

The Welfare Implications of Health Insurance

Sonia Jaffe and Anup Malani¹

Abstract. We analyze the financial value of insurance when individuals have access to credit markets. Loans allow consumers to smooth shocks across time, decreasing the value of the smoothing (across states of the world) provided by insurance. We derive a simple formula for the incremental value of insurance and show how it depends on individual characteristics and the features of available loans. Our central contribution is to derive formulas for aggregate welfare that can be taken to data from typical studies of health insurance. We provide both exact formulas that can be taken to data on the distribution of medical expenditures and income and an approximate formula for aggregate data on medical expenditure. Using the Medical Expenditure Panel Survey we illustrate how the incremental value of insurance is decreasing with access to loans. For consumers in the sickest decile, access to a five-year loan decreases the incremental value of insurance by \$338 (6%) on average and \$3,433 (36%) for the poorest consumers. We also find that our approximate formula is a reasonable proxy for the exact one in our data.

A central parameter of interest in health care policy is the value of health care insurance. Economic theory suggests health care insurance could improve both the health and the financial status of beneficiaries. Health care insurance lowers the price of medical care to beneficiaries on the margin and therefore may increase consumption of care. If increased care improves health, insurance may thereby improve health of beneficiaries. Moreover, health care insurance may be a more efficient means of raising capital for health care expenditures than alternatives such as current income, savings, or borrowing. Therefore, insurance may be able to reduce the overall cost of medical care – the price charged by the provider plus the cost of raising capital to pay that price – even for inframarginal consumption of care.

While there is some dispute about the impact of insurance on health in the U.S., there is a consensus that it has positive financial impacts on beneficiaries. There is consistent evidence that health care insurance increases health care utilization (Buchmueller, Grumbach, Kronick, & Kahn, 2005), but there is disagreement about impact of that health care on health. Reviews by Freeman, Kadiyala, Bell, and Martin (2008) and Levy and Meltzer (2008) found well-designed but non-randomized studies showing a positive effect on health. However, both the RAND and the Oregon health insurance experiments reported that increased utilization from insurance had

¹ University of Chicago. We thank Bartek Woda for excellent research assistance.

little effects on health, except for the poor² in the RAND study (perhaps because they were credit constrained) and for self-reported psychological ailments such as depression in the Oregon study (Finkelstein et al., 2012; Newhouse, 1993). By contrast, a wide array of studies – including the both the RAND and Oregon experiments – have noted that health insurance reduces out of pocket (OOP) payments and perhaps indebtedness.

The problem for economists, however, is that most studies of health insurance measure financial benefits in terms of OOP expenditures, but it is difficult to translate these into welfare.³ OOP payments imply foregone non-medical consumption, but the welfare effect depends on the amount of non-medical consumption sacrificed and the marginal utility of that consumption. These parameters in turn depend on the cost of financing those OOP payments, whether by saving, current income or borrowing. For example, if OOP payments are financed by borrowing, then the cost of financing OOP payments is lower when interest rates are low. Since insurance displaces OOP payments, this implies that the value of insurance falls as (real) interest rates fall.

We address this problem by offering a model in which consumers who might suffer a health shock have a choice of whether to finance medical care with insurance or with credit. We allow consumers to vary on medical risk, age, and income and we consider different models of insurance pricing. This model generates positive implications about who seeks insurance and who obtains loans. At its core, the choice depends on whether, given the ability to smooth foregone consumption over time, is it worth the added cost to also smooth over states of the world (i.e., other members of the pool). If so, then insurance is superior. If not, loans are preferred.

Moral hazard plays a complex role. Insurance creates the potential for moral hazard because it lowers the consumer's price for marginal care, perhaps inducing excess care. This inefficiency is reflected in a higher premium, which translates into a higher cost of financing care with insurance. Thus, moral hazard in insurance makes loans more attractive. Yet loans also face moral hazard. This comes from bankruptcy protection for debtors in case of default. The higher the risk of default, the higher the interest rate on loans to compensate creditors. Thus moral hazard in loans makes insurance more attractive.

Some implications of the model depend on the nature of insurance pricing. When insurance is community-rated, i.e., all member of an insurance pool pay the same premium regardless of their expected medical expenditures or "risk-rating", high risk individuals unsurprisingly prefer insurance to loans because their expenditures are cross-subsidized by the low risk. Less obviously, when insurance is experience-rated – an individual's premiums depends on her risk-rating – there are countervailing forces where high risk individuals may actually prefer loans

² Even these effects were not statistically significant. However, RAND could not rule out the possibility of medically significant benefits.

³ Knowing that insurance reduces indebtedness suffers the same problem. While reduction in indebtedness tells us how OOP were financed, it does not tell us the value of foregone consumption unless we know the cost of borrowed capital.

because the high risk implies less smoothing of costs across people/states of the world, whereas loans at least distribute those costs over time periods.

Regardless of how insurance is priced, older individuals tend to prefer insurance. The reason is that older individuals have fewer periods of life remaining in which to repay loans, so they have to concentrate their sacrifice of consumption for debt repayment in fewer years. The sacrifice of utility is thus greater for the old than the young for any given level of debt.

We then use our model to derive formulas for the welfare impacts of insurance. For the sake of realism in the US context, we focus on community-rated insurance and ask about the impact of switching from insurance to loans for one year in a population that is otherwise insured in other years. We present exact formulas and also approximations based upon Taylor expansions. The formulas derived from approximations depend only on the variance and average level of health expenditures and the coefficient of relative risk aversion rather than on the full distribution of potential shocks and a specified utility function. Both formulas require income, the interest rate and loan duration, and the markup on insurance.⁴

We apply our formulas to data from the Medical Expenditure Panel Survey (MEPS). For our exact formulas, we group households into bins of predicted health expenditures and use the distribution of actual shocks for each bin as an estimate of the distribution of possible shocks that a household in that bin faced ex-ante. Using CRRA utility we calculate the value of insurance for an individual for different loan lengths and cost assumptions. As expected, insurance is more valuable for sicker individuals, but only under community-rating. More interestingly, loans make insurance less valuable. In our baseline case, the ability to borrow decreases the incremental value of insurance by an average of \$65 (relative to no loans). This is relative to a baseline of \$1,147 for the average value of insurance to individuals without access to loans. For the sickest the average decrease is \$338 (6%) and the maximum is \$3,433 (36%). More generally, the value of insurance is decreasing in the length of time for which individuals can borrow, but increasing in the cost of financing loans.

