

Liquidity Constraints and the Value of Insurance

Keith Marzilli Ericson ^{*} Justin Sydnor ^{†‡}

September 21, 2022

Abstract

We explore how liquidity constraints affect the value of insurance using a model where insurance contracts span multiple consumption periods. Limits to borrowing create preferences over how insurance contracts move money across time as well as states. Rational liquidity-constrained individuals will demand insurance in ways that depart from the standard model's normative benchmarks. Liquidity-constrained individuals may not buy full insurance even at actuarially-fair prices. Yet because insurance can have a consumption-smoothing benefit, they may purchase some insurance even when premiums are so high that insurance is dominated in the standard model. Liquidity-constrained individuals will not treat all expected-out-of-pocket costs and premiums similarly, and can benefit from cost-sharing designs that are more complex than straight deductibles. Because some puzzling patterns of insurance demand can be consistent with optimal choices under liquidity constraints, the model highlights the importance of considering liquidity when studying insurance demand.

Keywords: Insurance, liquidity constraints, cost-sharing design, risk aversion

^{*}Boston University Questrom School of Business and NBER. kericson@bu.edu

[†]University of Wisconsin, Madison Wisconsin School of Business and NBER. justin.sydnor@wisc.edu

[‡]We thank Katsuhiko Komatsu for invaluable research assistance. We are also grateful to Jason Abaluck, Nikhil Agarwal, 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 workshop participants at Boston University, the BU-Harvard-MIT Health Seminar, the National Bureau of Economic Research Summer Institute, the University of Pennsylvania's Behavioral Economics and Health Symposium, the University of Texas, and the University of Wisconsin, Madison.

1 Introduction

The foundational theory on insurance uses 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 a majority of empirical studies of insurance markets, including on adverse selection (e.g. Einav et al. 2010; Handel 2013; Handel et al. 2015) and on evaluating how well people choose from insurance options (e.g. Abaluck and Gruber 2011; Handel and Kolstad 2015; Bhargava et al. 2017; Abaluck and Gruber 2020).

A limitation of the standard model of insurance is that it cannot address how insurance decisions translate to consumption shocks (Gollier 2003). Modeling insurance as lotteries over final wealth and ignoring financial flows within the contract term can be misleading, as many people face liquidity constraints due to low assets and limited or high-cost borrowing options. Around 40 percent of American adults are classified as living “hand-to-mouth” (Aguiar et al. 2020) and report that they would not be able to pay for an emergency \$400 expense using cash or its equivalents (Federal Reserve 2018).¹ Even for people with substantial assets, there can be difficult-to-observe frictions that make the effective adjustment costs for small uninsured shocks high (Chetty and Szeidl 2007). Moreover, many people with apparent access to assets and liquidity nevertheless act as if they consume out of whatever cash is available on hand (Olafsson and Pagel 2018; Hundtofte et al. 2019).

We investigate the demand for insurance in a dynamic consumption-utility framework that incorporates liquidity constraints.² We allow for insurance contracts that span multiple consumption periods: for example, a year-long insurance contract for an individual who has monthly consumption periods. This feature of our framework distinguishes it from prior

¹Periods of temporary liquidity constraints are common for lower and middle income Americans due to combinations of income volatility and spending shocks (Zeldes 1989; Morduch and Schneider 2017). Moreover, people show strong consumption responses around predictable timing of receipt of public benefits (Shapiro 2005; Gross et al. 2022), suggesting that low assets or liquidity constraints can be persistent states.

²While liquidity constraints are important for how insured people use insured services and thus make claims, we aim to contribute new insights about the baseline demand for insurance in the absence of moral hazard concerns. We leave extensions to moral hazard for future work.

literature that has examined risk aversion and insurance demand in lifecycle 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 would theoretically 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. These calibrations allow us to examine the quantitative importance of liquidity constraints for insurance demand in situations similar to those studied in recent literature.

Our first result is that the timing of when expenditures arrive within the insurance-policy period matters for those with liquidity constraints. For example, in some cases premiums must be paid up-front at the start of the policy period, but in other cases premiums are paid monthly throughout the policy period. We revisit Mossin's classic theorem (Mossin 1968), developed in the standard static framework, which states that 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. An implication of this result, 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. These results relate closely to Casaburi and Willis (2018) who find that take-up of crop-insurance in a developing country is dramatically higher if premiums are paid at harvest time when farmers have ample liquidity.³ The insight that timing of expenditures matters also extends to the types of losses being insured. For example, liquidity-constrained individuals will have lower demand for insurance against risks that are paid smoothly, such as prescription drugs refilled each month, than equivalent expenses that must be paid at a single time.

Our second result is 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, demand for insurance can be much higher for those with liquidity constraints than for unconstrained individuals. In fact, we show that because of the consumption-smoothing benefit of insurance, there are conditions under which liquidity constrained individuals will 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.

Our third result is that liquidity constraints affect the optimal design of insurance contracts. We revisit Arrow (1963)'s classic result that the optimal insurance contract in the absence of moral-hazard concerns will be a “straight deductible” in which the individual bears all losses below a deductible amount and is fully insured after the deductible is met. We prove that this result extends to the cash-on-hand case if the contract can be designed to match consumption periods (i.e., straight deductible for each period) or if there is at most one possible loss that can occur during the insurance term.⁴ However, in the common setting in which insurance contracts span multiple consumption periods and multiple loss events

³Liu and Myers (2016) have a related dynamic model of microinsurance demand where demand is low when premiums are due up front and high when due at the end of an earnings cycle.

⁴See also Hong and Mommaerts (2021) who examine health insurance and demonstrate that liquidity-constrained individuals benefit if deductibles are reset at more frequent intervals, and Remmerswaal et al. (2019) who examine a rebate-based cost-sharing design as an alternative to a deductible.

can arise, straight-deductible designs will not be optimal for the cash-on-hand individual. Liquidity-constrained individuals will benefit from alternative contract designs, such as the three-armed designs with coinsurance that are popular in health insurance, because they can reduce the potential for large uninsured losses early in the contract period. Complex non-linear 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, however, that even in the absence of moral-hazard concerns, liquidity constraints offer another reason for these insurance designs.

These results have implications for the normative benchmarks used in recent studies to evaluate the quality of insurance choices. Extreme willingness to pay for insurance and violations of dominance have been documented and interpreted as evidence of mistakes in insurance choice (Sydnor (2010); Handel (2013); Bhargava et al. (2017)). While prior research clearly shows that some of these decisions relate to confusion, liquidity constraints provide an alternative explanation for this seemingly extreme demand for insurance. For a liquidity-constrained individual high willingness-to-pay, and even in some cases dominated choice, is not per se irrational. Consistent with this point, we provide new survey evidence that controlling for age, health status, income, and education levels, people who report being liquidity constrained are more likely to say they would purchase dominated insurance contracts and find arguments in favor of the liquidity-smoothing benefits of doing so more compelling.

We also examine the implications of our model for common approaches to detecting insurance-choice anomalies. Abaluck and Gruber (2011) established a discrete-choice modeling approach that compares choices patterns to intuitive normative benchmarks for optimal choice, namely: a) individuals should value a one dollar reduction in premium equally to a one dollar reduction in expected out-of-pocket costs and b) they should not value particular contract features, such as deductibles, beyond their impact on the expected level and variance of total spending. Abaluck and Gruber (2011, 2016, 2020) show that on average people

violate both of these normative principles, and interpret this as evidence for mistakes. Drawing on a large dataset of health insurance claims (Truven Marketscan), we re-examine this normative benchmark for agents with liquidity constraints. We model the optimal choices of cash-on-hand individuals who vary in their ex-ante distribution of health expenditures. We then estimate the same choice models as used by Abaluck and Gruber (2011), and show that liquidity constrained individuals would systematically violate the normative benchmarks in the same ways that have been documented and interpreted as mistakes in past studies. These results show that it is important to account for liquidity constraints when establishing the normative benchmarks used to evaluate insurance choice quality.⁵ Furthermore, some caution may be warranted for interventions aimed at “nudging” individuals to make choices in line with standard normative benchmarks given that those choices may not be optimal for those with liquidity constraints.

Related Literature Our paper is motivated by recent work showing that liquidity constraints affect take up of insurance (Casaburi and Willis 2018), as well as utilization in insurance contracts (Gross et al. 2022). Our study contributes to a broader literature discussing how liquidity constraints affect the value of insurance. Gollier (1994, 2003) develops lifecycle models of consumption and savings incorporating insurance. These models highlight that insurance is most valuable when people have limited assets and savings and that precautionary savings should substitute for costly market insurance. A related stream of literature on optimal unemployment insurance examines how social insurance affects consumption flows (Hansen and Imrohoroglu 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 and

⁵To be clear, our results do not imply that prior observations of seeming mistakes are necessarily caused by optimal choice under liquidity constraints. In particular, there is substantial evidence that people are confused about insurance and that this confusion is part of the explanation for choice patterns (e.g., Johnson et al. (2013); Loewenstein et al. (2013); Handel and Kolstad (2015); Bhargava et al. (2017); Samek and Sydnor (2020)). However, 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.

save is a substitute for insurance (Handel et al. 2015; Malani and Jaffe 2018). 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.

There are naturally also limitations to our analysis. In our concluding section we discuss some of these limitations and give thoughts on directions that would be valuable for future research, including the micro-foundations and behavioral biases behind persistent liquidity constraints, the timing of when uninsured costs must truly be paid, and the link between liquidity constraints and how spending reacts to insurance (i.e., moral hazard).

2 Model

2.1 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, δ should be near 1 over the monthly time periods (Ericson and Laibson 2019) that we think is the most natural application of our model.

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. 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 often assume that the ex-ante probability of the loss is equal in any period ($\pi_t = \pi/N$).

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 premiums to be paid smoothly in equal 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 the “up-front premium case”.

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:

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

subject to the constraint that end-of-life assets are not negative ($a_T \geq 0$).

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:

$$W_t(a_t | \max_{\tau \in 0, \dots, t} l_\tau) = \max_{c_t} u(c_t) + \delta E W_{t+1}(a_{t+1} | \max_{\tau \in 0, \dots, t+1} l_\tau), \quad (3)$$

where $\max_{\tau \in 0, \dots, t} l_\tau$ 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 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.⁶

2.2 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:

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

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

⁶One 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.

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).⁷

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 EU representation). *An individual with consumption utility u can have their preferences represented with the static expected utility formulation with 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:

1. *perfect liquidity: $R_b = R_s = 1$*
2. *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 Appendix A.1. □

⁷For instance, an annual model, with annual income, cost-sharing, and premiums put in the utility function is used by both Finkelstein and McKnight (2008) when valuing the risk protection from Medicare, and Engelhardt and Gruber (2011) when evaluating Medicare Part D prescription drug insurance.

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 we lay out here only holds for the case where all uncertainty (over all future periods) is resolved before the time that consumption is chosen in the first period. When the possibility of loss arises over multiple periods, there is an additional impact on the consumption-utility process because the person cannot perfectly forecast lifetime resources because she does not know for sure if 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 $V(Z, L, \pi) = \pi u(\frac{w_L}{N}) + (1 - \pi)u(\frac{w_N}{N})$. 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 .

3 Cash-on-hand Individuals and Proportion 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 ($R_b = \infty, R_s = 0$). For a cash-on-hand individual, 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, $I(L) = \alpha L$). 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 λ (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:

1. Cash-on-hand individuals will appear “risk loving” and not purchase full insurance at actuarially-fair rates if premiums are paid up front.
2. 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.

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

1. Cash-on-hand individuals typically demand more insurance than the standard model would predict, especially when premiums are paid smoothly.
2. 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.
3. 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 4 shows with numeric examples that these insights apply more broadly, and they generalize to individuals who can borrow, but at high cost, and to other insurance designs.

3.1 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 upfront 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 when the loss does happen is $c_L = y - \frac{\lambda\alpha\pi L}{N} - (1 - \alpha)L$, and when the loss does not happen is $c_N = y - \frac{\lambda\alpha\pi L}{N}$. Their utility can then be represented by:

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

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: $c_{2,N} = y$.⁸

We begin by reconsidering one of the classic results in the theory of insurance demand, Mossin's (1968) Theorem. Mossin's 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.

⁸Then, utility in the up-front premium case is given by: $V_{Up} = (1 - \pi)[u(c_{1,N}) + (N - 1)u(c_{2,N})] + \frac{\pi}{N}[u(c_{1,L}) + (N - 1)u(c_{2,N})] + \frac{(N-1)}{N}\pi[u(c_{1,N}) + (N - 2)u(y) + u(c_{2,L})]$.

Proposition 2 (Mossin’s Theorem holds for cash-on-hand individuals with smooth but not upfront premiums). *A cash-on-hand individual facing a smooth premium schedule will purchase full insurance ($\alpha^* = 1$) if premiums are actuarially fair ($\lambda = 1$) and will purchase less than full insurance for premiums greater than the actuarially fair level ($\lambda > 1$). However, when facing up-front premiums the cash-on-hand individual will purchase less than full insurance for both actuarially fair and unfair loads ($\lambda \geq 1$).*

Proof. See 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 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 upfront. This will often be the case, and we prove in 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). We show an example under constant relative risk aversion (CRRA) utility where demand will be higher under upfront 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 upfront 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 upfront 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 the individual faces a probability π of a loss of size L that can either arrive 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 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.

3.2 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 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 also exists when premiums must be paid up-front.

We compare the insurance demand of a patient individual ($\delta = 1$) with perfect liquidity and foresight to that of a cash-on-hand individual facing smooth premiums. (Note that perfect liquidity and foresight generates the standard model of insurance demand, as shown by Proposition 1 and adding patience implies that the indirect utility function in the standard model matches the consumption utility.) 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 $Nu(y - \frac{\lambda\pi L}{N})$.

With no insurance, the cash-on-hand individual's utility is given by: $V_{COH}^{no} = (1 - \pi) Nu(y) +$

$\pi [(N - 1)u(y) + u(y - L)]$, 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 $V_{PLF}^{no} = N [(1 - \pi)u(y) + \pi u(y - \frac{L}{N})]$, 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:

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

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, which reflects the value they receive from being able to spread uninsured losses costlessly across consumption periods. That second term for the consumption-smoothing benefit has a simple equation:

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

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 sub-section:

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 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:

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

If we assume that the instantaneous utility function has the constant absolute risk aversion (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:

$$\alpha_{COH}^* \approx 1 - \frac{1}{rL} \frac{\lambda - 1}{1 - \frac{\lambda\pi}{N}}. \quad (10)$$

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

$$\alpha_{PLF}^* \approx 1 - \frac{N}{rL} \frac{\lambda - 1}{1 - \lambda\pi}. \quad (11)$$

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 $N > 1$.⁹

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 propor-

⁹In the case where insurance covers only one consumption period ($N = 1$) the level of demand converges in the two cases.

tional 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 $\lim_{c \rightarrow \underline{c}} u'(c) = \infty$, under dominated pricing ($\lambda = \frac{1}{\pi}$), 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 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 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.¹⁰ For dominated insurance ($\lambda = \frac{1}{\pi}$) and CARA utility with coefficient of absolute risk aversion r , the condition for positive insurance demand is:

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

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

¹⁰The 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.

