Frictions or Mental Gaps: What's Behind the Information We (Don't) Use and When Do We Care?

Benjamin Handel and Joshua Schwartzstein

I n a number of situations, there is strong evidence that people do not translate readily available information into the knowledge that would help them make better decisions. For example, people may choose a health insurance plan that costs $500 per year more in premiums in order to obtain a deductible that is $250 lower—despite having access to open enrollment booklets containing relevant information (Handel 2013; Bhargava, Loewenstein, and Sydnor 2017). People buy branded drugs over equivalent but less-expensive generics (Bronnenberg, Dubé, Gentzkow, and Shapiro 2015) even though information printed on the package reveals their equivalence. Investors pay a range of fees for investing in S&P 500 index funds—and index funds with higher fees have meaningful market shares (Hortaçsu and Syverson 2004). Consumers appear to demand the wrong cell phone plans given their previous usage patterns (Grubb and Osborne 2015).

Why don't people use available information? The many possibilities discussed in the research literature broadly fall into two camps, which we refer to as *frictions* and *mental gaps.* The frictions camp focuses on costs of acquiring and processing information. A consumer shopping in a health insurance exchange incurs a cost to explore more of the options in the choice set and to assess them. This camp, and the closely related framework of "rational inattention," maintains the neoclassical assumption that people form accurate beliefs using the information that is worth

■ *Benjamin Handel is Associate Professor of Economics, University of California, Berkeley, California. Joshua Schwartzstein is Assistant Professor of Business Administration, Harvard Business School, Boston, Massachusetts. Their email addresses are handel@berkeley.edu and jschwartzstein@hbs.edu.*

processing, but it incorporates realistic assumptions on how paying attention to or processing information is costly (Stigler 1961; McCall 1970; Caplin and Dean 2015; Sims 2003; Woodford 2012; Gabaix 2014).

The second camp deals with "mental gaps" or psychological distortions in information-gathering, attention, and processing. A consumer in the insurance exchange may neglect important information in selecting plans even if this information is readily available, perhaps from using an incorrect model, (for example, Schwartzstein 2014) or overweighting salient plan features (for example, Bordalo, Gennaioli, and Shleifer 2012, 2013; Kőszegi and Szeidl 2012). This camp emphasizes how, for a variety of reasons, there is a gap between what people think and what they should rationally think given costs. The categories of frictions and mental gaps are not mutally exclusive or exhaustive, but are intended as a broad classification of approaches that researchers take to studying poorly informed choice.

Most empirical research on frictions or mental gaps assumes that one mechanism dominates, without explicit consideration of possible alternatives, or doesn't try to specify the precise underlying mechanism. A primary reason is that, even with extensive data, it can be very difficult in a number of contexts to identify the source of apparent mistakes. For example, a researcher who observes that consumers of health insurance fail to switch to more valuable options over time may have a hard time distinguishing possible explanations including 1) high time costs of search and switching or 2) incorrect views of how likely product values are to change over time. When researchers assume that one specific mechanism underlies poorly informed choices, but cannot credibly distinguish between that mechanism and others, spurious conclusions often follow.

Beyond specifying the extent of and reason for poorly informed choices, a further goal of the literature is to investigate the consumer welfare (henceforth "welfare") impacts of policies in environments where consumers make such choices. When is it important for policy assessments to distinguish between the underlying mechanisms? We define two classes of policies. An *allocation policy* directly allocates (or strongly steers) consumers to specific actions. To assess the welfare impact of an allocation policy, it is sufficient to identify the combined effect of frictions and mental gaps empirically. A *mechanism policy* instead targets specific mechanisms, where policy predictions depend critically on understanding relative magnitudes of different frictions and mental gaps.[1] This classification may require some judgment to apply, but is intended to highlight factors to keep in mind. Our discussion largely focuses on *counterfactual policies*, by which we mean policies that are hard to evaluate empirically before implementing them, but we will also touch on the case of policies that can more easily be studied in action.

---

[1]The contrast between mechanism and allocation policies is not the same as the distinction between nudges and traditional policy instruments as introduced by Thaler and Sunstein (2009) and analyzed in detail by Farhi and Gabaix (2017). Many nudges, such as reminders, could be viewed as mechanism policies that target "behavioral" mechanisms, like forgetfulness, but we will view others, such as defaults, as allocation policies.

We begin by describing evidence across contexts in which consumers are not using important information. We then outline key frictions and mental gaps that could matter in these contexts. While many empirical papers describe their findings as related to one specific friction or mental gap, they typically provide little evidence to distinguish between mechanisms. After spelling out this key issue, we turn to three related questions. First, what can we say about the magnitude of frictions and mental gaps when we are uncertain about the mechanism? Second, how can we empirically distinguish the mechanisms? Third, for which policy questions is it sufficient to understand magnitudes and for which is it important to distinguish mechanisms?

Some Examples of Information that People Do Not Use

There is a substantial body of research documenting situations and consequences of people not using readily available information. Table 1 provides examples from the domain of health, and Table 2 provides a broader set of examples. Most of these papers do not attempt to distinguish, explicitly or implicitly, between reasons for not using information.

Consumer Ignorance and Misinformation in Health Markets

Consider a scenario: You have a headache, go to a pharmacy, and choose Advil over store-branded medication containing ibuprofen—which is the same active ingredient contained in Advil. This type of choice is common. Bronnenberg et al. (2015) find that the average consumer chooses national headache-remedy brands over chemically equivalent store-brand alternatives 26 percent of the time. What's going on? At a broad level, consumer misinformation appears to be a factor. Bronnenberg et al. find that pharmacists choose national headache-remedy brands over store-brand alternatives only 9 percent of the time, and nonexpert consumers are presumably less knowledgeable about active ingredients and relative safety. A subset of Nielson panelists were asked to name the active ingredient in national headache remedies. The average respondent answered 59 percent of these questions correctly, compared to over 85 percent for nurses, pharmacists, and doctors. Having this knowledge is highly positively associated with purchasing the store brand, as is reporting a belief that store brands are "just as safe" as national brands. This evidence strongly suggests that a lack of knowledge contributes to nonexpert consumers' demand for national brands. But the evidence has less to say about why consumers are misinformed.

Other papers documenting mistakes in the health treatment decisions of consumers likewise do not typically attempt to identify the causes or domains of misinformation. Pauly and Blavin (2008) and Baicker, Mullainathan, and Schwartzstein (2015) summarize evidence that people have a systematic propensity to under- or overuse certain treatments at the margin. For example, Choudhry et al. (2011) document that many recent heart attack victims do not adhere to drug

Table 1

Examples of Information People Don't Use in Health Markets

Paper	Findings	Potential explanations for not using information
Health insurance		
Handel and Kolstad (2015b)	"Uninformed" consumers leave substantial dollars on table when "over-choosing" generous insurance coverage, relative to "informed" consumers. Consumers who think (incorrectly) that more generous coverage gives them access to generous providers are willing to pay much more (~$2,300) for that coverage.	*Frictions:* Search costs lead to limited information; information processing costs lead to poor evaluations of plan characteristics. *Mental gaps:* Mistaken beliefs about important ways plans differ; neglect of key plan characteristics.
Bhargava, Loewenstein, and Sydnor (2017)	In active choices, consumers frequently choose dominated plans from menu of 48 insurance options at large employer, losing $300–$400 on average. Experiments show better choices in simplified choice environments and in environments with plan characteristics information.	*Frictions:* Search costs to find or explore plan options. *Mental gaps:* People have limited insurance competence, not understanding the mapping between plan characteristics (for example, deductibles) and payoff-relevant outcomes.
Handel (2013)	Consumer inertia leads to thousands of $ in financial losses (~$2,000) in insurance plan choice. Consumers choose dominated health plans with high frequency when possible to do so.	*Frictions:* Switching costs (from search, information processing, etc.); rational inattention to plan choice. *Mental gaps:* Consumers don't recognize potential benefits from switching, having wrong priors about plan changes over time (for example, not realizing that plans may become financially dominated); lack of competency in evaluating premiums relative to plan characteristics; neglect of certain key plan features.
Abaluck and Gruber (2011, 2016); Ho, Hogan, and Scott Morton (2017); Ketcham, Lucarelli, and Powers (2015)	Consumers leave money on the table in initial Medicare Part D choices, on average ~$300 per consumer. Consumers exhibit substantial inertia, leading to additional monetary losses.	
Health treatment		
Bronnenberg, Dubé, Gentzgow, and Shapiro (2015)	Experts (pharmacists and medical professionals) are less likely to pay extra for branded headache-remedy drugs relative to generic (bio-equivalent) alternatives.	*Frictions:* Information processing or search costs lead to unawareness of bio-equivalent alternatives. *Mental gaps:* People don't know which ingredients to focus on or realize that generic equivalents might be available; wrong priors about generic equivalence.
Pauly and Blavin (2008); Baicker, Mullainathan, and Schwartzstein (2015); Choudhry et al. (2011)	Health insurees seemingly underuse valuable treatments for chronic diseases. In such cases, health insurees' adherence is quite sensitive to copay changes.	*Frictions:* Information gathering and processing costs are too high for insurees to recognize the value of these treatments. *Mental gaps:* People do not know how to assess the value of treatments.

