Journal of Economic Perspectives—Volume 31, Number 4—Fall 2017—Pages 51–72

The Questionable Value of Having a Choice of Levels of Health Insurance Coverage

Keith Marzilli Ericson and Justin Sydnor

In most health insurance markets in the United States, consumers have substantial choice about their health insurance plan. In general, economists expect competition between insurance providers selling similar products to help hold down insurance costs, provided that people can make informed comparisons across insurers.

Health plans often differ both on the choice of insurance provider and also on a number of other dimensions. For example, plans differ in their *level of coverage,* like deductibles, limits on out-of-pocket payments, and the overall share of medical bills covered by insurance. The health insurance exchanges established by the Patient Protection and Affordable Care Act of 2010 feature four tiers of coverage—platinum, gold, silver, and bronze—which differ in actuarial value. The fraction of the population's medical bills that will be covered by insurance ranges from 90 percent for platinum plans to 60 percent for bronze plans. On HealthCare.gov, the average county has 46 health plans available across these tiers from five different insurers (Department of Health and Human Services 2016). This type of choice over coverage level is present in many other markets as well. For example, many employers offer a choice between a high- or low-deductible health plan, and over

■ *Keith Marzilli Ericson is Associate Professor of Markets, Public Policy, and Law, Questrom School of Business, Boston University, Boston, Massachusetts. Justin Sydnor is the Leslie P. Schultz Professor in Risk Management and Insurance, University of Wisconsin, School of Business, Madison, Wisconsin. Ericson is a Faculty Research Fellow and Sydnor is a Research Associate, National Bureau of Economic Research, Cambridge, Massachusetts. Their email addresses are kericson@bu.edu and justin.sydnor@ wisc.edu.*

† For supplementary materials such as appendices, datasets, and author disclosure statements, see the article page at
https://doi.org/10.1257/jep.31.4.51

doi=10.1257/jep.31.4.51

60 percent of people who get health insurance from their employer work in a firm that offers at least two different plan options (Claxton et al. 2016, p. 72).

Giving people choice over health insurance can have positive effects, but it also creates challenges related to both consumer confusion and adverse selection. There is mounting evidence that many people have difficulty understanding the value of insurance coverage, like evaluating the relative benefits of lower premiums versus lower deductibles. Also, in most US health insurance markets, people cannot be charged different prices for insurance based on their individual level of health risk. This creates the potential for well-known problems of adverse selection because people will often base the level of health insurance coverage they choose partly on their health status. As Geruso and Layton review in this issue, avoiding inefficiencies in these situations can require (sometimes imperfect) market regulations.

In this essay, we examine how the forces of consumer confusion and adverse selection interact with each other and with market institutions to affect how valuable it is to have multiple levels of health insurance coverage available in the market. We present an overview of how economists model the value of health insurance for a rational consumer and illustrate ideas using a set of simplified examples with parameters based on US data. We also review evidence on consumer confusion about health insurance and use an example simulation to illustrate the potential effect of consumer confusion on how the market functions.

We highlight a few key points. First, with fully informed consumers and no regulations to limit adverse selection, introducing plans with different levels of coverage can unsurprisingly lead to adverse selection and market unraveling that leaves consumers worse off (on average) than if only a single higher actuarial value plan were available. Having some confused and uninformed consumers can help prevent market unraveling in these situations, and thus raise average consumer welfare (Handel 2013). However, introducing choice across coverage levels can generate large transfers from uninformed consumers to sophisticated consumers. Market institutions, such as risk adjustment, that reduce how adverse selection affects plan prices can make offering choice over coverage levels beneficial to fully informed consumers. Yet these types of regulations are often imperfect and the gains to choice modest. Moreover, even with these regulations, confused consumers can erode any gains from choice.

We also briefly discuss how these issues can affect the value of choice on other dimensions of health insurance and health care. We discuss some policy options for addressing these issues of consumer choice, and some of the broader issues that the problem of consumer errors in making choices in health insurance and health care raises for economists and policymakers.

Modeling the Value of Health Insurance for Informed Consumers

We begin by reviewing the basic building blocks for how economists model the value of health insurance contracts: health risk, the level of coverage in the plans,

how those plans are priced, and how consumers value reductions in financial risk. A knowledgeable and informed consumer would take these factors into account in choosing a health insurance policy.

Building Block 1: Health Risk and Medical Spending

Some sources of variation in medical spending can be anticipated ahead of time, but some cannot. Anticipated variation comes from pre-existing conditions and other things that people (or insurers) know. We can think of this as variation between people who can be identified as more- or less-healthy at the time they choose insurance coverage. Unanticipated variation in medical spending is based on unexpected health shocks. For example, a young person with no chronic conditions can expect on average to have low health spending for the year, but might be involved in an accident that generates very large medical bills.

For our example, we want to capture realistic variation in both anticipated and unanticipated medical spending. We started with data from the Medical Expenditure Panel Survey (MEPS), which provides information about individuals' total medical expenditures (and other useful data) for a representative sample of Americans.[1] We focus on survey respondents of working age (18–64) who have health insurance coverage.

Survey respondents are asked, "In general, compared to other people of your age, would you say that your health is excellent, very good, good, fair, or poor?" For our simplified example, we group together people who choose excellent or very good and consider them "healthy" and those who answer good, fair, or poor as "unhealthy." Of course, dividing the population into these two types is a simplification. It ignores additional private information about health spending that an individual might have as well as demographic predictors of spending, such as age. While the question asks the respondents to report his or her health compared to other people of the same age, older individuals are still more likely to report lower health status.

As Table 1 shows, 68 percent of working-age adults are "healthy" and 32 percent are "unhealthy." These two groups differ substantially in their *anticipated* medical spending amounts. Healthy adults on average generate $3,045 in medical bills for the year, while unhealthy adults average more than twice as much.

Within health types, unexpected health shocks will still lead to substantial variation in the amount of medical bills a person has for the year. Table 1 gives a sense of the variability within each type—even unhealthy adults have a 9 percent chance of having no medical spending, and even healthy adults have a 7 percent chance of having over $10,000 in medical spending.

[1] The data is released publicly with a lag, and here we use data from the Agency for Healthcare Research and Quality (AHRQ) 2012 and 2013 Medical Expenditure Panel Survey yearly surveys. Because the MEPS data we use is a few years old and medical costs tend to rise over time, the estimates we provide here will underestimate current medical spending for the population. There is also some evidence (Aizcorbe et al. 2012) that the MEPS data tends to underestimate spending relative to data on claims from employer-sponsored insurance. The basic insights we gain from this exercise, however, are not affected by this issue.

Table 1
Medical Spending for Example of Population with Two Health-Risk Types

Population group	*Percent of population*	*Average medical spending*	*Probability of $0 medical spending*	*Probability of $10,000+ medical Spending*
All adults (18–64)	100%	$4,380	13%	10%
Healthy adults	68%	$3,045	15%	7%
Unhealthy adults	32%	$7,227	9%	18%

Note: Source is authors' calculations from 2012/2013 Medical Expenditure Panel Survey (MEPS) data. We limit the population to those with private insurance coverage for the full year. The split into "healthy" and "unhealthy" is described in the text. MEPS person weights are used to obtain the average spending amounts and the probability of spending thresholds.

