An Information-Costs Theory of Default Rules

Oren Bar-Gill and Omri Ben-Shahar*

February 2019

[Preliminary and Incomplete. Please do not cite without permission.]

Abstract

Policymakers and scholars – both lawyers and economists – have long been debating the optimal design of default rules. From the classic works on defaults for contracts and corporations to the modern, Nudge-inspired rush to use defaults for promoting social policies about issues as diverse as organ donations, retirement savings, mortgages, student loans, overdraft protection, privacy, and more. Much of the debate has focused on opt-out costs – how to minimize them, or how to increase them to secure the policymaker's desired outcome. We argue that this focus on what we call mechanical opt-out costs is misguided. The more important driver of opt-out decisions are information costs – the costs that must be incurred before an informed opt-out decision can be made. High information costs can make defaults stickier. But they can also make defaults slippery. This counterintuitive result is based on the phenomenon of uninformed opt-out, which we identify and characterize. Indeed, the importance of uninformed opt-out requires a reassessment of the conventional wisdom about Nudge and asymmetric or libertarian paternalism. We also show that different defaults provide different incentives to acquire the information necessary for informed opt out. With the ballooning use of default rules as a policy tool, our information-costs theory provides valuable guidance to policymakers.

* Harvard Law School and University of Chicago Law School, respectively. For helpful comments and suggestions, we thank Ian Ayres, Lucian Bebchuk, John Goldberg, Florencia Marotta-Wurgler, Mark Ramseyer. Haggai Porat provided outstanding research assistance.

I. Introduction

The design of default rules—provisions that govern unless actively negated—is one of the central techniques of lawmaking. Because of its centrality, the question how to design optimal default rules has been the subject of enormous commentary. Advancing this inquiry is also one of the greatest achievements of the economic approach to law.

The basic insight that ignited decades of interest in the theory of default rules is the Coase Theorem. Default rules should mimic the will of the parties, because otherwise they would force people to waste transactions costs in disclaiming them. If the population governed by a default rule is heterogeneous, continues this argument, better enact a "majoritarian" default so as to reduce the aggregate cost of opt out.

By recognizing the centrality of opt-out costs, the economic analysis of default rules opened a rich and influential agenda for the design of legal defaults. While early Coasian writers focused on the goal of minimizing opt-out costs,[1] later contributors adopted a markedly different approach asking, not how to reduce opt-out costs, but instead how to optimally exploit their presence. This more recent literature suggests that the "stickiness" of default rules—the tendency of people not to override defaults—is a blessed feature that could improve social welfare. With stickiness attributed to psychological or behavioral factors, and even deliberately "manufactured" by legally-required robust (and expensive) opt out procedures, writers have increasingly viewed sticky default rules as a major policy tool. Their popularity was further bolstered when relabeled as a "Nudge"—a soft regulatory arrangement that helps people secure superior outcomes.[2]

It is difficult to exaggerate the hopes that have been hung on sticky default rules. Devoted advocates view sticky defaults as a "one-size-fits-all" solution to contracting failures.[3] The idea that default rules are and should be sticky was optimistically applied to advocate a "plain vanilla" mortgage that all borrowers should receive from the banks unless they vigilantly choose otherwise.[4] The idea was also applied to reset the rules governing overdraft or student loans in the hope that borrowers will not disclaim them,[5] to reinforce digital privacy protection by requiring more robust forms of user consent,[6] and, famously, to design retirement plans that auto-enroll workers into saving rates which the savers, hopefully, would not undo.[7]

---

[1] See, e.g., Charles Goetz and Robert Scott, "The Mitigation Principle: Toward A General Theory of Contractual Obligation," 69 Va. L. Rev. 967 (1983); Lucian A. Bebchuk and Steven Shavell, "Information and the Scope of Liability for Breach of Contract: The Rule of "Hadley v. Baxendale"," 7 J.L. ECON. & ORG. 284-312 (1991); Richard A. Posner, Economic Analysis of Law (7th ed., Aspen Publishers, NY, 2007); Steven Shavell, Foundation of Economic Analysis of Law (2004); Cooter and Ulen, Law and Economics; Schwartz.
[2] [RA: Cite Nudge]
[3] See Barr, Mullainathan, and Shafir*, A One-Size-Fits-All Solution*, New York Times (December 26,2007).
[4] See Barr, Mullainathan, and Shafir, *An Opt-Out Home Mortgage System* (Brookings 2008).
[5] See Cox et al., Designed to Fail: Effects of the Default Option and Information Complexity on Student Loan Payment, NBER Working Paper #25258 (2018)
[6] [RA: Cite GDPR]
[7] [RA: Cite classic papers by Madrian and Shea (2001), Laibson et al; and also the critique, e.g., Bubb and Pildes.]