regimens aimed at preventing future heart attacks at regular copay levels, but show in a large-scale field experiment that eliminating copays for these drugs substantially boosts adherence and improves clinical outcomes. Baicker, Mullainathan, and Schwartzstein (2015) argue that it is difficult to rationalize such examples in a framework where consumers accurately trade off health benefits against the copay. Others argue that consumer misinformation is likely a key reason why consumers act as if they misweight treatment benefits (for example, Pauly and Blavin 2008). These findings have important policy implications: in many cases, when consumers make poor health choices it both increases long-run health costs and reduces consumer health, and everyone loses. How should policymakers use the evidence in these studies, or work to produce evidence in future studies, when considering different interventions to improve health care decisions?

Another set of examples comes from consumers' choices of health insurance plans. Handel (2013) analyzes this choice assuming that consumers have a bias toward inertia, modeled as costs from switching plans, but have rational expectations about their own health risk and full information about the plan options available. The paper estimates a switching cost of approximately $2,000 in the population. Many consumers leaving that much money on the table earn low incomes and have families, heightening the consequences. Handel acknowledges that the estimated switching cost likely reflects a range of underlying mechanisms, including true switching costs, search costs, and miscalibrated beliefs. Ho, Hogan, and Scott Morton (2017) model inertia using rational inattention as opposed to switching costs. They also find substantial inertia, modeled as a high cost of paying attention to the choice environment, with substantial negative consequences across the board for seniors. These two papers with similar data and identification assume distinct mechanisms underlying inertia, without teaching us which mechanism carries greater weight in the decision process.

Consumer Ignorance and Misinformation in Other Domains

Table 2 provides a few examples outside the health care arena. In one example, Hanna, Mullainathan, and Schwartzstein (2014) develop and test a model of technological learning. Focusing here on the empirical exercise, they study the knowledge, practices, and impact of a knowledge intervention on a community of Indonesian farmers who had a lot of experience: they farmed seaweed on average for 18 years with many cycles in each year. Seaweed is farmed by attaching strands of seaweed (or "pods") on lines submerged into the ocean, where many factors could affect yield. Local nongovernment organizations suggested that these farmers' practices tend to be far from the productivity frontier, a fact supported by Hanna et al.'s experimental estimates. Further, this appears to stem from farmers not understanding key relationships between input choices and yield. Farmers did precise things and had clear opinions on most dimensions: the length of their line, the distance between pods, the distance between lines, and the cycle length. But they did not have a clear opinion on their pod size (a truly important input dimension, according to Hanna et al.'s estimates): around 85 percent did not know the size they use and would not

Table 2
Examples of Information People Don't Use in Non-Health Markets

Paper	Facts	Potential explanations for not using information
Agriculture		
Hanna, Mullainathan, and Schwartzstein (2014)	Seaweed farmers persistently neglect an important input dimension (pod size).	*Frictions:* Learning potentially payoff-relevant relationships is costly.
	They respond to an intervention that filled in knowledge gaps.	*Mental gaps:* Farmers started with wrong beliefs about which inputs mattered.
Financial investments		
Hortaçsu and Syverson (2004)	There is significant price dispersion in S&P 500 index funds (financially undifferentiated products).	
	Higher-fee funds have meaningful market shares.	
Hastings, Hortaçsu, and Syverson (2017)	Consumers lose significant sums of money choosing among privatized, essentially homogeneous, mutual funds in Mexico's privatized social security.	*Frictions:* Large search costs to find prices or products; switching costs across firms.
	Advertising investment is associated with these poor choices.	*Mental gaps:* People don't realize that index funds differ only in fees; advertising or marketing of brands may reinforce or cause wrong beliefs; limited financial literacy; people don't think to check on their 401(k) contribution rate.
Choi, Laibson, and Madrian (2010)	Consumers leave significant sums of money on the table by choosing high-fee index funds. Experiment shows this is not because of nonportfolio features and also is not primarily the result of search costs. Consumers with lower financial literacy are more likely to make mistakes, and often even have a sense they are making mistakes.	
Madrian and Shea (2001)	Consumers exhibit substantial inertia in their choice of 401(k) investments and are highly sensitive to default investment settings.	
Cellular phones		
Grubb and Osborne (2015)	Consumers demand cell phone plans as if they underestimate the variance of future calling minutes.	*Frictions:* Keeping track of usage is costly; switching costs in plan choice.
	Consumers appear inattentive to past usage within a plan month, making usage alerts valuable.	*Mental gaps:* People underestimate the likelihood of using enough minutes to incur fees.
Energy		
Allcott and Taubinsky (2015)	Providing information on energy cost savings boosts demand for energy-efficient lightbulbs.	*Frictions:* Search costs for finding relevant product information.
		Mental gaps: People may be biased towards believing the upfront price is most important; people may focus too little on future costs.
Ito, Ida, and Tanaka (2016); Jessoe and Rapson (2014)	Consumers choose electricity tariffs that are bad for them as well as for society. Experiment shows that information provision helps reverse some of the poor decisions, but not a significant portion of them.	*Frictions:* Search costs of finding relevant electricity tariff information; switching costs of switching electricity plans; adjustment costs of changing electricity consumption in response to price fluctuations.
	High-frequency information provision makes consumers significantly better in responding to time-varying electricity tariffs and builds habits whereby consumers adjust behavior in the medium to long run even in the absence of information.	*Mental gaps:* People may have low literacy in evaluating complex multipart electricity tariffs, or real-time electricity pricing; people may believe that information is hard to obtain when it is in fact easy to obtain.

give an opinion on the optimal size. This lack of opinion appeared to translate into a lack of measurement: Each farmer had substantial variation in pod size within his own plot (which in theory he could learn from). The failure to optimize pod size appeared to meaningfully reduce farmers' output and income.

In household finance, Hortaçsu and Syverson (2004) show that consumers frequently purchase higher-fee S&P 500 index funds as if they do not know of the existence of lower-fee funds that will provide essentially equivalent returns. The authors pose a model with consumer search frictions and assume that these search costs are responsible for the low-value options consumers end up choosing. Madrian and Shea (2001) study 401(k) decisions of many employees at a large firm and show that a shift in the default policy for how contributions are matched and invested has a substantial impact on consumers' investment strategies. Choi, Laibson, and Madrian (2010) dive into the mechanisms behind why individuals invest in index funds that do not minimize fees and show that this continues to hold when search costs are removed and is not explained by nonportfolio services. Hastings, Hortaçsu, and Syverson (2017) show that consumers in Mexico are heavily persuaded by advertising and pay substantial fees since the public pension system was privatized in the 1990s. These papers show broadly that consumers often leave a lot of money on the table in this domain, arguably because of misinformation, but still only scratch the surface of determining precisely why.

Table 2 highlights several other examples. Consumers act as if they do not know the features of certain options, such as the energy cost savings associated with energy-efficient lightbulbs (Allcott and Taubinsky 2015). They act as if they do not know add-on prices such as the sales taxes and shipping costs associated with consumer products (Chetty, Looney, and Kroft 2009; Brown, Hossain, and Morgan 2010). They act as if they do not know basic features of income tax schedules, such as marginal tax rates at current income levels (for example, Rees-Jones and Taubinsky 2016). They act as if they do not know information about their own behavior, like the number of cell-phone plan minutes they have used within a plan month (Grubb and Osborne 2015).