We compare the results from the exact formulas to the results from our approximation formula to judge the accuracy of the latter. The approximation does reasonably well on average, but overestimates the value of insurance for the poorest individuals and the healthiest individuals; it underestimates it for the sickest. The approximation for the effect of the length of the borrowing period is also only accurate for those in the middle of the health distribution.

Our paper contributes to the sizable literature on the welfare implications of social insurance, particularly health insurance. As Einav, Finkelstein, and Levin (2010) note in their comprehensive review, nearly all existing papers in this literature focus on inefficiency in the market for insurance due to adverse selection (consumers' private information about their own

⁴ In a future draft we plan to adapt the formula to public insurance and account for the cost and benefits of insurance premium subsidies financed by public funds.

risk).⁵ The benchmark for these studies is either insurance with complete information or another insurance contract with private information.⁶

By contrast, our paper uses as a benchmark the case of no (or less) insurance, i.e., it seeks to calculate the value of pure insurance. Here its main contribution is elaborating on the idea that the better are other consumption smoothing tools, e.g., credit markets, the lower the value of insurance. Another contrast is that, to the extent we examine an inefficiency from insurance, it is due to moral hazard not adverse selection.⁷ (In a future draft we will account for adverse selection along the lines that Handel, Hendel, and Whinston (2015) do, though our aim will remain understanding how credit markets affect the value of insurance.)

This paper is closely related to Finkelstein, Hendren, and Luttmer (2016), which uses the approach of sufficient statistics (Chetty, 2009) to value Medicaid using data from the Oregon Health Insurance Experiment. Whereas the present draft of our paper focuses on the value of private insurance, Finkelstein et al. (2016) examine the net welfare benefits of subsidized (free) public insurance, i.e., the pure insurance plus the transfer benefit from Medicaid minus the cost of public funds to finance those benefits. There is nonetheless an overlap because both examine consumption with and without insurance to estimate the welfare value of pure insurance, particularly the value of reducing OOP medical expenditures.

There are multiple ways to estimate the impact of insurance on consumption. One approach Finkelstein et al. (2016) employs is to view OOP payments as lost consumption. This approach implicitly assumes that consumers have no access to credit markets and cannot smooth OOP costs over time. This may overestimate the value of insurance. Another approach that paper employs is to directly examine consumption. That consumption reflects the savings or borrowing that was enlisted to smooth consumption. However, if one cannot measure all consumption lost in past or future periods of due to health shocks, one may underestimate the value of insurance. We follow a third approach, which is to reallocate OOP payments over time assuming that consumers borrow to smooth consumption. We allow for the fact that because of borrowing, concurrent consumption is reduced by less than the full amount of the OOP payments (like the second approach in Finkelstein et al. (2016)), and we also capture the effect of that borrowing on consumption in later time periods when the loan taken out to cover OOP payments

⁵ For example, Einav, Finkelstein, and Cullen (2010) use variation in prices to identify consumer demand and marginal and average costs to estimate the welfare loss from adverse selection in private health insurance. See also Carlin and Town (2009).

⁶ For example, Bundorf, Levin, and Mahoney (2012) examine impact of private consumer information (or the equivalent, non-utilization by sellers of information about consumers in pricing) and compare welfare amongst employer plans with community or experience rating and plans with uniform and non-uniform employer contributions.

⁷ Whether that selection is on risk or propensity for moral hazard (Einav, Finkelstein, Ryan, Schrimpf, & Cullen, 2013).

is repaid.⁸ However, this approach is only as accurate as its assumptions about the type of loans consumers can and do access.

Our model is similar to the model Dobkin, Finkelstein, Kluender, and Notowidigdo (2016) employ to understand how consumers finance hospital bills. There are two important differences. First, Dobkin et al. (2016) emphasizes the impact of sickness on income. The paper examines how hospital shocks were financed using the Consumer Expenditure Survey and reported that borrowing declined after shocks among both the insured and uninsured, a result suggesting that shocks reduce income and thereby capacity for borrowing.⁹ Our model does not presently account for income shocks or the effect that those have on borrowing. That said, just as individuals are able to borrow less because of income shocks, they are also able to insure less as these shocks cannot be smoothed with insurance. Second, Dobkin et al. (2016) do not treat debt collection by providers as credit, though it is trade credit and an important consequence of the Emergency Medical Treatment and Labor Act of 1986 (EMTALA), which requires hospitals to treat first and collect payment later, i.e., to extend credit to their customers. Our assumptions about access to credit implicitly account for trade credit as they do not specify the source of the credit.

Our overall analysis is most closely related to Handel, Hendel, and Whinston (2015). That paper calculates the welfare value of one-year insurance contracts under community rating versus experience rating.¹⁰ It also compares these two pricing schemes when consumers have access to credit. It reports that the relative advantage of community rating declines with access to credit because saving reduces the relative cost of reclassification risk. Unfortunately, the paper provides an estimate of the difference in value between the two insurance contracts with credit and without credit, but not the absolute value of a given insurance contract with and without access to credit, the focus of this paper. Going beyond Handel et al. (2015), this paper also considers how different aspects of the credit contract (interest rate, repayment period) affect the value of each type of insurance contract.

⁸ As Finkelstein et al. (2016) say, their first approach over estimates the value of insurance because it assumes that consumption is reduced by the full OOP payments. Their second approach potentially understates the value of insurance because it does not account for the consumption loss in later periods due to borrowing. Our approach is in the middle because we allow for the consumption loss to be smaller than the OOP payments, but also include the consumption loss in later periods.

⁹ It should be noted that their evidence does not conclusively demonstrate that people cannot borrow because of income shocks following hospitalization. The uninsured have a smaller absolute reduction in borrowing than the insured after a health shock, a result consistent with the uninsured borrowing more relative to the insured. This result is unconditional on income, but the uninsured have less income in their data so their borrowing capacity should be less not more than the insured. While the authors conduct a difference-in-difference analysis on matched individuals from insured and uninsured groups (reported in their appendix) which shows borrowing declines more amongst the uninsured on a percentage rather than absolute basis, it is absolute levels rather than percentages that matter for borrowing as hospital bills reflect costs rather than the income of beneficiaries.