If opt out costs are so profoundly critical for regulatory design, you would think that a rich account would have developed to explain what exactly are these opt-out costs, these impediments to smooth contracting, that the Coasian tradition wants to minimize and the Nudge folklore wants to exploit? Surprisingly, despite the fundamental importance of opt-out costs, existing accounts are fairly thin in characterizing them. The typical view of opt-out costs focuses on what we call "mechanical costs": the process of negotiating, deciding upon, and drafting a tailored alternative to the default. In arms-length commercial contracts, such negotiation and redrafting may be time consuming.[8] In mass contracts, opt out requires lengthier standard-form drafting, more disclosures, and additional signatures and clicks. The cost of each opt out may be trivial, but when summed over vast transactions they add up.[9] Indeed, when deliberately designing sticky default rules, lawmakers have sought to increase these mechanical costs, for example by requiring comprehensive disclosures, segregated signatures, and personalized interactions—to make it harder to mindlessly (and cheaply) opt-out.

We think that mechanical costs provide a poor foundation for the theory of default rules, and we suggest a different, arguably more important, factor affecting opt-out – information costs. A decision to stick with or opt-out from a default is based on the parties' perceptions concerning the existence and relative value of the alternatives. People often do not readily know what the default is, how to value it, and how to compare it to any of the (sometimes unknown) alternatives.[10]

Consider the canonical example of retirement savings. The mechanical costs of opting out from the default savings rate are trivial, especially at the time of employment initiation, when the various benefits choices are made. Opting in or opting out of the default is just another brief click. But the information costs surrounding such decision could be substantial. People have to know how to value the default, which requires them to make financial projections about their lifetime income and evaluate their consumption and borrowing needs. An intensive session of financial counseling is required to make an informed opt out decision.

Information costs, we argue, are the key impediment to optimal opt out in a multitude of contexts. Consider contracting over data privacy. No matter how robust and meaningful the consent gesture needed to reverse the no-data-collection default and allow websites to collect personal information, the mechanical costs are still small. They are greatly

---

[8] [Add: A different type of opt out cost in arms-length contracts are "strategic costs" – the delay or failure to reach agreement in the presence of strategic bargaining behavior.]

[9] Even when the per-individual opt-out cost is small, massive opt outs by millions of individuals can be very costly to society. Opt-out costs are especially troubling, when firms have the incentive to deliberately increased such costs. See, e.g., https://apnews.com/f60bc112665b458cb6473d7ee9492932 (an example where Google increased opt-out costs).

[10] We began to lay the foundation for our information-costs theory in Oren Bar-Gill and Omri Ben-Shahar, "Optimal Defaults in Consumer Markets," 45 J. Legal Stud. S137 (2016). A different notion of information costs as an impediment to opt out has been studied extensively in the literature. It focused on the ways opt outs convey information in environments of asymmetric information. See, e.g., [RA: Add citations to Ayres and Gertner, Johnston, Yale L. J., Bernstein, S.Cal., Spier. See, generally, Ben-Shahar & Pottow, Fl. St. L. Rev.]

overshadowed by information costs: learning what data is collected and evaluating the costs and benefits of various permissions to collect the data. Or, consider mortgage borrowing and the proposal that the law establish a "plain vanilla" thirty-year fixed-rate mortgage as a sticky default. Lenders who offer other loan products may have to comply with "heightened" disclosure requirements and an assortment of opt out rituals. But these mechanical costs are nothing compared to the staggering costs of evaluating the financial implications of different loan provisions.

Shifting the focus away from mechanical costs and onto information costs as the primary impediment to optimal opt-outs leads to novel insights, including a fundamental rethinking of the notion of stickiness. Upon first reflection, it might be thought that because of information costs and of incomplete information people might not opt out of the default, even when the alternative to the default is better. In other words, uninformed people would stick to the default.[11] If this were the effect of information costs—to freeze people into the default—information costs would indeed be similar to mechanical costs, and their presence would make default rules all the more sticky.