When economists analyze health insurance markets, they typically assume that people are aware of the distribution of their possible medical bills for the year and choose their health plan with that information in mind. For our example simulation, when we consider fully informed consumers we assume that our two types of consumers—healthy types and unhealthy types—know the distribution of possible annual medical spending amounts their type could have. We use the observed variation in the data from the Medical Expenditure Panel Survey for each of the health types to create the distribution of possible medical spending each type might have for the year.

Building Block 2: Level of Financial Coverage and Cost-Sharing

The "actuarial value" of a health insurance plan is the percentage of the overall population's medical costs that would be covered by the plan. The person must pay the rest out of pocket, which is called "cost-sharing."

For our example, we consider two different coverage levels: a high and a low actuarial value plan. Each plan is defined by three typical cost-sharing features. The deductible is the amount of money you have to pay *for the year* in medical bills before insurance starts to (partially) cover additional bills. The coinsurance rate is the percentage of *each* bill the individual must pay out-of-pocket once the deductible has been met (the insurance covers the remaining portion). Finally, each plan has a maximum out-of-pocket limit. Once the combination of the deductible and coinsurance payments from the individual for that year hits this level, insurance fully covers all remaining bills. Table 2 describes our two example plans.

The high actuarial value plan, which covers 90 percent of medical costs, is similar to both a fairly generous employer-sponsored plan and to "platinum" plans on the health insurance exchange established by the Patient Protection and Affordable Care Act of 2010. The low actuarial value plan, covering 70 percent of medical costs, is a "high-deductible health plan," similar to the types of plans offered in the silver tier on private health insurance exchanges. High-deductible health plans have also become increasingly common in employer-sponsored plans over the past two decades.

Table 2
Two Examples of Health Plans

Plan type	*Actuarial value*	*Deductible*	*Coinsurance rate*	*Maximum out-of-pocket limit*
High actuarial value	90%	$250	10%	$1,250
Low actuarial value	70%	$2,000	10%	$4,500

Note: The actuarial value of our two example plans is calculated based on the sampling-weighted average amount of spending the plan would cover for people in the Medical Expenditure Panel Survey data.

Of course, real-world choices between insurance plans often include many different coverage levels (for reasons discussed by Geruso and Layton, in this issue). However, we believe that a simple example with two plans is quite helpful for thinking about how consumer choice and adverse selection patterns play out. In particular, the challenges of informed choice and adverse selection would likely be even more problematic with more choice between coverage options. National debates over health care reform during the first half of 2017 included proposals to lift restrictions on the range of coverage levels insurers could offer in the private health insurance market.

Building Block 3: How Insurance Premiums Are Set

Insurance premiums are largely determined by the expected amount of medical bills that will be covered by that insurance. Exactly how expected medical bills are linked to premiums, though, depends on how the market is regulated.

For this example, we focus on health insurance markets with two important regulations: *guaranteed issue*, which means that insurance plans have to sell to anyone who wants to buy, regardless of health status; and *community rating*, which means that insurance plans have to charge everyone in a given plan the same price, regardless of health status. These regulations are similar to circumstances in the employer-sponsored health insurance market and provisions of the 2010 Affordable Care Act.[2]

The premiums for insurance typically include a "load" on top of the medical costs covered by the insurer, which covers markups for insurer profit and administrative costs associated with processing claims. For our example, we assume a load factor of 1.25, implying that insurers charge $1.25 in premiums for each $1 in expected medical bills the insurance covers. This assumption is consistent with regulations from the Patient Protection and Affordable Care Act of 2010 that require insurers

[2] Technically, the Patient Protection and Affordable Care Act of 2010 has "modified community rating," as it permits premiums to vary based on a limited set of factors (for example, geography, age, smoking status) but not health status. Prior to this law, the individual insurance market typically did not have community rating or guaranteed issue, but employer-based insurance did.

to maintain a "medical loss ratio" (medical claim payments divided by premiums) of 80 percent or higher. The load factor is the reciprocal of the medical loss ratio.

For our example, assuming that plans covered the average population, the premiums for the high actuarial value plan would be set at $4,930, while the low actuarial plan would cost $3,810. These premiums come from taking the expected covered spending with each plan (the average medical spending multiplied by the plans' actuarial value) and then multiplying by the load factor of 1.25. These premiums are roughly in line with the types of premiums seen for real health insurance policies during this time period. They are somewhat lower, though, than is typical for employer-sponsored insurance, probably in part because the Medical Expenditure Panel Survey data tends to report lower total medical spending than data from insurance claims in employer-sponsored insurance (Aizcorbe, Liebman, Pack, Cutler, Chernew, and Rosen 2012).

These premiums are what we would expect if these plans enrolled the same average population. When people have choices over plans, the premiums for plans may be affected by the health status of people who enroll in the plan. We discuss this equilibrium process below.

Building Block 4: Risk Aversion and the Value of Insurance

In a standard model, more generous insurance is valuable primarily because it reduces the financial risk that people face (Einav, Finkelstein, and Levin 2010).[3] For someone who is risk averse, a health insurance plan that covers an additional $1,000 in expected medical bills will actually provide more than $1,000 in value because it also reduces the variation, or the risk, in spending the person faces. However, people differ in the optimal level of coverage they prefer.[4] People who are more risk averse will prefer plans with more coverage. Similarly, people who expect to have higher medical spending will see more value in plans with more coverage.

Economists typically model risk aversion using the idea of concave utility functions that capture the diminishing marginal utility of wealth. With a concave utility function, a person prefers more stable wealth over more variable wealth, even if it means giving up some wealth in expectation. Economists often turn to a few common mathematical functions for the concave utility functions that allow them to quantify the value people get from reducing risk. For our example, we use the constant absolute risk aversion utility function. With this utility function, the parameter r (known as the coefficient of absolute risk aversion) governs how risk averse a person is, with higher r implying more risk aversion. To put the coefficient of absolute risk aversion

[3] We are only discussing the choice between insurance plans of different coverage levels. Having insurance (as compared to being uninsured) may be valuable other reasons, such as access to care (for example, Nyman 1999) and access to negotiated rates for health care services.

[4] Individuals may also vary in how they value medical care. There is an interesting question of how variation in the price elasticity of medical care demand, which can be viewed as related to moral hazard, affects the value of choice (for example, Einav, Finkelstein, Ryan, Schrimpf, and Cullen 2013). There isn't a simple model we know of that can be used to determine this factor's effect on the value of choice, and for this essay, we set aside this issue.

into context, consider a person who holds a lottery ticket with a 50 percent chance of winning either $1,000 or nothing. A risk neutral person, with $r = 0$, values that lottery ticket at its expected value of $500. A risk-averse person, however, would be willing to sell the ticket for less than $500, to receive a certain gain and avoid the risk associated with the lottery. A person with $r = 0.001$, for example, would be willing to take as little as $380 for sure instead of the lottery.