Discussion

While we have focused on a subset of markets, the evidence suggests that researchers would find that consumers face similar challenges in markets that have not yet been studied empirically, whether because of a lack of data or because it is difficult for researchers to assess mistakes in a given context. For example, it is simpler as a researcher to study branded versus generic drugs, which are chemically equivalent, than it is to study decisions where consumer heterogeneity is more important. But the finding that consumers overpurchase branded drugs suggests that consumers make misinformed choices in a variety of similar contexts. Likewise, the documented difficulties consumers have in choosing health insurance and financial products suggest that they also likely experience similar difficulties in choosing other complex financial products, such as life insurance, car insurance, credit cards, or loans.

As Tables 1 and 2 illustrate, "not knowing" in many of these examples could arise from a range of mechanisms. To think about how we might go about trying to distinguish between them (and the situations in which doing so is more or less important), it is useful to elaborate on what these mechanisms might be.

Possible Mechanisms

To help spell out possible mechanisms, consider the following framework. A person wishes to choose an action that maximizes utility. For example, the person could be choosing between health insurance plans, or between branded or generic drugs, or inputs to production that yield utility-relevant outcomes. This person faces uncertainty about the optimal action, such as uncertainty about prices, attributes of options, or the relationship between the action and outcome. However, the person can gather and process information that helps resolve this uncertainty. We'll simplify this discussion by collectively referring to the process of gathering and processing data as "attending to data."

In this setting, as one example, the person chooses a health insurance plan given attended-to information on prices and features of plans. Any strategy for attending to information includes a probabilistic distribution over information the person ultimately processes, and induces some potential cost to the person in terms of time and effort. The person should trade off the expected benefits of attending to information, b, against the costs of attending, c, thereby attending if $b > c$. In a number of settings, the benefits b of attending appear to be large, but the person doesn't seem to be attending to information. What could be going on?

The cost frictions framework says the costs of attending, c, must be large as well. For example, a consumer shopping in an insurance exchange may have correct beliefs about the distribution of prices in the market but incur cost (time and hassle) in finding and exploring each option in the choice set. Or the consumer may have all information on the insurance choice easily available but may not want to do the full calculation on expected costs given the nonlinear contract or health risk because it is too complex or time consuming. Models focusing on cost frictions in gathering, attending to, and integrating information include McCall (1970), Sims (2003), Gabaix (2014), and Woodford (2012).

But this isn't the only possibility. In the alternative mental gaps view, the person may misweight the benefits to attending, using some $\hat{b} \neq b$ in evaluating whether to attend, because important features of the problem are not at the top of the mind. For example, the consumer in a health insurance exchange may mistakenly believe the benefits from searching or attending to information about different options is low when in fact there is substantial price dispersion (or there have been substantial changes to the market). Alternatively, in considering employer plans, the consumer may believe that it is important to focus on the size of provider networks when instead the focus should be on deductibles and premiums. Similarly, a seaweed farmer may not appreciate that pod size matters much for yield. Models focusing on mental

gaps in gathering, attending to, and integrating information include Schwartzstein (2014) and Gagnon-Bartsch, Rabin, and Schwartzstein (2017). Recent laboratory experiments by Enke and Zimmerman (2017) and Enke (2017) explore mental gaps in some detail, as well as de-biasing strategies. Closely related for our purposes are models where a person overreacts to certain salient features of a problem, such as differences in deductibles. Models focusing on systematic errors in integrating information include Bordalo, Gennaioli, and Shleifer (2013), Kőszegi and Szeidl (2012), and Bushong, Rabin, and Schwartzstein (2017).

While not a focus of this article, there are several other possibilities for why people might act as if they are not attending to important information, even when $\hat{b} = b$ and c is low. First, it is of course possible that we as analysts are mismeasuring the potential benefits of improved attention to information. Second, the person may be motivated not to attend to information in order to preserve optimistic beliefs, for example about their own health status (Caplin and Leahy 2001; Brunnermeier and Parker 2005; Kőszegi 2006; Karlsson, Loewenstein, and Seppi 2009; Oster, Shoulson, and Dorsey 2013; Bénabou and Tirole 2016). Third, the person may act on the "wrong" decision utility function, placing too little weight on future benefits (Laibson 1997; O'Donoghue and Rabin 1999) or mispredicting future utility (for example, Loewenstein, O'Donoghue, and Rabin 2003).

Table 3 presents a more detailed look at the frictions and mental gaps mechanisms together with examples from the literature, more carefully decomposing the choice process into stages when barriers to acquiring and optimally using information arise. Some of the examples discussed earlier arguably reflect either mostly frictions or mostly mental gaps. But turning back to Tables 1 and 2, the final column illustrates how many of the examples discussed earlier are consistent with both. Consider the Bronnenberg et al. (2015) branded versus generic drugs example. Cost frictions could be at play: it may be costly to find the generic alternative on the store shelf or to learn about active ingredients. Mental gaps may also play a role: people may not appreciate that generic alternatives to headache remedies are available or believe that chemical equivalence between the products is a possibility worth exploring. Distinguishing between mechanisms in examples such as these requires a more nuanced approach.

Empirical Approaches to the Magnitude of and Reasons for Error

This section discusses empirical approaches for studying environments where cost frictions and mental gaps are present. In this discussion, we will assume that we are considering situations where such frictions and gaps are the primary drivers of the wedge between choices people "should" make and choices they in fact make.

Total Impact on Demand

A range of empirical work seeks to identify the demand curve that represents consumers' actual choices separately from the demand curve in a frictionless

Table 3

Some "Whys" of Not Using Information

Frictions:		
When gathering	*When attending*	*When integrating*
Search costs + Rational expectations	**Rational inattention:**	**Costs of complex thinking, difficulty doing math**
Models: Stigler (1961); McCall (1970); Caplin and Dean (2015)	*Models:* Sims (2003); Gabaix (2014); Matějka and McKay (2015)	*Models:* Ortoleva (2013)
Examples: Hortaçsu and Syverson (2004)—Mutual funds; Cebul, Rebitzer, Taylor, and Votruba (2011)—Health insurance; Ellison and Ellison (2009)—Online markets	*Examples:* Bartoš, Bauer, Chytilová, and Matějka (2016)—Labor market discrimination	*Examples:* Handel and Kolstad (2015b)—Health insurance choice

Mental gaps:		
When gathering	*When attending*	*When integrating*
Search with subjective priors	**Noticing / Selective attention**	**Salience, focusing, relative thinking, limited financial literacy**
Models: Rothschild (1974); Rosenfield and Shapiro (1981)	*Models:* Schwartzstein (2014); Gagnon-Bartsch, Rabin, and Schwartzstein (2017)	*Models:* Bordalo, Gennaioli, and Shleifer (2012, 2013); Kőszegi and Szeidl (2012)—Salience and focusing; Bushong, Rabin, and Schwartzstein (2017)—Relative thinking
Examples: De los Santos, Hortaçsu, and Wildenbeest (2012)—Web browsing and purchasing	*Examples:* Hanna, Mullainathan, and Schwartzstein (2014)—Farming; Malmendier and Lee (2011)—eBay bidding	*Examples:* Bhargava, Loewenstein, and Sydnor (2017)—Health insurance choice

environment with fully rational consumers, which is called the "welfare-relevant" curve. Understanding and estimating the wedge between these two demand curves is sufficient for a variety of important policy questions (Mullainathan, Schwartzstein, and Congdon 2012). Again, we will equate "welfare" with consumer welfare throughout our discussion.

Figure 1 illustrates the demand curve and the welfare-relevant valuation curve for a hypothetical product. The welfare curve is defined conditional on the demand curve, such that the value of the welfare curve shown at any point reflects the average value for marginal consumers on the demand curve at a given quantity level.[2] The wedge between them represents the case where demand is higher than in a rational frictionless environment, leading to over-purchasing in an allocative sense. Each of

---

[2] For simplicity, we will refer to the "demand curve" and "welfare curve" as the key sufficient objects. For certain policy cases, discussed in more depth in the next section, the researcher will also want to understand heterogeneity conditional on a given level of demand in order to use these objects to evaluate policies where consumers may have heterogeneous responses—for example, to taxes or subsidies that may not be equally salient for everyone.