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.¹¹ Equation 8 shows that as π increases, the consumption smoothing benefit increases, and simulations presented in Figure 3 show that dominated insurance demand can arise with even moderate losses if the probability of a loss π is high enough.

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 EU 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 $\lim_{c \rightarrow \underline{c}} u'(c) = \infty$ 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 $L \geq \bar{L}$.*

Proof. See 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.

3.3 Numerical Examples for Proportional Insurance

A simple numerical example can be used to illustrate the propositions in this section. In Figure 1 we plot the optimal level of insurance (α) for an individual with a year-long insurance policy with monthly consumption periods ($N = 12$) and log utility ($u(c) = \ln(c)$) who has income of $y = 2,000$ each period and faces a potential loss of $L = 1,800$ across a range


¹¹This property holds for CRRA utility functions as consumption approaches zero. For CARA utility functions, this property holds if consumption is allowed to go negative

of insurance loads (λ). We show the optimal insurance curves as a function of loads for the three cases we have highlighted: 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 $\pi = 0.5$.

Figure 1: Optimal Insurance Demand by Load and Loss Probability

(a) Low probability $\pi = 0.1$

(b) High probability $\pi = 0.5$



Note: This figure shows the optimal level of insurance demanded (α) for the three liquidity and premium-payment scenarios discussed in this section as a function of the load for insurance. The periodic income is set at $y = \$2,000$ for this example and the loss size is $L = \$1,800$. The individual has log utility ($u(c) = \ln(c)$). In figure (a) the probability of a loss is set at 0.1, while in figure (b) it is set at 0.5. The vertical dashed lines show the load value at which insurance pricing becomes dominated given the probability of loss.

In both panels, demand for insurance for the cash-on-hand individual with smooth premiums is equal to that of the perfect-liquidity-and-foresight case when insurance is fairly priced ($\alpha = 1$ when $\lambda = 1$) and otherwise is always higher. Demand is also always higher for the cash-on-hand individual with smooth premiums than 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 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.

4 Generalization to Deductibles and Borrowing Costs

In this section, we aim to facilitate the empirical application of our model. We move beyond cash-on-hand individuals to consider individuals who face a range of different borrowing costs. We consider the willingness to pay to reduce a deductible, which is a frequently examined decision in many recent empirical papers (Cohen and Einav 2007; Einav et al. 2010; Sydnor 2010; Barseghyan et al. 2013; Handel 2013; Bhargava et al. 2017; Ericson and Sydnor 2017). Often these choices are used to make inferences about levels of risk aversion. For the example in this section we assume the individual can choose between insurance with a \$500 deductible or a \$1,000 deductible and that all losses beyond the deductible are fully covered.


We consider a relatively low-income individual with no existing assets who earns annual post-tax income of \$20,000 paid out equally in monthly installments; this is about 170% of the individual federal poverty level for 2018. The individual has log utility over monthly consumption, is perfectly patient ($\delta = 1$) and has no return to or constraint on saving ($R_s = 1$).

We assume the individual faces a 70% chance of a single large loss ($L > \$1,000$) per year.¹² The premium for the higher coverage option (i.e., \$500 deductible) is set at \$4,000, which together makes this situation roughly calibrated to the cost of employer-sponsored health insurance in the U.S. Our individual faces this insurance choice in the current year and then lives for another 20 years with full insurance in those later years. To calculate the value of an insurance contract, we solve a dynamic consumption optimization problem to determine the ex-ante expected value of insurance at different premiums. This is necessary

¹²For this simulation, we let there be a constant per-month probability of the loss conditional on the loss having not yet occurred, so if the annual probability is π , then per month probability is $1 - (1 - \pi)^{(1/12)}$.

because consumption decisions will depend on the timing of the arrival of losses (if any), about which the individual does not have perfect foresight.¹³

Figure 2: WTP for \$500 Lower Deductible by Borrowing Cost and Premium Timing



Note: This figure presents the value of lowering the deductible from \$1,000 to \$500. The example assumes a 70% chance of loss. Expected value = \$350. Assumes annual income=\$20,000, $\delta = 1, R_s = 1$, and $u(c) = \ln(c)$


Figure 2 shows the premium difference that would make an individual indifferent between policies with a \$500 deductible vs \$1,000 deductible for different levels of borrowing costs and for different ways premiums are paid (smoothly versus upfront). Borrowing interest rates plotted range from 0 (i.e., costless borrowing) to 200% 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 is close to the expected value (\$350) when the person

¹³The 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 willingness to pay 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 in for the smooth premium cases are less than \$3 different.

can borrow costlessly from future income, regardless of the premium structure. However, at higher borrowing costs, the willingness to pay for insurance diverges sharply depending on how premiums are paid. At subprime-level borrowing costs, the individual is willing to pay about 20% above the expected cost to reduce the deductible by \$500 when premiums are paid smoothly, a substantial “risk premium”. Yet that same individual is willing to pay only around 20% *below* expected cost to reduce the deductible 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. This leads the individual to appear risk loving for this additional insurance coverage on the margin.

Figure 3 plots the value of a \$500 deductible versus a \$1,000 deductible as a function of the probability of a loss. Each line in the graph represents a different liquidity condition, spanning from no borrowing costs (perfect liquidity but not perfect foresight) to 400% APR borrowing costs to the cash-on-hand situation. The figure shows that individuals facing sub-prime 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 paying more than \$500 to decrease the deductible by \$500. On the other hand, when it is costless to borrow the individual’s willingness to pay is close to expected value at all points and does not violate dominance.

Figure 3: Value of Lower Deductible by Probability of Loss and Liquidity Constraints



Note: This figure presents the value of lowering the deductible from \$1,000 to \$500 with smooth premiums for different levels of liquidity constraints. Assumes annual income is \$20,000, 12 monthly consumption periods, $u(c) = \ln(c)$, 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 ($R_s = 0$).

5 Effects on optimal insurance policy design

Liquidity constraints affect the optimal design of insurance contracts. The simplest design is a full-insurance contract that removes all risk for the individual. However, in practice we rarely observe full insurance, because loading costs and inefficiencies can make it rational for people to want less than full insurance. Full insurance can also create moral hazard problems. There are a number of ways that insurance products can be designed to offer partial coverage. How do liquidity constraints affect the desirability of these different designs from the perspective of the individual?

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.

In this section we provide two key results:

1. We extend Arrow’s result for cash-on-hand individuals and show that the optimal insurance design involves a straight deductible that resets at each consumption period.
2. When insurance contract features are constrained to cover multiple consumption periods, alternative designs that increase total spending risk but smooth shocks across periods, such as three-arm insurance designs, can be preferred for liquidity-constrained individuals over simple straight deductibles.

5.1 Formal Results

So far, we have considered a single loss that arrives all at once. To examine the design of contractual form, we will allow for losses that arrive in multiple periods.

For cash-on-hand individuals, per-period consumption in any insurance contract is a random variable that is given by

$$c_t(I) = y - p_t - L_t + I(t, L_t, \{L_\tau\}_{\tau \leq t}), \quad (13)$$

where p_t is the insurance premium in period t , L_t is a random loss that is i.i.d. over time, and $I(\cdot)$ is an indemnity paid by the insurer at time t that may depend on the history of prior losses ($\{L_\tau\}_{\tau \leq t}$) up to that time. We assume throughout this section that the insurance premium is paid smoothly in equal installments throughout the contract period so that $p_t = p$. Our goal here is to find the optimal functional form of I that maximizes expected

utility over potential claim realizations. That is, that maximizes

$$V(I) = \sum_{t=1}^N \mathbb{E}[u(y - p - L_t + I(t, L_t, \{L_\tau\}_{\tau \leq t}))] \quad (14)$$

subject to

$$p = \frac{\lambda}{N} \sum_{t=1}^N \mathbb{E}[I(t, L_t, \{L_\tau\}_{\tau \leq t})]. \quad (15)$$

where $\lambda \geq 1$ captures the insurance load.

The optimal plan design for the cash-on-hand individual will involve a simple straight deductible design with the deductible set at each consumption period. We first establish a Lemma showing that the optimal indemnity function does not depend on time.

Lemma 1. *For cash-on-hand individuals, The optimal indemnity function does not depend on t , i.e., for any t, t' , $I(t, L) = I(t', L)$ almost everywhere on the support of L .*

Proof. See Appendix A.7. □

The logic of the proof is that if the indemnity for a particular loss varies over time, there would exist small transfers from time periods with higher indemnities to those with lower indemnities that would hold fixed the expected cost of the contract and better smooth consumption for the cash-on-hand individual.

Since the optimal indemnity function does not depend on t , we focus on indemnity functions that do not depend on t and write it as $I(L)$. Since L_t is i.i.d. over time, the problem is now given by

$$\max_I V(I) = N\mathbb{E}[u(y - p - L + I(L))] \quad \text{s.t. Equation 15} \quad (16)$$

This is equivalent to the static problem considered in Arrow (1963) and hence we obtain the following proposition.

Proposition 7. *For a cash-on-hand individual, the optimal insurance contract takes the*

form of separate straight deductibles in each consumption period, i.e., there exists $d \geq 0$ such that $I(L) = \max\{0, L - d\}$.

Proof. From Lemma 1 we know that the optimal indemnity function does not depend on time. Then, since L_t is i.i.d. over time, the proposition is equivalent to Proposition 1 in Arrow (1963). \square

This result relates to Hong and Mommaerts (2021) and Diaz Campo (2021) who show that there can be sizeable liquidity benefits to resetting health insurance deductibles over shorter timespans than the common annual deductible structure. Our result highlights that for a cash-on-hand individual without moral hazard concerns, the optimal resetting matches consumption periods, implying that insurance contracts would be set at the level of consumption periods.

While insurance policies with deductibles that reset over short consumption-period time frames are not common, many policies come close to approximating this type of design. For example, in property insurance, such as automobile insurance, it is common for contracts to carry a deductible that applies to each loss that occurs. Since property losses are fairly rare, in practice this comes close to ensuring the optimal design for the cash-on-hand individual. In some settings, designs that focus on fixed dollar co-payments per incident (e.g. as sometimes seen in health insurance) may also come close to approximating the per-period straight-deductible design.

5.2 Quantitative Simulations

The prior subsection showed that cash-on-hand individuals will desire contract designs with straight deductibles, but where the deductible resets in conjunction with consumption periods. In this subsection we take up the question of how liquidity constraints affect the preferred design of insurance contracts in situations where that sort of consumption-period-based design is not available or infeasible for other reasons.

In particular, in some settings, such as health insurance, insurance contracts use deductibles that are set on annual terms. If contract terms are constrained to be set at longer frequencies than consumption periods, are straight deductible designs still desirable?

It is not possible to generally characterize the optimal policy for cash-on-hand individuals when there are constraints that prevent consumption-period-based deductibles, because the optimal constrained policy would depend on the exact nature of the potential loss distribution. In the special case where there can be at most one loss during the policy period, then we are back in the standard setting of Arrow (1963) and the straight deductible design over the policy period is optimal. When more than one loss can occur during the policy period, though, annual straight deductible designs will not be optimal for the liquidity constrained. Deductibles create larger uninsured shocks for losses early in the policy period and offer more indemnity to later losses. The liquidity-constrained individual will be willing to accept some increase in total spending risk to smooth out the indemnity schedule across loss possibilities.

To illustrate this point, we focus on a class of common contracts observed in health insurance markets, called “three armed” designs. These designs involve a deductible, some partial coverage once the deductible is met (coinsurance), and maximum limit on out-of-pocket costs for the insured. A policy with an annual straight deductible is a special case of the more general class of annual three-armed designs. We provide a simulation that demonstrates that liquidity constraints can lead to a preference for three-armed contracts that lower deductibles at the expense of increasing the size of total potential uninsured losses.

For our simulation we consider a baseline straight-deductible contract, 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 \$2000. 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 $m = \$2500$ in each case, and then for each value of $d < \$2000$, we solve for the coinsurance rate that delivers the same actuarial value

as the straight-deductible contract. As such, we consider a series of contracts with the same level of expected coverage and each contract other than the straight-deductible involves an increase in the coinsurance rate.

We examine how the desired contract from this set of possible contracts depends on the lumpiness of the risk that the individual faces. We consider two different distributions with the same probability distribution of total annual claims, but different distributions across months. In each distribution, the individual faces a $\frac{1}{4}$ chance each of total annual claims of \$30,000, \$2000, \$1500, or \$0. In the “smooth claims” distribution, annual claims are spread equally out across months, so the individual will either have 12 months of \$2500/month, \$167.67/month, \$125/month, or \$0/month in claims. This represents situations like health insurance for chronic or serious illnesses where losses can accumulate through a policy term. In the “lumpy claim” distribution, annual claims are located entirely within a single month (and each of the 12 months is equally likely to incur the claim). This represents something more like property losses or the chance of an accident for an otherwise healthy individual in health insurance. Note that we choose a distribution with more than 2 outcomes so that coinsurance rates will be relevant and that the actuarial value will change smoothly with changes in contractual form.¹⁴

Figure 4 shows the willingness to pay a liquidity constrained individual would have for contracts with deductibles lower than the \$2000 baseline under both the lumpy, single claim and smooth, multiple claims scenarios.¹⁵ We assume an individual with log monthly consumption utility (i.e., CRRA = 1) and income of \$20,000. We assume extremely high borrowing costs (900% APR) but no borrowing limit. This comes close to modeling a cash-on-hand


¹⁴With this distribution the actuarial value for all contracts we consider is 83.6%.

¹⁵Willingness to pay here is defined as a level of additional spending the individual would pay for sure spread evenly across consumption periods for the non-straight deductible option that would equalize utility between the alternative and the straight deductible. To make the two different cases of loss distributions more comparable and focus on the effect of liquidity constraints, we make a simplifying assumption of “perfect foresight” in the “single claim” scenario. We assume that at the beginning of the policy year the individual learns what loss size he will experience, if any, and the month in which it will occur. This allows the individual to choose consumption levels for the year without solving an evolving dynamic consumption problem. This makes the “single claim” similar to the “multiple claim” scenario where the information about losses is naturally revealed at the start of the year.

scenario but avoids technical issues that arise with strict borrowing limits. We see that in the lumpy, single-claim case, the classic Arrow (1963) result holds: the individual prefers the straight-deductible contract to any of the equivalent-actuarial-value options with lower deductibles. However, in the smooth, multiple claims case, the liquidity-constrained individual prefers a contract with a lower deductible but positive coinsurance coverage. Among these contracts with constant actuarial value, the individual prefers the contract with deductible just above \$1,500 and coinsurance around 3%. Even though lowering the deductible to this level increases the risk exposure, the liquidity-constrained person is willing to pay around \$100 to lower the deductible relative to the Arrow straight-deductible contract. It is worth noting, however, that liquidity constrained individuals do not necessarily want the lowest possible deductible level. There remains a trade-off with risk protection. In this case, the individual would prefer the straight-deductible to plans with deductibles much below \$1500, which have much higher co-insurance rates and hence higher chances of larger total spending amounts.