There is no agreed upon range of typical risk aversion in the population. For our example, we assume that people vary in their level of risk aversion, which we assume is uniformly distributed between 0 (risk neutral) and $r = 0.001$. This level of variation allows us to model a population with substantial variation in risk aversion and is broadly in line with the range of risk aversion in the health insurance literature (for example, see the distributions compared in Ericson, Kircher, Spinnewijn, and Starc 2015).

We can then calculate the expected utility for both types of consumers—healthy and unhealthy—at different levels of risk aversion for both the high and low actuarial value plans. The expected utility for each plan is the weighted average of their utility for each level of total spending the person might have (premium + out-of-pocket costs) given the distribution of medical spending for their health type. Once we have these expected utilities for each plan, we can then calculate the dollar value of the difference in consumer welfare each person would have for different plans. We do this by calculating how much money could be given to (taken from), a person for the case of welfare gain (loss), to make them indifferent between staying in a baseline plan option versus moving to an alternative plan.[5]

The value of the higher level of coverage in our example is strongly affected by a person's health status. Assuming that plans were available at the premiums set for the average population described above, a risk-neutral healthy type would perceive that moving from the high to low actuarial value plan would increase their welfare by $350. This gain comes because the low actuarial value plan reduces the premium by $1,110, but increases a healthy types' expected out-of-pocket medical costs by only $760. On the other hand, a risk-neutral unhealthy person moving from the high to low actuarial value plan sees a small decrease in welfare: their expected out-of-pocket health care costs rise by $1,160, which is more than the reduction in premiums.

With a higher level of risk aversion, the perceived value of insurance increases. Healthy types with higher levels of risk aversion are worse off in the low actuarial value plan due to the increased financial risk they face, even though their expected total spending will fall. Moreover, unhealthy types with high risk aversion experience large losses (for example, $750 or more) in the low actuarial value plan, even though their expected total spending would increase by only $47. For our example

[5] The constant absolute risk aversion utility function is defined as $u(w) = 1 - e^{-rw}$ for $r > 0$ and $u(w) = w$ for $r = 0$. With this function, the expected welfare of having the low actuarial value plan relative to a benchmark of the high actuarial value plan can be calculated by first calculating the expected utility for each plan and then calculating: $-\frac{1}{r}\ln\left(\frac{EU(low\ AV)}{EU\ (high\ AV)}\right)$.

simulation, averaging over the different people, we find that average consumer welfare would be $127 higher with everyone in the high actuarial value plan than if everyone were in the low actuarial value plan. Thus, if a social planner had to select from only one plan from these two options, the high actuarial value plan would be preferred.

Evidence on Consumer Confusion in Health Insurance Choices

In the workhorse model of insurance choices, people are active deciders with full information who make choices about insurance plans based on their risk-averse expected utility over final wealth outcomes. However, recent empirical research offers evidence against this characterization. We first review some of that evidence, and then examine what happens to the value of choice if consumers are confused when selecting plans.

First, few people have a strong understanding of health insurance plans. For example, Loewenstein et al. (2013) found that in a representative sample, fewer than 14 percent of people could correctly answer a series of multiple-choice questions about key health-insurance terms, including deductible, copay, coinsurance, and maximum out-of-pocket costs. People in this survey were also overconfident, thinking they understood the terms better than they really did. Many people cannot easily calculate how much money they would spend with different plans even if they knew exactly what medical spending they would have. For example, Johnson et al. (2013) found that even with incentives, a majority of people could not identify the cost-minimizing plan from a few choices when given a specific amount of anticipated medical spending and details about premiums and cost-sharing. Even people with high education and income have difficulty understanding health insurance options. In the Johnson et al. (2013) study, MBA students were more likely to select cost-minimizing plans, but even in that group a significant share got it wrong. Handel and Kolstad (2015) found that among employees at a high-paying firm (median income around $125,000), a significant share were confused about details of different plan options.

Second, it is challenging for people to both forecast their possible range of medical needs for the year and also then to know how those needs would map to medical bills. The price of health care services can vary dramatically (Cooper, Craig, Gaynor, and Van Reenen 2015) and providers often cannot tell patients what the price will be in advance (Rosenthal, Lu, and Cram 2013). People also often have distorted perceptions of risk in the context of insurance, sometimes ignoring the possibility of low-probability events altogether and sometimes overreacting to certain salient risks (for example, Johnson, Hershey, Meszaros, and Kunreuther 1993). Although we are not aware of research on how people perceive their health risks when making health insurance decisions specifically, it is likely that they may be subject to similar biases.

Third, individuals are often not active deciders. Many people stick with an initial choice of an insurance plan, even if premiums or other features change dramatically. For example, Handel (2013) finds this pattern in a study of employees selecting plans offered by their employer, and Ericson (2014a) finds it in a study on choices that seniors make between Medicare Part D prescription drug plans. With consumer inertia, even people who originally make an optimal choice may not later be enrolled in the plans that offer the highest expected utility.

Fourth, an abundance of research in psychology and economics shows that people often become overwhelmed when faced with many options, a phenomenon known as "choice overload" (for discussion, see Iyengar and Kamenica 2010). Choice overload can cause people to gravitate toward simple options or to focus on isolated features of products in ways they would not do if there were fewer options available. It can even cause people to disengage entirely and opt not to purchase a product at all. To date, direct evidence of choice overload in health insurance is limited. Bhargava, Lowenstein, and Sydnor (2017) find that people made seemingly suboptimal choices when employers offered them many possible plans, but provide some evidence the number of options per se was not the problem. However, choice overload could be one reason behind consumer inertia in health insurance markets

With these challenges, a number of people may not be able to make choices that maximize their expected utility. For example, Abaluck and Gruber (2011) document that seniors often choose Medicare Part D prescription-drug insurance plans that are not on the efficient frontier: that is, they could purchase cheaper expected cost plans that provide equally good risk protection. Bhargava, Loewenstein, and Sydnor (2017) studied employees' plan choices at a firm where many of the available plans were financially dominated by other options: in one case, a plan with a $500 lower deductible cost an additional $600 in yearly premiums! Yet the majority of employees ended up choosing a dominated plan. Sinaiko and Hirth (2011) also document violations of dominance with a situation where many employees selected a plan with more difficult access to specialists, even though it had the same cost as a more flexible option.

Finally, some people may have some other objective rather than reducing variation in annual health-care spending. Out-of-pocket costs may have different consequences for people than spending on premiums, due to self-control issues, loss aversion, or liquidity constraints. For example, people may recognize that if they face out-of-pocket costs like deductibles and co-pays, they may pull back on valuable health care, like treatment for chronic diseases (Baicker, Mullainathan, and Schwatzstein 2015; Brot-Goldberg, Chandra, Handel, and Kolstad 2017). Liquidity constraints also can increase the value of paying smooth regular premiums over possibly more lumpy out-of-pocket costs, which in some extreme cases can make it rational to purchase a plan that otherwise seems dominated by a different choice (Ericson and Sydnor 2017). Economists are just beginning to incorporate these ideas into their analysis of healthcare markets.