But there is another possibility. We argue that information costs do not necessarily make defaults sticky. High information costs may stop people from becoming informed, but may not stop them from opting out! We develop the idea of "uninformed opt out"—where people are uninformed about the relative value of the default and are unwilling or unable to incur the cost necessary to become informed, but nevertheless decide to affirmatively opt out, based on their uninformed perceptions about the relative value of the alternative. We call such default rules "slippery"—not only do they not stick, but they prompt people to descend from them without the traction of informed deliberation. Stickiness, in other words, is an artifact of high mechanical costs, not of high information costs. And if, in reality, opt out behavior is determined primarily by information costs, default rules are less sticky and more slippery than otherwise hoped.

Uninformed opt out has received surprisingly little, if any, attention. The identification and analysis of uninformed opt out as an important phenomenon in the world of default rules is one of the main contributions of our information-costs theory. We compare two versions of this phenomenon: rational and irrational. Rational uninformed opt-out occurs, when people correctly think that "on average" opt out is better than the default. The irrational version of uninformed opt out occurs when people misperceive the value of the opt out

[11] See, e.g., Cass R. Sunstein and Richard H. Thaler, *Libertarian Paternalism Is Not an Oxymoron*, 70 U. CHI. L. REV. 1159, 1197-98, 1201 (2003) ("In some domains, consumers and workers are highly informed - so much so that they will not even be influenced by default rules."; "… in many domains, people's preferences are labile and ill-formed, and hence starting points and default rules are likely to be quite sticky."); Jeffrey R. Brown, Anne M. Farrell, and Scott J. Weisbenner, *The Downside of Defaults* (Unpublished Manuscript, 2011) ("…lack of adequate information about decision alternatives is a significant driver of the likelihood of default… information problems are especially important. Notably, more than 50 percent of participants who defaulted cited at least one form of information problem as a reason behind their default."); Cass R. Sunstein, *Deciding by Default*, 162 U. PA. L. REV. 1, 20 (2013) ("… there is strong evidence that a lack of information on the part of choosers, including a lack of information about alternatives, helps to account for the power of defaults.")

option. The latter is particularly plausible in consumer markets where firms have an incentive to create such misperceptions.[12]

Shifting the focus to information costs reveals another set of insights: how the default rule shapes the amount of information people ultimately have. We show that the content of the default influences the incentives to acquire information. This, in turn, drives people's decisions to engage in informed or uninformed opt out. The effect of the default rule on the amount of information acquired, and indirectly on opt-out decisions, should play a central role in the design of default rules.

Let us summarize the gist of our analysis in the context of designing a default rule for retirement savings. Start with two placeholder cases—very low and very high information costs. When information costs are very low, people will become informed regardless of the default. Because people vary in their retirement goals, some of them will then opt out (unless the mechanical costs are prohibitive). The best default rule in this case is the majoritarian default, which gets people to sort out while incurring the least mechanical, opt-out costs. This is the familiar opt-out cost minimization principle, widely discussed in the literature to justify majoritarian defaults.

In contrast, when information costs are very high (a much more likely scenario in the context of retirement savings), people will not acquire the information, regardless of the default. Instead, they will evaluate the default and the alternative options based on their uninformed beliefs. They might opt out—an uninformed opt-out—if they deem the default to produce lower expected value (assuming, again, that mechanical costs are low). Here, regardless of the default, people end up with a retirement plan that maximizes the expected long-term value. High information costs do not in themselves cause the default to be sticky. Policymakers should therefore choose a default plan not according to the majoritarian principle; rather they should choose a default that maximizes the expected value.

Notice that these benchmark cases help us distinguish two different criteria for efficient design of default rules, both within the paradigm of reducing opt out costs. The Majoritarian principle fits well the case of informed heterogenous individuals who selectively opt out. It prescribes a majoritarian default to minimize the incidence and cost of informed opt out. The Expected Value Maximization principle (EVM) fits the case of uninformed individuals that act uniformly, based on similar beliefs, and are therefore best served by a retirement plan that fits their ex ante, uninformed beliefs. Such default retirement plan minimizes the cost of uninformed opt out.

