t} l_\tau\right) = \max_{c_t} u(c_t) + \delta E W_{t+1}\left(a_{t+1} \middle| \max_{\tau \in 0,\ldots,t+1} l_\tau\right),$$

subject to how assets evolve (equation (2)) and the constraint that end-of-life assets are not negative $(a_T \geq 0)$. Note that $\max_{\tau \in 0,\ldots,t} l_\tau$ indicates whether the loss has occurred by period t. The expectation for period $t+1$ is over whether the loss occurs in period $t+1$ since that affects assets in $t+1$.

This basic framework is flexible and allows us to explore different types of liquidity constraints. For example, we can consider a case where the individual can save but is not allowed to borrow by setting $R_b = \infty$. Setting both $R_b = \infty$ and $R_s = 0$ captures the cash-on-hand case that we consider extensively below, where the individual has no ability to borrow or accumulate assets and consumes whatever cash is on hand in that period.[8]

B. *Relation to the Classical Model of Insurance Demand*

Much of the insurance literature uses a static expected utility of wealth framework, in which insurance is modeled as determining lotteries over terminal wealth levels. Given our setup in the preceding subsection, in this classical model, the expected utility from an insurance contract can be written as

$$(4)\qquad V(Z, L, \pi) = \pi v(w_0 - P - L + I(L)) + (1 - \pi) v(w_0 - P).$$

[8] One might want to consider a limit on borrowing, such that $a_t \geq a_{min}$ in all periods, or a more general cost of borrowing that is convex in the amount borrowed. While these may be helpful in the empirical application of the model, they are not necessary for our results.

Notation is the same as the prior section, except in lieu of a consumption utility function u, we have v: the utility function over final wealth states, with $v' > 0$ and $v'' < 0$. We also need to define w_0, the initial wealth level. When this framework is applied, it is often unclear what the relevant level of wealth is, and many different assumptions are made in practice (e.g., annual or lifetime wealth).[9]

In general, it is not possible to map the consumption-maximization model into the classic static expected utility framework because the static formulation abstracts from issues of timing that are important for the dynamic problem. However, it can be instructive to consider some restrictive conditions under which the two frameworks coincide. Suppose for simplicity that the lifetime T is equal to the number of periods in an insurance contract N. The following proposition states that preferences for insurance in the consumption utility model can be represented with a static expected utility of wealth formulation under conditions we refer to as "perfect liquidity and perfect foresight".

PROPOSITION 1 (Perfect Liquidity and Perfect Foresight Necessary and Sufficient for Consumption Utility to Have Static Expected Utility Representation): *An individual with consumption utility u can have their preferences over all combinations of plans Z, loss sizes L, and probabilities π represented with the static expected utility formulation with an indirect utility function v*:

$$V(Z, L, \pi) = \pi v(w_L) + (1-\pi) v(w_N), \tag{5}$$

where $w_L = Ny - P - L + I(L)$ *and* $w_N = Ny - P$, *if and only if the following assumptions hold*:

(*i*) *Perfect liquidity*: $R_b = R_s = 1$.

(*ii*) *Perfect foresight*: *After the insurance decision is made, the individual gains perfect knowledge about the arrival and timing of losses prior to any consumption decisions.*

PROOF.

See Supplemental Appendix A.1.

If the assumptions of "perfect liquidity" are not met, it is clear that the static model cannot capture utility-relevant differences between consumption profiles induced by different insurance contracts that have the same final wealth states. More subtly, however, even if we assume perfect liquidity, the mapping that we lay out here only holds for the case where all uncertainty (over all future periods) is resolved before consumption is chosen in the first period. When the possibility of loss arises over multiple periods, there is an additional impact on consumption utility because the

[9] For instance, an annual model, with annual income, cost sharing, and premiums put in the utility function is used both by Finkelstein and McKnight (2008) when valuing the risk protection from Medicare and by Engelhardt and Gruber (2011) when evaluating Medicare Part D prescription drug insurance.

person cannot perfectly forecast lifetime resources, as they do not know whether the loss will occur at all. Due to that uncertainty, it will not generally be possible to perfectly smooth consumption over all periods, even under "perfect liquidity," unless one purchases a full-insurance contract.

With perfect liquidity and perfect foresight, after the individual learns whether the loss occurred, they choose consumption subject to the budget constraint that consumption must be less than w_L or w_N, and their utility function can be represented by some indirect utility of wealth function v. Since we assume perfect patience $(\delta = 1)$, an intuitive static expected utility formulation results. The individual smooths consumption equally over time and so consumes $\frac{w_L}{N}$ or $\frac{w_N}{N}$ when the loss does or does not happen, respectively. Then (rescaling by N), we have $V(Z, L, \pi) = \pi u\left(\frac{w_L}{N}\right) + (1 - \pi) u\left(\frac{w_N}{N}\right)$. However, the proposition does not rely on $\delta = 1$, and without perfect patience the indirect utility function v will not generally be the same as the consumption utility function u.

II. Cash-on-Hand Individuals and Proportional Insurance

This section develops formal results for cash-on-hand individuals choosing what level of proportional insurance to buy. The cash-on-hand individuals in our model can neither save nor borrow $(R_b = \infty, R_s = 0)$, so consumption decisions are completely dictated by available resources. This means that it is simple to represent the expected utility for the individual as a function of the insurance level selected.

We consider "proportional insurance" policies in which the insurance policy pays to the individual a share $\alpha \in [0, 1]$ of the loss L. That is, $I(L) = \alpha L$.[10] Proportional insurance with a fixed loss amount has been featured in the insurance literature dating back at least since Mossin (1968). The individual chooses their desired insurance level α^* from a menu of plans $\alpha \in [0, 1]$, where the premium for the chosen insurance contract covers expected costs times a proportional loading factor $\lambda \geq 1$ (constant across all levels of insurance). Thus, total premium is $P = \lambda \pi \alpha L$. Insurance is actuarially fair when $\lambda = 1$ and "unfair" when $\lambda > 1$.

First, we show a series of results on the timing of expenditures:

(i) Cash-on-hand individuals will appear "risk loving" and not purchase full insurance at actuarially fair rates if premiums are paid up front.

(ii) Cash-on-hand individuals will demand more insurance when losses arrive in a lumpy manner (e.g., all in a single month) than when the same annual loss occurs smoothly across consumption periods.

We then show that liquidity-constrained individuals derive a consumption-smoothing benefit from insurance, which gives additional results:

(i) Cash-on-hand individuals typically demand more insurance than the standard model would predict, especially when premiums are paid smoothly.

[10] With a single loss of fixed amount, this is equivalent to choosing a deductible of amount $(1 - \alpha)L$.

(ii) Cash-on-hand individuals can have positive demand for insurance even under dominated pricing, in which the price of insurance is at or above the potential payout for the insurance.

(iii) Cash-on-hand individuals can have positive demand for insurance even for losses that will happen for certain.

While we prove these results for cash-on-hand individuals, Section IIC shows with numeric examples that these insights apply more broadly and that they generalize to individuals who can borrow, but at high cost.

A. *The Importance of the Timing of Expenditures*

For a cash-on-hand individual, the utility of insurance will depend on how the premiums are paid. We consider two cases: smooth premiums and up-front premiums.

Smooth Premium Case.—Here, premiums are paid in equal installments $p = \frac{\lambda\pi\alpha L}{N}$ each consumption period. The cash-on-hand individual's consumption level is $c_L = y - \frac{\lambda\alpha\pi L}{N} - (1-\alpha)L$ when the loss does happen and is $c_N = y - \frac{\lambda\alpha\pi L}{N}$ when the loss does not happen. Their utility can then be represented by

$$V_{COH} = Nu(c_N) - \pi\left[u(c_N) - u(c_L)\right], \tag{6}$$

which highlights that in principle the individual consumes the no-loss consumption level every period but, with probability π, consumes the lower amount (accounting for losses) in a single period.