The Value of Choice of Coverage Levels

In this section, we build upon the parameters laid out earlier to illustrate the effect of adding choice over coverage levels under different market institutions and different levels of consumer errors about plan choice. We take a situation where only the high actuarial value plan is available as the benchmark and then explore what happens when the low actuarial value plan is introduced.

The results of our example simulation are summarized in Table 3, which we refer to throughout this section. As described above, we assume that there are healthy and sick people in the market and use the Medical Expenditure Panel Survey to create the distribution of medical spending each type might have for the year. We assume that each person in the market has a level of risk aversion, uniformly distributed from risk neutral ($r = 0$) to substantially risk averse ($r = 0.001$).

Fully Informed versus Uninformed Choice

Fully informed individuals select between the high and low actuarial value options based on which one gives them the higher expected utility, given their health type, risk aversion, and the premiums for the two plans. We assume that fully informed people correctly anticipate the distribution of possible spending, pay attention to the premiums for the plans, and understand how the two plan options will affect their distribution of possible out-of-pocket medical costs.

For uninformed or confused consumers, we consider two ways these consumers might choose: randomly and nonrandomly. In our example, random choosers select each plan with 50 percent probability regardless of that person's underlying risk aversion or health status. For a useful example of a nonrandom error, we will focus on the "inattention heuristic," in which confused consumers select the plan they would have selected if plan premiums were fixed at the levels appropriate for the average population. This error can be thought of as representing a form of inertia: some consumers initially choose plans optimally when they are offered at population-average premiums, but then fail to pay attention if premiums change over time. It can also represent a bias people may show when selecting plans, if those who struggle to understand and compare plan options might naively believe that plans are priced "fairly" for the average population. In this setting, people may select plans based on their relative levels of health risk and risk aversion (what Kamenica 2008 terms "contextual inference"), but without properly incorporating the real differences in premiums into their selection.

There are, of course, many other ways people might choose plans. These two examples of alternative choice processes—random and "inattention heuristic"—however, allow us to discuss some forces that arise when considering how consumer confusion interacts with adverse selection and health insurance market institutions.

In our example, we show the welfare effects of introducing choice for these confused consumers, basing our calculation of welfare on how they would perceive the value if they were fully informed. If someone selects a plan they would not have if they were fully informed, we refer to that as a choice "error." However, we do not

Table 3

Example Simulations of the Effect of Introducing Choice Relative to Baseline with Only the High Actuarial Value Plan Available

Market environment	*Share of uninformed choosers*	*How the uninformed choose*	*Average change in per-person consumer welfare*			*Empirical pattern in equilibrium*
			Overall	*Healthy types*	*Unhealthy types*	
No regulation: premiums reflect average costs of people who enroll in that plan.	0%	—	*–$127*	$25	*–$449*	Market fully unravels so all choose low actuarial value plan.
	50%	Randomly	*–$78*	$31	*–$310*	Plan selection by health type, mitigated in part by random errors. Moderately elevated premium differences between plans. Both plans selected by some informed and some uninformed choosers.
	50%	Nonrandom: inattention heuristic	*–$59*	$151	*–$505*	Plan selection strongly related to health type and large resulting premium differences between plans. Only uninformed select the high actuarial value plan.
Regulation and risk adjustment: premium differences between plans fixed for average population	0%	–	$9	$42	*–$60*	Plan selection strongly related to health type. Low actuarial value plan attracts the unhealthy types as well as healthy types who are highly risk averse.
	50%	Randomly	*–$27*	$27	*–$143*	Plan selection by health type, mitigated in part by random errors.
	50%	Nonrandom: inattention heuristic	$9	$42	*–$60*	Same as the fully informed case because the inattention heuristic is appropriate to the market environment in this case.

Note: Authors' calculations based on simulation described in the text.

model the costs of processing information about health plans, so choosing with a heuristic may be sensible for some people given the challenges of trying to make an informed choice.

The Effects of Choice in Markets with No Regulations to Address Selection Effects

In a competitive insurance market, with no additional regulations, the premiums for an insurance policy will reflect the average covered spending of the people who choose that plan. Plans that attract more unhealthy types will have higher prices. At the same time, people respond to the price of different plans and will tend to switch toward lower-cost plans. Eventually, the market reaches a *competitive equilibrium* in which premiums are equal to a plan's costs (plus load) and no one wants to switch plans. Einav and Finkelstein's (2011) overview article in this journal shows how competitive equilibrium is reached with this sort of pricing.

This equilibrium can "unravel," so that only the plan with the least coverage is actually purchased; for example, this outcome will arise if everyone were fully informed in our example simulation. Suppose plans started out priced for the average population, as described earlier. In that case, all of the unhealthy people want the high actuarial value plan, but among the healthy types, those with below-average risk aversion prefer the low actuarial value plan. Those choice patterns, though, mean that the low actuarial value plan will only have healthy types enrolled, while the high actuarial value plan will have a higher share of unhealthy types than are in the total population. That puts upward pressure on the premiums for the high actuarial value plan and downward pressure for the other plan. As the difference in premiums rises, more people will prefer the low actuarial value plan. Ultimately, premium differences rise to the point that the market fully unravels, with everyone ending up choosing the low actuarial value plan. In this case, introducing the low actuarial value option in this environment has the same end result as only offering that option.

The first row of Table 3 shows the results of this process on consumer welfare relative to the benchmark situations where only the high actuarial value plan was available. Average welfare *falls* by $127 per person when choice is introduced. On average, unhealthy types face losses of $449 per person from this shift, while the healthy on average gain $25. Those who are more risk averse lose more from the shift to low actuarial value plans, so that even healthy types with above average risk aversion lose from introducing choice. The basic tradeoffs at play here are that 1) healthy types would prefer to have less coverage in the market because then they have to do less cross-subsidizing of the unhealthy types and vice versa for unhealthy types (they would prefer to have more coverage, expecting it will be cross-subsidized by the healthy types) and 2) reducing coverage is especially costly for those who are more risk averse.

This stark example involves an insurance "death spiral." In other examples, the market may not unravel completely and there may be less of a welfare loss (or even a possible welfare gain, depending on other parameters) from having everyone enrolled in the low actuarial value plan. However, the basic insight of this example is relevant for considering the value of offering choice. Cutler and Reber (1998), for example, show how a milder form of adverse selection operated when additional health insurance choices were offered to employees of Harvard University. Handel, Hendel, and Whinston (2015) also simulate equilibrium in the state-level health insurance exchanges established by the 2010 Affordable Care Act,

and predict unraveling to the lowest level of coverage unless there are regulations and risk adjustments for premiums.