This simulation shows that liquidity constraints can be a reason individuals might prefer more complex “three arm” insurance designs even in the absence of moral hazard concerns. More generally, liquidity-constrained individuals will tend to want contract designs that do more to equalize spending across consumption periods. For example, in Appendix A.8 we show that relative to a straight deductible contract defined across multiple consumption periods, the cash-on-hand individual would prefer a design with the same expected coverage that spread coverage across claim events (by increasing coverage for the first claim event and decreasing it for a later claim event). The logic is that this modification helps to alleviate part of the discrepancy in indemnity across consumption periods in the policy term and helps the cash-on-hand individual better smooth consumption.

Figure 4: WTP For Alternative 3-Arm Design



Note: This figure shows the WTP for an alternative 3-arm design with the same actuarial value relative to the baseline contract with \$2,000 annual deductible and max-out-of-pocket limit (i.e., straight deductible). The WTP is calculated at the amount of additional spending (if positive) 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. For this simulation we assume borrowing costs ($R_b = 10$), annual income \$20,000, discount factor $\delta = 1$, savings gross interest ($R_s = 1$), and $u(c) = \ln(c)$. In the “lumpy claim” distribution, annual claims are located entirely within a single month. In “smooth claims” distribution, claims are spread evenly across months.

6 Reevaluating the Normative Benchmark

In a series of papers, Abaluck and Gruber (2011, 2016, 2020) pioneered a discrete-choice approach to evaluating choices in insurance markets. They lay out intuitive normative benchmarks that stem from the standard expected-utility model of insurance demand: a) people should place the same weight on premiums and expected out-of-pocket costs and b) people should not place weight on contract features (e.g., deductibles, copays) once premi-

ums and the expectation and variance of out-of-pocket costs are controlled for. They have demonstrated empirically in a number of health insurance markets that people routinely violate these normative benchmarks. People tend to place more weight on premiums than expected-out-of-pocket costs and place significant weight on contract features above and beyond their impact on the expectation and variance of out-of-pocket costs.

In this section, we examine how liquidity constraints change the normative benchmark for optimal plan selection. In prior sections we demonstrated that an individual selecting coverage optimally under liquidity constraints can violate dominance (Section 3) and will prefer plans with designs that help to minimize the variations in the flow of uninsured expenses over the course of the policy term (Section 5). Here we examine how choice patterns for a liquidity-constrained individual choosing optimally given those constraints compare to the normative benchmarks established by Abaluck and Gruber.

We simulate choices from menus of health insurance options for both standard expected-utility and cash-on-hand individuals who face distributions of health care expenditures based on their risk type. Individuals with standard expected utility preferences make choices consistent with the established normative benchmark. Cash-on-hand individuals, however, violate the established normative benchmarks in the same ways documented by Abaluck and Gruber. The cash-on-hand agents put significantly more weight on premiums than expected out-of-pocket costs and place significant weight on contract features, like deductibles, even after controlling for the expectation and variance of out-of-pocket costs.

Analyzing the quality of insurance choices for liquidity constrained individuals will require establishing normative benchmarks that account for the impact of liquidity constraints on the value of insurance. 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 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.

6.1 Model Setup

We simulate choices from a menu of health insurance plans for a large population of commercially insured individuals. Insurance plans in our context are comprised of annual premiums (paid smoothly each month), deductibles, coinsurance that applies after the deductible, and a maximum out-of-pocket amount. An insurance plan j translates health spending h to out-of-pocket spending $f_j(h)$. We do not have brand variation between insurance plans. We simulate choice from two different menus of options similar to those observed in health insurance markets in the United States. Menu A, with ten plan options, was based off of Ericson et al. (2021), while Menu B, with five plan options, was based off of Samek and Sydnor (2020). See Appendix Table B1 for plan menu details.

Our analysis sample comes from a large set of just over 11 million commercially insured individuals in the Truven MarketScan data base who were enrolled in coverage for the full 2018 calendar year.¹⁶ For each individual we obtain their total healthcare spending across inpatient, outpatient, and prescription drug coverage at the monthly level. This gives us one realized sequence of monthly health spending needs for each individual.

To account for ex-ante uncertainty in spending, we assign each individual a risk type, which will determine their expected distribution of healthcare cost. They expect to face the empirical distribution of spending across all individuals with that risk type. To assign risk types, we run the Centers for Medicare and Medicaid Services (CMS) hierarchical condition category (HCC) risk adjustment model for 2018, which is used for risk adjustment in the Health Insurance Marketplaces. See CMS (2018) for more details. The risk adjustment software assigns individuals to HCCs based on their diagnosis codes for the year.¹⁷ Within a disease category hierarchy, individuals are assigned to the most severe form of the disease

¹⁶We further clean the data by dropping individuals with days with negative recorded health spending.

¹⁷For example, HCC19 is “Diabetes without Complications”, while HCC19 is “Diabetes with Chronic Complications”.

that they manifest but individuals can have multiple HCCs across different disease categories. From these HCCs, the software then assigns a risk score to each individual.

Our baseline model then defines risk types as the combination of risk-score decile and age decile, giving us 100 potential risk types.¹⁸ Individuals expect to face a distribution of possible health spending flows across the year that come from the observed set of yearly spending profiles for people with that same risk type. Table 1 describes how spending varies by risk score and age. In the appendix we show that the results in this section are robust to an alternative risk-type classification procedure.

For each simulated individual we calculate the expected utility of each insurance contract under two models of behavior: a) the “standard expected utility” model in which the individual’s utility is defined over annual consumption net of health expenses and b) a cash-on-hand model in which the individual consumes income net of monthly health costs and can neither save nor borrow.

In the standard expected utility model, an individual i ’s expected utility for an insurance contract j is given by the expectation over all realizations of health spending among individuals with the same risk type k as individual i :

$$U_{ij}^{standard} = E [u(y_i - f_j(h_k))]$$

where y_i is the individual’s annual income, h_k the observed annual health spending for each person with risk type k , and f_j maps each annual health spending amount into annual out of pocket spending in insurance plan j . In our baseline simulation, we assume annual income $y = \$45,000$ and that u is log utility; the appendix shows that the results generalize to other utility functions.

The cash-on-hand model considers an individual who, at the monthly level, consumes

¹⁸Despite our large sample, data-sparsity is a challenge. Data sparsity is why we do not define risk types based on the unique combinations of HCCs a person is assigned to; in many cases, there is only one individual with this unique combination of HCCs. Appendix Table B2 shows that not all combinations of risk-score decile and age deciles are observed. For instance, none of the individuals in our data are both in the oldest decile and the lowest risk score decile.

Table 1: Descriptive Statistics of Commercially Insured Population

	Age		Annual Spending (\$)			
	Mean	Mean	Std. Dev	p25	Median	p75
<i>Full Sample:</i>	47	8,555	29,254	687	2,040	6,408
<i>Age Decile:</i>						
1	27	5,278	22,986	412	1,204	3,708
2	32	6,041	25,856	455	1,375	4,681
3	36	6,414	22,596	478	1,424	4,794
4	41	6,699	22,843	540	1,540	4,823
5	45	7,190	25,178	616	1,699	5,173
6	49	8,175	27,899	705	1,963	5,959
7	53	9,390	30,482	890	2,574	7,157
8	57	10,480	33,241	959	2,725	7,856
9	60	11,648	35,374	1,087	3,089	8,904
10	63	13,159	38,528	1,271	3,610	10,226
<i>Risk Decile:</i>						
1	32	2,293	8,169	252	637	1,852
2	35	2,734	7,319	394	993	2,533
3	40	3,107	7,982	430	1,088	2,840
4	44	3,571	8,287	490	1,286	3,401
5	49	3,991	9,418	580	1,451	3,725
6	54	4,443	10,082	686	1,669	4,218
7	55	4,591	9,825	849	1,982	4,662
8	55	66,53	13,120	1,079	2,755	7,079
9	51	10,668	18,920	2,109	5,212	11,914
10	49	35,517	72,715	5,518	15,167	37,192

Notes: Total N = 11,196,057 observations from the 2018 Truven MarketScan Data.

that month's income minus that month's out-of-pocket spending. The expected utility for an insurance plan j for a cash-on-hand individual is:

$$U_{ij}^{COH} = E \left[\sum_{m=1}^{12} u \left(\frac{y_i}{12} - f_j(h_{km}|h_{k,t < m}) \right) \right]$$

where h_{km} are the observed health spending levels for the individuals in risk-group k in the MarketScan data for month m , $h_{k,t < m}$ denotes each individual's sum of monthly health

spending in prior months, and f_j maps monthly health spending h_{km} into monthly out of pocket spending (conditional on the prior spending up to that month).¹⁹ We use the same annual income, divided into equally monthly income levels, and same utility function as in the standard model.

Finally, to simulate choices consistent with a logit choice model, we add a noise term ε_{ij} distributed type-1 extreme value to the expected utility, so that $V_{ij} = U_{ij} + \varepsilon_{ij}$. An individual's simulated choice is the plan with the highest V_{ij} .

Once we have simulated choices for each individual, we then analyze the choices using the discrete-choice framework adapted from Abaluck and Gruber (2020). They estimate the following equation from individual's choices:

$$U_{ij} = \beta_1 \text{premium}_j + \beta_2 E(OOP)_{ij} + \beta_3 \text{Var}(OOP)_{ij} + \gamma X_j \quad (17)$$

where X_j contains plan j 's cost-sharing characteristics (deductible, coinsurance, etc), and $E(OOP)_{ij}$ and $\text{Var}(OOP)_{ij}$ are the expectation of and variance of annual out-of-pocket costs that individual i will experience in plan j . Since individuals vary in their health status, $E(OOP)_{ij}$ and $\text{Var}(OOP)_{ij}$ vary at the level of plan by risk type.

As described by Abaluck and Gruber (2020), the standard model predicts that $\beta_1 \approx \beta_2$. That is, a dollar of premium costs should be treated equally to a dollar of expected out-of-pocket costs, holding fixed variance of out-of-pocket costs. The standard model also predicts that $\gamma = 0$: once accounting for expected costs and variance of costs, individuals should not attach any additional value to contract features.²⁰

¹⁹The function $f_j(h_k)$ in the standard model assumes no spending prior to this point, since in the standard model, timing of realization within the year does not matter. In the cash-on-hand model, the history of spending up to this point matters to determine the cost-sharing this month.

²⁰This relation is exact for CARA utility and a normal distribution of out-of-pocket costs, which is rarely the case. We focus on cases where this approximation works well in the standard model. However, utility is not always well approximated by this form and we can generate cases where the approximation breaks down in the standard case as well.

6.2 Results

Table 2 shows the results of the discrete-choice model estimated on choices simulated by models for the two plan menus. Columns (1) and (3) show the results under the simulation of the standard expected utility model, while Columns (2) and (4) show the results for the cash-on-hand model. The top part of the table shows the coefficient estimates from the discrete-choice model, while the bottom part of the table shows an interpretation of the coefficients in terms of equivalence to a one dollar increase in premium.

Table 2: Discrete Choice Model Results

	Plan Menu A		Plan Menu B	
	Standard (1)	Cash on Hand (2)	Standard (3)	Cash on Hand (4)
Premium	-0.002 (2.0.E-06)	-0.007 (2.4.E-06)	-0.002 (9.E-06)	-0.004 (1.E-05)
E(OOP)	-0.002 (1.6.E-06)	-0.002 (2.4.E-06)	-0.002 (3.E-06)	-0.004 (3.E-06)
Var(OOP) (\$10,000s)	-0.0002 (1.3.E-05)	0.00004 (1.7.E-05)	-0.0001 (5.E-05)	0.008 (7.E-05)
Deductible (\$1,000s)	-0.002 (1.1.E-03)	-1.764 (1.5.E-03)	-0.003 (3.E-03)	-0.450 (3.E-03)
Coinsurance (10 p.p.)	-0.003 (8.4.E-04)	-1.017 (1.E-03)	0.001 (2.E-03)	-0.017 (2.E-03)
Max(OOP) (\$1,000s)			-0.005 (3.E-03)	-0.347 (3.E-03)
\$ of Premium Equivalent to:				
Increase \$1 in E(OOP)	\$1.00	\$0.30	\$1.00	\$0.92
Increase \$1,000 Deductible	\$0.68	\$263.46	\$1.30	\$115.76
Increase 10 p.p. in Coinsurance	\$1.10	\$151.99	-\$0.56	\$4.28
Increase \$1,000 in Max(OOP)			\$2.24	\$89.18

Notes: This table displays the results of the estimation of the discrete choice model described in Equation 17. The top part of the table shows the coefficient estimates and standard errors (in parentheses). The bottom part of the table provides an interpretation of the other coefficients in terms of equivalent premium dollar increases by dividing each coefficient by the coefficient on premium. The odd numbered columns use data from the simulation of choices using the standard model of insurance demand, while the even numbers use data from the simulation of choices for the cash-on-hand model. The estimation sample ($N = 11,196,057$) is based on the commercially insured population in the Truven Marektscan data base for 2018 as described in the text.

When choices come from the standard model, we find that the estimates are consistent with the normative benchmarks for both plan menus. The response to premiums and expected out-of-pocket costs is virtually identical in both plan menus. Moreover, there is little response to contract features once we have conditioned on expected costs and variance of costs. A \$1000 increase in deductible is equivalent to a \$0.68 increase in annual premiums in Menu A (deductibles in this menu range from \$0 to \$3000) and a \$1.30 increase in annual premiums in Menu B (deductibles in this menu range from \$250 to \$2500). Similarly, there is little direct impact of a 10 percentage point increase in coinsurance in either menu or of a \$1,000 increase in maximum-out-of-pocket limit in Menu B beyond how these features affect the expectation and variance of out-of-pocket costs.²¹

In contrast, estimates from decisions generated by the cash-on-hand model deviate significantly from the normative benchmark in the same ways documented by Abaluck and Gruber (2020). Liquidity constrained individuals severely underweight expected out-of-pocket costs as compared to premiums: it takes \$1 in expected out-of-pocket costs to have the same impact on choice as just \$0.30 in premiums in Menu A. In menu B, a \$1 increase in expected out-of-pocket costs is equivalent to a \$0.92 increase in premiums, which is closer but still lower. In addition, the estimates show that cash-on-hand individuals place a strong weight on contract features: for example, increasing the annual deductible by \$1000 is equivalent to an increase in annual premium of \$263 in Menu A and of \$116 in Menu B.