Between these two polar cases, we consider the "intermediate" information costs case. Here, information about retirement plans is acquired and acted upon only under some

[12] To elaborate: When will we see uninformed opt-out? The answer depends on a comparison between (i) the expected net benefit of the default (relative to the opt-out alternative), and (ii) the individual's opt-out cost. When factor (i) is positive, there will be no uninformed opt-out. When factor (i) is negative, namely, when the default generates a lower expected net benefit that the opt-out alternative, there is a possibility of uninformed opt-out. The uninformed party would still stick with the default, as long as the disadvantage of the default, in terms of negative expected net benefits, is dominated by the opt-out cost (factor (ii)). It is when the negative expected net benefits dominates the opt-out cost that we get uninformed opt out.

default provisions, not others. This case helps us demonstrate that the value of information varies by the default. Within this region of intermediate information costs, we distinguish two possibilities. First, when information costs are relatively low, the default plan that induces more information acquisition is superior, as it results in more tailored retirement choices. Second, when information costs are relatively high, the default plan that does not induce information acquisition is superior. In this range, the value of tailored choice is outweighed by the cost of information. Getting people to become informed is not necessarily better!

This framework—of examining how the default rules affect information acquisition and how opt-out occurs with or without information—applies to rational as well as irrational choice. We examine how people make opt out choices when they have inaccurate beliefs or straight-up misperceptions about the value of the default and its alternatives. In the retirement savings context, misperceptions are prevalent because it is hard to make unbiased assessments of household finances. In other contexts, misperceptions are prevalent because self-interested, well-informed counter-parties work hard to create misperceptions (e.g., banks luring customers to unnecessarily opt into costly overdraft protection). Mis-information may lead people to poorly estimate the ex-ante relative value of the default and commit uninformed opt out that is inefficient (or, vice versa, not to commit uninformed opt out that is efficient). Mis-information may also distort people's ex ante assessment of the value of information. Such second-order error may lead individuals to either acquire too much or too little information.

Apart from laying a foundation for the concept of opt out, which is so central to the design of default rules, our analysis sheds new light, and forces a reevaluation, of some of the popular applications of default rule theory, specifically Nudge-type regulation through sticky defaults. Such regulation has been hailed as a "libertarian paternalistic" sorting device—rules that help less sophisticated individuals, who stick with the default, without harming the more sophisticated individuals who make a deliberate decision whether to opt out.[13] But what (or who) counts as "sophisticated"? If sophistication is a proxy for low information costs, our analysis shows that libertarian paternalistic sorting does not always work. When uninformed opt out takes place, sorting fails: even unsophisticated people, with high information costs, may opt out of the default.

---

[13] See Colin Camerer, Samuel Issacharoff, George Loewenstein and Ted O'Donoghue, "Regulation for Conservatives: Behavioral Economics and the Case for Asymmetric Paternalism," 151 U. PA. L. REV. 1211, 1225 (2003) ("As long as actively making a choice requires very little effort, the choice of defaults has essentially no effect on fully rational consumers. But for boundedly rational people who have a status quo bias, the choice of defaults is important."); Cass R. Sunstein, "Boundedly Rational Borrowing," 73 U. CHI. L. REV. 249, 257 (2006) ("…interventions that are choice-preserving (and hence libertarian) are generally asymmetrical, because they are not likely to impose significant costs on people who do not suffer from bounded rationality."); Richard H. Thaler and Cass R. Sunstein, Nudge: Improving Decisions About Health, Wealth, and Happiness, p. 242 (Yale University Press, 2008) ("Most of the time, nudging helps those who need help while imposing minimal costs on those who do not."); Ryan Bubb and Richard H. Pildes, "How Behavioral Economics Trims Its Sails and Why," 127 HARV. L. REV. 1593, 1598-99 (2014) ("…the default is designed to put those who stay with the default in the best position but to enable those with different preferences, more sophistication, greater resources, or other appropriate bases to opt out and choose whatever they prefer.")

In addition, our information-costs theory sheds new light on the conventional wisdom about optimal default design and, in some cases, rejects this conventional wisdom. One strand of the conventional wisdom supports majoritarian defaults, namely, defaults that are optimal for most individuals or, equivalently, optimal for the median individual. The assumption is that the majority – whether informed or uninformed – will stick with the default, and the minority will opt out if informed. Our theory shows that such a majoritarian default can lead to uninformed opt-out, when it is not the expected value maximizing rule. In such cases, we show, a minoritarian default can be better.[14]

The analysis in this Article helps bridge an uncomfortable gap between the academic appetite for sticky default rules and the reality of slippery defaults. The few examples in which default rules managed to stick have led to widespread hopes for this technique. But, as writers in the area of contract law have long recognized, opt out is exceedingly common, especially when consumers are prompted by a self-interested seller. People opt out without sophistication and without information. It is the overlooked phenomenon of uninformed opt-out that accounts for the unintended slipperiness of so many default rules.