Up-Front Premium Case.—When premiums are instead paid in full during the first consumption period, there are four different levels of consumption. When the loss does not happen in the first period but premiums still must be paid, consumption is $c_{1,N} = y - \lambda\pi\alpha L$. When the loss does happen in the first period (along with premiums being due), consumption is $c_{1,L} = y - \lambda\pi\alpha L - (1-\alpha)L$. After the first period, premiums do not have to be paid. So consumption in any later period where there is a loss is $c_{2,L} = y - (1-\alpha)L$, while later periods without a loss have consumption equal to income: $c_{2,N} = y$.[11]

We begin by reconsidering one of the classic results in the theory of insurance demand, Mossin's (1968) theorem. Mossin's (1968) theorem states that a risk-averse individual will demand full insurance if and only if its price is actuarially fair. The following proposition shows that while this holds under liquidity constraints when premiums are paid smoothly, it does not hold when premiums are paid up front. We also refer readers to Liu and Myers (2016), who have an equivalent result in their Proposition 2 with a slightly different modeling setup.

[11] Then utility in the up-front premium case is given by $V_{Up} = (1-\pi)\left[u(c_{1,N}) + (N-1)u(c_{2,N})\right] + \frac{\pi}{N}\left[u(c_{1,L}) + (N-1)u(c_{2,N})\right] + \frac{(N-1)}{N}\pi\left[u(c_{1,N}) + (N-2)u(y) + u(c_{2,L})\right]$.

PROPOSITION 2 (Mossin's Theorem Holds for Cash-on-Hand Individuals with Smooth but Not Up-Front Premiums): *A cash-on-hand individual facing a smooth premium schedule will purchase full insurance* $\left(\alpha^* = 1\right)$ *if premiums are actuarially fair* $\left(\lambda = 1\right)$ *and will purchase less than full insurance for premiums greater than the actuarially fair level* $\left(\lambda > 1\right)$. *However*, *when facing up-front premiums*, *the cash-on-hand individual will purchase less than full insurance for both actuarially fair and unfair loads* $\left(\lambda \geq 1\right)$.

PROOF:

See Supplemental Appendix A.2.

The intuition for only purchasing partial insurance when premiums are paid up front is that if full insurance is purchased, then consumption will be lower in the first period $(y - \lambda\pi L)$ than in the remaining periods (y). As a result, reducing α from 1 will help to better smooth consumption. This is not to say that cash-on-hand individuals facing up-front premiums have a low value for insurance. As we will show later, they have very high demand for having *some* level of insurance. The point, though, is that they will no longer demand *full* insurance at fair prices.

A direct corollary of this result is that when cash-on-hand individuals face up-front premiums, they will sometimes appear "risk loving." As the individual approaches full insurance coverage, on the margin, additional insurance will be worth less to the individual than the expected cost of that insurance.

One might reasonably conjecture from Proposition 2 that the cash-on-hand individual will demand more insurance when premiums are paid smoothly than when they are paid up front. This will often be the case, and we prove in Supplemental Appendix A.2.1 that it will be true for individuals with constant absolute risk aversion (CARA) utility. However, it is not generally true because demand for proportional insurance can be non-monotonic when agents have decreasing absolute risk aversion (Schlesinger 2000). The Supplemental Appendix A.2.1 also shows an example with constant relative risk aversion (CRRA) utility where demand will be higher under up-front premiums than smooth premiums for a very risk-averse individual facing high insurance loads.

Although the predictions on the effect of premium timing on the level of demand for insurance can be somewhat subtle, there is a more basic point that being able to pay premiums smoothly rather than up front should benefit liquidity-constrained individuals. We discuss the value of insurance with smooth premiums more in the next subsection. In practice, many insurance markets provide opportunities for paying premiums smoothly. For example, this is a feature of most health insurance programs in the United States. In property insurance markets, e.g., automobile and home, it is common for insurers to offer both up front and monthly premium options and many insurers charge an additional fee for the monthly premium option.

Beyond our results on the timing of premium payments, the issue of expenditure timing is also relevant when considering the nature of losses that are being insured. The value of insurance for a liquidity-constrained individual will be higher when potential losses create concentrated shocks in a single consumption period than if they arrive in a more spread-out fashion. The following proposition formalizes this point.

PROPOSITION 3 (Higher Insurance Demand for "Lumpy" Loss Processes): *Suppose that the individual faces a probability π of a loss of size L that can arrive either in a single consumption period* (*lumpy losses*) *or as a series of equal losses in all consumption periods $l = \frac{L}{N}$* (*smooth losses*). *The cash-on-hand individual* (*facing either smooth or up-front premiums*) *will choose a strictly higher level of insurance when faced with lumpy losses than with smooth losses unless they choose a corner solution* ($\alpha^* = 0$ *or* 1) *in both cases.*

PROOF:
See Supplemental Appendix A.3.

The simple intuition for this result is that cash-on-hand individuals are not able to spread consumption shocks from uninsured losses across consumption periods. When losses arrive in a "lumpy" fashion, this creates an additional financing benefit to insurance. The practical impact of this point is that liquidity-constrained individuals should have a stronger demand for insurance against a more sudden loss process, like automobile accident damage, than against costs that are likely to be more ongoing and spread out, like some types of chronic medical diseases. More generally, the ability to pay for uninsured losses over time will be a partial substitute for the value of insurance to the liquidity-constrained individual.

B. *Consumption-Smoothing Benefit Creates Demand for Insurance*

In this subsection we show that for those with liquidity constraints, insurance provides an additional consumption-smoothing benefit on top of the traditional risk-transfer benefit. Insurance not only transfers resources from states of the world with no losses to states with losses, but it also helps to transfer resources across consumption periods in the event that a loss happens. This consumption-smoothing benefit is especially strong when insurance can be paid for with smooth premium schedules, but it also exists when premiums must be paid up front.

We compare the insurance demand of a cash-on-hand individual facing smooth premiums to that of an individual with perfect liquidity, foresight, and patience (that is, an individual with the standard model of insurance demand whose indirect utility function matches the consumption utility function.) With full insurance, both types of individuals have smooth consumption across periods and states. With partial insurance, an individual with perfect liquidity and foresight smooths the uninsured loss over all periods, while a cash-on-hand individual absorbs the uninsured loss all in a single period.

For the cash-on-hand individual, we decompose the value of going from uninsured ($\alpha = 0$) to fully insured ($\alpha = 1$) into the classic risk-transfer benefit plus a new benefit of consumption smoothing across time. Under full insurance, both types of individuals completely smooth consumption across periods and states, and the utility for the cash-on-hand individual (V_{COH}^{full}) is the same as that of the individual with perfect liquidity and foresight (call this V_{PLF}^{full}): Both are $Nu\left(y - \frac{\lambda \pi L}{N}\right)$.

With no insurance, the cash-on-hand individual's utility is given by $V_{COH}^{no} = (1-\pi)Nu(y) + \pi\left[(N-1)u(y) + u(y-L)\right]$ as, with probability π, they face the loss and bear it all in a single period. In contrast, the utility from no insurance with perfect liquidity and foresight is $V_{PLF}^{no} = N\left[(1-\pi)u(y) + \pi u\left(y - \frac{L}{N}\right)\right]$ as, with probability π, they face the loss and bear it distributed across N periods.

Then we can write the difference in utility under full insurance for the cash-on-hand individual as

$$V_{COH}^{full} - V_{COH}^{no} = \underbrace{\left(V_{PLF}^{full} - V_{PLF}^{no}\right)}_{\text{risk transfer}} + \underbrace{\left(V_{PLF}^{no} - V_{COH}^{no}\right)}_{\text{consumption smoothing}}. \tag{7}$$

The first term is the classic risk-transfer benefit of full insurance versus no insurance in the standard model of insurance demand. The second term is the consumption-smoothing benefit of insurance for the cash-on-hand individual. This reflects the value from insurance supplying a source of implicit financing that lets them spread uninsured losses costlessly across consumption periods. That second term for the consumption-smoothing benefit has a simple equation:

$$V_{PLF}^{no} - V_{COH}^{no} = \pi\left[\left(u(y) - u(y-L)\right) - N\left(u(y) - u\left(y - \frac{L}{N}\right)\right)\right]. \tag{8}$$

Since u is concave, equation (8) for the consumption-smoothing benefit of full insurance for the cash-on-hand individual is positive and linearly increasing in the probability of loss π.