¹⁰ Their main insight is that, whereas community-rated insurance suffers adverse selection, experience-rated insurance suffers reclassification risk and the welfare cost of reclassification risk outweighs the welfare cost of adverse selection.

The remainder of the paper has three parts. Section I presents a simplified version of our model of consumer choice between health care insurance and loans to cover treatment for health shocks. (A more realistic model is presented in the appendix.) We discuss the positive implications of the choice model. Section II derives formulas and approximations for the welfare value of health care insurance. Section III estimates formulas with data from MEPS.¹¹

I. Theory

To illustrate the basic tradeoffs, this section presents a simplified model of consumer choice between health care insurance and loans for financing the medical care required to treat a health shock. Simplicity is achieved mainly via two assumptions. First, we assume that individuals face only one potential shock per period. Second, we focus the model on the choice of insurance or loans in a single period of the model, assuming that the consumer has insurance in all other periods. In the appendix we present a somewhat more general model. In addition, our numerical analysis uses a model with a distribution of shocks, relaxing the first assumption.

To provide a conservative estimate of the value of credit markets we also assume that the individual only has access to loans to cover the cost of medical care in case she suffers a health shock; she cannot use the credit markets to more generally smooth consumption across time. Because we hobble credit markets in this manner and the value of insurance falls with the value of credit markets, the model here provides an upper bound on the value of insurance. The appendix also presents results from a model that allows credit markets to be used more generally to smooth consumption, providing a lower bound on the value of insurance.

A. Simplified model

We will model an individual from a larger population of heterogeneous consumers. For simplicity, all consumers have an identical per period utility function $u(c)$ that draws value exclusively from consumption, not from health. To abstract from the issue of state-dependent utility, health shocks are modelled as health expenditure shocks that reduce overall consumption. All consumers share a discount factor of β .

Consumers vary in age, income, and the health risk that they face. The individual we model will live for T more periods. If T is large, we have a younger individual and, if T is small, an older individual. In each period, the individual earns income y . Consumption is equal to income net of medical expenses and any credit market transactions. Each period, with probability π , the individual incurs a medical expense that reduces her income available for consumption by p . In the population at large, consumers may vary in T , y , and π .

The individual must decide how to smooth consumption given this shock. She has two options – purchase health care insurance or borrow money. We assume that in periods 2, ..., T , the individual chooses to purchase health care insurance. We focus on the consumer's choice in

¹¹ In future drafts we will extend the analysis to other settings.

period 1 of whether to buy insurance or take out loans to cover the medical shock. Naturally, this choice must be made before the period 1 shock is realized.

Insurance. If the consumer chooses insurance, she pays a premium at the start of the period equal to $m^I q$, where q is the component of the premium that reflects expected costs absent insurance and $m^I(L^I, \epsilon)$ is the markup on insurance.

The expected costs q depends on how insurance is priced. If insurance is community-rated, then it depends on the average risk of the people in the insurance pool, i.e., $q^C = \pi^C p$, where π^C is the average π in the pool. Conversely, if insurance is experience-rated, then the expected cost depends only on the consumer's own risk, i.e., $q^E = \pi p$. For the main analysis, we will assume insurance is community-rated. Later we discuss how the results differ under experience-rated health care insurance.

The markup on insurance embeds the load L^I on insurance and moral hazard from insurance. The load includes administrative costs. The moral hazard is attributable to the excess consumption of care due to insured individuals facing, on the margin, a price that is less than the full marginal cost of care; it therefore depends on the price elasticity of demand for medical care, ϵ . The moral hazard cost is not simply the cost of additional care purchased by consumers under insurance because this excess care does have value. Rather it is the cost of the additional care (care beyond the level of consumption the individual would have chosen if she faced the full marginal costs of medical care), minus the dollar value of that care to the consumer. Obviously, the markup is increasing in the elasticity of demand for care: $\partial m^I / \partial \epsilon > 0$.

Loans. If the consumer chooses a loan, she must pay back $m^L p$ over n periods at an interest rate r . We require that $n \leq T$, so people with more remaining periods (younger people) have access to loans with a longer repayment period. The amount borrowed is p . The total amount to be repaid includes a factor m^L that captures the administrative costs of lending (e.g., underwriting), which for symmetry we shall call the load on loans.¹² The interest rate and repayment period affect the amount of consumption the consumer must forego each period for repayment of the loan. The repayment amount is $m^L \alpha p$ per period during periods 1 through n , where α satisfies

$$\alpha + \frac{\alpha}{1+r} + \frac{\alpha}{(1+r)^2} + \dots + \frac{\alpha}{(1+r)^{n-1}} = 1.$$

The αp component of the repayment is in some sense the fair component of the repayment which is multiplied by the load, m^L . Consistent with this idea, we assume that the interest rate corresponds to the individual's time preference, that is $\beta = 1/(1+r)$; this implies that $\alpha = (1-\beta)/(1-\beta^n)$.

¹² We don't explicitly allow for the possibility of default on a loan, but if existed this "moral hazard" on loans would be reflected in m^L . As with insurance, the cost of default would be the added borrowing cost minus the benefit the individual would receive from the ability to default. Since the focus of this paper is the value of insurance rather than the value of credit access, we keep credit pricing simple and do not allow for community-rated loans.

Choice. If the individual chooses community-rated insurance in period 1, then her utility is

$$V^I = \frac{1 - \beta^T}{1 - \beta} u(y - m^I \pi^c p).$$

The individual must pay a premium each period, which reduces her consumption by $m^I \pi^c p$.