Indeed, if information costs are the primary factor determining opt out behavior, the idea of sticky defaults must be reevaluated.[15] As a descriptive matter, we said, high information costs may *reduce* stickiness—a counterintuitive effect that prior accounts of stickiness did not recognize. Normatively, our analysis maps the relationship between stickiness and welfare: when a default is sticky, this stickiness may either increase or decrease welfare. Sticky defaults increase welfare, when they prevent inefficient, uninformed opt-out. Conversely, sticky default rules reduce welfare when high information costs prevent efficient, informed opt-out.

The remainder of the Article is organized as follows. Part II works through a numeric example that illustrates the main positive and normative implications of our information-costs theory of default rules. Part III develops a formal model that generalizes and extends the results from the numeric example. Part IV highlights policy implications. Part V considers several extensions. Part VI offers a more detailed analysis of several applications – policy domains where default rules play a key role. A brief Conclusion follows.

II. A Simple Example

We present here a theoretical illustration of the relationship between information costs and default rules. The analysis is presented through a numeric example. Section A presents the conventional perfect information case, as a benchmark for the imperfect information

[14] [Compare and contrast with other justifications for minoritarian defaults: Ian Ayres and Robert Gertner, "Filling Gaps in Incomplete Contracts: An Economic Theory of Default Rules," 99 YALE L.J. 87, 93, 116 (1989) (discussing reasons to adopt minoritarian defaults); Ayres and Gertner (1992)] [While an OCM default targets the median individual, an EVM default targets the average individual.]

[15] Existing accounts of stickiness include Bernheim et al (2011); Bernheim and Taubinsky (2018); Blumenstock et al (2017) [RA: Add citations to these papers].

analysis in the remaining Sections. Section B assumes that all uninformed parties have the same, accurate beliefs about the distribution of types. Section C relaxes the accurate beliefs assumption.

A. Benchmark: Perfect Information

There are two possible arrangements, "High" and "Low" ("H" and "L"). High could denote, for example, a commitment to be an organ donor, a broader warranty, a higher pension savings rate, or greater privacy protection, relative to Low. The legal default rule could be set at either High or Low. Because the choice High v. Low could affect other aspects of the transaction (like price), people vary in how they value the two arrangements. We assume that 60% of the population prefer Low, and we call them "Type 1." 40% of the population prefers High, and are called "Type 2." Let's use numbers to reflect the valuations assigned:

	Type 1 (60%)	Type 2 (40%)
High	$v = -10$	$v = 20$
Low	$v = 0$	$v = 0$

Table 1: The Example – Setup

To make the example simple, we assumed that everyone assigns value $v = 0$ to Low. For Type 1, High is worse; they assign a value of $v = -10$ to High. And for Type 2, High is better; they assigning a value of $v = 20$ to High. In the perfect information benchmark, everyone knows their type. If the default rule is unattractive to them, they opt out. The mechanical cost of opt out is 1 (assumed to be sufficiently low that any party would opt out from an unattractive default; this allows us to focus attention on information costs).

With a Low default, Type 2 individuals opt out and average social welfare is: $W_{Low} = 0.4 \cdot (20 - 1) + 0.6 \cdot 0 = 7.6$. With a High default, Type 1 individuals will opt out and average social welfare will be: $W_{High} = 0.4 \cdot 20 + 0.6 \cdot (-1) = 7.4$. We have calculated welfare for a representative individual.[16]

Comparing social welfare under the two defaults, we see that Low is the better default, because it generates fewer costly opt outs. This is the standard result that, with perfect information, the majoritarian default is the best. We now turn to the imperfect information case, where this standard result will be qualified.

[16] To asses overall social welfare we would need to multiply these values by the number of individuals who are subject to the default rule.

B. Accurate Uninformed Beliefs

Assume that individuals do not know their type, but can invest some fixed cost to find out. What they do know, even without investing, is the distribution of types: they know that 60% are Type 1 and 40% are Type 2. They can therefore calculate the average value of the two provisions, Low and High. The expected value of Low is 0. The expected value of High is $0.6 \cdot (-10) + 0.4 \cdot 20 = 2$. Accordingly, based solely on such average valuations, people prefer High.[17] And if the default is set at Low, they will want to opt out to High. Given the assumption that the mechanical opt-out cost is 1 (< 2), when the default is Low such "uninformed opt out" would occur.