Demand versus Welfare-Relevant Valuation
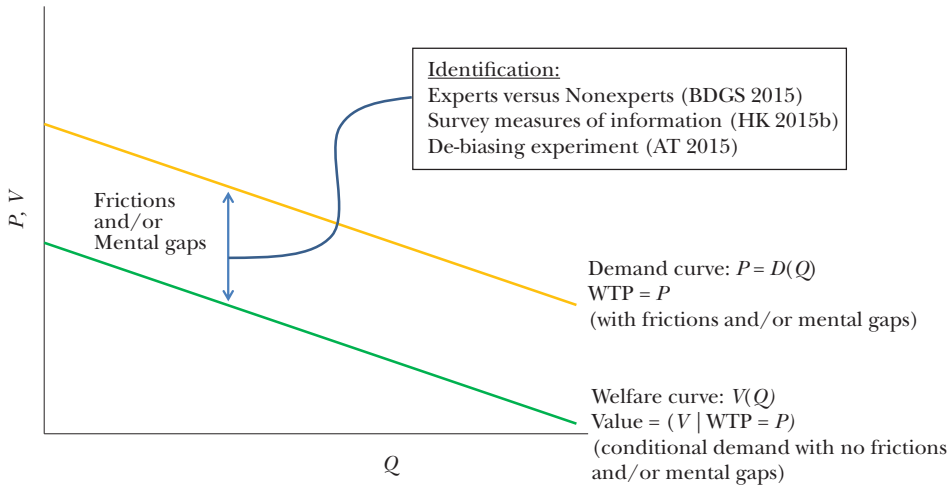
Notes: This figure illustrates empirical approaches that seek to identify observed demand, including frictions and mental gaps, from the welfare-relevant valuation curve, which in some contexts is equivalent to the demand curve for fully informed, frictionless, and bias-free consumers. The welfare-relevant valuation curve gives what true experienced product values would be for consumers at a given level of demand. The wedge between the demand and welfare curve can be due to a range of underlying mechanisms (for example frictions and/or mental gaps) and there are several identification approaches used in the literature to identify this gap. BDGS stands for Bronnenberg et al., HK for Handel and Kolstad, and AT for Allcott and Taubinsky. WTP is "willingness to pay."

the frictions or mental gaps described in the previous section could contribute to this wedge. Recent empirical research highlights several different options for identifying both the demand and welfare-relevant valuation curve in a given environment.

A first empirical strategy estimates a demand curve for experts and a separate demand curve for nonexperts based on the assumption that the demand curve for experts represents the demand curve in a rational frictionless world for experts *and* nonexperts, conditional on a range of observables. In a study mentioned earlier, Bronnenberg et al. (2015) take this approach in studying demand for generic drugs relative to their branded counterparts. When they have quantified the wedge between true demand (of nonexperts) and the welfare-relevant valuation (of experts) for branded versus generic drugs, they can then use this calculation as an input into a welfare analysis of various policies that shift consumers towards generic drugs.

A second approach to identifying this wedge, based on a similar intuition, uses a survey that separates informed from uninformed consumers. The underlying assumption is that informed consumers as measured by the survey make rational full information choices in the context of a neoclassical expected utility model.

One can then quantify the wedge between the demand for informed and uninformed consumers. Handel and Kolstad (2015b) take this approach in seeking to understand why consumers under-purchase high-deductible health plans in a large-employer health insurance context.

A third approach involves using a randomized trial to create a class of well-informed consumers, who can then be compared to others. Allcott and Taubinsky (2015) take this approach in studying the demand for energy-efficient lightbulbs, which seem to be under-purchased relative to both their value for a given individual and relative to their social value (given the externalities imposed by inefficient energy consumption). They assume that consumers in the treatment group are "fully de-biased"—that is, equivalent to the rational frictionless experts and fully informed consumers in the previous two methods. Under this assumption, the demand curve for treated consumers represents the welfare-relevant value curve for all consumers conditional on key observable factors, while actual demand including mental gaps and frictions can be estimated using the control group.

These three approaches differ in the assumptions required to identify welfare-relevant valuation separately from demand.[3] The first strategy (comparing acknowledged experts to nonexperts) is probably the most robust approach of the three, assuming that experts can be appropriately differentiated. Here, the assumptions are that for experts the cost of attention c is relatively low and the perceived benefits to attention are similar to the actual benefits, $\hat{b} \approx b$.

The second approach (using a survey to identify informed and uninformed consumers) presumes that informed consumers have similar preferences to uninformed consumers (conditional on detailed observables), but because they are better informed, they are able to accurately link those preferences to choices. One weakness of this approach relative to the first approach is that eliciting preferences and information sets via survey can introduce well-known issues of measurement error (for discussion, see Bertrand and Mullainathan 2001). A weakness of both approaches is that experts (or informed consumers) who look similar to nonexperts (or uninformed consumers) on observable characteristics may be different on unobservable characteristics.

The third, "de-biasing experiment," approach assumes that the treatment gives a consumer the expertise necessary to operate as a rational frictionless agent (through better calibrating their estimates of benefits $\hat{b}$ or by reducing costs c). The assumptions in this approach are likely the strongest of those needed across the three approaches; indeed, in some cases, the "de-biasing" may even overshoot the true demand curve for reasons argued by Bordalo, Gennaioli, and Shleifer (2015). This approach assumes that the intervention *causes* expertise in a domain, rather than measuring it (survey) or verifying it (occupation data). Of course, experiments can be combined with detailed surveys to assess the level of information or biases

---

[3]A fourth approach, explored by Baicker, Mullainathan, and Schwartzstein (2015), is to estimate (or bound) the welfare curve by directly measuring proxies for inputs to welfare, such as health outcomes.

a consumer has, potentially improving on approaches that use one method or the other.

All three of these approaches assume that a constellation of cost frictions and mental gaps drive the wedge between choices people "should" make and choices they do make. Along with biases specifically related to information, other biases may also be at play in some of the decisions studied, such as present-bias (Laibson 1997; O'Donoghue and Rabin 1999). Greater knowledge of the mechanism(s) driving the wedge in a particular application can bolster confidence in estimates of its size.

Empirical Identification of Specific Mechanisms

The majority of papers that seek to estimate a wedge between demand and welfare curves suggest a specific mechanism that may have caused the wedge, but rarely test their suggested explanation against other possible explanations. For example, a paper that estimates search, switching, or attentional costs typically models a consumer with beliefs closely tied to a rational beliefs framework who incurs costs to acquire key information and improve choices (for example, Hortaçsu and Syverson 2004; Handel 2013). While such papers acknowledge that other factors could also drive poorly informed choices, typically these other factors are not explicitly included in the model. A range of other papers, which are typically less-structural, alternatively allow consumers to make mistakes but largely abstract from more traditional search or processing costs that a social planner might not want consumers to incur (for example, Baicker, Mullainathan, and Schwartzstein 2015).

Distinguishing between the potential mechanisms can provide a more precise characterization of demand versus welfare-relevant value in environments and enable more accurate predictions of policy impacts, but may also require additional data or empirical assumptions. Researchers have used several approaches to differentiate empirically between competing mechanisms. The first uses theoretically motivated assumptions in the context of structural models to test hypotheses about underlying mechanisms: for example, Malmendier and Lee (2011) use this approach to study why some consumers pay more for an item in an eBay auction than they would in a simultaneous fixed-price offering. They test for whether a combination of price uncertainty and switching costs can explain these patterns and argue that the data are inconsistent with this mechanism, implying that some additional mental gap must be a partial cause of these mistakes.[4] One feature of these and other structural approaches (for example, Grubb and Osborne 2015; Barseghyan, Molinari, O'Donoghue, and Teitelbaum 2013) is that they can distinguish between a specific set of potential mental gaps or frictions but must maintain assumptions about other gaps and frictions to do so. Ultimately, the credibility of

---

[4]Schneider (2016) comments on the Malmendier and Lee (2011) paper, suggesting that, under some assumptions, adding traditional search costs into the model can rationalize what might otherwise appear to be bidder mistakes on eBay.

each approach rests on how reasonable these assumptions are in the context of the specific environment being studied.