If the individual instead chooses to borrow to smooth her period 1 shock, her utility is

$$V^L = (1 - \pi) \left(u(y) + \frac{\beta - \beta^T}{1 - \beta} u(y - m^I \pi^c p) \right) + \pi \left(\frac{(1 - \beta^n)}{1 - \beta} u(y - m^I \pi^c p - m^L \alpha (1 - m^I \pi^c) p) + \frac{\beta^n - \beta^T}{1 - \beta} u(y - m^I \pi^c p) \right).$$

If there is no shock, the individual loses no consumption in period 1, and must pay insurance premiums in periods 2 through T . If there is a shock, the individual pays a part of the shock (specifically, $m^I \pi^c p$) directly and borrows the remainder of the shock ($(1 - m^I \pi^c) p$) in order to smooth consumption across the first n periods. In periods 2 through n , she pays that period's premium and her per period loan repayment. Thus in periods 1 through n , the individual pays $m^I \pi^c p + m^L \alpha p (1 - m^I \pi^c)$. From period $n + 1$ onwards, the individual only pays her insurance premium.

The incremental value of insurance over the value of loans, $V^{UB} = V^I - V^L$, is

$$V^{UB} = (1 - \pi) (u(y - m^I \pi^c p) - u(y)) + \pi \frac{(1 - \beta^n)}{1 - \beta} (u(y - m^I \pi^c p) - u(y - m^I \pi^c p - m^L \alpha p (1 - m^I \pi^c))). \quad (1)$$

If the individual takes out insurance and there is no shock, she loses the premium $m^I \pi^c p$ in period 1. However, if there is a shock, then the individual gains for the first n periods because she avoids having to repay $m^L \alpha (1 - m^I \pi^c) p$ each period on a loan. We superscript the incremental value with UB because we have hobbled credit markets by allowing them to be used only for smoothing the health expenditure shock when it occurs, not for saving absent a shock. Since the value of insurance is falling in the value of credit markets, this provides some semblance of an upper bound on the value of insurance.

In general, loans only have value over insurance if loans have a cost advantage over insurance, meaning they have lower load or community rating makes insurance a bad deal for a low risk individual. To see this, consider the special case where there is no load on insurance or loans (i.e., $m^I = m^L = 1$) and we focus on the ‘‘average’’ individual whose actual risk is identical to average risk (i.e., $\pi = \pi^c$). Using $(1 - \beta^n)/(1 - \beta) = 1/\alpha$, the incremental value of insurance becomes

$$V^{UB} = (1 - \pi) (u(y - \pi p) - u(y)) + \frac{\pi}{\alpha} (u(y - \pi p) - u(y - \pi p - \alpha p (1 - \pi))), \quad (2)$$

which is positive by Jensen's inequality.¹³ The consumer already smooths across time by paying a constant premium each period so, when insurance is fair (which the assumptions for this special case guarantee), loans offer no cost advantage over insurance.

That being said, even if insurance is more valuable than loans, the incremental value of insurance is lower when consumers have access to loans than when they do not. To see this, note that no loans is equivalent to reducing n to 1, which in turn increases α to 1, and that $\partial V^{UB}/\partial \alpha > 0$. Intuitively, with no loans, consumers have worse options for smoothing in the absence of insurance.

B. Positive implications

Effect of risk. Risk is affected by both the probability of shock and the price of medical care. As the probability π rises, the incremental value of insurance rises: $\partial V^{UB}/\partial \pi > 0$. This is driven by the well-known point that community rating cross-subsidizes higher-than-average risk members of the pool.

With experience rating, things are more interesting. First, consider changing an individual's probability of a shock *only in the first period*. We show in the appendix that V^{UB} is concave in π_1 , the first period probability of a shock, and that $V^{UB} = 0$ at $\pi_1 = 1$. For the marginal person¹⁴ who is indifferent between insurance and loans, higher first period risk actually pushes them towards loans (because of the concavity). To see the intuition for this, consider what happens as π approaches 1, e.g., for chronic conditions. In that case experience rated insurance loses its smoothing advantage because the premium rises to the point where the individual pays the cost of the near-certain shock.

In reality, risks are strongly positively correlated across periods, so if we want to compare people with different risk levels, we need to consider the implication for higher risk in future periods as well. Changing an individual's risk *in future periods* creates a countervailing effect. The increased premium payments in future periods raises the marginal utility of income in those periods, meaning loan repayment is more costly in those periods. As a result, it is costlier to take out a loan in the first period, so insurance is more valuable. The net effect of changing the risk probability in all periods is ambiguous.

¹³ To see this, multiply through by α and rearrange Equation (2) to obtain

$$(\alpha(1 - \pi) + \pi)u(y - \pi p) - (\alpha(1 - \pi)u(y) + \pi u(y - p[\alpha(1 - \pi) + \pi])).$$

Note that $u(y - \pi p)$ in the first term is utility of the average consumption in the second and third terms, i.e., $u(y)$ and $u(y - p[\alpha(1 - \pi) + \pi])$, where the weights are $\alpha(1 - \pi)/(\alpha(1 - \pi) + \pi)$ and $\pi/(\alpha(1 - \pi) + \pi)$, the same relative weights as those on the second and third terms.

¹⁴ That is, the marginal person with $\pi > 0$. Someone with $\pi = 0$ is also indifferent between loans and insurance, but that case seems less relevant; the effect of increasing π for them depends on the loads on loans and insurance. If there is a fixed cost of insurance (rather than a multiplicative load), then $[V^{UB}]_{\pi_1=0} < 0$ and there could be another $\pi_1 > 0$ at which an individual is indifferent between insurance and loans and an increase in π_1 pushes them towards insurance. Concavity means that this will be at a lower π_1 than the one referenced in the main text.

These experience-rated dynamics are not purely hypothetical. While health care insurance in the U.S. is largely community-rated, in some pockets it is meaningfully experience-rated. For example, in pools that suffer a lot of adverse selection or in the small group market, individual experience can affect premiums. Moreover, even under the ACA, insurers are allowed to discriminate to some extent on age, family size, and health behaviors such as smoking.

Medical price inflation, reflected in a higher p also affects the risks people face. There is not a substantial difference between the effect of price under community-rated and experience-rated insurance, so we focus on the former. First consider changing the price only in the first period. The increased price is more costly in terms of utility when the person takes loans than when they chose insurance (since consumption is lower), so if loans are “costlier” ($\pi m^L > \pi^c m^I$) a higher price in the first period unambiguously increases the incremental value of insurance. However if insurance is “costlier” ($\pi^c m^I > \pi m^L$) there is a countervailing effect because the load is multiplicative. The higher the price, the more individuals have to pay for insurance above the “fair” price.