Table B3 shows that these general patterns of results are robust to an alternative way of assigning individuals to risk types, to different levels of risk aversion for CRRA and CARA preferences, and to a very different specification of the variances in the logit error terms. We re-estimate the specifications from Columns (1) and (2) for Menu A in Table 2 for each of these modifications and find patterns similar to the main results: weights on plan attributes

²¹The extremely large sample size for the simulation leads to very small standard errors. Despite the small standard errors, most of the coefficients on plan features for the standard model are not statistically significantly different than zero at conventional levels. Two exceptions are the coinsurance rate in Menu A and the maximum-out-of-pocket limit in Menu B. In both cases the estimated coefficients are economically small, in line with the normative benchmark expectations, but in these cases, we do observe that controlling for mean and variance of costs does not perfectly proxy for expected utility given the log-utility specification.

close to the normative benchmark for choices generated from the standard expected utility model but substantial deviations from those generated by the cash-on-hand model.²²

7 Survey Evidence

To investigate the links between liquidity constraints and insurance demand, we fielded a survey using a Qualtrics panel. We recruited 206 adults between 18 and 65 years old and targeted specific enrollment percentages by gender, age, and household income in order to get a sample that was similar to the overall U.S. working-age population on those characteristics. See Appendix B.2 for more details on the survey population and survey design. Appendix Table B4 gives summary statistics for this sample. While not a fully representative U.S. sample, this survey sample has substantial diversity in age, income and other characteristics.

For this survey we designed a primary measure of liquidity constraints based on how an individual would finance an uninsured medical bill. We asked subjects the following question:

“Suppose you had to go to the emergency room because of an accident and just got a bill from the hospital for \$1,000 that is not covered by insurance and is due within a month. What percent of the \$1,000 hospital bill would you cover from each of these sources (total must add to 100 percent)?”

We asked participants to consider these sources of funds: money you already have (e.g., savings/checking account); extra money you save by pulling back on spending; extra money you earn by working more; borrowing from friends/family; borrowing using credit cards or home equity lines; borrowing using payday or pawn-shop loans; selling things you own; and other sources.”

We find that only 31% say they would pay the medical bill fully from money they already have, suggesting that the majority would have to engage in some sort of borrowing or

²²In some of these specifications, the standard model simulation does not as closely match the normative benchmarks. This is a result of the fact that the Abaluck and Gruber discrete-choice approach is not a perfect approximation for the standard model under certain conditions.

consumption response to finance the bill. We use the share of the bill the person says they would pay from existing funds as a simple primary indicator of liquidity constraints. This measure of liquidity constraints is, unsurprisingly, strongly but not perfectly correlated with household income. Those reporting household income below \$15,000 estimate they would pay on average 21% of the bill with cash. Among the top two income groups (income of \$50,000 or more), the average was 58%. We also fielded a more traditional question from prior research (Lusardi, Annamaria et al. 2011) that asks people “How confident are you that you could come up with \$2,000 if an unexpected need arose within the next month?” The answers to that question are closely correlated with our primary measure of liquidity constraints and also match the prior findings by Lusardi et al.

The survey then asked questions to assess whether respondents placed value on lower deductibles and smooth premium payments in ways that would be difficult to reconcile with frictionless models of consumption smoothing. Two measures related to a willingness to pay for dominated lower-deductible health plans. We replicated a menu of four hypothetical plan options from Bhargava et al. (2017), in which three lower-deductible options are dominated by an option with a \$1000 deductible. The second measure asked participants to rate which of two arguments they found more persuasive about the benefits of choosing either a \$500 or \$1000 deductible in a situation where the \$1000 deductible cost \$650 less in premium. One argument highlighted that the high deductible’s premium was so much lower it more than covered the deductible difference (i.e., dominance argument) and the other highlighted that it might be difficult to set aside money to pay for higher deductibles (i.e., budgeting argument favoring low deductibles). A final question asked participants about their preference for a “rebate plan,” motivated by previous work suggesting this idea (Johnson et al. 1993; Remmerswaal et al. 2019). This question asked respondents to consider either a standard health insurance plan with a \$1,500 annual deductible and an annual premium of \$2,000 or an equivalent “prepay with rebate option”. The rebate plan had a premium that was \$1,500 higher for the year and no deductible. However, this plan would give a rebate at the end of

the year equal to the difference between \$1,500 and their medical spending if their spending came in under \$1,500 for the year.

Table 3: Survey Results: Liquidity Constraints and Demand for Insurance

	Measures of desire for insurance to smooth consumption			
	(1) Chose dominated health plan	(2) Find argument for dominated plan persuasive	(3) Chose and agree w/ dominated (Combo 1 + 2)	(4) Prefer rebate to deductible
Overall mean of dependent var:	0.54	0.33	0.27	0.34
Estimated effect of going from 0 to 100% share of \$1k bill paid with available money:	-0.23	-0.19	-0.26	-0.16
	(0.11)	(0.10)	(0.10)	(0.10)
	[-0.44, -0.03]	[-0.39, -0.01]	[-0.45, -0.07]	[-0.35, 0.03]
Number of observations	206	206	206	206

Notes: Linear regression with heteroscedasticity-robust standard errors in parentheses and 95% confidence intervals in square brackets. Each regression includes controls for household income in discrete income bands (<\$15k, \$15k-\$25k, \$25k-\$50k, \$50k-\$100k, >\$100k), discrete age bins (18-24, 25-34, 35-44, 45-54, 55-64, 65+), discrete education levels (less HS, HS/GED, some college, 2-year college degree, 4-year college degree, masters degree, doctoral degree, professional degree), and self-reported current health status (excellent, very good, good, fair, poor). The key independent variable measure is the share of a hypothetical \$1k medical bill due within one month that the individual states they would pay with money they already have available (e.g., cash, checking, savings account), divided by 100 so that the estimated effect coefficient reflects the effect of going from 0 to 100% share paid with liquid money. The dependent variables are as follows: (1) indicator for selecting a deductible of less than \$1,000 in a 4-option hypothetical menu of health plans (dominated options); (2) Indicator for finding an argument in favor of choosing dominated plans for budgeting reasons persuasive; (3) An indicator for both choosing a dominated option and finding the argument for it more persuasive, (4) An indicator for stating a preference for a hypothetical health insurance plan with higher premiums and a rebate at the end of the year over an equivalent plan with a deductible.

Table 3 presents regression results showing the relationship between our measure of liquidity constraints and participants' answers on these questions. The regressions control for household income, respondent's level of education, respondent's age, and respondent's self-reported health status. There is a strong negative relationship between paying for the bill out of liquid money and a preference for dominated lower deductibles. Moving from 0 to 100% of the bill paid with cash is associated with a 23 percentage point reduction in the likelihood of selecting the dominated option, relative to a baseline likelihood of 54%. The estimates for finding the argument for choosing dominated plans and preferring the rebate plan are also substantial, though the result on preferring the rebate plan is not statistically significant at conventional levels. In Appendix Section B.2 we further explore the sources

that respondents state they might use for paying a bill other than money on hand. We find that, with the exception of selling goods and borrowing from friends and family, all of the categories predict a higher likelihood of choosing dominated plans relative to paying with cash. The strongest effect, and the only one that is consistently statistically significant, is the share that would be borrowed on credit cards.

8 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 survey evidence suggests that there may be value in future empirical work to collecting measures on individual liquidity constraints when assessing insurance-market dynamics. 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 such 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 upfront. 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 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 we are considering (Andrews 2016). For many other services, such as emergency room visits and hospitalizations, though, 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 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, but we can highlight a few initial thoughts on how liquidity constraints may interact with moral hazard. Liquidity constraints partially explain why individuals respond to the “spot price” of medical care (the cost sharing they must pay today), not merely their effective end-of-year price (Aron-Dine et al. 2015). We also conjecture that, all else equal, a liquidity constrained person’s medical utilization will often be more responsive to cost-sharing than a fully liquid person’s. For the same cost-sharing level, the liquidity-constrained person faces an additional financing or consumption-distortion cost for medical services than a person with perfect liquidity. As a result, liquidity constraints may provide an explanation for Einav et al.’s (2013) finding that those whose medical utilization would fall most with higher deductible plans are least likely to choose them when given an option. It may also be valuable to explore how liquidity constraints interact with other forms of moral hazard such as how people decide to time when they incur claims (Cabral 2017; Diamond et al. 2018; Gross et al. 2022).

To the extent that liquidity constraints affect service utilization under insurance, they may change some of the welfare implications associated with those responses. It may be that increased utilization when insurance coverage is high represents increased social efficiency if it solves a liquidity-constraint problem rather than the usual assumption that it represents

inefficient waste, which is similar to arguments in Nyman (1999) and Baicker et al. (2015).

Our consumption utility 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; Jaffe and Malani, 2017). As one example, improving access to and awareness of payment plans for medical bills may make people less averse to high-deductible health plans. There may then be important interactions between financing opportunities and the extent to which high-deductible plans can be used effectively to address overutilization of some medical services. As another example, our survey provided some evidence that people with liquidity constraints might benefit from slightly altered insurance arrangements that, for example, substitute rebates for deductibles.

Exploring the connection between liquidity constraints and behavioral biases is likely to be a fruitful direction for research. In our model, individuals are fully optimizing and make no mistakes. However, individuals may have incorrect beliefs (e.g. overconfidence), present-bias, or self-control issues. These behavioral biases may explain why individuals are liquidity constrained, as optimizing models suggest 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 willingness to pay 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

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. Further exploring these biases would enrich our framework.

Finally, we have limited our analysis to formal 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.

References

- Abaluck, J. and J. Gruber (2011, June). Choice Inconsistencies among the Elderly: Evidence from Plan Choice in the Medicare Part D Program. *American Economic Review* 101(4), 1180–1210.
- Abaluck, J. and J. Gruber (2016). Evolving choice inconsistencies in choice of prescription drug insurance. *American Economic Review* 106(8), 2145–84.
- Abaluck, J. and J. Gruber (2020). When Less is More: Improving Choices in Health Insurance Markets.
- Aguiar, M. A., M. Bils, and C. Boar (2020, January). Who Are the Hand-to-Mouth? Working Paper 26643, National Bureau of Economic Research. Series: Working Paper Series.

- Andrews, M. (2016, December). Doctors And Hospitals Tell Patients: Show Us The Money Before Treatment. *NPR*.
- Aron-Dine, A., L. Einav, A. Finkelstein, and M. Cullen (2015, October). Moral Hazard in Health Insurance: Do Dynamic Incentives Matter? *The Review of Economics and Statistics* 97(4), 725–741.
- Arrow, K. J. (1963). Uncertainty and the Welfare Economics of Medical Care. *The American Economic Review* 53(5), 941–973. Publisher: American Economic Association.
- Baicker, K., S. Mullainathan, and J. Schwartzstein (2015, November). Behavioral Hazard in Health Insurance. *The Quarterly Journal of Economics* 130(4), 1623–1667.
- Barseghyan, L., F. Molinari, T. O’Donoghue, and J. C. Teitelbaum (2013, October). The Nature of Risk Preferences: Evidence from Insurance Choices. *American Economic Review* 103(6), 2499–2529.
- Bhargava, S., G. Loewenstein, and J. Sydnor (2017, August). Choose to Lose: Health Plan Choices from a Menu with Dominated Options. *The Quarterly Journal of Economics* 132(3), 1319–1372.
- Cabral, M. (2017, January). Claim Timing and Ex Post Adverse Selection. *The Review of Economic Studies* 84(1), 1–44.
- Casaburi, L. and J. Willis (2018, December). Time versus State in Insurance: Experimental Evidence from Contract Farming in Kenya. *American Economic Review* 108(12), 3778–3813.
- Chetty, R. (2006, November). A general formula for the optimal level of social insurance. *Journal of Public Economics* 90(10), 1879–1901.

- Chetty, R. (2008, April). Moral Hazard versus Liquidity and Optimal Unemployment Insurance. *Journal of Political Economy* 116(2), 173–234. Publisher: The University of Chicago Press.
- Chetty, R. and A. Szeidl (2007, May). Consumption Commitments and Risk Preferences*. *The Quarterly Journal of Economics* 122(2), 831–877.
- CMS. Hhs-developed risk adjustment model algorithm “do it yourself (diy)” software instructions for the 2018 benefit year. Technical report.
- Cohen, A. and L. Einav (2007, June). Estimating Risk Preferences from Deductible Choice. *American Economic Review* 97(3), 745–788.
- Diamond, R., M. J. Dickstein, T. McQuade, and P. Persson (2018, May). Insurance without Commitment: Evidence from the ACA Marketplaces. Working Paper 24668, National Bureau of Economic Research. Series: Working Paper Series.
- Diaz Campo, C. (2021). Dynamic Moral Hazard in Nonlinear Health Insurance Contracts. Working Paper.
- Einav, L., A. Finkelstein, and M. R. Cullen (2010, August). Estimating Welfare in Insurance Markets Using Variation in Prices*. *The Quarterly Journal of Economics* 125(3), 877–921.
- Einav, L., A. Finkelstein, and J. Levin (2010). Beyond Testing: Empirical Models of Insurance Markets. *Annual Review of Economics* 2(1), 311–336. eprint: <https://doi.org/10.1146/annurev.economics.050708.143254>.
- Einav, L., A. Finkelstein, S. P. Ryan, P. Schrimpf, and M. R. Cullen (2013, February). Selection on Moral Hazard in Health Insurance. *American Economic Review* 103(1), 178–219.
- Engelhardt, G. V. and J. Gruber (2011). Medicare part d and the financial protection of the elderly. *American Economic Journal: Economic Policy* 3(4), 77–102.

- Ericson, K. M., P. Kircher, J. Spinnewijn, and A. Starc (2021, February). Inferring Risk Perceptions and Preferences Using Choice from Insurance Menus: Theory and Evidence. *The Economic Journal* 131(634), 713–744.
- Ericson, K. M. and D. Laibson (2019). Intertemporal choice. In *Handbook of Behavioral Economics: Applications and Foundations 1*, Volume 2, pp. 1–67. Elsevier.
- Ericson, K. M. and J. Sydnor (2017, November). The Questionable Value of Having a Choice of Levels of Health Insurance Coverage. *Journal of Economic Perspectives* 31(4), 51–72.
- Federal Reserve (2018). Report on the Economic Well-Being of U.S. Households in 2017. Technical report, Board of Governors of the Federal Reserve System.
- Finkelstein, A. and R. McKnight (2008). What did Medicare do? The initial impact of Medicare on mortality and out of pocket medical spending. *Journal of Public Economics* 92(7), 1644–1668.
- Gollier, C. (1994). Insurance and Precautionary Capital Accumulation in a Continuous-Time Model. *The Journal of Risk and Insurance* 61(1), 78–95. Publisher: [American Risk and Insurance Association, Wiley].
- Gollier, C. (2003). To Insure or Not to Insure?: An Insurance Puzzle. *The Geneva Papers on Risk and Insurance Theory* 28(1), 5–24. Publisher: Palgrave Macmillan Journals.
- Gollier, C. and H. Schlesinger (1996). Arrow’s Theorem on the Optimality of Deductibles: A Stochastic Dominance Approach. *Economic Theory* 7(2), 359–363. Publisher: Springer.
- Gross, T., T. J. Layton, and D. Prinz (2022). The Liquidity Sensitivity of Healthcare Consumption: Evidence from Social Security Payments. *American Economic Review: Insights*, forthcoming.
- Gruber, J. (1997). The Consumption Smoothing Benefits of Unemployment Insurance. *The American Economic Review* 87(1), 192–205. Publisher: American Economic Association.