Uninformed opt out decisions (or uninformed decisions not to opt out) will determine outcomes when the cost of becoming informed is high. When information costs are lower, individuals may decide to acquire information and thereby make the opt out decision in a fully informed manner, knowing their actual types. Let us examine how they would behave, and the resulting social welfare, under the two defaults.

1. Low Default

The first decision people face is whether or not to acquire information. The value of information is

$$I_{Low} = [0.4 \cdot (20 - 1)] - [2 - 1] = 6.6$$

The first component represents a typical payoff with perfect information: the individual has a 40% chance of learning that she is Type 2, and she will then opt out of the Low default, incurring an opt-out cost of 1, and earn a payoff of 20. (There is a 60% chance that the individual will learn that she is Type 1 and stick with the Low default, earning 0.) The second component represents the payoff without perfect information: in that case, uninformed opt-out occurs; the individual earns an expected payoff of 2 while incurring an opt-out cost of 1. The value of acquiring information is the difference between the payoff with versus without information, and it equals 6.6. When information costs are below 6.6, the individual will acquire information.

It is interesting to note that, under the Low default, high information costs *reduce* both stickiness and welfare. Higher information costs make it less likely that people will become informed and selectively opt out. Instead, they opt-out non-selectively, uninformed. At the same time, high information costs reduce welfare, because they eliminate a "separating" outcome in which different people choose the outcome most suitable to them.

2. High Default

Under a High default, the value of information is

[17] We assume that individuals are risk neutral and thus focus on expected values.

$$I_{High} = [0.6 \cdot (-1) + 0.4 \cdot 20] - [2] = 5.4$$

The first component represents a typical payoff with perfect information: the individual has a 60% chance of learning that she is Type 1, and she will then opt out of the High default, incurring an opt out cost of 1, and earning 0 instead of -10; and she has a 40% chance of learning that she is Type 2, and she will then stick with the High default and earn 20. The second component represents the payoff of an uninformed individual: Since the expected value of High (2) is larger than the expected value of Low (0), there will be no (uninformed) opt-out, and the individual will earn an expected payoff of 2. Now, when information costs are below 5.4, people will acquire information.

3. Welfare Comparison

First, notice that the value of information is greater with Low default ($I_{Low} > I_{High}$), which means that a Low default leads to more acquisition of information. Why? Primarily because the expected value of the Low default is lower and even without acquiring information people gain by opting out. Information is more valuable with Low default, because it saves (some of) the costs of these uninformed opt outs. To amplify, we rewrite I_{Low} and I_{High} as follows:

$$I_{Low} = (0.4 \cdot 20 - 2) - 0.4 + 1 = 6.6$$

$$I_{High} = (0.4 \cdot 20 - 2) - 0.6 = 5.4$$

With both defaults, informed parties get $0.4 \cdot 20$ and uninformed parties get 2. The difference lies in the opt-out costs. With Low default, informed opt out occurs 40% of the times (when information is acquired) and uninformed opt-out occurs 100% of the times (when information is not acquired). Information acquisition thus saves 0.6 in opt-out costs. With High default, informed opt out occurs 60% of the times (when information is acquired) and there is no uninformed opt-out. Information acquisition thus *costs* 0.6 in opt-out costs.

It may be thought that a default that induces more information acquisition is necessarily better. But, we show, this turns out to be false. Since information is costly to acquire, it may be better to remain uninformed.

We proceed by distinguishing between three ranges of information costs:

(a) Upper range of Information Costs (Larger than 6.6)

When information costs are above a certain threshold, information is not acquired – under either the Low or the High default. Under Low default, all parties opt-out; there is a 100% opt-out rate – all uninformed opt out. Under High default, no one opts out; the opt out rate is zero. Regardless of the default, all parties end up at the High outcome. Low default is slippery, whereas High default is sticky. The High default happens to be more efficient, because it saves the mechanical costs of uninformed opt out. Generalizing, when

information costs are at the upper range and opt-out costs are small, the optimal default is the one that maximizes the expected value for uninformed parties.