While the cash-on-hand individual will not always demand full insurance, this decomposition illustrates the broader point that a cash-on-hand individual values insurance in part from the ability it provides to finance consumption smoothing across periods. This leads to our first key formal proposition in this subsection:

PROPOSITION 4 (Demand Is Higher for Cash-on-Hand Individuals with Smooth Premiums than in the Standard Model): *When insurance is priced with a positive load* ($\lambda > 1$), *an individual with perfect liquidity and foresight will optimally select a weakly lower level of proportional insurance* (α) *than a cash-on-hand individual who faces the smooth premium schedule. The inequality is strict whenever the cash-on-hand individual demands a positive level of insurance.*

PROOF:

See Supplemental Appendix A.4.

To illustrate this result, we can consider the first-order conditions in each situation and the associated linear approximations for optimal demand under a specific utility function. For the cash-on-hand case, the first-order condition from differentiating the indirect utility function from a contract with insurance level α (equation (6)) can be written as

$$\text{(9)} \qquad \frac{u'\big(c_L(\alpha^*)\big) - u'\big(c_N(\alpha^*)\big)}{u'\big(c_N(\alpha^*)\big)} = \frac{\lambda - 1}{1 - \frac{\lambda\pi}{N}}.$$

If we assume that the instantaneous utility function has the CARA form with absolute risk aversion parameter r, equation (9) yields the following linear approximation for the interior solution to the optimal level of insurance in the cash-on-hand case with smooth premiums:

$$\text{(10)} \qquad \alpha^*_{COH} \approx 1 - \frac{1}{rL}\frac{\lambda - 1}{1 - \frac{\lambda\pi}{N}}.$$

We see here that optimal insurance is increasing in risk aversion r, loss size L, and the number of consumption periods in the insurance term N. The equivalent first-order condition and CARA approximation for the perfect-liquidity-and-foresight case gives

$$\text{(11)} \qquad \alpha^*_{PLF} \approx 1 - \frac{N}{rL}\frac{\lambda - 1}{1 - \lambda\pi}.$$

Comparing equations (10) and (11) illustrates the result that the optimal level of insurance is smaller in the perfect-liquidity-and-foresight case since $N\frac{\lambda - 1}{1 - \lambda\pi} > \frac{\lambda - 1}{1 - \frac{\lambda\pi}{N}}$ when $N > 1$.

In fact, the consumption-smoothing benefits of insurance are significant enough for the cash-on-hand individuals that when losses are large enough, they will demand insurance even at "dominated" price levels. In our context, dominated pricing occurs when the proportional load on insurance equals or exceeds the reciprocal of the probability of a loss occurring $(\lambda \geq \frac{1}{\pi})$. When loads reach this level $(\lambda = \frac{1}{\pi})$, the premium for insurance $P_\alpha = \lambda\pi\alpha L = \alpha L$ becomes equal to the indemnity that insurance provides in the case of a loss (αL). It is trivial to show that in the standard model, there is no demand for insurance at these dominated prices. However, for the cash-on-hand individuals, when losses are large enough, there will be a demand for some level of insurance.

We thus have the following proposition.

PROPOSITION 5 (Cash-on-Hand Individuals Demand Insurance Even at Dominated Prices for Large Losses): *Assuming that there exists a consumption level* $\underline{c}$ *such that* $\lim_{c \to \underline{c}} u'(c) = \infty$, *under dominated pricing* $\left(\lambda = \frac{1}{\pi}\right)$, *there exists a loss* $\bar{L}$ *such that a cash-on-hand individual (facing either the smooth or up-front premium schedule) has positive demand for insurance* $(\alpha > 0)$ *if and only if* $L \geq \bar{L}$.

PROOF:

See Supplemental Appendix A.5.

The basic intuition of the proof starts by noting that if the individual purchases no insurance, they absorb the full loss in a single consumption period. If losses are large enough, the marginal utility will be much higher in the consumption period with the loss than it will be in other consumption periods. Shifting some consumption toward the period with the loss via insurance is valuable enough to justify even very high premium costs borne by other consumption periods.[12] For dominated insurance ($\lambda = \frac{1}{\pi}$) and CARA utility with coefficient of absolute risk aversion r, the condition for positive insurance demand is

$$L > \frac{1}{r}\ln\left(1 + \frac{\frac{1}{\pi} - 1}{1 - \frac{1}{N}}\right), \tag{12}$$

which highlights that the individual is more likely to have demand for insurance under dominated pricing when they have higher risk aversion r, have a greater number of consumption periods in the insurance term (N), or face a higher probability of loss (π).

The proof holds for arbitrarily low probability of loss π, so long as there is a loss large enough for the utility function to have arbitrarily high marginal utility.[13] Equation (8) shows that as π increases, the consumption smoothing benefit increases. Later, in numerical simulations, we will see that that dominated insurance demand can arise with even moderate losses if the probability of a loss π is high enough.

Proposition 5 shows that observations of dominance violations interpreted in prior literature as evidence for choice frictions and mistakes could also be consistent with optimizing behavior for those with liquidity constraints. We note, however, that the dominance here is from the perspective of annual wealth levels. Hand-to-mouth individuals in our model will still not violate dominance from the perspective of the paths of consumption flows. This provides an empirical way to identify choices that could not be consistent with our model but instead likely result from choice frictions: first-order stochastic dominance violations for paths of consumption flows.

A closely related proposition is that cash-on-hand individuals will have demand for insurance even against losses that happen with certainty. In contrast, we know that individuals in the standard static expected utility model (and, hence, individuals with perfect liquidity and foresight) will never purchase actuarially unfair insurance for events that are certain to occur.

PROPOSITION 6 (Cash-on-Hand Individuals Demand Insurance against Large Certain Losses): *Assume that there exists a consumption level* $\underline{c}$ *such that* $\lim_{c \to \underline{c}} u'(c) = \infty$ *and suppose that* $\pi = 1$ *and* $\lambda < N$. *There exists* $\bar{L}$ *such that cash-on-hand individuals have a positive demand for insurance* $(\alpha > 0)$ *if and only if* $L \geq \bar{L}$.

[12] The logic is especially clear for the smooth-premium case because the indemnity benefit accrues to a single consumption period, while the premium costs can be financed across periods. When premiums are paid up front, purchasing dominated insurance is essentially a transfer from first-period consumption to potentially support another period when the loss might occur.

[13] This property holds for CRRA utility functions as consumption approaches zero. For CARA utility functions, this property holds if consumption is allowed to go negative.

PROOF:

See Supplemental Appendix A.6.

While the risk-transfer benefits of insurance fall to zero as the probability of loss goes to one, the consumption smoothing benefit of insurance for the cash-on-hand individual is increasing in the probability of loss. As long as loads are not too high, the cash-on-hand individual will value insurance simply for this consumption smoothing benefit.

C. *Numerical Examples*

In this subsection, we use simple numerical examples to illustrate some of the theoretical results. Throughout, we assume an individual with a yearlong insurance policy with monthly consumption periods $(N = 12)$ and log utility $(u(c) = \ln(c))$. We assume that the individual has yearly income of $20,000, implying constant income of $y = \$1,667$ per period. This level of income is slightly above the 2024 US federal poverty level. We further assume that the individual faces a risk of a single potential loss $L = \$1,200$ that could occur during one period during the insured year. If uninsured, this loss size is equivalent to a loss of 72 percent of a month's income. As in our theoretical analysis above, we continue to assume no time discounting $(\delta = 1)$. For consistency with our theoretical results, in these numerical examples, we continue to explore proportional insurance contracts. We note, however, that given our assumption of a simple single possible loss size, all of the results below can also be recast as an analysis of insurance contracts with a straight deductible design.[14]

Optimal Proportional Insurance.—In Figure 1 we plot the optimal level of proportional insurance (α) for this numerical example across a range of insurance loads (λ). We show the optimal insurance curves as a function of loads for the three cases we have highlighted in our theoretical analysis: perfect liquidity and foresight (i.e., standard model), cash-on-hand with smooth premiums, and cash-on-hand with up-front premiums. In panel A, we assume the loss occurs with probability $\pi = 0.1$, while, in panel B, $\pi = 0.5$.