- Handel, B., I. Hendel, and M. D. Whinston (2015). Equilibria in Health Exchanges: Adverse Selection versus Reclassification Risk. *Econometrica* 83(4), 1261–1313.
- Handel, B. R. (2013). Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts. *American Economic Review* 103(7), 2643–2682.
- Handel, B. R. and J. T. Kolstad (2015, August). Health Insurance for "Humans": Information Frictions, Plan Choice, and Consumer Welfare. *American Economic Review* 105(8), 2449–2500.
- Hansen, G. D. and A. Imrohoroğlu (1992). The Role of Unemployment Insurance in an Economy with Liquidity Constraints and Moral Hazard. *Journal of Political Economy* 100(1), 118–142. Publisher: University of Chicago Press.
- Hastings, J. S. and J. M. Shapiro (2013, November). Fungibility and Consumer Choice: Evidence from Commodity Price Shocks*. *The Quarterly Journal of Economics* 128(4), 1449–1498.
- Hong, L. and C. Mommaerts (2021, February). Time Aggregation in Health Insurance Deductibles. Working Paper 28430, National Bureau of Economic Research. Series: Working Paper Series.
- Hundtofte, S., A. Olafsson, and M. Pagel (2019, October). Credit Smoothing. Working Paper 26354, National Bureau of Economic Research. Series: Working Paper Series.
- Jaffe, S. and A. Malani (2017). The welfare implications of health insurance. *Human Capital and Economic Opportunity Working Group Working Paper 45*.
- Johnson, E. J., R. Hassin, T. Baker, A. T. Bajger, and G. Treuer (2013, December). Can Consumers Make Affordable Care Affordable? The Value of Choice Architecture. *PLOS ONE* 8(12), e81521. Publisher: Public Library of Science.

- Johnson, E. J., J. Hershey, J. Meszaros, and H. Kunreuther (1993, August). Framing, probability distortions, and insurance decisions. *Journal of Risk and Uncertainty* 7(1), 35–51.
- Liu, Y. and R. J. Myers (2016). The Dynamics of Microinsurance Demand in Developing Countries Under Liquidity Constraints and Insurer Default Risk. *The Journal of Risk and Insurance* 83(1), 121–138. Publisher: [American Risk and Insurance Association, Wiley].
- Loewenstein, G., J. Y. Friedman, B. McGill, S. Ahmad, S. Linck, S. Sinkula, J. Beshears, J. J. Choi, J. Kolstad, D. Laibson, B. C. Madrian, J. A. List, and K. G. Volpp (2013, September). Consumers’ misunderstanding of health insurance. *Journal of Health Economics* 32(5), 850–862.
- Lusardi, Annamaria, Schneider, Daniel, and Tufano, Peter (2011). Financially Fragile Households: Evidence and Implications. Technical report.
- Malani, A. and S. P. Jaffe (2018, July). The Welfare Implications of Health Insurance. Working Paper 24851, National Bureau of Economic Research. Series: Working Paper Series.
- Morduch, J. and R. Schneider (2017, April). *The Financial Diaries: How American Families Cope in a World of Uncertainty* (Illustrated edition ed.). Princeton, New Jersey: Princeton University Press.
- Mossin, J. (1968). Aspects of Rational Insurance Purchasing. *Journal of Political Economy* 76(4), 553–568. Publisher: University of Chicago Press.
- Nyman, J. A. (1999, April). The value of health insurance: the access motive. *Journal of Health Economics* 18(2), 141–152.
- O’Donoghue, T. and M. Rabin (1999, March). Doing It Now or Later. *American Economic Review* 89(1), 103–124.

- Olafsson, A. and M. Pagel (2018, November). The Liquid Hand-to-Mouth: Evidence from Personal Finance Management Software. *The Review of Financial Studies* 31(11), 4398–4446.
- Remmerswaal, M., J. Boone, M. Bijlsma, and R. Douven (2019). Cost-sharing design matters: A comparison of the rebate and deductible in healthcare. *Journal of Public Economics* 170, 83–97.
- Rothschild, M. and J. Stiglitz (1976, November). Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information*. *The Quarterly Journal of Economics* 90(4), 629–649.
- Samek, A. and J. R. Sydnor (2020). Impact of consequence information on insurance choice. Technical report, National Bureau of Economic Research.
- Schlesinger, H. (2000). *The Theory of Insurance Demand*, pp. 131–151. Dordrecht: Springer Netherlands.
- Shapiro, J. M. (2005, February). Is there a daily discount rate? Evidence from the food stamp nutrition cycle. *Journal of Public Economics* 89(2), 303–325.
- Sydnor, J. (2010, October). (Over)insuring Modest Risks. *American Economic Journal: Applied Economics* 2(4), 177–199.
- Thaler, R. (1985). Mental Accounting and Consumer Choice. *Marketing Science* 4(3), 199–214. Publisher: INFORMS.
- Zeckhauser, R. (1970, March). Medical insurance: A case study of the tradeoff between risk spreading and appropriate incentives. *Journal of Economic Theory* 2(1), 10–26.
- Zeldes, S. P. (1989). Consumption and Liquidity Constraints: An Empirical Investigation. *Journal of Political Economy* 97(2), 305–346. Publisher: University of Chicago Press.

Online Appendix For: Liquidity Constraints and the Value of Insurance

Keith Marzilli Ericson and Justin Sydnor

Appendix A Proofs

A.1 Proof of Proposition 1

Proof. First, we show these assumptions are sufficient for a representation. With perfect liquidity and perfect foresight, the timing of when the loss occurs during the insurance contract period does not matter. Then, if the loss happens, there will be some vector of optimal consumption utility $\{c_{Lt}^*\}_{t=1}^N = \operatorname{argmax} \sum_{t=1}^N \delta^t u(c_t)$ subject to the budget constraint $\sum_{t=1}^N c_t \leq w_L$. Similarly, if the loss does not happen, $\{c_{Nt}^*\}_{t=1}^N$ will solve the maximization problem subject to the budget constraint $\sum_{t=1}^N c_t \leq w_N$. Then, define the indirect utility function $v(w_L) = \sum_{t=1}^N \delta^t u(c_{Lt}^*)$ and $v(w_N) = \sum_{t=1}^N \delta^t u(c_{Nt}^*)$. Since the loss happens with probability π , we have the expected utility formulation $V(Z, L, \pi) = \pi v(w_L) + (1 - \pi)v(w_N)$.

Now, we show these assumptions are necessary for a representation. Suppose there was not perfect liquidity. Then, consider two different risks, A and B, with vectors of ex ante per-period probabilities (mutually exclusive) in which the one loss may occur: π_t^A and π_t^B , respectively, where $\sum_{t=1}^N \pi_t^A = \sum_{t=1}^N \pi_t^B = \pi$.

However, the budget constraint for the utility maximization problem will differ depending on what period t the loss arises in. When $R_b = R_s = R > 1$, the effective loss amount is larger if it happens earlier than it if happens later.

There are $N + 1$ possible budget constraints that will be realized (each of the N periods that the loss might happen in, plus the case where the loss does not happen). If $R_b \neq 1$ and the individual wants to borrow at any point, or $R_s \neq 1$ and the individual wants to save at any point, then these budget constraints will yield different utilities. As a result, the average utility when the loss happens (over all the periods it might happen in) with risk A will not always be the same as risk B, as π_t^A and π_t^B differ. (Utility when the loss does not happen with risk A will be the same as risk B.) However, $V(Z, L, \pi)$ must be identical for both risk A and B, a contradiction. Hence, without perfect liquidity, an individual with consumption utility cannot have their preference represented with the static expected utility formulation.

Now, suppose there was not perfect foresight. Then, the individual solves a dynamic programming problem to choose consumption each period, with two state variables: current assets and an indicator for whether the loss has happened yet or not. As in paragraph above, $V(Z, L, \pi)$ must be identical for both risk A and B. However, the consumption solution vector for risk A will not in general yield the same solution and the same utility as risk B, and the individual will strictly prefer one risk to the other.

□

A.2 Proof of Proposition 2

Proof. Recall that the indirect utility function of the cash-on-hand individual with smooth premium schedule is given by

$$V_{COH}(\alpha) = Nu(c_N(\alpha)) - \pi [u(c_N(\alpha)) - u(c_L(\alpha))], \quad (6)$$

where $c_L(\alpha) = y - \frac{\lambda\alpha\pi L}{N} - (1 - \alpha)L$ and $c_N(\alpha) = y - \frac{\lambda\alpha\pi L}{N}$. Differentiating the indirect utility with respect to α and arranging terms yield

$$V'_{COH}(\alpha) = \pi \left(1 - \frac{\pi}{N}\right) \lambda L \{u'(c_L) - u'(c_N)\} - (\lambda - 1)\pi L u'(c_L). \quad (18)$$

Case 1: Smooth premiums & $\lambda = 1$

Substituting $\lambda = 1$ and arranging the terms gives us

$$V'_{COH}(\alpha) = \pi \left(1 - \frac{\pi}{N}\right) L \{u'(c_L) - u'(c_N)\}. \quad (19)$$

Note that $c_L < c_N$ for $\alpha \in [0, 1)$ and $c_L = c_N$ for $\alpha = 1$. Together with the assumption that u is strictly concave, we have $V'_{COH}(\alpha) > 0$ for $\alpha \in [0, 1)$ and $V'_{COH}(1) = 0$. Therefore $\alpha = 1$ maximizes utility.

Case 2: Smooth premiums & $\lambda > 1$

If $\alpha = 1$, then $c_L = c_N$. Therefore, we have

$$V'_{COH}(1) = -(\lambda - 1)\pi Lu'(c_L) < 0. \quad (20)$$

This implies that decreasing the coverage from full insurance increases utility, i.e., full insurance is not optimal.

Case 3: Up-front premiums & $\lambda \geq 1$ Recall that the indirect utility function of the cash-on-hand individual with up-front premium schedule is given by

$$\begin{aligned} V_{Up}(\alpha) = & (1 - \pi)[u(c_{1,N}) + (N - 1)u(c_{2,N})] + \frac{\pi}{N}[u(c_{1,L}) + (N - 1)u(c_{2,N})] \\ & + \frac{(N - 1)}{N}\pi[u(c_{1,N}) + (N - 2)u(c_{2,N}) + u(c_{2,L})], \end{aligned} \quad (21)$$

where $c_{1,N} = y - \lambda\alpha\pi L$, $c_{2,N} = y$, $c_{1,L} = y - \lambda\alpha\pi L - (1 - \alpha)L$, and $c_{2,L} = y - (1 - \alpha)L$.

Differentiating the indirect utility function with respect to α yields

$$V'_{Up}(\alpha) = -\left(1 - \frac{1}{N}\pi\right)\lambda\pi Lu'(c_{1,N}) + \frac{\pi}{N}(1 - \lambda\pi)Lu'(c_{1,L}) + \frac{N - 1}{N}\pi Lu'(c_{2,L}). \quad (22)$$

Arranging terms yields

$$V'_{Up}(\alpha) = \pi L \left[\left\{ \frac{1}{N}u'(c_{1,L}) + \frac{N - 1}{N}u'(c_{2,L}) \right\} - \lambda \left\{ \left(1 - \frac{\pi}{N}\right)u'(c_{1,N}) + \frac{\pi}{N}u'(c_{1,L}) \right\} \right]. \quad (23)$$

Full insurance ($\alpha = 1$) implies $c_{2,L} = y$ and $c_{1,N} = c_{1,L} = y - \pi L$. Therefore, the marginal value at full insurance is given by

$$\begin{aligned} V'_{Up}(1) = & \pi L \left[\left\{ \frac{1}{N}u'(y - \pi L) + \frac{N - 1}{N}u'(y) \right\} - \lambda u'(y - \pi L) \right] \\ & < -(\lambda - 1)\pi Lu'(y - \pi L) \\ & \leq 0. \end{aligned} \quad (24)$$

The inequality in the second line is obtained by replacing $u'(y)$ in the first line with $u'(y - \pi L)$

which is strictly larger than $u'(y)$. The expression in the second line is nonpositive since $\lambda \geq 1$. Since $V'_{Up}(1) < 0$, the full insurance is not optimal.

□

A.2.1 Demand under smooth vs up-front premiums

As discussed in the text, Proposition 2 might initially suggest that optimal demand for proportional insurance will be higher under smooth premiums than upfront premiums. Here we show that while this can be proven to be true for constant absolute risk aversion (CARA), it is not always true under other preferences. We show a specific numerical example for the constant relative risk aversion CRRA case.

First we can prove that an individual who has CARA preferences will purchase a strictly higher level of insurance α^* with smooth premiums than with upfront premiums for all premium loads at or above the actuarially fair level ($\lambda \geq 1$), unless both types choose $\alpha^* = 0$.