A second approach used more informally in the literature is to choose one specific mechanism to represent the set of frictions and biases, but then to use calibration arguments to argue that this mechanism is unlikely to explain the entire wedge between demand and welfare-relevant value. For example, Handel's (2013) model of health plan choice assumes that inertia—in which consumers stay with their previous plan even after the elements of the plan have shifted—might result from consumer switching costs. But the size of the switching costs needed to produce this result are estimated to be approximately $2,000. Based on typical values of time costs and some intuition about consumer valuation, this cost seems "too large." The author uses this observation to discuss other potential explanations for inertia in switching between insurance plans, such as biased beliefs and inattention.

A third option is to use survey data to elicit responses about different frictions or mental gaps. For example, Handel and Kolstad (2015b) ask questions about information on a range of dimensions to assess the contribution of different kinds of limited information to demand for health plans. They show, for example, that a lack of information about provider networks has a large impact on demand for high-deductible plans. Their primary structural framework includes indicators for limited information in a reduced-form way, and an alternative framework (presented in an appendix) structurally links indicators of limited information to biased beliefs about certain plan dimensions. Hanna, Mullainathan, and Schwartzstein (2014) similarly combine survey and behavioral data to differentiate between some reasons why seaweed farmers' practices are seemingly off the production possibilities frontier. While classical explanations would likely involve frictions to information-gathering—for example, perhaps due to costs of experimentation—the data instead suggest that farmers were not paying attention to key input dimensions in their own activities. As discussed above, a vast majority could not answer questions about their own practices with regard to key inputs.

A fourth option is to use "mechanism experiments" (Ludwig, Kling, and Mullainathan 2011) to understand the relative impacts of different frictions or biases. Bhargava, Loewenstein, and Sydnor (2017) explore the quality of individuals' health insurance decisions, and reasons for what appear to be mistakes, by analyzing data from an employer where employees choose from large menus of insurance plan options. They find that a *majority* of employees choose health insurance plans that are financially dominated: For example, an employee might pay $500 more in annual premiums to reduce the deductible from $1,000 to $750. One natural hypothesis is that consumers choose financially dominated options because search is difficult and consumers do not know that financially dominating options are available. But evidence from follow-up experiments suggests a basic error may be even more important: many consumers do not appear to know how to map insurance plan features into final wealth outcomes. In a follow-up survey done using the Qualtrics online survey platform, 66 percent of participants choose a financially dominated plan even when the presentation of options was highly simplified to

include four options that only varied in deductible and premium. On the other hand, in another follow-up experiment, this time on Amazon Mechanical Turk, clarifying the relationship between various premium and deductible combinations and total health costs reduced the fraction of participants choosing dominated plans from 48 to 18 percent. Further, those with higher measured understanding of health insurance concepts in this experiment were less likely to choose dominated plans.

When Do We Care Why?

In contexts where consumers appear to leave a lot of money on the table, an obvious accompaniment to looking at the welfare losses is to consider counterfactual public policies, which by definition are out-of-sample. For example, in the health insurance exchanges set up under the Patient Protection and Affordable Care Act of 2010, it is very costly to change regulations related to consumer choice environments (for example, specifying a set of allowable contracts, web designs, or ways in which benefits are represented) and useful to predict impacts of potential new policies.

As you recall, *allocation policies* directly allocate (or strongly steer) consumers to specific actions, and so the underlying cause of the error is unlikely to matter much for policy analysis. *Mechanism policies* instead target specific mechanisms, and so the underlying cause of the consumer error will matter for analysis of that type of policy. While these definitions are not intended to be mutually exclusive or exhaustive, they are intended to broadly frame policies as those that either do or do not strongly interact with the mechanism underlying poorly informed choices.

Allocation Policies

Regulations that remove specific poor options from choice sets, force or nudge consumers into certain better products, or use targeted default options are all examples of allocation policies. Traditional price instruments, such as taxes and subsidies (assuming consumer awareness of those taxes and subsidies) can also be allocation policies. For allocation policies, knowing the precise mechanism behind poorly informed choices is arguably less important than knowing the existence and magnitude of the consumer error.

Table 4 provides examples of some allocation policies. One example in health insurance markets is plan regulation that restricts the actuarial value of plans that insurers can offer in the market. Exchanges set up under the Patient Protection and Affordable Care Act of 2010 allow insurers to offer plans in four tiers of actuarial value: 60, 70, 80, and 90 percent of expected healthcare costs. Consider potential policies that either 1) raise the minimum allowable coverage to 70 percent actuarial value or 2) reduce the maximum allowable coverage to 80 percent actuarial value. Though there are some potential equilibrium pricing consequences that result from such regulation, the first-order effect is likely to shift an entire population of

Table 4
Allocation versus Mechanism Policies

Allocation Policies	*Mechanism Policies*
Health insurance	
Market regulation in Affordable Care Act regarding plan design, like minimum cost-sharing, or structure of cost-sharing.	Choice-framing in insurance markets through web design and information display.
Regulation of minimum networks and covered services.	Education campaigns to promote insurance literacy.
Changes to default insurance options or processes (for example, targeted defaults).	Availability of aggregate and disaggregate information on insurer networks.
	Standardized representation of insurance plans.
Health care services	
Mandatory generic substitution for drugs.	Information pamphlets and posting about equivalence of brand and generic drugs.
Changing medical guidelines to induce changes in default treatments for patients.	Choice framing for brand versus generic drugs.
Value-based cost-sharing.	Shared decision making for difficult medical decisions.
	Information provision on costs or outcomes of medical services.
Financial investment	
Fee regulation eliminating plans with certain types of hidden fees.	Education campaigns promoting financial literacy.
Default options in 401(k) choices.	Standardized display of key fund features.
	Improvements to search tools.
Energy-efficient products	
Regulation on level of energy efficiency required for products.	Education about the value energy efficiency can provide financially.
Taxing energy-inefficient products or subsidizing efficient products.	Education about the impacts of energy efficiency on the environment.
Agricultural production	
Subsidizing or directly distributing inputs like fertilizer.	Agricultural extension, outreach, and education.
School choice	
Changes to the default options or the steering inherent in the choice mechanism.	Information provision about school-choice mechanism, or school options.
Limiting the set of available schools.	Changes to the complexity of the mechanism.

consumers either up or down in terms of coverage generosity. The welfare implications could be assessed if the demand curve and welfare curve for one coverage tier relative to another are identified, without appealing to the specific mechanisms driving the wedge between these curves. Handel and Kolstad (2015b) study a similar example where a large employer shifts from offering multiple insurance options to just one option, a high-deductible plan. The authors are able to assess the welfare implications of such a move after identifying the relevant demand and

welfare curves prior to this forced switch. Because most employers offer only one insurance plan, and many switch their plans over time, forced choice is especially relevant in that market.

In health care services more broadly, a number of states have implemented mandatory generic substitution laws, which essentially require mandatory substitution from brand drugs to generic equivalents except in certain exceptional circumstances. Work like Bronnenberg et al. (2015) that identifies the demand curve and welfare curve for purchases of brand versus equivalent generic drugs can help predict the welfare effects of such a policy. In the domain of over-the-counter drugs, where consumers may have more discretion than for prescribed drugs, the estimates of Bronnenberg et al. also inform how we might want to tax branded drugs or subsidize generic drugs. In health treatment markets, Baicker, Mullainathan, and Schwartzstein (2015) argue that knowing the extent to which people on the price margin are underusing certain treatments (for example, drugs to prevent future heart attacks) is enough to conclude that it would be welfare-enhancing to reduce prices, even without knowing exactly what leads to such underuse.