The second two rows of Table 3 show what happens in equilibrium if instead of fully informed consumers, half of consumers are confused and choose either randomly or with the inattention heuristic. Consumer mistakes in selecting health plans can have two effects, as discussed in Handel (2013) and formalized in Handel, Kolstad, and Spinnewijn (2015). On the one side, mistakes lead people to sort less optimally between plans, which can lower welfare. However, random mistakes blunt the force of adverse selection and improve average welfare relative to the case where everyone is fully informed. Even a small fraction of random choosers helps to stabilize premium differences and prevents the market from unraveling. As the fraction of random choosers grows, premium differences will fall toward the population-average level, and average welfare improves.

In Table 3, we see that if half of people choose randomly, the welfare loss from introducing choice is $78 per person, about two-thirds the size of the loss when everyone chooses rationally. With half the people choosing randomly, the difference in premiums between the plans is around $1,760, which is around $600 higher than the premium difference would be with premiums set for the average population, but still low enough for some informed risk-averse unhealthy types to prefer the high actuarial value plan. However, no matter how many random choosers there are, we still find in this example that choice lowers welfare as compared to the benchmark case with only the high actuarial value plan available.

This result highlights a paradox of consumer confusion in health insurance markets where premiums can be affected by adverse selection (first highlighted clearly by Handel 2013). A decision aid that helps people select an optimal insurance plan could have a big return for them. For example, in our simulation, helping a single highly risk-averse unhealthy person avoid the mistake of randomly selecting the low actuarial value plan could provide that person with a few hundred dollars of consumer welfare value at the equilibrium prices that prevail when half of people are confused. However, if everyone started to choose optimally, premiums would change and the equilibrium would unravel, and only the low actuarial value plan could be purchased. The people who benefit the most from having some confused consumers in the market are the informed choosers who are unhealthy and highly risk averse. Those people benefit from having the high actuarial value available and in particular benefit from the healthy types who randomly select that plan with them and help keep its premiums down. Even among the confused consumers, the only ones who really benefit in equilibrium from having everyone become informed are the healthy types who are not very risk-averse. Only those types like the equilibrium where everyone is in the low actuarial value plan.

The third row of Table 3 shows the results if the uninformed consumers follow the nonrandom "inattention heuristic." Again, confused consumers help to ensure that the market does not completely unravel from adverse selection. However, strong selection effects arise. Among those with the inattention heuristic, all of the unhealthy types (along with the more risk-averse healthy types) select the high

actuarial value plan. That tends to push up the premium difference between the high and low actuarial value plans. In fact, with half the people choosing based on the heuristic, the equilibrium premium differences between the two plans would be nearly $2,500, which is more than the difference in deductibles between the plans and nearly as large as the difference in maximum out-of-pocket limits. Only the uninformed consumers would be selecting the high coverage plan in this situation. In effect, the premium differences look like the market has unraveled, but uniformed consumers do not realize it, and select the high actuarial value plan even though it is extremely expensive.

This pattern appears to be at play in some employer-sponsored insurance settings in which employees can choose between different levels of coverage. Handel (2013) and Bhargava, Loewenstein, and Sydnor (2017) analyzed data from firms where the difference in premiums between plans with higher and lower levels of coverage had risen to the point where the higher coverage plans were so expensive that they were dominated. Yet a substantial share of employees selected the higher coverage plans, due in part to inattention and inertia (Handel 2013), but also due to active choice processes like the naive sorting by health type we simulate with our inattention heuristic (Bhargava, Loewenstein, and Sydnor 2017).

With these nonrandom errors, the average welfare loss from introducing choice is again lower when there are uninformed consumers relative to the case with only fully informed consumers. With nonrandom errors, however, there is a clearer transfer of welfare from uninformed to informed consumers. In this case, many uninformed consumers are selecting the high actuarial value plan and paying premiums that reflect the fact that many of them are unhealthy types. The informed consumers, especially the healthy types, get a large benefit from selecting the low actuarial value plan, which has an advantageous mix of more healthy types. In effect, the option to choose the low coverage plan allows these informed healthy types to avoid pooling with many of the unhealthy types who naively select the higher coverage plan. We see, for example, in our simulation with half the people choosing with the heuristic, that the average welfare for healthy types is $151 higher per person with choice, while the loss for unhealthy types is over $500 and worse than the case where everyone is fully informed and the market unravels.

Thus, consumer confusion has distributional consequences. When mistakes help to stabilize the market, confused consumers are in effect subsidizing savvier consumers. The individual welfare losses for those making the mistakes may be large, especially when these errors are nonrandom. Moreover, economically vulnerable populations, including those with less education, lower incomes, the elderly, and those with health problems, are all more likely to have problems selecting optimal health insurance plans (for example, Loewenstein et al. 2013; Bhargava, Loewenstein, and Sydnor 2017). The type of consumer welfare calculations we have shown in our example, and which are the norm in much of the economic analysis of health insurance markets, treat a dollar of value the same for all people. However, a dollar has greater utility for someone with lower income, which means the average

social welfare effects of offering choice might be worse even on average when a substantial share of consumers are confused about their choices.

The Effects of Choice in Markets with Regulations and Risk Adjustments

Many insurance markets have regulations and policies in place that seek to prevent the problems with selection by health type and the consequent market unraveling. Many insurance markets with a consumer choice component, including the Affordable Care Act health insurance exchanges, Medicare Part D, and Medicare Advantage, use a system of "risk adjustments" that transfers money from plans that enroll healthier individuals to plans that enroll sicker individuals. While the details vary by market, risk adjustments tend to target equalizing the average cost of enrollees between plans (for a review, see Van den Ven and Ellis 2000; Geruso and Layton, this issue). For our simulation, we model risk adjustment as perfectly equalizing costs. However, actual risk adjustment tends to be imperfect, which means that the insights from markets with adverse selection remain important even when the market has some risk adjustment.

The bottom panel of Table 3 shows what happens in our simulations when effective risk adjustments are in place. In particular, we assume that the risk-adjustment process stabilizes the difference in premiums between plans at the level appropriate for the population on average (roughly $1,110). The level of the premiums adjusts in equilibrium to be high enough to ensure that the plans collect enough money overall to cover the expected covered medical spending plus loads.

When all the people choosing plans in the market are fully informed, risk adjustment helps make offering choice welfare-enhancing on average, relative to the benchmark case with only the high actuarial value plan available. The risk adjustments keep the premiums between plans stable even though the pool of enrollees for the low actuarial value plan is much healthier (in our example, only healthy individuals choose that plan). In our simulation, we see an average increase in consumer welfare from having choice of $9 per person, with healthy people gaining around $42 per person on average and unhealthy people losing $60 per person.