In both panels, demand for insurance for the cash-on-hand individual with smooth premiums is always higher than that of the perfect-liquidity-and-foresight case, except when insurance is fairly priced $(\alpha = 1$ when $\lambda = 1)$ and both types purchase full insurance. Demand is also always higher for the cash-on-hand individual with smooth premiums than the individual with up-front premiums in these examples. As loads increase, demand for insurance falls sharply for the individual with perfect liquidity and foresight but decreases only gradually for cash-on-hand individuals. For the lower probability of loss in panel A, the optimal demand falls to zero before loads are high enough such that pricing is dominated from the annual

[14] To be precise, with a fixed loss size L and a proportional insurance indemnity level α, the equivalent deductible d is $d = (1 - \alpha)L$. Thus, an insurance plan with $\alpha = 0.8$ is equivalent to having a deductible of $240 for this loss size of $1,200.

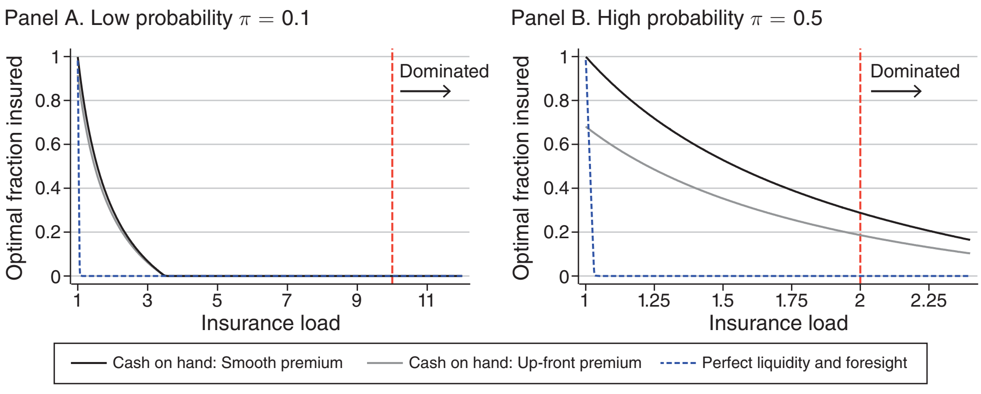

Figure 1. Optimal Insurance Demand by Load and Loss Probability

Notes: This figure shows the optimal level of insurance demanded (α) for three liquidity and premium-payment scenarios as a function of the load for insurance. The vertical dashed lines show the load value at which insurance pricing becomes dominated given the probability of loss. Parameters: Periodic income $y = \$1,667$, loss size $L = \$1,200$, and the individual has log utility $(u(c) = \ln(c))$. In panel A, the probability of a loss is set at 0.1, while in panel B it is set at 0.5.

perspective (denoted with the vertical dashed line). For the higher probability of loss in panel B, however, the cash-on-hand individuals desire positive levels of insurance at and above the dominated pricing level due to the strong consumption-smoothing benefit of insurance.

Allowing for Borrowing.—Our main theoretical results and the numerical example in the prior subsection focus on the cash-on-hand case. Here we move beyond cash on hand to consider individuals who face a range of different borrowing costs. We continue to assume the individuals enter the insurance policy period with zero assets. Except when we discuss the cash-on-hand case, we assume throughout that the agent faces no barrier or return to savings (gross interest rate on savings $(R_s = 1)$).

Now that we are analyzing borrowing at different interest rates, we make three modifications to our baseline assumptions. First, since we need to solve a dynamic programming problem to incorporate borrowing, we relax the simplifying assumption of "perfect foresight" used above and incorporate the fact that the timing of when a loss occurs (if at all) is stochastic. Second, to allow for costly borrowing to be spread over time in a natural way, we assume the agent lives for multiple years, setting the number of years to 20. To focus the analysis on the value of insurance for a single policy period, we assume that the individual faces the insurance choice in the first year and is fully insured in later years.[15] Third, we use parameter values

[15] The assumption of full insurance in later years simplifies the analysis. The value of assets in the current year is naturally affected by assumptions about the nature of risk and insurance exposure in future years. We have explored adding income risk, and the impact on our estimated WTP is small. For instance, we consider a scenario in which individuals face equal chances of having future income that is either 1.25 times or 0.75 times current income. WTP amounts for the smooth premium cases are different by less than \$3.

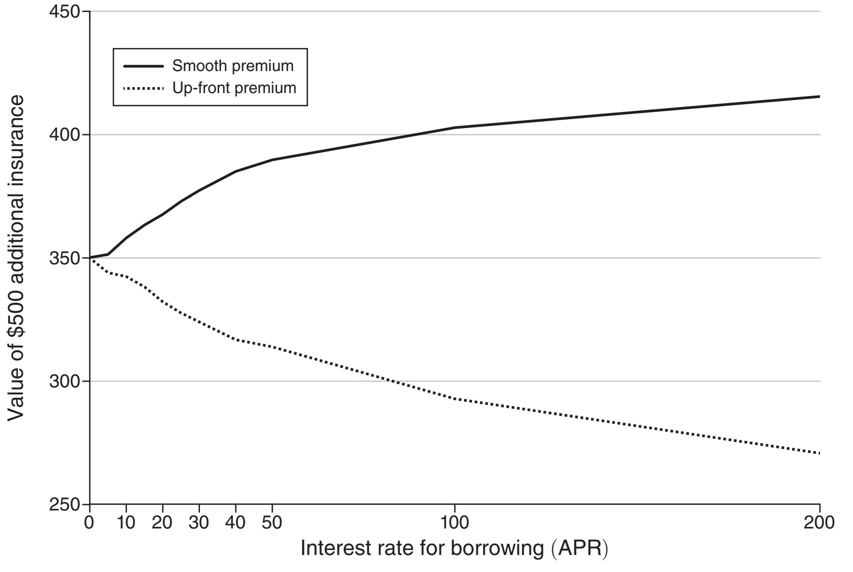

Figure 2. WTP for $500 Additional Indemnity by Borrowing Cost and Premium Timing

Notes: This figure presents the value of increasing the insurance indemnity by $500 (from $4,000 to $4,500). Parameters: Loss size $L = \$5,000$, 70 percent chance of loss (giving expected value of coverage difference $= \$350$), annual income $= \$20,000$, premium for the higher-coverage contract set at $4,000. $\delta = 1$, $R_s = 1$, and $u(c) = \ln(c)$.

that make this example more closely calibrated to the magnitude of costs for choices over health insurance contracts that have been studied in recent literature. To do so, we assume that the uninsured loss size is $L = \$5,000$.

We compare the value of a high-coverage proportional insurance contract $(\alpha = 90\%)$ to a contract with more moderate coverage $(\alpha = 80\%)$. Given our assumption that the loss size $L = \$5,000$, this equates to estimating the value of increasing the insurance indemnity by $500. To calculate the value of an insurance contract, we solve a dynamic consumption optimization problem to determine the ex ante expected utility value of insurance at different premiums given the agent's ability to borrow. We assume that the premium for the higher-coverage contract is set at $4,000. We then find the premium differential for the lower insurance coverage that equalizes the ex ante expected utility under the two contracts.

Figure 2 shows the premium difference that would make an individual indifferent between the two insurance policies if there is a 70 percent chance that the loss will happen.[16] The *x*-axis varies the level of borrowing costs, and the lines show the estimates for different ways that premiums are paid (smoothly versus up front). Borrowing interest rates plotted range from 0 (i.e., costless borrowing) to 200 percent annualized percentage rates, which starts to approach the level for those relying on subprime borrowing such as payday loans. The value of this additional insurance

[16] We choose 70 percent for this example because it is similar to the likelihood of experiencing losses in employer-sponsored health insurance settings.