Table 4 lists examples of allocation policies related to financial investments, energy-efficient products, agricultural production, and school choice. Across these sectors, and the others already discussed, we include default policies that strongly influence the allocation of consumers to products as a borderline case of allocation policies. For example, Madrian and Shea (2001) and follow-up work illustrate how changing the default choice of whether one is automatically enrolled in a retirement savings program powerfully affects the extent of consumer savings. While the effect of defaults of course depends somewhat on the mechanism that drives a wedge between the demand curve and the welfare curve, arguably it is broadly independent of precise details of this mechanism. Table 4 includes a number of other contexts where default policies are likely to be close in spirit to allocation policies.
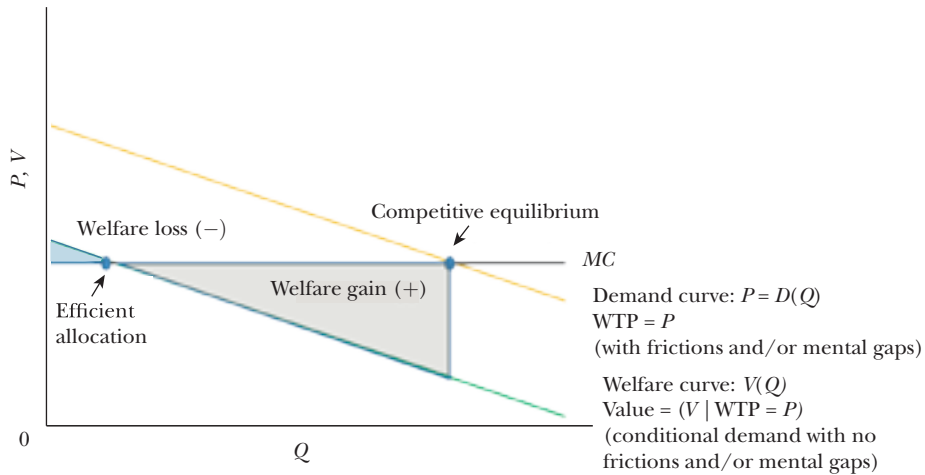
Figure 2 illustrates the welfare implications of an allocation policy in the context of a market for a commodity. For example, assume that the purchase decision studied is the case where consumers are considering whether to buy a brand drug or a generic drug, but that consumers on average are biased towards purchasing a brand drug, relative to actual benefits. The demand curve represents the *relative* revealed preference for a brand drug relative to a generic drug, as a function of price, while the welfare curve represents the distribution of the *actual welfare-relevant relative value* for fully informed, frictionless, and bias-free consumers. The cost curve represents the higher social marginal cost of the brand drug, perhaps in this case because of advertising.

The figure illustrates the welfare effects of an allocation policy that allocates all consumers to the generic counterpart of a branded drug. Consumers who had been purchasing branded drugs, but for whom the actual relative value of the branded drug is much lower, have a large welfare gain. But the figure also allows for the possibility that some subset of consumers loses from this allocation policy: even if all consumers have the same bias towards purchasing branded drugs, as the figure posits, some subset might still value the branded drug above its relative marginal cost.

Figure 2
Welfare Impact of an Allocation Policy
(for instance, forcing consumers to buy a generic drug rather than a brand drug)

Notes: This figure illustrates the welfare impact of an allocation policy that restricts the quantity consumed to zero in a market where there is a wedge between demand and welfare-relevant valuation, as a result of frictions and/or mental gaps. The figure applies to the simple case of a competitive market for two products with constant marginal costs, for example, as in the Bronnenberg et al. (2015) case of consumers who consider whether to purchase a brand drug or a generic equivalent. In that case, the demand and welfare curves reflect the relative willingness-to-pay and valuation for a brand drug compared to its generic counterpart, and quantity reflects the amount of the branded drug consumed.

This could, for example, be a case where consumers have a placebo effect induced by advertising for a branded drug, or just gain utility from purchasing a heavily advertised product. This figure underscores that in order to predict the welfare effect of an allocation policy, assessing heterogeneity in perceived value (demand), actual value, and the overall extent of frictions or mental gaps are crucial, while differentiating between specific frictions and mental gaps may be less important.

The distinctions in Figure 2 should be viewed as approximations. In some empirical contexts the data and identification strategy for determining the wedge between demand and welfare-relevant demand do not allow for tight estimates. Relatedly, the definition of an allocation policy is one of degree. Clearly, mandatory generic substitution moves consumers to generic drugs in a way that leaves little room for consumer mental gaps or frictions to affect the outcome of the policy. However, the identified demand and value curves may not be sufficient for studying a tax or subsidy policy if, for example, taxes and subsidies are not particularly "salient" for some consumers (Chetty, Looney, and Kroft 2009). To evaluate such policies, the researcher needs to analyze heterogeneous consumer responses, which may be a function of underlying mechanisms (Taubinsky and Rees-Jones forthcoming).

Figure 3
Welfare Impact of a Mechanism Policy
(for instance, providing consumers with information about the relative value of branded drugs and generic drugs)
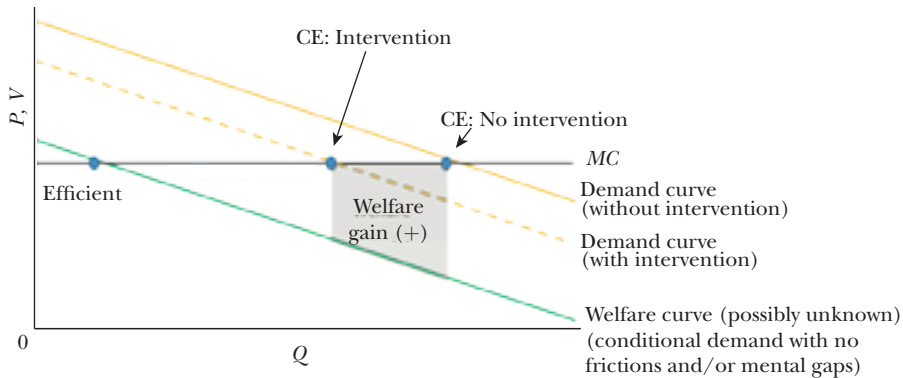
Notes: This figure shows the impact of a mechanism policy, for example one that provides information to consumers about the relative value of branded drugs compared to generic drugs. It shows the case where the policy has a homogeneous impact on all consumers, shifting them part-way towards the true welfare curve from the demand curve. CE stands for competitive equilibrium.

Mechanism Policies

Mechanism policies target specific frictions or mental gaps. For example, sending a consumer a targeted message with choice-relevant information may effectively promote better outcomes if information availability or search costs were the first-order problem, but will be ineffective if the information were always readily available and mental gaps having to do with using or processing that information are more material.

Table 4 also lists some examples of mechanism policies. In health insurance markets, for example, these include standardized representation of insurance plans (Ericson and Starc 2016), education campaigns to promote insurance literacy, choice-framing via specific choice orderings and web designs, or intensive targeted information provision (Kling, Mullainathan, Shafir, Vermeulen, and Wrobel 2012; Handel and Kolstad 2015a). Table 4 also lists some policies related to energy, school choice, and agricultural production. In order to predict the effects of policies that target a specific information-related friction or mental gap, it is necessary to identify the role that mechanism plays in driving choices and potential mistakes. As discussed earlier, this can be quite difficult, usually requiring either multi-arm experiments, comprehensive linked surveys that target information acquisition and processing issues, or natural experiments linked with structural assumptions about these microfoundations.

Figure 3 illustrates the welfare impact of a mechanism policy. For simplicity, the figure assumes that the mechanism policy impacts all consumers evenly, though this

is unlikely to be the case in reality. One can imagine this example as representing the case of information provision for the relative quality of branded drugs versus generics, which would likely reduce the relative demand of some consumers for branded drugs. The figure illustrates the demand impact of this policy, assuming that limited information is one reason, but not the entire explanation, for the wedge between demand and welfare-relevant valuation.

The figure shows the potential pitfalls of using a mechanism policy, like a helicopter drop of information, without having a good sense of the mechanism beforehand. First, the policy may not be very effective: in this case, if information frictions are but one of several frictions and mental gaps, then the drop in demand from the policy is small relative to the drop if the policy truly eliminated all frictions and mental gaps. Second, if the policy used to *remove* frictions and fill in mental gaps was also used in earlier research to *measure* the magnitude of frictions and mental gaps, then the results could seriously understate the benefits from trying hard to eliminate all frictions and gaps. Third, if the demand curve under the policy is mistakenly viewed as the welfare curve, then not only will we understate the potential welfare gains from an ideal policy, we will understate the welfare gains from *this* policy. A given fall in demand from an information drop may appear to barely raise welfare not only because the fall is small, but because this small effect could mislead researchers to infer that people were making good choices to begin with.