It may seem surprising that the unhealthy people would lose anything when the low actuarial value plan is introduced, since they all select the benchmark high actuarial value plan and premium differences between plans do not adjust. However, risk adjustments that stabilize premium differences based on the appropriate difference for the average person are not fully efficient. At these premium differences, the healthy types get a larger discount from selecting the high actuarial value plan than is warranted by their reduction in covered spending from reducing coverage. Equivalently, the unhealthy types pay less for higher coverage than the increase in covered spending they receive. This remaining inefficiency with risk adjustment is nearly impossible to avoid and results in the level of premiums (for both plans) rising above the level they would be if everyone were enrolled in a single plan. Ultimately, risk adjustment technology can generate positive welfare gain from choice for informed consumers, but since risk adjustment is not fully efficient, those gains are not guaranteed, may be modest, and may favor the healthy types.

In contrast to the situation with completely unregulated premiums, random mistakes tend to lower consumer welfare when there is risk adjustment. Once risk adjustment reduces the problem of adverse selection, the primary effect of random choice is that people sort less well into the plan that is best for them. In our simulation, with 50 percent of the population choosing randomly, there is now an average loss of $27 per person from introducing choice. Because the gains to choice are so small even in the best case with fully informed consumers, it only takes a small fraction of people making mistakes (about 15 percent in our simulation) for the value of introducing choice to be negative.

As the fraction of those who choose their insurance policy randomly rises, welfare continues to fall, both on average and for each health type. However, random choice will tend to lower welfare most quickly for risk-averse unhealthy types. This group receives an especially strong benefit from additional insurance coverage, and conversely, the welfare costs for this group of wrongly choosing the low actuarial value plan is especially harmful. Thus, expanding choice of coverage levels creates a distributional tradeoff between relatively small gains for many of those who use their added choice wisely (again, given that they are paying a risk-adjusted premium for the health group to which they belong) but the risk of large losses for those in the risk-averse unhealthy group who make random errors.

The final row of Table 3 highlights that unlike random choosing, there are no adverse consequences to uninformed consumers using the inattention heuristic when premiums are controlled by risk adjustment. In this case, the heuristic happens to be perfectly matched to the market environment and people are choosing in the same way they would if they were fully informed. Of course, other forms of nonrandom choice errors could lead to welfare losses even when there are risk adjustments. However, in general, effective risk adjustments will help dampen the negative consequences of nonrandom choice errors in which selection is partly related to one's health status.

Nudges and Other Interventions to Improve Consumer Choice

Behaviorally informed policies can seek to address the consumer confusion and biases that are widespread in health insurance choice. Initial research suggests that such policies can sometimes be effective, but also suggests that they face substantial challenges.

First, *standardizing policy health insurance plan options* within a level of actuarial value can make it easier to compare plans. For example, Ericson and Starc (2016) examined an earlier natural experiment in which health plans on the Massachusetts health insurance exchange were standardized within each tier. Standardization led consumers to choose more generous health insurance plans and to substantial shifts in brands' market shares. However, seemingly small details about the design of choice platforms also affect consumer choice, such as the labels attached to tiers (like calling one level "bronze") and the order in which plans are sorted (Ubel, Comerford, and

Johnson 2015). The HealthCare.gov website which provides information for the state-based health insurance exchanges recently introduced standardized options for each coverage tier called "simple choice" plans, which all have the same deductible and co-pay levels, although nonstandardized options do still remain available. Future studies will likely investigate how this change affects choices.

Second, *providing personalized information* to health insurance shoppers may help their decision-making. There is substantial variation in what is included in these consumer decision support tools (for discussion, see Wong, Polsky, Jones, Wiener, Town, and Baker 2016). Many markets provide out-of-pocket cost calculators that help people estimate how plan choices will affect their expected spending, based on demographics, health status, and/or past claims history. For instance, Medicare Plan Compare sorts Medicare Part D prescription drug insurance plans based on expected costs given the drugs an individual is currently taking. Similarly, for health insurance bought on HealthCare.gov, the site presents expected annual spending amounts for a few representative spending scenarios.

Research on whether out-of-pocket calculators meaningfully affect plan choices is limited, and the results are mixed. Earlier work found that providing personalized information about out-of-pocket costs did induce people to switch plans (Kling, Mullainathan, Shafir, Vermeulen, and Wrobel 2012). However, Abaluck and Gruber (2016) find that providing an out-of-pocket cost predictor at a large employer had little effect on plan choices. Similarly, Ericson, Kingsdale, Layton, and Sacarny (2017) find that a randomized experiment providing people with personalized information about the potential premium savings they could have in the state-level health insurance market induced more people to shop more actively—but did not lead people to switch plans.

Third, *smart defaults* offer a more aggressive approach to nudging consumers, in which the decision-assistance software actually chooses a default plan based on the consumer's information, but the consumer can override that default if desired. Smart defaults have been explored in Medicare Part D (Hoadley, Thompson, Hargrave, and Merrell 2007), and can be used either at the point of initial enrollment or to switch inattentive consumers (Ericson 2014b). Johnson et al. (2013) suggest that it may be necessary to couple calculators with smartly chosen defaults or recommendations to meaningfully improve consumer choices.

Efforts to nudge consumer choice will create tradeoffs of their own. The example of out-of-pocket cost calculators in health insurance can be used to illustrate the concerns. They are typically implemented as an "expected value nudge" that recommends a plan that minimizes a person's total expected health costs (premiums plus out-of-pocket costs for cost sharing). Importantly, most out-of-pocket calculators do not include a measure of risk aversion (although Picwell.com offers a counterexample). As a result, healthy people will typically be nudged to choose low actuarial value plans, even though a healthy person with high risk aversion might prefer a plan with high actuarial value. Such out-of-pocket cost calculators encourage sorting by health status. As noted above, nudges that reduce random errors can worsen adverse selection and lead to market unraveling. Nudges to improve consumer choice are

more likely to lead to an overall social gain if accompanied by well-designed risk-adjustment payments to limit the effects of adverse selection.

Finally, policymakers could turn toward explicitly *limiting the amount of choice available* in the market. This is an ongoing debate in many areas of health care policy. For example, the Patient Protection and Affordable Care Act of 2010 regulates the range of options for coverage levels available in the private health insurance exchanges. Plans are required to cover a common set of essential health benefits and there are four allowed levels of actuarial value. These regulations are contentious and there have been proposals to relax them. As our essay has sought to highlight, there are real tradeoffs involved with offering additional choices of health insurance. On one side, a greater range of choice creates opportunity to raise consumer welfare as buyers sort themselves into the options they prefer. On the other side, a greater variety of choice also raises the possibility of adverse selection dynamics as well as losses due to consumer mistakes.

Other Dimensions of Choice: Value and Pitfalls

Choice over coverage levels, in which insurance plans can be ranked by their financial generosity and actuarial value, can be modeled very tractably. This dimension of choice has been most extensively studied. However, health insurance and health care more broadly pose many other choices. Features of insurance plans that are not reflected in actuarial value create many of the same tradeoffs and potentials for interactions between consumer confusion and adverse selection that we highlighted here for choice over coverage levels.