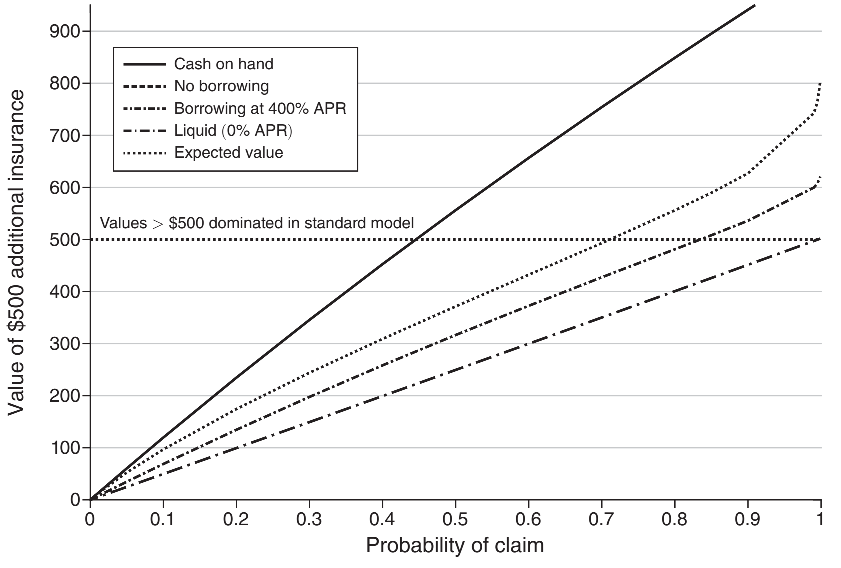

Figure 3. WTP for $500 Additional Indemnity by Loss Probability and Liquidity Constraints

Notes: This figure presents the value of increasing the insurance indemnity by $500 (from $4,000 to $4,500) with smooth premiums for different levels of liquidity constraints. Parameters: Loss size $L = \$5,000$, annual income is $20,000, 12 monthly consumption periods, premiums set at $4,000, $u(c) = \ln(c)$, no time discounting ($\delta = 1$), and no return or barrier to saving ($R_s = 1$), except in the cash-on-hand case where saving is not allowed ($R_s = 0$). Note that the "Expected value" line lies very close to the "Liquid (0% APR)" line.

is close to the expected value ($350) when the person can borrow costlessly from future income, regardless of the premium structure. However, at higher borrowing costs, the WTP for insurance diverges sharply depending on how premiums are paid. At subprime-level borrowing costs, the individual is willing to pay about 20 percent above the expected cost for the additional $500 insurance indemnity when premiums are paid smoothly, a substantial "risk premium". Yet that same individual is willing to pay only around 20 percent below expected cost for the additional coverage when premiums must be paid up front. In the up-front premium case, the premium generates a strong potential consumption shock in the first period that can only be smoothed out at high costs. As such, the additional insurance is unattractive, and it leads the individual to appear risk loving by having a WTP below the expected cost for the additional insurance coverage.

Figure 3 plots the value of the additional $500 insurance indemnity as a function of the probability of a loss under the assumption that premiums can be paid smoothly. Each line in the graph represents a different liquidity condition, spanning from no borrowing costs (perfect liquidity but not perfect foresight) to 400 percent APR borrowing costs to the cash-on-hand situation. The figure shows that individuals facing subprime-level borrowing costs (APR $=$ 400%) will pay significantly more than the expected value for the additional $500 of insurance and that this risk premium rises with the probability of a loss. For loss probabilities of around 0.80 and higher, the individual seemingly violates dominance by being willing to pay more than \$500 for an additional \$500 insurance indemnity. On the other hand, when it is costless to borrow, the individual's WTP is close to expected value at all points and does not violate dominance.

III. Valuing Insurance Contract Designs

A. *Introduction*

In this section we will examine how liquidity constraints affect individuals' valuation of the attributes of insurance contracts. While the simplest insurance design is a full-insurance contract that removes all risk for the individual, we rarely observe full insurance. Loading costs and inefficiencies can make it rational for people to buy only partial insurance. Moreover, full insurance can also create moral hazard problems.

There are a variety of ways that insurance products can be designed to offer partial coverage, including proportional insurance, deductibles with full coverage afterward, and "three arm" designs that give partial coverage after the deductible up to a maximum loss amount. Arrow (1963) derived a classic result that, in the absence of moral hazard concerns, the optimal insurance contract for a risk-averse individual will take the form of a "straight deductible" with full coverage above the deductible. Gollier and Schlesinger (1996) demonstrate the intuition for the optimality of deductibles by establishing that for the same level of expected coverage any other contract design will create a mean-preserving spread of total uninsured losses.

The Arrow (1963) result implies a normative benchmark for how individuals should value features of insurance design: They should be willing to pay more for a straight deductible contract than for a contract of the same actuarial value (percent of costs covered) but of a different design. However, empirical evidence from multiple studies suggests that people do not consistently value insurance features in accordance with the normative ideal, violating dominance (e.g., Bhargava, Loewenstein, and Sydnor 2017), purchasing too much insurance (e.g., Sydnor 2010), and valuing contract attributes over and above their impact on the distribution of out-of-pocket spending (e.g., Abaluck and Gruber 2011, 2016, 2023).[17]

In this section, we examine how liquidity constraints affect how insurance plan design impacts individual utility. We do not attempt to characterize a general set of normative benchmarks for liquidity-constrained choices, as these are likely to depend on the precise menu of insurance options and the nature of both underlying risk distributions and liquidity constraints in the population. The analysis in this section, though, provides a framework for the type of simulation that can be used to simulate the appropriate normative benchmark for a given setting.

Our analysis provides a few insights. A direct consequence of Proposition 1 is that the result of Arrow (1963)—that individuals without moral hazard concerns should prefer straight deductible plans to other designs—holds for individuals with

[17] Specifically, Abaluck and Gruber (2011, 2016, 2023) argued that people should not place weight on contract features (e.g., deductibles, copays) once premiums and the expectation and variance of out-of-pocket costs are controlled for. This result is exact for CARA utility and a normal distribution of out-of-pocket costs; it is otherwise an approximation. They find that people place significant weight on contract features above and beyond their impact on the expectation and variance of out-of-pocket costs.

perfect liquidity and foresight in our setup. However, this is not generally true for cash-on-hand individuals: They often will prefer an actuarially equivalent contract with a lower deductible, even though it will have higher total potential uninsured losses. Individuals value contract attributes that smooth the cost sharing across time. Compared to coinsurance, deductibles create larger uninsured shocks for losses early in the policy period (when there is no coverage) and offer more indemnity to later losses. The liquidity-constrained individual will be willing to accept some increase in total spending risk for these smoothing benefits.

Second, because cash-on-hand individuals are constrained, the precise design of the contract matters much more to them. While individuals with access to liquidity prefer different plans than cash-on-hand individuals, the losses from choosing the "wrong" plan are much larger for cash-on-hand individuals than for perfect-liquidity individuals (at least in the regions we consider).

Third, the distribution of losses matter. When all losses within the policy period occur within a single consumption period (i.e., all annual medical claims paid for in a single month), then the straight-deductible result of Arrow (1963) holds for a cash-on-hand individual. With such lumpy losses, the insurance design merely affects their distribution of consumption in the period in which the loss occurs. The gains to alternative insurance plan designs come when losses can happen in multiple consumption periods. With this distribution of losses, alternative insurance plan designs produce value for the liquidity constrained by smoothing losses across consumption periods.

B. *Setup*

We examine how individuals value plan design within the space of "three-armed" designs, a class of contracts commonly observed in health insurance markets. These designs involve a deductible, some partial coverage once the deductible is met (coinsurance), and a maximum limit on annual out-of-pocket costs for the insured. A policy with a straight deductible is a special case of the more general class of three-armed designs.

Our simulations consider an array of contracts, each with the same actuarial value (the fraction of losses covered by insurance), to focus on the valuation of contract attributes rather than the total amount of insurance. Our baseline contract has a straight deductible in which the individual is responsible for total losses up to the deductible and then is fully covered for losses in excess of the deductible. For this example, we set the straight-deductible at $3,000. We then consider a set of insurance contracts with a "three-arm design" with a deductible, a coinsurance rate of partial coverage after the deductible up to a maximum-out-of-pocket limit (d, c, m). We fix the out-of-pocket maximum $m = \$3,500$ in each case, and then for a grid of 20 values of d between $2,500 and $3,000, we solve for the coinsurance rate that delivers the same actuarial value as the straight-deductible contract. As such, each contract has the same level of expected coverage and so would be valued equivalently by a risk-neutral individual.