In many cases—such as with providing information, making an interface simpler, or encouraging the consumer to make an active choice—it is useful to remember that even policies that seem blunt or obvious may not necessarily target the relevant mechanisms.

An additional key issue in mechanism policies is the extent to which changing the nature of consumer engagement with the choice process causes them to incur additional costs. For example, a policy that reduces consumer information processing costs—for example, via standardized presentation of product attributes—may have multiple effects: 1) help consumers make better choices; 2) cause them to devote more time to the choice process; and 3) incur *more* processing costs than before as a result of this increased engagement. A more straightforward example is a policy that encourages active choices (Carroll, Choi, Laibson, Madrian, and Metrick 2009). For such policies, it is important not only to understand how those policies might improve outcomes (as shown in Figure 3) but also to understand how the costs incurred during the choice process change (and Figure 3 abstracts from that change).

The above policy discussion comes from the perspective of an analyst who seeks to evaluate the likely impacts of a counterfactual policy, whether that policy is an allocation or mechanism policy. It is also possible to evaluate the welfare impacts of a mechanism policy without identifying the exact underlying mechanisms in the case where the empirical analyst can both separate true consumer value from willingness-to-pay and also evaluate the positive impacts of a mechanism policy on choices. In this case, the researcher can use the techniques described (like comparisons of

experts versus nonexperts) to identify value from demand, and can use these fundamentals together with the empirical implementation of a policy to assess its welfare impacts. This can be an efficient way to evaluate mechanism policies when testing these policies is simple and cheap, and when there is a clear way to identify true value from willingness-to-pay in the empirical environment.

Discussion

Rapid improvements in the depth and scope of data available to empirical research have fueled a wave of recent research on the extent to which consumers leave meaningful value on the table as a result of frictions and mental gaps. Policymakers have used this research to motivate a wide range of policies, including setting default options, influencing or constraining choice sets, providing information, standardizing products, and promoting active choices.

Yet there is much weaker evidence on which mechanisms are most important in given contexts. Many research articles explicitly model one mechanism as *the* key friction or mental gap and assume away all other potential explanations. This is typically done for simplicity and for exposition: it is often useful for researchers to act as if one mechanism were the true mechanism even if there is little in the data to distinguish it from other potential mechanisms. Such articles often discuss other potential mechanisms as alternative explanations but don't consider them in depth.

Our main goal in this article is to highlight this issue and investigate how to deal with it in empirical work and policy analysis. Economists sometimes have an intuition that nudges are more conservative than traditional policy instruments when we are uncertain about the mechanism underlying poor choices (Thaler and Sunstein 2009). In contrast, we emphasize a way that blunter allocation policies may actually be conservative: for allocation policies, it is less important to understand the precise mechanisms leading to consumer mistakes than to estimate the wedge between demand and welfare. A growing literature uses survey data, data on experts, and "de-biasing" experiments to identify this wedge and to illustrate its implications for different policies.

The ability to characterize the impacts of allocation policies more easily means that policymakers may have a more precise assessment of those policies, not necessarily that those policies are preferable. The direct intervention of allocation policies is a blunt instrument that may ignore heterogeneity in consumer preferences and the valuable role that informed consumers play in causing the market to provide the best possible products at the lowest possible prices.

More targeted mechanism policies may be better or more politically palatable. However, to evaluate the potential effects of these policies, it is crucial to understand which specific mechanisms lead to consumer mistakes in the first place. While noting the paucity of such evidence across important contexts, we highlighted some promising approaches. As data depth and scope improve, empirically disentangling mechanisms in a given context will become increasingly viable.

References

Abaluck, Jason, and Jon Gruber. 2011. "Choice Inconsistencies among the Elderly: Evidence from Plan Choice in the Medicare Part D Program." *American Economic Review* 101(4): 1180–1210.

Abaluck, Jason, and Jon Gruber. 2016. "Evolving Choice Inconsistencies in Choice of Prescription Drug Insurance." *American Economic Review* 106(3): 2145–84.

Allcott, Hunt, and Dmitry Taubinsky. 2015. "Evaluating Behaviorally Motivated Policy: Experimental Evidence from the Lightbulb Market." *American Economic Review* 105(8): 2501–38.

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

Barseghyan, Levon, Francesca Molinari, Ted O'Donoghue, and Joshua C Teitelbaum. 2013. "The Nature of Risk Preferences: Evidence from Insurance Choices." *American Economic Review* 103(6): 2499–2529.

Bartoš, Vojtěch, Michal Bauer, Julie Chytilová, and Filip Matějka. 2016. "Attention Discrimination: Theory and Field Experiments with Monitoring Information Acquisition." *American Economic Review* 106(6): 1437–75.

Beaman, Lori, Jeremy Magruder, and Jonathan Robinson. 2014. "Minding Small Change among Small Firms in Kenya." *Journal of Development Economics* 108: 69–86.

Bénabou, Roland, and Jean Tirole. 2016. "Mindful Economics: The Production, Consumption, and Value of Beliefs." *Journal of Economic Perspectives* 30(3): 141–64.

Bertrand, Marianne, and Sendhil Mullainathan. 2001. "Do People Mean What They Say? Implications for Subjective Survey Data." *American Economic Review* 91(2): 67–72.

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

1319–72.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2012. "Salience Theory of Choice under Risk." *Quarterly Journal of Economics* 127(3): 1243–85.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2013. "Salience and Consumer Choice." *Journal of Political Economy* 121(5): 803–43.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2015. "Memory, Attention and Choice." https://scholar.harvard.edu/files/shleifer/files/evokedsets_march18_0.pdf.

Bronnenberg, Bart J., Jean-Pierre Dubé, Matthew Gentzkow, and Jesse M. Shapiro. 2015. "Do Pharmacists Buy Bayer? Informed Shoppers and the Brand Premium." *Quarterly Journal of Economics* 130(4): 1669–1726.

Brown, Jennifer, Tanjim Hossain, and John Morgan. 2010. "Shrouded Attributes and Information Suppression: Evidence from the Field." *Quarterly Journal of Economics* 125(2): 859–76.

Brunnermeier, Markus K., and Jonathan A. Parker. 2005. "Optimal Expectations." *American Economic Review* 95(4): 1092–1118.

Bushong, Benjamin, Matthew Rabin, and Joshua Schwartzstein. 2017. "A Model of Relative Thinking." http://www.hbs.edu/faculty/Pages/download.aspx?name=RelativeThinking.pdf.

Caplin, Andrew, and Mark Dean. 2015. "Revealed Preference, Rational Inattention, and Costly Information Acquisition." *American Economic Review* 105(7): 2183–2203.

Caplin, Andrew, and John Leahy. 2001. "Psychological Expected Utility Theory and Anticipatory Feelings." *Quarterly Journal of Economics* 116(1): 55–79.

Carroll, Gabriel D., James J. Choi, David Laibson, Brigitte C. Madrian, and Andrew Metrick. 2009. "Optimal Defaults and Active Decisions." *Quarterly Journal of Economics* 124(4): 1639–74.

Cebul, Randall D., James B. Rebitzer, Lowell J. Taylor, and Mark E. Votruba. 2011. "Unhealthy

Insurance Markets: Search Frictions and the Cost and Quality of Health Insurance." *American Economic Review* 101(5): 1842–71.

Chetty, Raj, Adam Looney, and Kory Kroft. 2009. "Salience and Taxation: Theory and Evidence." *American Economic Review* 99(4): 1145–77.

Choi, James, David Laibson, and Brigitte C. Madrian. 2010. "Why Does the Law of One Price Fail? An Experiment on Index Mutual Funds." *Review of Financial Studies* 23(4): 1405–32.