Proof. Let α^*_{smooth} be the optimal level of insurance when the agent is facing the smooth premium schedule and α^*_{up} be the optimal level of insurance when the agent is facing the upfront premium schedule. Consider a CARA utility function $u(c) = -\frac{1}{\gamma}e^{-\gamma c}$. If $\alpha^*_{smooth} = 1$, we always have $\alpha^*_{smooth} > \alpha^*_{up}$. Suppose $\alpha^*_{smooth} < 1$. Then, α^*_{smooth} satisfies

$$V'_{COH}(\alpha^*) = -\left(1 - \frac{\pi}{N}\right) \lambda \pi L e^{-\gamma c_N(\alpha^*_{smooth})} + \pi \left(1 - \frac{\lambda \pi}{N}\right) L e^{-\gamma c_L(\alpha^*_{smooth})} \leq 0, \quad (25)$$

where we get $\alpha^*_{smooth} = 0$ if the inequality is strict. Dividing both sides with $e^{-\gamma c_L(\alpha^*_{smooth})}$ yields

$$-\left(1 - \frac{\pi}{N}\right) \lambda \pi L e^{-\gamma [c_N(\alpha^*_{smooth}) - c_L(\alpha^*_{smooth})]} + \pi \left(1 - \frac{\lambda \pi}{N}\right) L \leq 0. \quad (26)$$

Since $c_N(\alpha) - c_L(\alpha) = (1 - \alpha)L$, we get

$$-\left(1 - \frac{\pi}{N}\right) \lambda \pi L e^{-\gamma (1 - \alpha^*_{smooth})L} + \pi \left(1 - \frac{\lambda \pi}{N}\right) L \leq 0. \quad (27)$$

The marginal indirect utility under the upfront premium schedule is

$$\begin{aligned}
V'_{Up}(\alpha) &= - \left(1 - \frac{1}{N}\pi\right) \lambda\pi Lu'(c_{1,N}) + \frac{\pi}{N}(1 - \lambda\pi)Lu'(c_{1,L}) + \frac{N-1}{N}\pi Lu'(c_{2,L}) \\
&= - \left(1 - \frac{1}{N}\pi\right) \lambda\pi Lu'(c_{1,N}) + \pi \left(1 - \frac{\lambda\pi}{N}\right) Lu'(c_{1L}) \\
&\quad - \pi \left(1 - \frac{\lambda\pi}{N}\right) Lu'(c_{1L}) + \frac{\pi}{N}(1 - \lambda\pi)Lu'(c_{1,L}) + \frac{N-1}{N}\pi Lu'(c_{2,L}) \\
&= e^{-\gamma c_{1L}} \left[- \left(1 - \frac{\pi}{N}\right) \lambda\pi L e^{-\gamma(1-\alpha)L} + \pi \left(1 - \frac{\lambda\pi}{N}\right) L \right] + \left(1 - \frac{1}{N}\right) \pi L [u'(c_{2L}) - u'(c_{1L})].
\end{aligned} \tag{28}$$

If we evaluate this at α^*_{smooth} , the first term in the last line is nonpositive because the expression in the bracket equals the first-order condition for α^*_{smooth} . Since $c_{2L} > c_{1L}$, we have

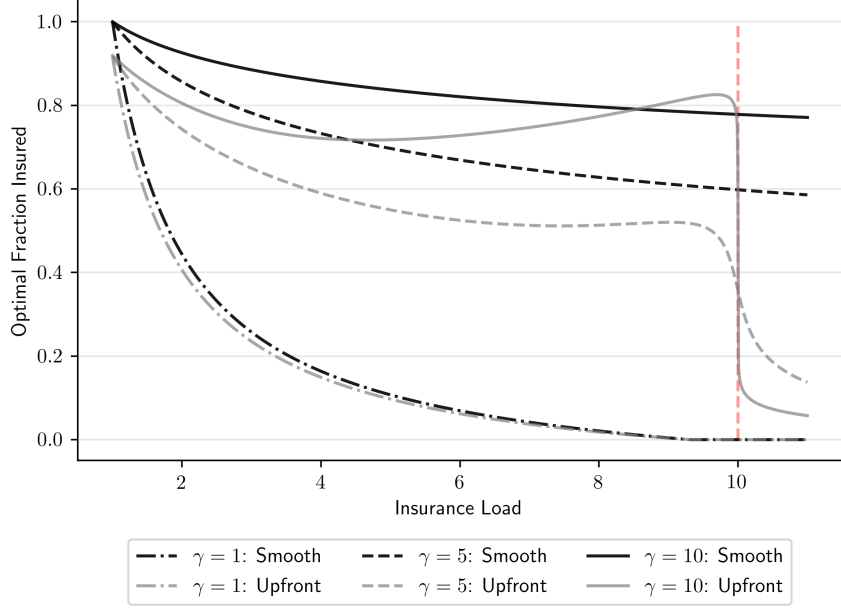
$$V'_{Up}(\alpha^*_{smooth}) \leq \left(1 - \frac{1}{N}\right) \pi L [u'(c_{2L}) - u'(c_{1L})] < 0. \tag{29}$$

This implies that an agent with the upfront premium schedule prefers decreasing α from α^*_{smooth} . In other words, demand for insurance is smaller under the upfront premium schedule than the smooth premium schedule.

□

However, this will not always be true. Figure A1 shows a numerical example of the optimal fraction insured under CRRA preferences for different insurance loads and risk aversion levels. The individual would select higher levels of insurance with smooth premiums than upfront premiums when insurance loads are fairly low or for low levels of risk aversion. However, when the individual is very risk averse (and hence has a strong degree of decreasing absolute risk aversion), there can be a non-monotonicity in insurance demand at high insurance loads. This property occurs more easily for the upfront premium case than the smooth premium case. As such, as the example for CRRA with risk aversion parameter of 10 shows in the figure, it can be the case that the level of insurance demanded will be higher with upfront than smooth premiums.

Figure A1: Optimal Insurance Demand by Load and CRRA Parameter



Note: This figure shows the optimal level of insurance demanded (α) by insurance load and the risk aversion parameter γ of the CRRA utility function $u(c) = \frac{c^{1-\gamma}}{1-\gamma}$ under the smooth premium schedule (black lines) and the upfront premium schedule (gray lines). The number of periods is $N = 12$. The periodic income is set at $y = \$2,000$ for this example and the loss size is $L = \$1,800$. The probability of a loss is $\pi = 0.1$. The vertical dashed line shows the load value at which insurance pricing becomes dominated given the probability of loss.

The intuition for this result is that the highly risk averse agent will want insurance even if loads are very high. Yet that purchase of costly insurance will mean that they are at a low consumption level and extremely sensitive to uninsured loss amounts. For the upfront case, the premium charge occurs in a single period and heightens this sensitivity to uninsured losses.

A.3 Proof of Proposition 3

Proof. First, consider what happens when premiums are paid smoothly. For our original case of lumpy losses, Equation (6) gives the indirect utility: $V_{COH} = Nu(c_N) - \pi [u(c_N) - u(c_L)]$ where $c_N = y - \frac{\lambda\alpha\pi L}{N}$ and $c_L = y - \frac{\lambda\alpha\pi L}{N} - (1 - \alpha)L$. For smooth losses, the indirect utility is given by

$$\hat{V}_{COH} = (1 - \pi)Nu(c_N) + \pi Nu(c_{L,smooth}), \tag{30}$$

where

$$c_{L,smooth} = y - \frac{\lambda\alpha\pi L}{N} - (1 - \alpha)\frac{L}{N}. \quad (31)$$

The marginal indirect utilities are given by

$$\begin{aligned} V'_{COH}(\alpha) &= -\left(1 - \frac{\pi}{N}\right)\lambda\pi Lu'(c_N) + \left(1 - \frac{\lambda\pi}{N}\right)\pi Lu'(c_L) \\ V'_{COH}(\alpha) &= \left(1 - \frac{\lambda\pi}{N}\right)\pi L[u'(c_L) - u'(c_N)] - (\lambda - 1)\pi Lu'(c_N) \end{aligned} \quad (32)$$

and

$$\begin{aligned} \hat{V}'_{COH}(\alpha) &= -(1 - \pi)\lambda\pi Lu'(c_N) + (1 - \lambda\pi)\pi Lu'(c_{L,smooth}) \\ &= (1 - \lambda\pi)\pi L[u'(c_{L,smooth}) - u'(c_N)] - (\lambda - 1)\pi Lu'(c_N). \end{aligned} \quad (33)$$

The second term is the same in both equations. The first term in the first equation is larger than the second equation since $1 - \frac{\lambda\pi}{N} > 1 - \lambda\pi$ and $c_L < c_{L,smooth}$. The marginal indirect utility is therefore larger in the case of lumpy losses for the same level of α .

Since the marginal indirect utility is larger under lumpy losses for the same level of α , if there is an interior solution for the smooth-losses case, then the first order condition holds for that case and the optimal α will be higher than that level for the lumpy-loss case. At a corner solution for high λ , the agent may demand zero insurance ($\alpha = 0$) under both loss-arrival cases. For the other corner solution, when $\lambda = 1$, the logic of the proof for Proposition 2 under smooth premiums holds and the agent will demand full insurance $\alpha = 1$ for both loss-arrival cases. This implies that when premiums are paid smoothly the demand for insurance is weakly larger with lumpy losses than with smooth losses and is strictly larger when the solution $\alpha^*_{COH} \in (0, 1)$.

Now, we consider up-front premiums. For our original case of lumpy losses, the indirect

utility is given by

$$\begin{aligned}
V_{Up} &= (1 - \pi)[u(c_{1,N}) + (N - 1)u(c_{2,N})] \\
&\quad + \frac{\pi}{N}[u(c_{1,L}) + (N - 1)u(c_{2,N})] \\
&\quad + \frac{(N - 1)}{N}\pi[u(c_{1,N}) + (N - 2)u(c_{2,N}) + u(c_{2,L})]
\end{aligned} \tag{34}$$

where $c_{1,N} = y - \lambda\alpha\pi L$, $c_{2,N} = y$, $c_{1,L} = y - \lambda\alpha\pi L - (1 - \alpha)L$, and $c_{2,L} = y - (1 - \alpha)L$.

For smooth losses, the indirect utility is given by

$$\begin{aligned}
\hat{V}_{Up} &= (1 - \pi)[u(c_{1,N}) + (N - 1)u(c_{2,N})] \\
&\quad + \pi[u(c_{1,L,smooth}) + (N - 1)u(c_{2,L,smooth})],
\end{aligned} \tag{35}$$

where

$$c_{1,L,smooth} = y - \lambda\alpha\pi L - (1 - \alpha)\frac{L}{N} \tag{36}$$

$$c_{2,L,smooth} = y - (1 - \alpha)\frac{L}{N}. \tag{37}$$

The marginal indirect utilities are given by

$$\begin{aligned}
V'_{Up}(\alpha) &= -\left(1 - \frac{1}{N}\pi\right)\lambda\pi Lu'(c_{1,N}) + \frac{\pi}{N}(1 - \lambda\pi)Lu'(c_{1,L}) + \frac{N - 1}{N}\pi Lu'(c_{2,L}), \\
&= \frac{\pi}{N}(1 - \lambda\pi)L[u'(c_{1,L}) - u'(c_{1,N})] - \left(\lambda - \frac{1}{N}\right)\pi Lu'(c_{1,N}) + \frac{N - 1}{N}\pi Lu'(c_{2,L}).
\end{aligned} \tag{38}$$

and

$$\begin{aligned}
\hat{V}'_{Up}(\alpha) &= -(1 - \pi)\lambda\pi Lu'(c_{1,N}) + \frac{\pi}{N}(1 - N\lambda\pi)Lu'(c_{1,L,smooth}) + \frac{N - 1}{N}\pi Lu'(c_{2,L,smooth}) \\
&= \frac{\pi}{N}(1 - N\lambda\pi)L[u'(c_{1,L,smooth}) - u'(c_{1,N})] - \left(\lambda - \frac{1}{N}\right)\pi Lu'(c_{1,N}) + \frac{N - 1}{N}\pi Lu'(c_{2,L,smooth}).
\end{aligned} \tag{39}$$

The second term is the same in both equations while the first and the last terms are larger in the first equation since $c_{1,L} < c_{1,L,smooth}$ and $c_{2,L} < c_{2,L,smooth}$.

By the same argument as above for the smooth-premiums scenario, this implies that the demand for insurance is weakly larger with lumpy losses than with smooth losses and is strictly larger when the solution $\alpha_{Up}^* \in (0, 1)$. \square

A.4 Proof of Proposition 4

Proof. Let α_{PLF} be the demand for insurance by an agent with perfect liquidity and perfect foresight that solves the first-order condition:

$$V'_{PLF}(\alpha_{PLF}) = \pi u'(c_L(\alpha_{PLF}))(1 - \lambda\pi)\frac{L}{N} - (1 - \pi)u'(c_N(\alpha_{PLF}))\frac{\lambda\pi L}{N} = 0. \quad (40)$$

This can be rewritten as

$$\frac{u'(c_L(\alpha_{PLF}))}{u'(c_N(\alpha_{PLF}))} = 1 + \frac{\lambda - 1}{1 - \lambda\pi}. \quad (41)$$

Let α_{COH} be the demand for insurance by a cash-on-hand individual facing smooth premium schedules that solves the first-order condition (Recall Equation 32): :

$$V'_{COH}(\alpha_{COH}) = \pi\{u'(c_L(\alpha_{COH})) - u'(c_N(\alpha_{COH}))\}\left(1 - \frac{\pi}{N}\right)\lambda L - (\lambda - 1)\pi u'(c_L(\alpha_{COH}))L = 0. \quad (42)$$

This can be written as

$$\frac{u'(c_L(\alpha_{COH}))}{u'(c_N(\alpha_{COH}))} = 1 + \frac{\lambda - 1}{1 - \frac{\lambda\pi}{N}}. \quad (43)$$

Now compare the two first-order conditions (40) and (43). Since we assume $\lambda\pi < 1$, we have $\frac{\lambda-1}{1-\lambda\pi} > \frac{\lambda-1}{1-\frac{\lambda\pi}{N}}$, which implies $\frac{u'(c_L(\alpha_{PLF}))}{u'(c_N(\alpha_{PLF}))} > \frac{u'(c_L(\alpha_{COH}))}{u'(c_N(\alpha_{COH}))}$. Therefore, we obtain

$$\frac{u'\left(y - \frac{\lambda\alpha_{PLF}\pi L}{N} - (1 - \alpha_{PLF})\frac{L}{N}\right)}{u'\left(y - \frac{\lambda\alpha_{PLF}\pi L}{N}\right)} > \frac{u'\left(y - \frac{\lambda\alpha_{COH}\pi L}{N} - (1 - \alpha_{COH})L\right)}{u'\left(y - \frac{\lambda\alpha_{COH}\pi L}{N}\right)}, \quad (44)$$

which can be rewritten as

$$\frac{u' \left(y - \frac{\lambda \alpha_{PLF} \pi L}{N} - (1 - \alpha_{PLF}) \frac{L}{N} \right)}{u' \left(y - \frac{\lambda \alpha_{COH} \pi L}{N} - (1 - \alpha_{COH}) L \right)} > \frac{u' \left(y - \frac{\lambda \alpha_{PLF} \pi L}{N} \right)}{u' \left(y - \frac{\lambda \alpha_{COH} \pi L}{N} \right)}. \quad (45)$$

Suppose that $\alpha_{PLF} \geq \alpha_{COH}$. Then, since $u'' < 0$, the right-hand side is greater than or equal to 1. This implies

$$y - \frac{\lambda \alpha_{PLF} \pi L}{N} - (1 - \alpha_{PLF}) \frac{L}{N} < y - \frac{\lambda \alpha_{COH} \pi L}{N} - (1 - \alpha_{COH}) L, \quad (46)$$

which in turn implies

$$\lambda \pi (\alpha_{PLF} - \alpha_{COH}) > N(1 - \alpha_{COH}) - (1 - \alpha_{PLF}) > \alpha_{PLF} - \alpha_{COH}. \quad (47)$$

Since we assume $\lambda < \frac{1}{\pi}$, or $\lambda \pi < 1$, this inequality implies $\alpha_{PLF} - \alpha_{COH} < 0$, but this contradicts $\alpha_{PLF} \geq \alpha_{COH}$. Therefore, we must have $\alpha_{PLF} < \alpha_{COH}$. \square

A.5 Proof of Proposition 5

Proof. Recall from Equation 32, we have:

$$V'_{COH}(\alpha) = \pi \{u'(c_L) - u'(c_N)\} \left(1 - \frac{\lambda \pi}{N}\right) L - (\lambda - 1) \pi L u'(c_N). \quad (48)$$

Since $\lambda = \frac{1}{\pi}$, we get

$$V'_{COH}(\alpha) = \{u'(c_L) - u'(c_N)\} \left(1 - \frac{1}{N}\right) L - (\lambda - 1) L u'(c_N). \quad (49)$$

The agent has a positive demand for insurance if and only if $V'_{COH}(0) > 0$. The derivative at $\alpha = 0$ is positive if and only if

$$\{u'(y - L) - u'(y)\} \left(1 - \frac{1}{N}\right) > (\lambda - 1) u'(y), \quad (50)$$

or equivalently,

$$\frac{u'(y-L) - u'(y)}{u'(y)} > \frac{\lambda - 1}{1 - \frac{1}{N}}. \quad (51)$$

Recall that we assume that there exists \underline{c} such that $\lim_{c \rightarrow \underline{c}} u'(c) = \infty$. Therefore, as L gets larger so that $y - L \rightarrow \underline{c}$, the left-hand side goes to infinity and the inequality is satisfied. Since we assume that u' is continuous, this implies that there exists $\bar{L}_S \in (0, y - \underline{c})$ such that this inequality holds if and only if $L \geq \bar{L}_S$. □

A.6 Proof of Proposition 6

Proof. Inserting $\pi = 1$ into the first order condition for choice of α from Equation 32, we have

$$V'_{COH}(\alpha) = \{u'(c_L) - u'(c_N)\} \left(1 - \frac{\lambda}{N}\right) L - (\lambda - 1)Lu'(c_N). \quad (52)$$

The agent has a positive demand for insurance if and only if $V'_{COH}(0) > 0$. Consider $\lambda < N$. Then, the derivative at $\alpha = 0$ is positive if and only if

$$\frac{u'(y-L) - u'(y)}{u'(y)} > \frac{\lambda - 1}{1 - \frac{\lambda}{N}}. \quad (53)$$

Then, the same argument as in the proof of Proposition 5 shows that there exists a threshold \bar{L}_S for a positive insurance demand. □

A.7 Proof of Lemma 1

Proof. Let S be a subset on the support of L satisfying $\Pr(L \in S) > 0$. Suppose that there exist t and t' with $t \neq t'$ such that $I(t, L) - I(t', L) > \epsilon > 0$ for any $L \in S$.²³ Then, we can find a transfer from t to t' that can improve $V(I)$ while satisfying the budget constraint.