Our analysis sample comes from a large dataset of just over 11 million commercially insured individuals aged 24–64 in the Truven Marketscan (Merative 2018) database who were enrolled in coverage for the full 2018 calendar year.[18] For each individual, we obtain their total health care spending across inpatient, outpatient, and prescription drug coverage at the monthly level. We take a random 1 percent sample for computational reasons. This gives us one realized sequence of monthly health spending needs for each individual. Descriptive statistics for our 129,952 individuals in the analysis sample are presented in Supplemental Appendix Table B1.

We analyze the expected utility of an individual who faces the empirical distribution in this dataset of possible health spending flows over the course of a year. In this "ex ante" framework, the individual is behind a veil of ignorance, meaning that they do not know in advance what specific health risks or spending they will face. This allows us to assess how they would value different insurance contract designs before knowing their actual health outcomes.

For each simulated individual we calculate the expected utility of each insurance contract under two models of behavior developed previously: (i) our "perfect liquidity and foresight" model, in which the individual can costlessly borrow and save across periods within the year, and (ii) our cash-on-hand model, in which the individual consumes income net of monthly health care costs and can neither save nor borrow.

We also consider how the concentration or smoothness of expenditure across consumption periods affects preferences for contract attributes. In addition to the empirical distribution, we generate a "smoothed claims" distribution in which we take the observed total annual expenditure for each individual and treat it as though it were paid in equal installments each month. In the "lumpy claims" distribution, we take the observed total annual expenditure for each individual and treat it as though it were paid in one single month.[19]

Finally, while the data do not provide information on individuals' incomes, commercially insured individuals generally have higher income than individuals on government subsidized insurance, such as Medicaid. We thus assume that individuals have an annual income received equally across months that is three times the federal poverty limit for a household of two, which in 2018 was a value of \$49,380. We additionally assume they have CRRA of 3.

C. *Results*

Figure 4 shows the value individuals would have for contracts with deductibles lower than the \$3,000 baseline in different scenarios.[20] First, consider the results for the perfect-liquidity-and-foresight case, which is the nearly horizontal line in the graph. We see that the classic Arrow (1963) result holds: The individual prefers the straight-deductible contract to any of the equivalent-actuarial-value options with lower deductibles. Utility is monotonically decreasing in the deductible in this space

[18] We further clean the data by dropping individuals with negative recorded spending for a given day (for inpatient or outpatient spending) or month (for prescription drug spending).

[19] We implement it as a randomly chosen month, but since there is no time preference or dynamic choice in this model, the results are identical regardless of which single month it is assigned to.

[20] WTP here is defined the amount the individual would need to receive (or lose) in the base contract (paid in equal installments each month) to make their utility equal to the alternative contract.

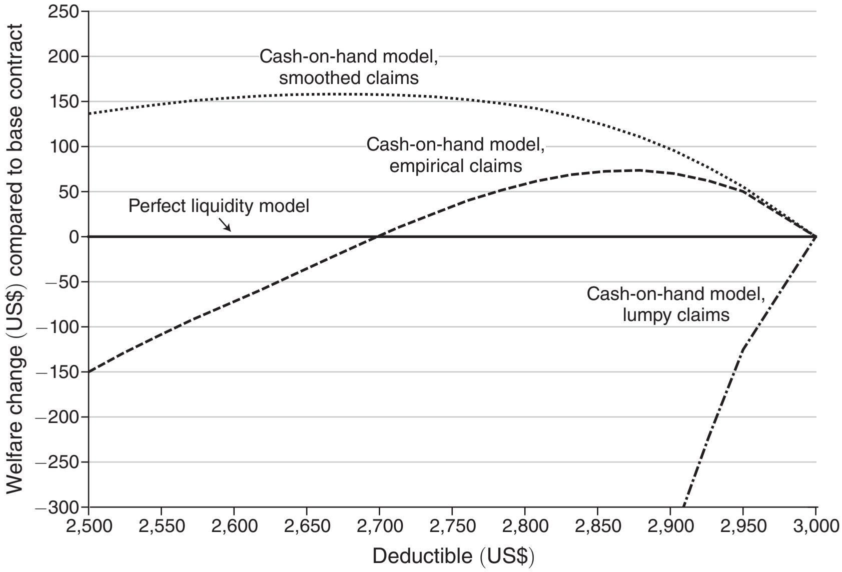

Figure 4. Value of Alternative Three-Arm Contract Designs

Notes: Shows WTP for the alternative three-arm design compared to straight deductible, as described in text. The WTP is calculated as the amount of additional spending or reduction in spending (if negative) spread out evenly across consumption periods that the individual would need to equalize utility between the alternative and straight-deductible designs. Assumed parameters: annual income \$49,380, discount factor $\delta = 1$, and CRRA $= 3$ utility. Source of claims data: 2018 Truven Marketscan (Merative 2018), limited to individuals 24–64 with commercial insurance, continuously enrolled for 12 months, 1 percent random sample. In the "lumpy claim" distribution, total annual claims are treated as though they were paid in a single month. In "smooth claims" distribution, claims are spread evenly across months.

of actuarially equivalent contracts. This is because, in this case, the Arrow (1963) theorem applies: With these assumptions, the problem can be represented as a decision over final wealth states (see Proposition 1). The scale of the loss from alternative plan designs is small, however—equivalent to less than a \$5 loss in annual income from choosing a deductible that is \$500 lower.[21]

In contrast, the cash-on-hand individual (with the empirical distribution of claims) prefers a contract with a lower deductible but positive coinsurance coverage. Among these contracts with constant actuarial value, the individual would pay about \$73 to have a deductible that is about \$100 lower than the base contract. Even though lowering the deductible to this level increases the risk exposure via a higher maximum out-of-pocket expense, the liquidity-constrained individual benefits from smoothing the expense across consumption periods.

Note also that the preferences of the cash-on-hand individual are nonmonotonic: They only benefit from a lower deductible up to a certain point. Afterward, the benefits of consumption smoothing are outweighed by the costs of the risk they are exposed to by higher coinsurance rates in these actuarially equivalent plans.

[21] The small differences in equivalent wealth values for those with perfect liquidity relates to the calibration theorem in Rabin (2000) since the scale of risk exposure over these different contract designs is modest.

Additionally, we see that cash-on-hand individuals are much more sensitive to the choice of deductible than individuals with perfect liquidity. While the worst contract in this space lowers the welfare of an individual with perfect liquidity by less than \$5, the impact of lower deductibles quickly grows large for cash-on-hand individuals, with a \$2,500 deductible lowering utility by about \$150.

Next, consider the alternative distributions of claims. The figure presents them only for the cash-on-hand individuals, as the results for individuals with perfect liquidity and foresight are invariant to the distribution of claims in the year.

When the claims distribution is smoothed so that claims come in equal amounts in every month, we see that cash-on-hand individuals receive greater benefits from lower deductibles. This is because the smaller but more frequent claims result in more smoothing of expenses across periods when the deductible is lowered, and the maximum out-of-pocket amount is less likely to be borne in a single consumption period.

Finally, for the lumpy claims distribution, in which all claims happen in a single consumption period (here, a month), the classic Arrow (1963) result holds: The individual again prefers the straight-deductible contract to any of the equivalent-actuarial-value options with lower deductibles. This is because, with this claims distribution and the cash-on-hand assumption, the decision problem is again reducible to a choice over wealth states in a single period. Moreover, because the cash-on-hand must bear this loss in a single consumption period (as opposed to spreading it across the year), utility very quickly falls as the deductible gets lower.