The fact that higher prices this period are associated with higher prices in the future again complicates things. Higher expected insurance premiums in years 2 through n raise the marginal utility of consumption in those periods, which makes taking out a loan in the first period more costly. Since higher prices in future periods increase V^{UB} the net effect of changing p in all periods is unambiguously positive if $\pi m^L > \pi^c m^I$, but otherwise could go either way.

Our model of financing care captures the idea from Nyman (1999) that insurance has “access value” when the individual cannot access very expensive care using loans because the price of that care exceeds her lifetime income and thus her ability to repay. To see this, note that, the individual lacks access to loans when $y < p(m^I \pi^c + m^L \alpha(1 - m^I \pi^c))$ as the individual cannot cover loan servicing. This inability to pay for the loan can be captured by assuming an infinitely negative utility for negative consumption. So if $p > y/(m^L \alpha + m^I \pi^c(1 - m^L \alpha))$, then V^{UB} is infinite as long as $p < y/(m^I \pi^c)$.¹⁵

Effect of age. Age operates to increase demand for insurance through two mechanisms. Most directly it decreases T . This in turn may lower n since necessarily $n \leq T$. As the loan period falls, the repayment rate α rises, which reduces the smoothing benefit of loans and thereby increases the incremental value of insurance.

A second mechanism is that age is empirically associated with greater π . Under community rating, older individuals are more likely to be above average risk and thus prefer insurance because their insurance premiums are subsidized by younger, lower risk individuals.

Effect of credit market access. The robustness of credit markets can be reflected both in the amount of time individuals can borrow for, n , and in the interest rate, r . For a fair interest rate,

¹⁵ These conditions are not mutually exclusive so long as $m^I \pi^c < m^I \pi^c + m^L \alpha(1 - m^I \pi^c)$, which holds so long as $m^I \pi^c < 1$, which is a necessary assumption for this question to be interesting since otherwise everyone would be better off paying p when they experienced a shock than paying the insurance premium $m^I \pi^c p > p$.

i.e., $1 + r = 1/\beta$, increasing the borrowing period allows individuals to smooth across more periods thereby increasing the value of loans and decreasing the incremental value of insurance. If interest rates are “unfair” so $1 + r > 1/\beta$, increasing r while holding n constant increases the cost of borrowing in that it increases the per period repayment on the loan, am^Lp . This in turn reduces utility under loans, increasing the incremental value of insurance.

These points are particularly important for low income countries where a central implication of under-developed credit markets is that loans are only available for short lengths of time, such as one year, and interest rates are high (Duflo, 2005). Improving credit markets such that n rises or r falls could substitute for increasing access to insurance.

These results are also relevant for high income countries, though less for future policy and more for understanding the demand of insurance. Over the last 50 years, consumer access to credit has improved – largely due to the proliferation of credit cards (Federal Reserve Bank of Philadelphia, 2015). Our model suggests, therefore, that the value of insurance – holding prices constant – should have fallen. (In aggregate the effect is likely swamped by high medical price inflation having increased demand for insurance during the same period.)

Effect of administrative costs and price elasticity. The value of insurance – however priced – is declining in the markup on insurance, i.e., in load and moral hazard, a well-known point. Since moral hazard is increasing in the price elasticity of demand for medical care, the value of insurance declines with that elasticity. Conversely, it is increasing in the markup on loans, i.e., administrative costs of banks. One thing that makes the latter somewhat less important than the former (even if $m^I = m^L$) is that, the markup on loans only has to be paid if loans are taken out, i.e., if there is shock, whereas the markup on insurance has to be paid in all states.

C. Extensions

Coinsurance and moral hazard. The model above assumes that insurance pays for the entire health shock. Many insurance plans require the individual to pay a fraction of the health expenditures as “coinsurance.” If insurance is “fair” (experience rating and no markup), this decreases the value of insurance since it only imperfectly smooths the shocks. However, when there is a markup on insurance ($m^I > 1$), individuals may prefer incomplete insurance since they do not have to pay the markup on the part of the price that they pay out of pocket.

One of the motivations for coinsurance is to reduce moral hazard. The more price-elastic consumers are, the more moral hazard there will be and the larger the effect of coinsurance on moral hazard. If ϵ is the arc-elasticity¹⁶ and δ is the coinsurance rate (often around 20%) then the relationship between the quantities demanded with and without insurance is

$$\frac{Q_{ins}}{Q} = \frac{1 + \delta - \epsilon(1 - \delta)}{1 + \delta + \epsilon(1 - \delta)}$$

¹⁶ The arc elasticity is $\frac{q_1 - q_2}{q_1 + q_2} / \frac{p_1 - p_2}{p_1 + p_2}$ (a negative number). Unlike an elasticity, it can be estimated even when one of the prices is zero.

We do not have the exogenous variation in our data to estimate ϵ so we use a range of values from the literature. In considering different co-insurance rates we account for both the direct effect of the copayments on consumption in different states and the effect, via moral hazard, on the price of insurance.

Income shocks. Health shocks can also affect individual's income. Dobkin et al. (2016) find that for those admitted to the hospital, the income declines by an average of 20% of the health expenditures. Income shocks are similar to co-insurance in being a risk that is not covered by insurance, however, we do not require that they be proportional to the health expenditure shock. In MEPS the income shocks are smaller, but still potentially important. Income shocks make loans less valuable relative to insurance because they increase the marginal utility of consumption (by lowering income) exactly in the states of the world where one must make loan payments.