²³The same proof applies to the case where the inequality is reversed.

Indeed, consider a function $\hat{I}(t, L)$ that equals $I(t, L)$ any t for $L \notin S$, and for $L \in S$,

$$\hat{I}(t, L) = I(t, L) - \frac{\epsilon}{2}, \quad (54)$$

$$\hat{I}(t', L) = I(t', L) + \frac{\epsilon}{2}, \quad (55)$$

$$\hat{I}(t'', L) = I(t'', L), \quad (56)$$

where $t'' \neq t, t'$. Note that $\sum_{t=1}^N \mathbb{E}[I(t, L_t)] = \sum_{t=1}^N \mathbb{E}[\hat{I}(t, L_t)]$ and therefore the constraint is satisfied. Also, we have (i) I and \hat{I} have the same total consumption; $c_t(\hat{I}) + c_{t'}(\hat{I}) = c_t(I) + c_{t'}(I)$, and (ii) \hat{I} has smaller consumption fluctuation; $|c_t(\hat{I}) - c_{t'}(\hat{I})| \leq |c_t(I) - c_{t'}(I)|$ where the inequality is strict with positive probability. Since u is strictly concave, it follows that $V(\hat{I}) > V(I)$. \square

A.8 Existence of a better policy than a straight deductible policy

We show that we can create a better policy than a straight deductible policy by marginally increasing an indemnity for the first claim and marginally reducing the indemnity for a second claim.

Consider an insurance contract $I(t, \{L_\tau\}_{\tau \leq t})$ that has a straight deductible. Now suppose we consider an alternative policy $I'(t, \{L_\tau\}_{\tau \leq t})$ that has the same design but for the following modifications:

1. Marginally increase $I(t, \{L_\tau\}_{\tau \leq t})$ by $d\epsilon > 0$ at the first period with a loss (if any).
2. Marginally decrease $I(t, \{L_\tau\}_{\tau \leq t})$ by $d\delta > 0$ at the second period with a loss (if any).

The modified contract I' provides a larger indemnity at the first period with a loss and a smaller indemnity at the second period with a loss.

Let π be the probability having a positive loss amount $L > 0$ in any period. To hold fixed the expected cost of the contract, we need to make sure the expected value of first-period additional indemnities equals the expected value of the second-period additional surcharges.

That is:

$$(1 - (1 - \pi)^N)d\epsilon = (1 - (1 - \pi)^N - N\pi(1 - \pi)^{N-1})d\delta, \quad (57)$$

where $1 - (1 - \pi)^N$ is the probability of a loss in at least one period and $1 - (1 - \pi)^N - N\pi(1 - \pi)^{N-1}$ is the probability of losses in two or more periods.

Now consider the utility gains and losses that come from the new transfer policy I' relative to the original policy I . Denote the loss size at the first period with a loss as L_t and the loss size at the second period with a loss as $L_{t'}$ where $t < t'$. The expected gain in utility for the first period with losses is:

$$\begin{aligned} & (1 - (1 - \pi)^N) \left[\mathbb{E}[u(y - p - L_t + I'(t, \{L_\tau\}_{\tau \leq t}))] - \mathbb{E}[u(y - p - L_t + I(t, \{L_\tau\}_{\tau \leq t}))] \right] \\ & = (1 - (1 - \pi)^N) \mathbb{E}[u'(y - p - L_t + I(t, \{L_\tau\}_{\tau \leq t}))] d\epsilon \end{aligned} \quad (58)$$

The loss in utility for the second period with losses is:

$$\begin{aligned} & (1 - (1 - \pi)^N - N\pi(1 - \pi)^{N-1}) \left[\mathbb{E}[u(y - p - L_{t'} + I'(t', \{L_\tau\}_{\tau \leq t'})] - \mathbb{E}[u(y - p - L_{t'} + I(t', \{L_\tau\}_{\tau \leq t'})] \right] \\ & = (1 - (1 - \pi)^N - N\pi(1 - \pi)^{N-1}) \mathbb{E}[u'(y - p - L_{t'} + I(t', \{L_\tau\}_{\tau \leq t'}))] d\delta \\ & = (1 - (1 - \pi)^N) \mathbb{E}[u'(y - p - L_{t'} + I(t', \{L_\tau\}_{\tau \leq t'}))] d\epsilon. \end{aligned} \quad (59)$$

Combining these, we have

$$V(I') - V(I) = (1 - (1 - \pi)^N) \left[\mathbb{E}[u'(y - p - L_t + I(t, \{L_\tau\}_{\tau \leq t}))] - \mathbb{E}[u'(y - p - L_{t'} + I(t', \{L_\tau\}_{\tau \leq t'}))] \right] d\epsilon. \quad (60)$$

Since I is a straight deductible policy, we have

$$-L_t + I(t, \{L_\tau\}_{\tau \leq t}) \leq -L_{t'} + I(t', \{L_\tau\}_{\tau \leq t'}) \quad (61)$$

for any realization of losses (L_1, \dots, L_N) . This implies $\mathbb{E}[u'(y - p - L_t + I(t, \{L_\tau\}_{\tau \leq t}))] >$

$\mathbb{E}[u'(y - p - L_{t'} + I(t', \{L_\tau\}_{\tau \leq t'}))]$ and hence $V(I') > V(I)$.

Appendix B Empirical Appendix

B.1 Data for Section 6

Table B1: Plan Menus Used in Discrete Choice Models Estimation

Menu A					
Plan	Deductible	Coinsurance	Maximum OOP	Premium	
1	2000	0.052	3500	2016	
2	2000	0.02	3500	2340	
3	1000	0.052	3500	2324	
4	1000	0.125	3500	2576	
5	250	0.152	3500	3000	
6	0	0.163	3500	3000	
7	0	0.03	3500	3300	
8	3000	0.052	3500	1800	
9	3000	0.163	3500	1700	
10	2000	0.163	3500	1900	

Menu B					
Plan	Deductible	Coinsurance	Maximum OOP	Premium	
1	250	0.2	2000	1536	
2	1300	0.2	3000	950	
3	500	0.1	1500	1450	
4	750	0.1	1200	1390	
5	2500	0.25	2750	875	

Notes: Menu A was based off of Ericson, Kircher, Spinnewijn, and Starc (2021), while Menu B was based off of Samek and Sydnor (2020).

Table B2: Average Annual Spending (\$) of Commercially Insured Population, by Risk Type

Risk Decile	Age Decile										Total	
	1	2	3	4	5	6	7	8	9	10		
1	2096	2216	2434	2603								2293
2	2701	5211	6880	2722	2823							2734
3	6247	3095	3240	5600	2999	3142						3107
4	6571	10008	3444	3537	10763	3377	3760					3571
5	27300	7824	9339	3821	3892			4031	4253			3991
6		29932	31286	10713	4061	4176	5917		4548	4839		4443
7			7042	4120		4295	4598	4620	4745	9729		4591
8	6032	6694	8129	8074	8424	9055	9612	9787	5410	5330		6653
9	11306	10062	10475	10631	10437	10380	10855	10958	10618	10744		10668
10	26396	24654	27982	33813	37261	38978	39155	39204	38541	38966		35517
Total	5278	6041	6414	6699	7190	8175	9390	10480	11648	13159		8555

Notes: Total N = 11,196,057 commercially insured individuals from the 2018 Truven MarketScan Database. Risk decile based on the decile of risk score generated by the 2018 CMS-HCC risk adjustment software. Cells are blank when there are no patients with that decile of risk score and age.

Table B3: Robustness Analysis for Additional Discrete Choice Models

	Standard (1)	Cash on Hand (2)	Standard (3)	Cash on Hand (4)	Standard (5)	Cash on Hand (6)	Standard (7)	Cash on Hand (8)	Standard (9)	Cash on Hand (10)
Premium	-0.002 (2.E-06)	-0.006 (2.E-06)	-0.007 (3.E-06)	-0.010 (3.E-06)	-0.00001 (2.E-06)	-0.0021306 (2.E-06)	2.19E-06 (2.E-06)	-0.0008241 (2.E-06)	-0.000243 (2.E-06)	-0.052513 (1.E-04)
E(OOP)	-0.002 (2.E-06)	-0.002 (3.E-06)	-0.007 (3.E-06)	-0.004 (3.E-06)	-0.00001 (2.E-06)	0.0007953 (2.E-06)	3.17E-06 (2.E-06)	-0.0007235 (2.E-06)	-0.000244 (2.E-06)	-0.0290081 (7.E-05)
Var(OOP) (\$10,000s)	-0.0003 (2.E-05)	-0.011 (3.E-05)	-0.0004 (2.E-06)	0.005 (3.E-05)	-0.000001 (1.E-05)	-4.24E-03 (1.E-05)	-1.11E-05 (1.E-05)	1.22E-03 (1.E-05)	-3.74E-05 (1.E-05)	-4.99E-02 (6.E-04)
Deductible (\$1,000s)	0.002 (2.E-03)	-0.841 (2.E-03)	-0.020 (2.E-03)	-3.348 (2.E-03)	0.00045 (1.E-03)	-0.5881 (1.E-03)	6.02E-06 (1.E-03)	-0.2803 (1.E-03)	-1.45E-04 (1.E-03)	-8.111 (4.E-02)
Coinsurance (10 p.p.)	0.001 (1.E-03)	-0.492 (1.E-03)	-0.012 (1.E-03)	-0.892 (1.E-03)	0.00061 (8.E-04)	-0.4430923 (8.E-04)	-0.00023 (8.E-04)	-0.1993357 (9.E-04)	0.0000921 (8.E-04)	-3.724035 (2.E-02)
\$ of Premium Equivalent to:										
Increase \$1 in E(OOP)	\$1.00	\$0.40	\$1.00	\$0.36	\$1.10	-\$0.37	\$1.45	\$0.88	\$1.00	\$0.55
Increase \$1,000 Deductible	-\$1.00	\$143.11	\$2.81	\$323.39	-\$35.97	\$276.03	\$2.75	\$340.13	\$0.60	\$154.46
Increase 10 p.p. in Coinsurance	-\$0.40	\$83.73	\$1.68	\$86.15	-\$49.38	\$207.97	-\$102.91	\$241.88	-\$0.38	\$70.92
Variation from Baseline Model	Risk types defined by combination of common HCCs		CRRA $r = 0.9$ (less risk averse)		CRRA $r = 1.5$ (more risk averse)		CARA $r = 0.001$		Alternative Error Variances	

Notes: This table replicates the estimation procedure from the results in Table 2. All of the results in this table use Plan Menu A. The first two columns deviate from the benchmark analysis by using an alternative rule for assigning individuals to risk types based on grouping individuals into sets of common CMS hierarchical condition categories (HCCs). Columns (3)-(6) produce results returning to the baseline assignment to risk types but now changing the utility function from the baseline log utility to a CRRA utility function with coefficient of relative risk aversion set at 0.9 (Columns 3-4) or 1.5 (Columns 5-6). Columns 7 and 8 use a constant absolute risk aversion (CARA) utility function with the absolute risk aversion parameter set at 0.001. Finally, in Columns 9 and 10 we return to the baseline risk-group assignment and log utility but now use a very different specification of the variances for the choice model error term ε . For this robustness check we increase the variance for the standard model dramatically by making it 10x larger, while for the cash-on-hand model we reduce it dramatically making it 0.1x smaller. These levels of error variances are arbitrary but are used here to demonstrate that even substantial changes to our model of error variances does not have a substantive impact on the qualitative conclusions of this estimation exercise.

B.2 Additional Detail for Section 7

We fielded a survey using a Qualtrics panel in November and December of 2015 in order to investigate the links between liquidity constraints, insurance demand and more generally preferences for smoothing spending shocks through paycheck withdrawals. The full text of the survey is included at the end of this Appendix section.

We contracted with Qualtrics to provide a 200-person sample and ended up with a final sample of 206 respondents. We limited the sample to those between 18 and 65 years old. We also targeted specific enrollment percentages by gender, age, and household income in order to get a sample that was similar to the overall U.S. working-age population on those characteristics. In order to ensure valid data, we also included two aggressive attention screeners in the survey and only those who passed both of those screeners and who took at least one third of the median time for the survey from a controlled pre-test (11 minutes) were included in the final sample.

Appendix Table B4 gives summary statistics on the self-reported demographic characteristics of the survey respondents. The balance targeting was successful, as age and household income breakdowns are close to those reported in the 2013 American Community Survey. Most importantly, the survey provides a sample with substantial diversity in age, income and other characteristics. While not a fully representative sample, this gives us increased confidence that the results from this survey are likely to be more broadly applicable to the U.S. population. Of course, it is important to remember that the Qualtrics Panels are volunteer panels and as such the respondents are a somewhat selected group even after attempts to obtain balance on a few target variables. For example, we find that 55% of our sample has an associate degree or higher, which is around 10% higher than we would expect from 2015 Census reports. On the other hand, we find that just under 70% of our sample reports private health insurance coverage (employer sponsored and exchange markets) and 11% are uninsured in 2015, which are both close to official statistics for the U.S. population in 2015.