(b) Bottom Range of Information Costs (Smaller than 5.4)

When information costs are below a certain (different) threshold, information is acquired – under both the High and the Low default. Under Low default, people who learn that they are Type 2 opt-out; there is a 40% opt-out rate. Under High default, people who learn that they are Type 1 opt-out; there is a 60% opt-out rate. All opt out is informed. Regardless of the default, people end up with the optimal outcome – Type 1 with Low and Type 2 with High. Here, Low default is stickier. It is also the more efficient rule, because it reduces the cost of informed opt out. Generalizing, when information costs are at the bottom range and opt-out costs are small, the optimal default is the majoritarian one, because it minimizes the costs of informed opt out.

(c) Intermediate Range of Information Costs (Between 5.4 and 6.6)

When information costs are intermediate, information is acquired under Low default, but not under High default. Under Low default, people who learn that they are Type 2 opt-out; there is a 40% opt-out rate. Under High default, the uninformed individuals stick with the default; the opt-out rate is zero. High default is stickier. And either rule may be more efficient. Low default leads to optimal matching (while incurring some opt-out costs), which generates a value of $0.4 \cdot (20 - 1) = 7.6$, but requires costly investment in information acquisition. High default provides the higher average value (of 2) and avoids costly investment in information; social welfare equals the average value of 2. Therefore, Low default is more efficient when information costs are below 5.6 (= 7.6 – 2). To summarize: At the lower end of the intermediate range of information costs, between 5.4 and 5.6, Low default is the more efficient rule; and at the higher end of that range, between 5.6 and 6.6, High default is the more efficient rule.

For expositional purposes, the preceding analysis considers the relative stickiness of the two defaults and the welfare outcomes that they produce for any given level of information costs. Of course, different individuals will have different information costs – some can acquire information quickly and cheaply, whereas others would need to invest more time and money to become informed. When we allow for heterogeneity in information costs, the optimal default will need to balance the different considerations listed above.

The preceding analysis suggests that when information costs are at the upper range, the opt-out rate is greater with Low default, because the uninformed prefer High (on average). When information costs are at the bottom range, information will be acquired and informed opt out will occur. Here, the opt-out rate would still be greater under Low default, if there are more Type 2 individuals (those who benefit from High default) in the population. In our example, there were more Type 1 individuals, and thus High default led to more informed opt out in this range.

C. Inaccurate Uninformed Beliefs

So far, people's beliefs were accurate "on average" – they correctly anticipated the share of Type 1's (60%) and the valuations each type would have under either the High or Low options. We saw that with such accurate beliefs, uninformed opt out guarantees that people can do no worse than to maximize the expected payoff, minus opt-out costs, and sometimes can do better.[18] But this is a lot to assume, and we now illustrate how things change when people's beliefs are inaccurate. Specifically, we examine a special case in which some people overestimate the likelihood that they are Type 1 to be 80%.[19] To be sure, there are many other ways in which beliefs could be inaccurate, and we examine the effects of other distortions more thoroughly in Section III. Here, we merely illustrate how inaccurate beliefs can change the effects of the different defaults and thus the optimal default choice.

Upper range of information costs. When information costs are prohibitive, people act solely on the basis of their beliefs, according to the perceived expected value of each option, which is now 0 for Low and -4 for High. (The perceived value of High is $0.8 \cdot (-10) + 0.2 \cdot 20 = -4$.) If the default is Low, people stick with it—better to keep 0 than pay an opt out cost to get -4. If the default is High, people act on their inaccurate beliefs and opt out uninformed, ending up with a net payoff of -1 (the value of Low, which is 0, minus the mechanical opt-out cost). Either way, the result is inefficient: either sticking with an inefficient default when it is Low, or opting out of an efficient default uninformed when it is High.[20]

Here, inaccurate beliefs change the welfare ordering of the two defaults and thus alter our policy prescriptions. With accurate beliefs, all parties ended up with High, regardless of the default, and the policy preference for High default was based on the avoidance of opt-out costs. With inaccurate beliefs, all parties end up with the Low option, regardless of the default, and it would be better to set the Low default—again, to avoid the waste of opt out costs. While the prescription changes, the principle remains the same: When information costs are at the upper range, set the default that maximizes the expected value that people believe they will get – even if this belief is inaccurate.

Bottom range of information costs. When information costs are easily affordable, such that all parties become informed, inaccurate beliefs about the share of Type 1 individuals do not matter, because people acquire information and do not act upon their beliefs. In other words, when parties acquire individualized information, beliefs about averages are irrelevant.

---

[18] With accurate beliefs, uninformed opt out, by definition, maximizes the expected payoff of the individual who decides to opt out. And, similarly, a decision not to opt out maximizes expected payoffs. The introduction of inaccurate beliefs allows for inefficient, uninformed opt out; and for inefficient decisions not to opt out – to stick with the default.