II. Welfare formulas

In this section we derive formulas for welfare that can be taken to data. Because individuals in real life – and the data – face multiple health expenditure risks each period, we now allow individuals in the model to face a more general distribution of shocks. Specifically, we assume an individual faces many possible shocks, indexed by i , that occur with probability π_i and cost p_i . The value from Equation (1) of insurance relative to loans becomes

$$V^{UB} = \sum_{i: p_i < m^l E_c[p]} \pi_i (u(y - m^l E_c[p]) - u(y - p_i)) \quad (3)$$

$$+ \frac{1}{\alpha} \sum_{i: p_i > m^l E_c[p]} \pi_i (u(y - m^l E_c[p]) - u(y - m^l E_c[p] - m^l \alpha (p_i - m^l E_c[p]))),$$

where $E_c[p]$ indicates the expected medical expenditure for the community-rated risk pool. As before, the individual must choose whether to insure all her shocks or none of them. Even if an individual chooses not to purchase any insurance, she does not necessarily smooth all her shocks by borrowing. If an individual has a shock that is less than what the premium would have been (and the premium will be in later years) she just pays for it directly out of present consumption. In that case she is better off by $u(y - p_i) - u(y - m^l E_c[p])$ relative to buying insurance. If the shock is higher, she will pay the initial $m^l E_c[p]$ directly and borrow the remaining $p_i - m^l E_c[p]$. In this case she is worse off by $u(y - m^l E_c[p]) - u(y - m^l E_c[p] - m^l \alpha (p_i - m^l E_c[p]))$ relative to buying insurance. No health shock corresponds to $p_i = 0$; for that term in the first summation, we get back the $u(y)$ term seen in Equation (1).

If individual data on medical expenditures for a population is available, this quantity can be calculated directly (with data or assumptions on m^l , m^l , and α and an assumption on the functional form of u) as we do in the next section. The value of insurance can also be approximated using data on the mean and variance of medical expenditures for a population. One way to approximate the value of insurance is to take a second order Taylor approximation of

$u(y - p_i)$ and $u(y - m^l E_c[p] - m^l \alpha(p_i - m^l E_c[p]))$ around $\tilde{y} = y - m^l E_c[p]$. However, since with limited credit markets (for the upper bound on insurance value) the individual only borrows after large shocks and does not save after small ones, this would yield a somewhat messy formula with conditional means and variance (conditioned on whether price is above or below the community-rated premium). Those conditional moments are hard to measure from the data as they depend on knowing the full distribution of shocks (or assuming a convenient functional form).

Two assumptions that simplify the approximation are that (1) credit markets allow the individual to save (for n periods) as well as borrow to smooth the shock in period 1 and (2) $m^l = 1$, so that, the rate the individual pays when borrowing is equal to what they get when saving. Because the banking system is doing a bit more work, namely it allows both borrowing and saving, the relative value of insurance falls, so this valuation is below the upper bound in Equation (3). In this case the value of insurance relative to loans is approximately

$$V \approx \frac{1}{\alpha} \sum_i \pi_i \left(u'(\tilde{y}) [\alpha(p_i - m^l E_c[p])] - \frac{u''(\tilde{y})}{2} [\alpha(m^l E_c[p] - p_i)]^2 \right).$$

This weighted sum is an expectation from the perspective of an individual. If we then take the expectation across all individuals in the risk pool, we obtain

$$V \approx \frac{1}{\alpha} \left(u'(\tilde{y}) \alpha (1 - m^l) E_c[p] - \frac{u''(\tilde{y})}{2} \alpha^2 (V_c[p] + (1 - m^l)^2 E_c^2[p]) \right).$$

Pulling out $\alpha u'(\tilde{y})$ and normalizing by $u'(\tilde{y})$ yields

$$\frac{V}{u'(\tilde{y})} \approx -(m^l - 1) E_c[p] + \frac{1}{2} \frac{\gamma}{\tilde{y}} \alpha (V_c[p] + (m^l - 1)^2 E_c^2[p]), \quad (4)$$

where γ is the coefficient of relative risk aversion (and γ/\tilde{y} is the coefficient of absolute risk aversion). Intuitively, the value of insurance, in dollar terms, is declining in the expected cost (because the effect of load and moral hazard are multiplicative) and increasing in the variance of shocks and relative risk aversion. Income reduces the impact of risk but also reduces the marginal utility of consumption.

Estimating this requires much less information than the full formula does. One only needs data on r or β and the loan period (to calculate α), the markup m^l on insurance (including moral hazard), the community-rated expectation $E_c[p]$ and variance $V_c[p]$ of health expenditure shocks, and the coefficient of absolute risk aversion (or the coefficient of relative risk aversion γ and income net of the premium \tilde{y}). If one wanted an apples-to-apples comparison of the incremental value of insurance using individual data and the approximation, note that one would want to convert the former into a measure of dollars by dividing by $u'(\tilde{y})$.

From Equation 5 we can easily derive how much ignoring loans leads one to overestimate the welfare value of insurance. Without loans, the full shock must be borne in the present period regardless of interest rate; with loans the repayment rate is extended and the full shock can be

smoothed over n periods. Hence, the overestimate is Equation (4) when $\alpha = 1$ (for $n = 1$) versus α for larger n :

$$\Delta_n \left(\frac{V}{u'(\tilde{y})} \right) \approx \frac{1}{2} \frac{\gamma}{\tilde{y}} (1 - \alpha) (V_c[p] + (1 - m')^2 E_c^2[p]). \quad (5)$$

Because α is declining in the loan repayment period and rising in the interest rate, this overestimate is increasing in the repayment period and falling in the interest rate. Intuitively, the overestimate is proportional to the variance of the shocks, the level of risk aversion and, critically, the ability of loans to reallocate OOP payments to future periods at low capital cost.

III. Welfare estimates

A. Data

We use data from the Medical Expenditures Panel Survey from 1996 to 2014. For the present analysis we focus on the value of private insurance so we drop households with a member who is 65 or older at the end of the year, since those individuals are eligible for Medicare, and households with income below 138% of the federal poverty line who are eligible for Medicaid.¹⁷ We are left with 335,870 individuals in 136,299 households.

Since welfare is naturally very sensitive to consumption at very low levels, we censor the distribution of health shocks to ensure consumption above a certain level. Importantly, we also use this censored distribution in calculating the expected costs for the insurance premiums. In limiting the right tail risk, censoring inevitable decreases the value of insurance somewhat, but it would decrease it dramatically if we did not correspondingly lower the premiums. Also, since the effect of censoring is larger for sicker consumers, not adjusting the premiums would make insurance particularly unattractive for high risk consumers.¹⁸

Table 1 shows the summary statistics for the sample. It is about half female; the average age is 32. The average (median) individual income is \$33,889 (\$24,894) and for households it is \$81,770 (\$66,116). For privacy reasons, diagnoses are aggregated to Clinical Classification Codes. The average person has diagnoses in 1.38 of these 252 categories; the median person has 1. On average there are 2.44 people (including children) in a household.