For example, consider two insurance plans. One is a "managed care" plan with strong limitations on the network of doctors and hospitals a patient can see, and a low degree of cost-sharing. Another is a "consumer-directed" plan with a broad network of doctors and hospitals, few limits on access to specialists, but with a high degree of cost-sharing. Many hybrids of these approaches exist in US health insurance markets. On one hand, people can benefit from choosing the breadth of provider network they prefer. On the other hand, such choices create potential for problems with consumer confusion and adverse selection. However, economists know less about and have less of a well-established framework for modeling how people value access to different networks of providers (for discussion, see Ericson and Starc 2015).

Even within types of insurance arrangements—managed care versus consumer directed—consumer confusion is likely to play an important and varied role. Researching the networks included in any insurance plan can be exceedingly difficult. Anticipating how network access might matter in the future is even more difficult. Recent research has also highlighted the potential for "surprise billing" for medical services where patients get charged for out-of-network doctors at hospitals that are in network for their insurance (Cooper, Scott-Morton, and Shekita 2017).

Most people are not well positioned to assess the value of different medical procedures, which can create a conflict of interest between providers and patients for some services that are of questionable value. It is also difficult for patients to assess both the quality of services and the prices they will be charged from different providers. Consumer-driven health plans are increasingly providing people with portals that facilitate "shopping" between plans by providing information on the prices for services and quality ratings for providers. However, the evidence so far suggests that many of these tools are not very effective (for example, Brot-Goldberg et al. 2017).

Finally, a recent literature, not specific to health care, looks at how firms that have the ability to design contracts may seek to exploit consumers by obfuscating the true price or generosity of the contract (Akerlof and Shiller 2015). In some cases, this obfuscation leads to redistribution from less-sophisticated to more-sophisticated households. Gabaix and Laibson (2006) refer to these cases as situations where products have "shrouded attributes" and highlight the example of credit cards, in which consumers who end up carrying costly debt and paying fees partly subsidize savvier and more financially secure consumers who get the convenience benefits and rewards from credit cards. In other cases, firms can make larger profits on socially wasteful products (Heidhues, Kőszegi, and Murooka 2017; for a review of the theoretical literature, see Grubb 2015).

While the literature on firms exploiting consumer confusion in health insurance has been limited to date, some evidence suggests that these insights may apply in the health care setting. For example, Ericson (2014a) showed that insurers offering Medicare Part D plans raised their prices over time on existing plans to take advantage of consumer inertia while simultaneously introducing new plans into the market at lower prices to attract attentive new customers. Similarly, our discussion in this essay highlights the possibility that when employers offer their employees multiple coverage level options, a situation can arise in which healthier and more sophisticated employees can take advantage of cheaper plans that allow them to avoid pooling together with unhealthier employees who are not sophisticated about their plan choice. Moreover, more sophisticated individuals might be able to follow plan rules more closely and get more value out of a given plan (for instance, by making sure to get pre-authorization or by effectively appealing a denied claim).

Discussion

Offering choice over the type of insurance policy or allowing people to select into plans with different networks of doctors may be beneficial. But as this essay has suggested, the additional choice is not an unmixed blessing.

The many and expanding dimensions of consumer choice in health insurance and health care in the US present challenges not only to individuals but also to economic modeling. For example, the Congressional Budget Office (CBO) is tasked with simulating how people will select individual health insurance plans to help inform legislators. For its simulations, the CBO does not use a model based

on the classical expected utility approach, in part because this standard economic framework does not capture many of the forces driving actual choice behavior (see for example https://www.cbo.gov/publication/45427, slide 9). Instead, the CBO bases its projections on elasticities measured from data in actual markets, which at least partially capture the effects of inertia, inattention, and consumer biases. Of course, a limitation of this elasticity approach is that elasticities observed in the past may change in different settings or when changes in regulations affect the extent of consumer confusion.

Policymakers and economists also need to wrestle with the fundamental challenge of how we judge what policies make people better off. Economists typically try to answer these questions by observing the choices people make, and then drawing inferences about peoples' preferences. However, consumer confusion means that the choices people make about health plans may not be directly informative about their underlying preferences. Studies that directly measure consumer confusion and use that information to map choice to welfare (for example, Handel and Kolstad 2015) are an important step forward. More work is needed in that direction.

Even the basics of how economists should evaluate welfare in environments where decision-makers have biases and confusion is a contentious issue (Beshears, Choi, Laibson, and Madrian 2008; Bernheim and Rangel 2009). For example, individuals with loss aversion may try to avoid being exposed to out-of-pocket costs. Such people may desire a high actuarial value plan even if the premiums are very expensive, or they may steer away from the combination of a high-deductible insurance plan with a health savings account—even if these options will save them money. Should this form of loss aversion be considered a "mistake" and nudges implemented to guide those with high loss aversion in another direction even if it makes the person miserable? Economists have not yet been able to offer clear guidance to policymakers on this issue.

Both economists and policymakers should pay more attention to how the complexity on many dimensions of modern health insurance in the United States creates confusion for consumers and can erode the benefits of competition. Given the complexity of healthcare and health insurance markets, health economists and health policy experts must in part also be behavioral economists with an eye toward understanding how people process information and decide, and how those forces shape health care markets.

References

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

Abaluck, Jason, and Jonathan Gruber. 2016. "Improving the Quality of Choices in Health Insurance Markets." NBER Working Paper 22917.

Agency for Healthcare Research and Quality (AHRQ). 2012 and 2013. Medical Expenditure Panel Survey. https://meps.ahrq.gov/mepsweb/.

Aizcorbe, Ana, Eli Liebman, Sarah Pack, David M. Cutler, Michael E. Chernew, and Allison B. Rosen. 2012. "Measuring Health Care Costs of Individuals with Employer-Sponsored Health Insurance in the U.S.: A Comparison of Survey and Claims Data." *Statistical Journal of the IAOS* 28(1–2): 43–51.

Akerlof, George A., and Robert J. Shiller. 2015. *Phishing for Phools: The Economics of Manipulation and Deception*. Princeton University Press.

Baicker, Katherine, Sendhil Mullainathan, and Joshua Schwartzstein. 2015. "Behavioral Hazard in Health Insurance." *Quarterly Journal of Economics* 130(4): 1623–67.

Bernheim, B. Douglas, and Antonio Rangel. 2009. "Beyond Revealed Preference: Choice-Theoretic Foundations for Behavioral Welfare Economics." *Quarterly Journal of Economics* 124(1): 51–104.

Beshears, John, James J. Choi, David Laibson, and Brigitte C. Madrian. 2008. "How Are Preferences Revealed?" *Journal of Public Economics* 92(8–9): 1787–94.

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

Brot-Goldberg, Zarek C., Amitabh Chandra, Benjamin R. Handel, and Jonathan T. Kolstad. 2017. "What Does a Deductible Do? The Impact of Cost-Sharing on Health Care Prices, Quantities, and Spending Dynamics." *Quarterly Journal of Economics* 132(3): 1261–318.