This simulation shows that liquidity constraints change how individuals value "three-arm" insurance designs. More generally, liquidity-constrained individuals will tend to want contract designs that do more to equalize spending across consumption periods. The results on how the value depends on the concentration or smoothness of the underlying loss distribution also highlights that developing different cost-sharing designs for different types of diseases can be valuable to the liquidity constrained.[22]

IV. Conclusion

This paper establishes the importance of accounting for liquidity constraints when evaluating insurance choices and assessing the value of insurance contracts. Insights about insurance emerge from our consumption-utility model that are not captured in standard expected utility of wealth models. Our model highlights new insurance market interventions that may be useful directions for future research. For instance, part of the demand for insurance when premiums can be paid smoothly under liquidity constraints comes from the consumption-smoothing benefit of insurance. Providing improved access to credit may reduce insurance demand in some settings (see also Handel, Hendel, and Whinston 2015; Malani and Jaffe 2018), and

[22] For example, many types of prescription drugs are taken routinely, and the results here suggest that three-arm designs with lower deductibles could be valuable in those cases. More acute needs that arise more rarely and in concentrated fashion, such as emergency room visits, may be better covered with deductibles or equivalently fixed copayments.

the design of savings vehicles, such as health savings accounts, may affect both utilization and insurance demand (Leive 2022). Improving access to and awareness of payment plans for medical bills may make people less averse to high-deductible health plans. Financing opportunities may enable high-deductible plans to be used more effectively to address overutilization of some medical services.

There are, however, some limitations to our analysis and some important areas for future research to better understand the links between liquidity constraints and insurance.

One issue is that in the consumption utility model, it matters when bills become due within the year. A natural question in practical applications is "when are bills actually paid?" We assume here that consumption reductions and borrowing coincide with when bills are generated. That assumption is likely reasonable for some types of insurance, where cost sharing must be paid before the service is rendered. For instance, for both home and auto insurance, contractors and mechanics typically won't make repairs without some payment up front. Yet the timing of payments is more complicated for some other insurance markets, such as health insurance. Some medical services require cost-sharing payment in advance of receiving care and likely fit our model assumptions well. A classic example is prescription drugs, which are typically paid for at the time the individual acquires the drug. Many physicians' offices require cost-sharing payment at the time that services are delivered, and this may be more strictly enforced in areas where patients are more likely to be a payment risk—precisely the liquidity constrained population that we are considering (Andrews 2016). For many other services, though—such as emergency room visits and hospitalizations—there may be more flexibility in how quickly bills must be paid. That flexibility creates an additional degree of freedom for a liquidity-constrained individual. Empirical applications of the consumption-utility model for these situations would ideally be paired with more information about the realities of bill payments and, most importantly, the beliefs that people have about their bill payment options.

Another direction for future research is to explore the link between liquidity constraints and moral hazard. Our analysis has abstracted from moral hazard to allow us to isolate important insights about how liquidity constraints interact with the ex ante value of risk protection. An analysis of moral hazard is beyond the scope of this paper. To the extent that liquidity constraints affect service utilization under insurance, they may change some of the welfare implications associated with those responses, as explored in Ericson, Jaspersen, and Sydnor (2025).

We have limited our analysis to optimizing choice in response to literal liquidity constraints. However, evidence suggests that people are prone to mental accounting (Thaler 1985) and treat assets as not fungible across accounts (e.g., Hastings and Shapiro 2013). That is, people may have access to a savings or retirement account but act *as though* it were not available to smooth unpredictable shocks and thereby reduce consumption rather than assets in response to shocks. Thus, mental accounting and related heuristics may lead people to act as if they are liquidity constrained, even if they could smooth consumption (Olafsson and Pagel 2018). A promising direction for future research is to examine the impact of mental accounting on consumption responses to insurance cost sharing.

Exploring the connection between liquidity constraints and behavioral biases—and distinguishing between the two—is likely to be a fruitful direction for research. In our model, individuals are fully optimizing and make no mistakes. However, observed patterns of choice may also result from misunderstanding and confusion. For instance, choosing a seemingly dominated plan may be a deliberate choice in response to liquidity constraints, may be a mistake, or may be a combination of both. Many empirical techniques exist to separate out liquidity constraints from other explanations of behavior, such as examining response to a liquidity shock (as in Casaburi and Willis 2018) or by providing decision aids to improve understanding (as in Samek and Sydnor forthcoming). In any empirical application, it will be valuable to distinguish between these two plausible explanations that have different implications for welfare. Moreover, even when decisions are partially a result of mistakes, it is important that welfare judgments accurately account for the liquidity constraints individuals face.

Finally, individuals may have incorrect beliefs (e.g., overconfidence) or present-bias or self-control issues. These behavioral biases may explain why individuals are liquidity constrained, as optimizing models suggest that they would strongly benefit from saving or reducing debt. Our model may provide a reasonable approach to modeling the demand for insurance for an agent with naïve present bias (O'Donoghue and Rabin 1999) who fails to accumulate assets because they perpetually delay saving. Naïve agents wrongly believe they will act like fully rational agents in the future. So the ex ante WTP for insurance should be the same for a temporarily liquidity-constrained rational agent and for a potentially persistently liquidity-constrained naïve-present-biased agent. For example, our model and the simulations in Figure 3 provides a guide to how extreme present bias would interact with liquidity constraints. Consider a fully naïve $\beta\delta$ discounter who had $\beta = 0$ but believed $\beta = 1$ and who had reached their borrowing limit. This individual would think that they would behave like someone who could save but not borrow (and so use that model of themselves to choose how much to pay for insurance) but would actually behave like a cash-on-hand individual once they made their choice. However, if behavioral biases rather than credit market constraints are leading people to act as if they are liquidity constrained, this would have implications for the impact of credit market interventions on insurance demand, which our theory predicts would be quite important. Testing the two would be valuable, as would incorporating models of behavioral biases into our framework with liquidity constraints.

REFERENCES

Abaluck, Jason, and Abi Adams-Prassl. 2021. "What Do Consumers Consider before They Choose? Identification from Asymmetric Demand Responses." *Quarterly Journal of Economics* 136 (3): 1611–63.

Abaluck, Jason, and Jonathan Gruber. 2011. "Choice Inconsistencies among the Elderly: Evidence from Plan Choice in the Medicare Part D Program." *American Economic Review* 101 (4): 1180–210.

Abaluck, Jason, and Jonathan Gruber. 2016. "Evolving Choice Inconsistencies in Choice of Prescription Drug Insurance." *American Economic Review* 106 (8): 2145–84.

Abaluck, Jason, and Jonathan Gruber. 2023. "When Less Is More: Improving Choices in Health Insurance Markets." *Review of Economic Studies* 90 (3): 1011–40.

Aguiar, Mark, Mark Bils, and Corina Boar. 2025. "Who Are the Hand-to-Mouth?" *Review of Economic Studies* 92 (3): 1293–340.

Andrews, Michelle. 2016. "Doctors and Hospitals Tell Patients: Show Us the Money before Treatment." *NPR,* December 7. https://www.npr.org/sections/health-shots/2016/12/07/504589131/doctors-and-hospitals-tell-patients-show-us-the-money-before-treatment.

Arrow, Kenneth J. 1963. "Uncertainty and the Welfare Economics of Medical Care." *American Economic Review* 53 (5): 941–73.

Barseghyan, Levon, Jeffrey Prince, and Joshua C. Teitelbaum. 2011. "Are Risk Preferences Stable across Contexts? Evidence from Insurance Data." *American Economic Review* 101 (2): 591–631.

Bhargava, Saurabh, George Loewenstein, and Justin Sydnor. 2017. "Choose to Lose: Health Plan Choices from a Menu with Dominated Options." *Quarterly Journal of Economics* 132 (3): 1319–72.

Casaburi, Lorenzo, and Jack Willis. 2018. "Time versus State in Insurance: Experimental Evidence from Contract Farming in Kenya." *American Economic Review* 108 (12): 3778–813.

Chetty, Raj. 2006. "A General Formula for the Optimal Level of Social Insurance." *Journal of Public Economics* 90 (10–11): 1879–901.

Chetty, Raj. 2008. "Moral Hazard versus Liquidity and Optimal Unemployment Insurance." *Journal of Political Economy* 116 (2): 173–234.

Chetty, Raj, and Adam Szeidl. 2007. "Consumption Commitments and Risk Preferences." *Quarterly Journal of Economics* 122 (2): 831–77.

Cole, Shawn, Xavier Giné, Jeremy Tobacman, Petia Topalova, Robert Townsend, and James Vickery. 2013. "Barriers to Household Risk Management: Evidence from India." *American Economic Journal: Applied Economics* 5 (1): 104–35.

Dionne, Georges, and Louis Eeckhoudt. 1984. "Insurance and Saving: Some Further Results." *Insurance: Mathematics and Economics* 3 (2): 101–10.