B. Calculation using individual level data on shocks

In addition to the characteristics of the insurance and loan markets, the value of insurance relative to loans depends on the utility function, the distribution of health risks and income. We use a utility function with constant relative risk aversion

¹⁷ The interplay between back-stop government insurance and the value of private insurance is an important interaction we want to explore.

¹⁸ This is not quite equivalent to setting a consumption floor because the censoring of the shocks is not income dependent. The concern with using a consumption floor is that the distribution of shocks is then income dependent and the community-rated premiums become income dependent, but that is an alternative approach we intend to explore.

$$u = \frac{c^{1-\gamma}}{1-\gamma}.$$

We assume a coefficient of relative risk aversion, $\gamma = 2$.¹⁹

To estimate the part of an income change that is a shock, we first try to remove the lifecycle component for income changes. We regress the change in income on dummies for each age groups g and initial income

$$\Delta y = \alpha + \sum_g \beta_g + \beta_2 \cdot y + \epsilon.$$

We then calculate an adjusted income change

$$\widetilde{\Delta y} = \Delta y - \left(\hat{\alpha} + \sum_g \hat{\beta}_g + \hat{\beta}_2 \cdot y \right).$$

To estimate the distribution of medical expenditures and income shocks that each household faces,²⁰ we first regress, at the individual level, medical expenditures on age, sex, and dummies for diagnoses they had at the *beginning of the year* in which we observe expenditures.²¹ We then calculate predicted expenditures for each individual and aggregate them to the household level. For each household size, we divide households into deciles of predicted medical expenditures. We use the distribution of medical expenditures and income shocks in a given decile for the distribution of potential shocks that each household in that decile faced ex-ante.

Within a group of households that are all facing the same distribution of potential shocks, the only part of the value calculation that varies by household is income. Since calculating V^{UB} is somewhat computationally intensive we calculate it at the mid-point of each income ventile for each health expenditure risk group instead of for every household. To have the values in dollar units, which are more intuitive than utility units, we divide by the marginal utility of income and report $V^{UB}/u'(y - m^l E_c[p])$ in all the graphs and calculations. For our baseline we use $n = 5$, $m^l = 1.1$, and $m^l = 1.3$ as plausible values and we show how the results change as these vary.

Figure 1 shows the average across income ventiles of the dollar value of insurance relative to loans for different size households for each decile of health expenditure risk. It shows the value for both community and experience rating. As expected, the value is increasing in sickness for community rating. It is fairly flat for experience rating. For a family of two with community rating, the average incremental value of insurance ranges from -\$4813 for the healthiest decile to \$6,208 for the sickest.

¹⁹ Chetty (2006) argues that the coefficient of relative risk aversion is at most 2. Using this upper bound is favorable to the value of insurance.

²⁰ Though some of the calculations would be simpler at the individual level, members of a household naturally co-insure each other's risk, so we think this the appropriate unit of analysis.

²¹ The MEPS surveys each household in two consecutive years, asking about past health events and spending. We use conditions reported at the end of the first year and the expenditures reported at the end of the second year.

To avoid showing 4 graphs for every result, we focus on households with 2 people for the rest of the analysis. Results for other household sizes are similar. Also, consistent with the focus in the text, we use the values for community-rated insurance.

Our model suggests that ignoring loans leads to an over-valuing of insurance. Our calculations suggest that the value of insurance falls by \$65 on average if one goes from not allowing loans to allowing 5 year loans with a load of 1.1. For the sickest decile the decrease is \$338 (6% of the utility of insurance without loans) and the maximum (for the poorest consumers) is \$3,433 (36%). Going from no loans to loans means going from an effective m^L of 1 to an m^L of 1.1. To see the effect of loan length without this conflating factor we compare longer loans ($n = 5, 10$) to a two-year loan. Figure 2 shows the average change in the incremental value of insurance across sickness bins and income. As expected, increasing the length of time for which households can borrow decreases the value of insurance, particularly for the sickest people – who have the highest initial value of insurance under community rating – and for the poorest – who most value the decrease in annual payments from the longer loan.

Unsurprisingly, the value of insurance is decreasing in m^I and increasing in m^L . The left panel of Figure 3 shows the change in the value of insurance relative to $m^I = 1$ for $m^L = 1.1, 1.3$. The right panel shows the change relative to $m^L = 1$ for $m^I = 1.1, 1.3$. Under community rating, the effect of the insurance load does not vary much by health since everyone pays the same premium.²² The effect of m^L is largest for the sickest households since they have to pay the additional load on a larger (on average) loan payment.

C. Calculation using Taylor approximation

To see how accurate the Taylor approximation is, we compare it to the results using the complete data. To compare apples-to-apples we redo the above calculation allowing individuals to optimally save for n periods when they don't have a shock or have a shock smaller than the premium. We also set $m^L = 1$, to be comparable to the approximation.

Figure 4 shows the approximate value and the exact value for each health decile across income. The right panel shows the averages across health bins for each ventile of income. This averaging is not ideal because different health groups have different cutoffs for their income bins, but it allows us to see how the approximation does on average. The approximation overestimates the value of insurance for the poorest individuals. These are the ones for which the next term in the Taylor expansion would be most negative.

²² The effect is slightly increasing with health bin because those in the sicker bins tend to be wealthier on average and we normalize V by $u'(\tilde{y})$. Equal shifts in V generate a larger shift in $V/u'(\tilde{y})$ for those with lower marginal utility – those with higher income.

We can also compare the approximation of the effect of loan length, based on Equation (7), with the calculated value. Figure 5 shows this comparison.²³ Similar to the levels approximation, the approximation over estimates (in absolute value) the change for healthy groups and underestimates it for sick groups. The change in α resulting from the changed loan length is multiplied by the variance, which is much higher for the sicker groups than for the population as a whole, so an approximation using group specific aggregate statistics would be more accurate.