Choudhry, Niteesh K., Jerry Avorn, Robert J. Glynn, Elliott M. Antman, Sebastian Schneeweiss, Michele Toscano, Lonny Reisman, Joaquim Fernandes, Claire Spettell, Joy L. Lee, Raisa Leven, Troyen Brennan, and William H. Shrank for the Post-Myocardial Infarction Free Rx Event and Economic Evaluation (MI FREEE) Trial. 2011. "Full Coverage for Preventive Medications after Myocardial Infarction." *New England Journal of Medicine.* Special article posted December 1. http://www.nejm.org/doi/full/10.1056/nejmsa1107913#t=article.

Cunningham, Tom. 2013. "Comparisons and Choice." Unpublished paper. May 29.

De los Santos, Babur, Ali Hortaçsu, and Matthijs R. Wildenbeest. 2012. "Testing Models of Consumer Search Using Data on Web Browsing and Purchasing Behavior." *American Economic Review* 102(6): 2955–80.

Ellison, Glenn, and Sara Fisher Ellison. 2009. "Search, Obfuscation, and Price Elasticities on the Internet." *Econometrica* 77(2): 427–52.

Enke, Benjamin. 2017. "Complexity, Mental Frames, and Neglect." Available at SSRN: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2922274.

Enke, Benjamin, and Florian Zimmermann. 2017. "Correlation Neglect in Belief Formation." Unpublished paper, August 8.

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.

Farhi, Emmanuel, and Xavier Gabaix. 2017. "Optimal Taxation with Behavioral Agents." NBER Working Paper 21524.

Gabaix, Xavier. 2014. "A Sparsity-Based Model of Bounded Rationality." *Quarterly Journal of Economics* 129(4): 1661–1710.

Gagnon-Bartsch, Tristan, Matthew Rabin, and Joshua Schwartzstein. 2017. "Channeled Attention and Stable Errors." Preliminary, December 10. https://scholar.harvard.edu/files/gagnonbartsch/files/channeledattentioncleandec10.pdf.

Grubb, Michael D., and Matthew Osborne. 2015. "Cellular Service Demand: Biased Beliefs, Learning, and Bill Shock." *American Economic Review* 105(1): 234–71.

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

Handel, Ben, and Jonathan Kolstad. 2015a. "Getting the Most From Marketplaces: Smart Policies on Health Insurance Choice." Brookings Hamilton Project Discussion Paper 2015-8. October. http://www.hamiltonproject.org/assets/files/smart_policies_on_health_insurance_choice_final_proposal.pdf.

Handel, Benjamin R., and Jonathan T. Kolstad. 2015b. "Health Insurance for 'Humans': Information Frictions, Plan Choice, and Consumer Welfare." *American Economic Review* 105(8): 2449–2500.

Hanna, Rema, Sendhil Mullainathan, and Joshua Schwartzstein. 2014. "Learning through Noticing: Theory and Evidence from a Field Experiment." *Quarterly Journal of Economics* 129(3): 1311–53.

Hastings, Justine S., Ali Hortaçsu, and Chad Syverson. 2017. "Sales Force and Competition in Financial Product Markets: The Case of Mexico's Social Security Privatization." *Econometrica* 85(6): 1723–61.

Ho, Kate, Joseph Hogan, and Fiona Scott Morton. 2017. "The Impact of Consumer Inattention on Insurer Pricing in the Medicare Part D Program." NBER Working Paper 21028.

Hortaçsu, Ali, and Chad Syverson. 2004. "Product Differentiation, Search Costs, and Competition in the Mutual Fund Industry: A Case Study of S&P 500 Index Funds." *Quarterly Journal of Economics* 119(2): 403–56.

Ito, Koichiro, Takanori Ida, and Makoto Tanaka. 2016. "Information Frictions, Inertia, and Selection on Elasticity: A Field Experiment on Electricity Tariff Choice." July 24. https://cenrep.ncsu.edu/cenrep/wp-content/uploads/2016/12/K-Ito.pdf.

Jessoe, Katrina, and David Rapson. 2014. "Knowledge is (Less) Power: Evidence from Residential Energy Use." *American Economic Review* 104(4): 1417–38.

Karlsson, Niklas, George Loewenstein, and Duane Seppi. 2009. "The Ostrich Effect: Selective Attention to Information." *Journal of Risk and Uncertainty* 38(2): 95–115.

Ketcham, Jonathan D., Claudio Lucarelli, and Christopher A. Powers. 2015. "Paying Attention or Paying Too Much in Medicare Part D." *American Economic Review* 105(1): 204–33.

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.

Kőszegi, Botond. 2006. "Ego Utility, Overconfidence, and Task Choice." *Journal of the European Economic Association* 4(4): 673–707.

Kőszegi, Botond, and Adam Szeidl. 2012. "A Model of Focusing in Economic Choice." *Quarterly Journal of Economics* 128(1): 53–104.

Laibson, David. 1997. "Golden Eggs and Hyperbolic Discounting." *Quarterly Journal of Economics* 112(2): 443–78.

Loewenstein, George, Ted O'Donoghue, and Matthew Rabin. 2003. "Projection Bias in Predicting Future Utility." *Quarterly Journal of Economics* 118(4): 1209–48.

Ludwig, Jens, Jeffrey R. Kling, and Sendhil Mullainathan. 2011. "Mechanism Experiments and Policy Evaluations." *Journal of Economic Perspectives* 25(3): 17–38.

Madrian, Brigitte C., and Dennis F. Shea. 2001. "The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior." *Quarterly Journal of Economics* 116(4): 1149–87.

Malmendier, Ulrike, and Young Han Lee. 2011. "The Bidder's Curse." *American Economic Review* 101(2): 749–87.

Matějka, Filip, and Alisdair McKay. 2015. "Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model." *American Economic Review* 105(1): 272–98.

McCall, John Joseph. 1970. "Economics of Information and Job Search." *Quarterly Journal of Economics* 81(1): 113–26.

Mullainathan, Sendhil, Joshua Schwartzstein, and William J. Congdon. 2012. "A Reduced-Form Approach to Behavioral Public Finance." *Annual Review of Economics* 4: 511–40.

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

Ortoleva, Pietro. 2013. "The Price of Flexibility: Towards a Theory of Thinking Aversion." *Journal of Economic Theory* 148(3): 903–34.

Oster, Emily, Ira Shoulson, and E. Ray Dorsey. 2013. "Optimal Expectations and Limited Medical Testing: Evidence from Huntington Disease." *American Economic Review* 103(2): 804–30.

Pauly, Mark V., and Fredric E. Blavin. 2008. "Moral Hazard in Insurance, Value-Based Cost Sharing, and the Benefits of Blissful Ignorance." *Journal of Health Economics* 27(6): 1407–17.

Rees-Jones, Alex, and Dmitry Taubinsky. 2016. "Heuristic Perspectives of the Income Tax: Evidence and Implications for Debiasing." NBER Working Paper 22884.

Rosenfield, Donald B., and Roy D. Shapiro. 1981. "Optimal Adaptive Price Search." *Journal of Economic Theory* 25(1): 1–20.

Rothschild, Michael. 1974. "Searching for the Lowest Price When the Distribution of Prices is Unknown." *Journal of Political Economy* 82(4): 689–711.

Schneider, Henry S. 2016. "The Bidder's Curse: Comment." *American Economic Review* 106(4): 1182–94.

Schwartzstein, Joshua. 2014. "Selective Attention and Learning." *Journal of the European Economic Association* 12(6): 1423–52.

Sims, Christopher A. 2003. "Implications of Rational Inattention." *Journal of Monetary Economics* 50(3): 665–90.

Stigler, George J. 1961. "The Economics of Information." *Journal of Political Economy* 69(3): 213–25.

Taubinsky, Dmitry, and Alex Rees-Jones. Forthcoming. "Attention Variation and Welfare: Theory and Evidence from a Tax Salience Experiment." *Review of Economic Studies.*

Thaler, Richard H., and Cass R. Sunstein. 2009. *Nudge: Improving Decisions about Health, Wealth, and Happiness.* Penguin Books.

Woodford, Michael. 2012. "Inattentive Valuation and Reference-Dependent Choice." May 2. http://www.columbia.edu/~mw2230/InattentiveValue%20Harvard%20Seminar.pdf.