Einav, Liran, Amy Finkelstein, and Mark R. Cullen. 2010. "Estimating Welfare in Insurance Markets Using Variation in Prices." *Quarterly Journal of Economics* 125 (3): 877–921.

Einav, Liran, Amy Finkelstein, Iuliana Pascu, and Mark R. Cullen. 2012. "How General Are Risk Preferences? Choices under Uncertainty in Different Domains." *American Economic Review* 102 (6): 2606–38.

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

Ericson, Keith Marzilli, Johannes G. Jaspersen, and Justin R. Sydnor. 2025. "Moral Hazard under Liquidity Constraints." NBER Working Paper 33648.

Ericson, Keith Marzilli, and David Laibson. 2019. "Intertemporal Choice." In *Handbook of Behavioral Economics: Applications and Foundations*, Vol. 2, edited by B. Douglas Bernheim, Stefano DellaVigna, and David Laibson, 1–67. Elsevier.

Ericson, Keith Marzilli, and Justin Sydnor. 2026. *Data and Code for: "Liquidity Constraints and the Value of Insurance."* Nashville, TN: American Economic Association; distributed by Inter-university Consortium for Political and Social Research, Ann Arbor, MI. https://doi.org/10.3886/E237008V1.

Federal Reserve. 2018. *Report on the Economic Well-Being of U.S. Households in 2017.* Board of Governors of the Federal Reserve System.

Finkelstein, Amy, and Kathleen McGarry. 2006. "Multiple Dimensions of Private Information: Evidence from the Long-Term Care Insurance Market." *American Economic Review* 96 (4): 938–58.

Finkelstein, Amy, and Robin McKnight. 2008. "What Did Medicare Do? The Initial Impact of Medicare on Mortality and Out of Pocket Medical Spending." *Journal of Public Economics* 92 (7): 1644–68.

Geruso, Michael. 2017. "Demand Heterogeneity in Insurance Markets: Implications for Equity and Efficiency." *Quantitative Economics* 8 (3): 929–75.

Giné, Xavier, Robert Townsend, and James Vickery. 2008. "Patterns of Rainfall Insurance Participation in Rural India." *World Bank Economic Review* 22 (3): 539–66.

Gollier, Christian. 1994. "Insurance and Precautionary Capital Accumulation in a Continuous-Time Model." *Journal of Risk and Insurance* 61 (1): 78–95.

Gollier, Christian. 2003. "To Insure or Not to Insure?: An Insurance Puzzle." *Geneva Papers on Risk and Insurance Theory* 28 (1): 5–24.

Gollier, Christian, and Harris Schlesinger. 1996. "Arrow's Theorem on the Optimality of Deductibles: A Stochastic Dominance Approach." *Economic Theory* 7 (2): 359–63.

Gottlieb, Daniel, and Kent Smetters. 2021. "Lapse-Based Insurance." *American Economic Review* 111 (8): 2377–416.

Gross, Tal, Timothy J. Layton, and Daniel Prinz. 2022. "The Liquidity Sensitivity of Healthcare Consumption: Evidence from Social Security Payments." *American Economic Review: Insights* 4 (2): 175–90.

Gruber, Jonathan. 1997. "The Consumption Smoothing Benefits of Unemployment Insurance." *American Economic Review* 87 (1): 192–205.

Handel, Ben, Igal Hendel, and Michael D. Whinston. 2015. "Equilibria in Health Exchanges: Adverse Selection versus Reclassification Risk." *Econometrica* 83 (4): 1261–313.

Handel, Benjamin R. 2013. "Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts." *American Economic Review* 103 (7): 2643–82.

Handel, Benjamin R., and Jonathan T. Kolstad. 2015. "Health Insurance for 'Humans': Information Frictions, Plan Choice, and Consumer Welfare." *American Economic Review* 105 (8): 2449–500.

Hansen, Gary D., and Ayşe Imrohoroğlu. 1992. "The Role of Unemployment Insurance in an Economy with Liquidity Constraints and Moral Hazard." *Journal of Political Economy* 100 (1): 118–42.

Hastings, Justine S., and Jesse M. Shapiro. 2013. "Fungibility and Consumer Choice: Evidence from Commodity Price Shocks." *Quarterly Journal of Economics* 128 (4): 1449–98.

Hofmann, Annette, and Richard Peter. 2016. "Self-Insurance, Self-Protection, and Saving: On Consumption Smoothing and Risk Management." *Journal of Risk and Insurance* 83 (3): 719–34.

Huber, Tobias. 2022. "Comparative Risk Aversion in Two Periods: An Application to Self-Insurance and Self-Protection." *Journal of Risk and Insurance* 89 (1): 97–130.

Hundtofte, Sean, Arna Olafsson, and Michaela Pagel. 2019. "Credit Smoothing." NBER Working Paper 26354.

Jaffe, Sonia, Anup Malani, and Julian Reif. 2026. "Access to Credit Reduces the Value of Insurance." *Journal of Public Economics* 258: 105633.

Johnson, Eric J., Ran Hassin, Tom Baker, Allison T. Bajger, and Galen Treuer. 2013. "Can Consumers Make Affordable Care Affordable? The Value of Choice Architecture." *PLOS ONE* 8 (12): e81521.

Kaplan, Greg, Giovanni L. Violante, and Justin Weidner. 2014. "The Wealthy Hand-to-Mouth." *Brookings Papers on Economic Activity* 45: 77–153.

Leive, Adam. 2022. "Health Insurance Design Meets Saving Incentives: Consumer Responses to Complex Contracts." *American Economic Journal: Applied Economics* 14 (2): 200–27.

Liu, Yanyan, and Robert J. Myers. 2016. "The Dynamics of Microinsurance Demand in Developing Countries under Liquidity Constraints and Insurer Default Risk." *Journal of Risk and Insurance* 83 (1): 121–38.

Loewenstein, George, Joelle Y. Friedman, Barbara McGill, Sarah Ahmad, Suzanne Linck, Stacey Sinkula, John Beshears, et al. 2013. "Consumers' Misunderstanding of Health Insurance." *Journal of Health Economics* 32 (5): 850–62.

Malani, Anup, and Sonia P. Jaffe. 2018. "The Welfare Implications of Health Insurance." NBER Working Paper 24851.

Merative. 2018. *MarketScan® Research Databases: Proprietary Health Insurance Claims Data.* Merative. https://www.merative.com/healthcare-analytics.

Mossin, Jan. 1968. "Aspects of Rational Insurance Purchasing." *Journal of Political Economy* 76 (4): 553–68.

O'Donoghue, Ted, and Matthew Rabin. 1999. "Doing It Now or Later." *American Economic Review* 89 (1): 103–24.

Olafsson, Arna, and Michaela Pagel. 2018. "The Liquid Hand-to-Mouth: Evidence from Personal Finance Management Software." *Review of Financial Studies* 31 (11): 4398–446.

Peter, Richard. 2017. "Optimal Self-Protection in Two Periods: On the Role of Endogenous Saving." *Journal of Economic Behavior and Organization* 137: 19–36.

Rabin, Matthew. 2000. "Risk Aversion and Expected-Utility Theory: A Calibration Theorem." *Econometrica* 68 (5): 1281–92.

Rothschild, Michael, and Joseph Stiglitz. 1976. "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information." *Quarterly Journal of Economics* 90 (4): 629–49.

Samek, Anya, and Justin Sydnor. Forthcoming. "The Impact of Consequence Information on Insurance Choice." *Review of Economics and Statistics*.

Schlesinger, Harris. 2000. "The Theory of Insurance Demand." In *Handbook of Insurance*, edited by Georges Dionne, 131–51. Springer Netherlands.

Sydnor, Justin. 2010. "(Over)insuring Modest Risks." *American Economic Journal: Applied Economics* 2 (4): 177–99.

Thaler, Richard. 1985. "Mental Accounting and Consumer Choice." *Marketing Science* 4 (3): 199–214.

Zeckhauser, Richard. 1970. "Medical Insurance: A Case Study of the Tradeoff between Risk Spreading and Appropriate Incentives." *Journal of Economic Theory* 2 (1): 10–26.