²³ Since we've set $m^L = 1$, for the approximation, we can now compare to 1 year loans. The graphs are truncated at the 8th decile so the relative values can be seen when they are close. The calculated change is roughly twice as large in the 10th as in the 8th decile.

References

- Buchmueller, T. C., Grumbach, K., Kronick, R., & Kahn, J. G. (2005). Book Review: The Effect of Health Insurance on Medical Care Utilization and Implications for Insurance Expansion: A Review of the Literature. *Medical Care Research and Review*, 62(1), 3-30. doi:doi:10.1177/1077558704271718
- Bundorf, M. K., Levin, J., & Mahoney, N. (2012). Pricing and Welfare in Health Plan Choice. *THE AMERICAN ECONOMIC REVIEW*, 102(7), 3214-3248.
- Carlin, C., & Town, R. (2009). Adverse selection, welfare and optimal pricing of employer-sponsored health plans. *U. Minnesota Working Paper*.
- Chetty, R. (2009). Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-Form Methods. *Annual Review of Economics*, 1(1), 451-488. doi:10.1146/annurev.economics.050708.142910
- Dobkin, C., Finkelstein, A., Kluender, R., & Notowidigdo, M. J. (2016). *The economic consequences of hospital admissions*. Retrieved from
- Duflo, A. (2005). Health Shocks and economic vulnerability in rural India: break the vicious circle. *Centre for Micro Finance Research Working Paper, Center for Micro Finance, IFMR*.
- Einav, L., Finkelstein, A., & Cullen, M. R. (2010). Estimating Welfare in Insurance Markets Using Variation in Prices*. *The Quarterly Journal of Economics*, 125(3), 877-921. doi:10.1162/qjec.2010.125.3.877
- Einav, L., Finkelstein, A., & Levin, J. (2010). Beyond Testing: Empirical Models of Insurance Markets. *Annual Review of Economics*, 2, 311-336. doi:10.1146/annurev.economics.050708.143254
- Einav, L., Finkelstein, A., Ryan, S. P., Schrimpf, P., & Cullen, M. R. (2013). Selection on moral hazard in health insurance. *THE AMERICAN ECONOMIC REVIEW*, 103(1), 178-219.
- Federal Reserve Bank of Philadelphia. (2015). Consumer Credit & Payments Statistics. Retrieved from <https://www.philadelphiafed.org/consumer-credit-and-payments/statistics>
- Finkelstein, A., Hendren, N., & Luttmer, E. F. (2016). *The value of medicaid: Interpreting results from the oregon health insurance experiment*. Retrieved from
- Finkelstein, A., Taubman, S., Wright, B., Bernstein, M., Gruber, J., Newhouse, J. P., . . . Group, O. H. S. (2012). The Oregon Health Insurance Experiment: Evidence from the First Year*. *The Quarterly Journal of Economics*, 127(3), 1057-1106. doi:10.1093/qje/qjs020
- Freeman, J. D., Kadiyala, S., Bell, J. F., & Martin, D. P. (2008). The causal effect of health insurance on utilization and outcomes in adults: a systematic review of US studies. *Med Care*, 46(10), 1023-1032.
- Handel, B., Hendel, I., & Whinston, M. D. (2015). Equilibria in health exchanges: Adverse selection versus reclassification risk. *Econometrica*, 83(4), 1261-1313.
- Levy, H., & Meltzer, D. (2008). The Impact of Health Insurance on Health. *Annual Review of Public Health*, 29(1), 399-409. doi:doi:10.1146/annurev.publhealth.28.021406.144042
- Newhouse, J. P. (1993). *Free for all?: lessons from the RAND health insurance experiment*: Harvard Univ Pr.
- Nyman, J. A. (1999). The value of health insurance: the access motive. *Journal of Health Economics*, 18(2), 141-152.

Table 1: Summary Statistics

Individuals N=335,870			
	Mean	Std Dev	Median
Age	31.81	17.86	33
Female	0.49	0.50	0
Income	33,889	39,810	24,894
# of conditions	1.38	1.76	1
Households N=136,299			
	Mean	Std Dev	Median
Family size	2.44	1.44	2
Income	81,770	58,192	66,116
Income per person	39,894	30,023	31,587
Medical expenditures	5,719	6,351	3,188
Medical Expenditures per person	2,514	2,780	1,432

Note: Incomes and expenditures are in 2014 dollars. Income is adjusted using the standard CPI and medical expenditures are adjusted using the “Medical care services” CPI. Averages are across all individuals, including children; “per person” includes children in the household.




Figure 1: Incremental value of insurance, relative to loans, under community and experience rating

Note: This figure shows, for each family size the value of insurance relative to a 5 year loan when the load on insurance is 1.3 and the load on loans is 1.1. The value is the average across the twenty income ventiles for each decile of health risk. Health risk is determined by predicting individual medical expenditures and aggregating to the household level. Households are grouped by predicted total expenditure and a household is assumed to face the distribution of medical expenditure shocks observed in its decile. Experience rating is based on the expected risk in each decile. Community rating is based on expenditures in the whole population.




Figure 2: How varying loan length changes the incremental value of insurance relative to loans
 Note: This shows how changing loan length changes the average incremental value of community-rated insurance over loans. See the note in Figure 1 for an explanation of the health risk deciles.




Figure 3: The effect of m^l and m^L on the incremental value of insurance relative to loans
 Note: This shows the effect of the load on insurance or loans on the average incremental value of community-rated insurance over loans. See the note in Figure 1 for an explanation of the health risk deciles.




Figure 4: Calculated and approximated value of insurance relative to loans

Note: This shows the value of insurance relative to loans calculated from Equation (3) and estimated from Equation 5. The left panel shows the calculated value separately for each health risk decile (see the note in Figure 1). The right shows the average across health risk for each income ventile.




Figure 5: Calculated and approximated effect of loan length on the incremental value of insurance, relative to loans

Note: Similar to Figure 2, this shows how changing loan length changes the incremental value of community-rated insurance over loans. The right panel shows the approximation from Equation (7). See the note in Figure 1 for an explanation of the health risk deciles. Both panels are truncated at the 8th decile.