Table B4: Demographic Characteristics of Survey Respondents

Balance-targeted Variables		Non-targeted Variables	
Male	49%	Employment	
Age		Full time	53%
18-24	15%	Part time	17%
25-34	21%	Not employed	29%
35-44	21%	Education	
45-54	23%	HS or less	22%
55-64	19%	Some college	22%
		2 or 4-year degree	48%
<\$15,000	13%	Advanced degree	7%
15,000–24,999	11%	Married	52%
25,000–49,999	24%	Has children under 24	53%
50,000–99,999	31%	Health Insurance Coverage	
\$100,000 +	21%	Private coverage	68%
		Public coverage	13%
		Other coverage	8%
		Uninsured	11%

Notes: N = 206.

Table B5: Survey Results: Sources Other than Liquid Money

	Measures of desire for insurance to smooth consumption			
	(1) Chose dominated health plan	(2) Find argument for dominated plan persuasive	(3) Chose and agree w/ dominated (Combo 1 + 2)	(4) Prefer rebate to deductible
Overall mean of dependent var:	0.54	0.33	0.27	0.34
Estimated effect of going from 0 to 100% share of \$1k bill paid with:				
Consumption cuts	0.28 [-0.09, 0.65]	0.09 [-0.29, 0.48]	0.23 [-0.15, 0.61]	0.11 [-0.28, 0.50]
Borrowing on credit cards	0.41 [0.14, 0.68]	0.49 [0.17, 0.81]	0.56 [0.25, 0.86]	0.36 [0.10, 0.63]
Borrowing from family	0.04 [-0.27, 0.35]	0.02 [-0.28, 0.32]	0.06 [-0.20, 0.33]	-0.03 [-0.34, 0.28]
Increasing work	0.33 [-0.11, 0.77]	0.11 [-0.29, 0.50]	0.17 [-0.24, 0.57]	0.18 [-0.25, 0.62]
Other sources	0.31 [-0.09, 0.71]	0.18 [-0.20, 0.57]	0.23 [-0.14, 0.60]	-0.15 [-0.42, 0.13]
Selling property	-0.26 [-1.04, 0.52]	-0.13 [-0.71, 0.46]	-0.11 [-0.67, 0.45]	0.49 [-0.22, 1.21]
Borrowing on payday loans	-0.19 [-0.82, 0.44]	0.02 [-0.68, 0.73]	0.07 [-0.58, 0.71]	0.39 [-0.37, 1.15]
Number of observations	206	206	206	206

Notes: Linear regression with heteroscedasticity-robust standard errors in parentheses and 95% confidence intervals in square brackets. Each regression includes controls for household income in discrete income bands (<\$15k, \$15k-\$25k, \$25k-\$50k, \$50k-\$100k, >\$100k), discrete age bins (18-24, 25-34, 35-44, 45-54, 55-64, 65+), discrete education levels (less HS, HS/GED, some college, 2-year college degree, 4-year college degree, masters degree, doctoral degree, professional degree), and self-reported current health status (excellent, very good, good, fair, poor). The independent variables included in the regression are the share of a hypothetical bill due within one month that the individual states they would pay using the listed source, with an omitted category of the share the individual states they would pay with money on hand. The shares are divided by 100 so that the estimated effect coefficient reflects the effect of going from 0 to 100% share paid with that source (holding fixed other sources other than liquid money). The dependent variables are as follows: (1) indicator for selecting a deductible of less than \$1,000 in a 4-option hypothetical menu of health plans (dominated options); (2) Indicator for finding an argument in favor of choosing dominated plans for budgeting reasons persuasive; (3) An indicator for both choosing a dominated option and finding the argument for it more persuasive, (4) An indicator for stating a preference for a hypothetical health insurance plan with higher premiums and a rebate at the end of the year over an equivalent plan with a deductible.

Survey Instrument Used In Section 7

Welcome to our survey.

Please first answer a few short demographic questions to see if you qualify to take the full survey.

What is your gender?

- Male
- Female
- Other/Prefer not to state

What is your age?

- Under 18
- 18-24
- 25-34
- 35-44
- 45-54
- 55-64
- 65+

What is your yearly household income level?

- Less than \$15,000
- \$15,000 - \$24,999
- \$25,000 - \$49,999
- \$50,000 - \$99,999
- \$100,000 and above

What is your current employment status?

- Working full-time
- Working part-time
- Looking for work
- Not currently working or looking for work

What is the highest level of education you have completed?

- Less than High School
- High School / GED
- Some College
- 2-year College Degree
- 4-year College Degree
- Masters Degree
- Doctoral Degree
- Professional Degree (JD, MD)

What is your current marital status?

- Married
- Unmarried

How many children do you have under the age of 24?

- 0
- 1
- 2
- 3
- 4+

Based on your answers, you qualify to participate in the study. The next page shows you a brief university informed consent statement and after that you can complete the remainder of the survey.

INFORMED CONSENT

You are being asked to participate in a research study. Please read this form before agreeing to the research. If you have questions, please direct them to the Principal Investigator, Dr. Justin Sydnor of the University of Wisconsin, Madison via email at (Justin.sydnor@wisc.edu)

Purpose: The purpose of the research is to learn about how people select health-insurance plans.

Procedures: The survey should take you around 10 minutes to complete.

Risks and Benefits: This research has no foreseeable risks or direct benefits to you.

Compensation: If you complete the entire survey, you will be compensated for your time as communicated to you through the standard Qualtrics Panel agreement.

Confidentiality: The records of this research will be kept private and the survey does not ask for information that could be used to identify you.

Voluntary nature of study: Your participation is voluntary and you may withdraw from the study at any time. If you choose not to participate, it will not affect your current or future relationship with the University of Wisconsin, Madison. In order to fully participate and receive compensation it is necessary to complete the survey and it is generally not possible to skip or decline to answer the questions.

Contacts and questions: The researcher conducting this study is Dr. Justin Sydnor. If you have any questions, concerns or complaints about the study, you may contact him at Justin.sydnor@wisc.edu. If you are not satisfied with the response of the research team, please contact the University of Wisconsin, Madison Education and Social/Behavioral IRB Office at 608-263-2320.

Statement of consent: By clicking next to continue with the survey you are attesting that: 1) you read this informed-consent document, 2) you are at least 18-years old, and 3) you voluntarily consent to participate in the study.

Welcome to our survey about financial decisions.

We appreciate your participation. Please read each of the questions carefully.

Many of the questions ask you what you prefer between a few options.

Also, please note that there are attention checks in the survey, so please read carefully.

The first part of the survey asks you about which health insurance plans you would prefer from different possible options.

Before we get to those questions, we want to define a few terms we will use throughout. Please read these definitions and then answer two questions below to check that you understand.

Insurance terms we use in this survey:

Deductible: Total amount of medical bills you have to pay for the year before insurance coverage starts. For the plans in this survey, after you have hit the deductible, insurance pays for all medical bills for the rest of the year.

Premium: Amount of money you pay for the insurance. The premium is withdrawn in even amounts from your monthly paychecks.

Based on the information above, which of these is true about the deductible for the health plans we ask you to consider in this survey?

- Every time you go to the doctor you have to pay the amount of the deductible
- The deductible is how much you pay from your paychecks for the insurance
- After you have paid medical bills equaling the deductible, the rest of your medical bills for the year are fully covered by insurance.
- None of these are true
- All of these are true

Insurance terms we use in this survey:

Deductible: Total amount of medical bills you have to pay for the year before insurance coverage starts. For the plans in this survey, after you have hit the deductible, insurance pays for all medical bills for the rest of the year.


Premium: Amount of money you pay for the insurance. The premium is withdrawn in even amounts from your monthly paychecks.

Based on the information above, which of these is true about the premium for the health plans we ask you to consider in this survey?

- The premium is what you pay for the insurance and is withdrawn monthly from paychecks
 - The premium is the amount of medical bills you have to pay for the year
 - The premium is what you pay for the insurance and is paid in full in January
 - None of these are true
 - All of these are true
-

Suppose you had the choice of the following health insurance plans from your employer. All plans have the same access to doctors and hospitals and differ only on annual deductible and annual premium. Which plan would you choose?

- \$350 annual deductible with a cost of \$1,957 in annual premium (\$163/month)
 - \$500 annual deductible with a cost of \$1,419 in annual premium (\$118/month)
 - \$750 annual deductible with a cost of \$1,321 in annual premium (\$110/month)
 - \$1,000 annual deductible with a cost of \$817 in annual premium (\$68/month)
-

(Randomized Order of Arguments) 

Sam has a choice between two plans:

Plan A: \$500 deductible and \$1,500 annual premium (\$125/month)

Plan B: \$1,000 deductible and \$850 annual premium (\$71/month)

Which argument do you find more persuasive?

Argument A: Sam should choose Plan A because it is difficult to budget for out-of-pocket medical bills. It can be difficult to set aside money for unexpected bills. With Plan B, Sam might be hit with an extra \$500 in medical bills he cannot pay (\$1,000 instead of \$500 deductible).


Argument B: Sam should choose Plan B because the premium is so much lower. With Plan B, he will pay \$650 less in premium for the year (\$850 instead of \$1,500). That more than covers the extra \$500 deductible if he had to pay it.

- Argument A is much more persuasive
 - Argument A is somewhat more persuasive
 - Argument B is somewhat more persuasive
 - Argument B is much more persuasive
-

People are busy these days and do not always have time to research insurance options. While some have time to pay attention to their options, others may not even have time to read survey questions carefully. To show that you have read carefully, please select "Other approach" as your choice option. That's right, there is no question here -- just select "Other approach" to show you were reading carefully.

Which best describes your approach?

- Recommendation from a friend
- Recommendation from employer
- Recommendation from a third party source
- Research options myself
- Other approach

(Randomized Order of Deductible and Prepay Option) 

Suppose your employer is considering two different health insurance options. These options give you the same access to doctors and hospitals.

Deductible Option: \$1,500 annual deductible and annual premium of \$2,000 (\$167/month from paycheck withdrawals).

Prepay with Rebate Option: This plan is similar but you prepay for the \$1,500 in medical expenses and get a rebate if you are under that amount for the year. You pay annual premium of \$3,500 (\$292/month). There is no deductible, so you do not pay anything when you go to the doctor or hospital. However, if you have total medical spending below \$1,500 at the end of the year, you will get a rebate equal to the difference between \$1,500 and your medical spending for the year.

Which option do you prefer?

- Strongly prefer the deductible option
- Somewhat prefer the deductible option
- Neutral/indifferent between options
- Somewhat prefer the prepay with rebate option
- Strongly prefer the prepay with rebate option

If you were sick and deciding between going to the doctor or not, do you think you would be more likely to go to doctor if you had one of these plans?

- More likely to go with deductible option
- Equally likely to go with either option
- More likely to go with the prepay and rebate option

How would you rate your current health status?

- Excellent
- Very good
- Good
- Fair
- Poor

Do you currently have health-insurance coverage?

- Yes
- No

From what source do you get your health-insurance coverage?

- your employer
- the employer of a spouse/partner
- through a state or federal insurance marketplace
- public program (e.g., Medicare, Medicaid)
- another source

How confident are you that you could come up with \$2,000 if an unexpected expense arose and had to be paid within the next month?

- I am certain I could come up with the full \$2,000
- I could probably come up with \$2,000
- I could probably not come up with \$2,000
- I am certain I could not come up with \$2,000

If you had to go to the emergency room and the hospital billed you for charges that were not covered by insurance, how long do you think you would have before you had to pay the bill in full?

- Less than a month
- 1 month
- 2 months
- More than 2 months

Suppose you had to go to the emergency room because of an accident and just got a bill from the hospital for \$1,000 that is not covered by insurance and is due within a month.

What percent of the \$1,000 hospital bill would you cover from each of these sources?

(Total must add to 100 percent)

Money you already have (e.g., savings/checking account): _____

Extra money you save by pulling back on spending: _____

Extra money you earn by working more: _____

Borrowing from friends/family: _____

Borrowing using credit cards or home equity lines: _____

Borrowing using payday or pawn-shop loans: _____

Selling things you own: _____

Money you get from other sources: _____

Total: _____

Which of these actions would be your most likely approach if you received this \$1,000 hospital bill?

- I would pay the full bill by the due date
- I would pay at least half, but not the full bill, by the due date
- I would pay some of the bill, but less than half, by the due date
- I would not pay any of the bill by the due date

Imagine you found out today that because of a mistake in your taxes, the government was going to force you to make a payment of \$500 at some point next year. There is no way to know when they would force you to make the payment, but it would occur at some point during 2016 and you would have to pay in full within a week.

If you had the option to instead pay by having your employer make equal withdrawals from your paycheck over the full year next year, which would you prefer?

	1	2	
Paycheck withdrawals totaling \$500	<input type="radio"/>	<input type="radio"/>	Pay the full \$500 when told to
Paycheck withdrawals totaling \$520	<input type="radio"/>	<input type="radio"/>	Pay the full \$500 when told to
Paycheck withdrawals totaling \$540	<input type="radio"/>	<input type="radio"/>	Pay the full \$500 when told to
Paycheck withdrawals totaling \$560	<input type="radio"/>	<input type="radio"/>	Pay the full \$500 when told to
Paycheck withdrawals totaling \$580	<input type="radio"/>	<input type="radio"/>	Pay the full \$500 when told to
Paycheck withdrawals totaling \$600	<input type="radio"/>	<input type="radio"/>	Pay the full \$500 when told to

Suppose now that you are considering a trip to the doctor for a routine procedure, but then the researchers want to know if you are still reading these questions carefully. If you are still reading carefully, please select "other" for the choice below. That's right, this is just an attention check question. Please select "other" if you are reading this.

What is your preference?

- Prefer the emergency room
- Prefer an outpatient clinic
- Prefer a visit to my primary-care doctor
- Other

On the next few pages we will ask you to answer a series of questions about experiences that could happen either sooner or later. For each question we want you to think about whether you would prefer the sooner or later experience.

Please pay attention to the timing of these options as you answer the questions.

If you had the choice between these two prizes that would be delivered at different times, which would you prefer?

- \$20 guaranteed today
- \$30 guaranteed one month from today

If you had the choice between these two prizes that would be delivered at different times, which would you prefer?

- \$20 guaranteed in 6 months
- \$30 guaranteed in 7 months

If you had the choice between these two prizes that would be delivered at different times, which would you prefer?

- \$50 guaranteed today
- \$60 guaranteed one month from today

If you had the choice between these two prizes that would be delivered at different times, which would you prefer?

- \$50 guaranteed in 4 months
- \$60 guaranteed in 5 months

Please let us know if you have any comments about the survey, especially if anything was unclear or confusing. _____

Thank you for completing this survey.