Claxton, Gary, Matthew Rae, Michelle Long, Anthony Damico, Bradley Sawyer, Gregory Foster, Heidi Whitmore, and Lindsey Schapiro. 2016. *Employer Health Benefits: 2016 Annual Survey*. Menlo Park, CA*:* Kaiser Family Foundation.

Cooper, Zack, Stuart V. Craig, Martin Gaynor, and John Van Reenen. 2015. "The Price Ain't Right? Hospital Prices and Health Spending on the Privately Insured." NBER Working Paper 21815.

Cooper, Zack, Fiona Scott Morton, and Nathan Shekita. 2017. "Surprise! Out-of-Network Billing for Emergency Care in the United States." NBER Working Paper 23623.

Cutler, David M., and Sarah J. Reber. 1998. "Paying for Health Insurance: The Trade-off between Competition and Adverse Selection." *Quarterly Journal of Economics* 113(2): 433–66.

Department of Health and Human Services. 2016. "Health Insurance Marketplace Premiums after Shopping, Switching, and Premium Tax Credits, 2015–2016." ASPE Issue Brief, April 1, Office of the Assistant Secretary for Planning and Evaluation. https://aspe.hhs.gov/system/files/pdf/198636/MarketplaceRate.pdf.

Einav, Liran, Amy Finkelstein, and Jonathan Levin. 2010. "Beyond Testing: Empirical Models of Insurance Markets." *Annual Review of Economics* 2(1): 311–36.

Einav, Liran, Amy Finkelstein, Stephen P. Ryan, Paul Schrimpf, and Mark R. Cullen. 2013. "Selection on Moral Hazard in Health Insurance." *American Economic Review* 103(1): 178–219.

Einav, Liran, and Amy Finkelstein. 2011. "Selection in Insurance Markets: Theory and Empirics in Pictures." *Journal of Economic Perspectives* 25(1): 115–38.

Ericson, Keith M. Marzilli. 2014a. "Consumer Inertia and Firm Pricing in the Medicare Part D Prescription Drug Insurance Exchange." *American Economic Journal: Economic Policy* 6(1): 38–64.

Ericson, Keith M. Marzilli. 2014b. "When Consumers Do Not Make an Active Decision: Dynamic Default Rules and their Equilibrium Effects." NBER Working Paper 20127.

Ericson, Keith M. Marzilli, Jon Kingsdale, Tim Layton, and Adam Sacarny. 2017. "Nudging Leads Consumers in Colorado to Shop But Not Switch ACA Marketplace Plans." *Health Affairs* 36(2): 311–19.

Ericson, Keith Marzilli, Philipp Kircher, Johannes Spinnewijn, and Amanda Starc. 2015. "Inferring Risk Perceptions and Preferences using Choice from Insurance Menus: Theory and Evidence." NBER Working Paper 21797.

Ericson, Keith Marzilli, and Amanda Starc. 2015. "Measuring Consumer Valuation of Limited Provider Networks." *American Economic Review* 105(5): 115–19.

Ericson, Keith M. Marzilli, and Amanda Starc. 2016. "How Product Standardization Affects Choice: Evidence from the Massachusetts Health Insurance Exchange." *Journal of Health Economics* 50: 71–85.

Ericson, Keith M., and Justin Sydnor. 2017.

"Liquidity Constraints and the Demand for Insurance." Unpublished paper.

Gabaix, Xavier, and David Laibson. 2006. "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets." *Quarterly Journal of Economics* 121(2): 505–40.

Grubb, Michael D. 2015. "Overconfident Consumers in the Marketplace." *Journal of Economic Perspectives* 29(4): 9–36.

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

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

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

Handel, Benjamin R., Jonathan T. Kolstad, and Johannes Spinnewijn. 2015. "Information Frictions and Adverse Selection: Policy Interventions in Health Insurance Markets." NBER Working Paper 21759.

Heidhues, Paul, Botond Köszegi, and Takeshi Murooka. 2017. "Inferior Products and Profitable Deception." *Review of Economic Studies* 84(1): 323–56.

Hoadley, Jack, Laura Summer, Jennifer Thompson, Elizabeth Hargrave, and Katie Merrell. 2007. *The Role of Beneficiary-Centered Assignment for Medicare Part D.* Washington, DC: Medicare Payment Advisory Commission.

Iyengar, Sheena S., and Emir Kamenica. 2010. "Choice Proliferation, Simplicity Seeking, and Asset Allocation." *Journal of Public Economics* 94(7–8): 530–39.

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

Johnson, Eric J., John Hersehy, Jaqueline Meszaros, and Howard Kunreuther. 1993. "Framing, Probability Distortions, and Insurance Decisions." *Journal of Risk and Uncertainty* 7(1): 35–51.

Kamenica, Emir. 2008. "Contextual Inference in Markets: On the Informational Content of Product Lines." *American Economic Review* 98(5): 2127–49.

Kling, Jeffrey R., Sendhil Mullainathan, Eldar Shafir, Lee C. Vermeulen, and Marian V. Wrobel. 2012. "Comparison Friction: Experimental Evidence from Medicare Drug Plans." *Quarterly Journal of Economics* 127(1): 199–235.

Loewenstein, George, Joelle Y. Friedman, Barbara McGill, Sarah Ahmad, Suzanne Linck, Stacey Sinkula, John Beshears, James J. Choi, Jonathan Kolstad, David Laibson, Bigitte C. Madrian, John A. List, and Kevin G. Volpp. 2013. "Consumers' Misunderstanding of Health Insurance." *Journal of Health Economics* 32(5): 850–62.

Nyman, John A. 1999. "The Value of Health Insurance: The Access Motive." *Journal of Health Economics* 18(2): 141–52.

Rosenthal, Jaime A., Xin Lu, and Peter Cram. 2013. "Availability of Consumer Prices from US Hospitals for a Common Surgical Procedure." *JAMA Internal Medicine* 173(6): 427–32.

Sinaiko, Anna D., and Richard A. Hirth. 2011. "Consumers, Health Insurance and Dominated Choices." *Journal of Health Economics* 30(2): 450–57.

Ubel, Peter A., David A. Comerford, and Eric Johnson. 2015. "Healthcare.gov 3.0—Behavioral Economics and Insurance Exchanges." *New England Journal of Medicine* 372(8): 695–98.

van de Ven, Wynand P. M. M., and Randall P. Ellis. 2000. "Risk Adjustment in Competitive Health Plan Markets." In *Handbook of Health Economics,* vol. 1A, edited by Anthony J. Culyer and Joseph P. Newhouse, 755–845. Amsterdam: Elsevier.

Wong, Charlene A., Daniel E. Polsky, Arthur T. Jones, Janet Weiner, Robert J. Town, and Tom Baker. 2016. "For Third Enrollment Period, Marketplaces Expand Decision Support Tools to Assist Consumers." *Health Affairs* 35(4): 680–87.