[19] We could assume inaccurate yet unbiased beliefs, such that some parties overestimate the share of Type 1 individuals while others underestimate this share. Or we could allow for inaccurate, biased beliefs, e.g., when a business invests in manipulating consumer beliefs in order to maximize its profits.

[20] With accurate beliefs, High default was sticky and Low default was slippery. With the inaccurate beliefs, High default is slippery and Low default is sticky.

Intermediate information costs. When information costs are intermediate, the key difference between the two defaults has to do with the relative incentives that they provide for information acquisition. With accurate beliefs, information was acquired with Low default, but not with High default. Inaccurate beliefs affect the perceived value of information under the two default rules. With Low default, the perceived value of information is

$$I_{Low} = [0.2 \cdot (20 - 1)] - [0] = 3.8$$

The perceived expected payoff if information is acquired decreases relative to the accurate beliefs case (from $0.4 \cdot (20 - 1)$ to $0.2 \cdot (20 - 1)$). The perceived payoff if information is not acquired changes, relative to the accurate beliefs case, from the uninformed opt-out payoff of 1 (= 2 – 1) to the no opt-out payoff, zero.

With High default, the perceived value of information is

$$I_{High} = [0.8 \cdot (-1) + 0.2 \cdot 20] - [-1] = 4.2$$

The perceived expected payoff if information is acquired decreases relative to the accurate beliefs case (from $0.6 \cdot (-1) + 0.4 \cdot 20$ to $0.8 \cdot (-1) + 0.2 \cdot 20$). The perceived payoff if information is not acquired changes, relative to the accurate beliefs case, from the no opt-out payoff, 2, to the uninformed opt-out payoff of -1.

Now it is High default that provides stronger incentives to acquire information – incentives that can be particularly important, since individualized information alleviates the need to rely on inaccurate beliefs about average values. (Still, given the cost of acquiring information, the stronger incentive to become informed is not necessarily efficient.) [21]

It is worth noting that, with inaccurate beliefs, a policy that reduces information costs might have the counterintuitive effect of reducing social welfare. This effect could happen under either default rule, and it is due to people's overestimation of the value of information. Specifically, lower information costs cause more people to acquire information. With accurate beliefs, this increased tendency to acquire information is good, because information is acquired only when the value of the subsequent informed opt-out exceeds the cost of information. But with inaccurate beliefs people might imagine a benefit from information that is not real and will then acquire too much information. For example, if the actual benefit from information is 10 and the misperceived benefit is 20, then people would inefficiently acquire information that costs anywhere between 10 and 20. If information costs are high, say 25, the greater tendency to acquire information will not result in inefficient acquisition of information. But if information costs are reduced to, say, 15, people will inefficiently invest in information.[22]

[21] We also observe that the very definition of low vs. intermediate vs. high information costs changes, as the cutoff values of information are affected by the inaccurate beliefs.

[22] This result does not obtain in our example, where the inaccurate beliefs are limited to an underestimation of the share of Type 2 individuals (who benefit from High), which leads to a false belief that Low default is

III. Model

We now generalize and extend the numerical example of Part II, using a formal model. In Section A, we present our framework of analysis. In Section B, we analyze outcomes and welfare with Low default. In Section C, we analyze outcomes and welfare with High default. In Section D, we compare the two defaults and provide guidance to policymakers about optimal default design. We initially assume that uninformed individuals hold accurate beliefs about the relevant parameters and can accurately assess the expected values of the different options. The implications of inaccurate beliefs are explored in Section E.

A. Framework

Consider a binary choice between two options that we will call Low and High. We normalize the net benefit from Low to zero. The benefits and costs generated by High differ across individuals. Specifically, a share $\alpha \in [0,1]$ of individuals enjoy a net benefit $B > 0$, whereas the remaining $1 - \alpha$ incur a net cost of $C > 0$. We call individuals who prefer Low type 1, and call individuals who prefer High type 2.

We consider two possible default rules: Low default (or L default), which corresponds to Low, and High default (or H default), which corresponds to High. Parties can opt out of either default at a cost *k*. We assume that *k* is distributed across contracting pairs according to $F(\cdot)$.

Initially, individuals do not know whether they are type 1 or type 2. Individuals can invest x and learn their type. The investment x varies among individuals, according to the cumulative distribution function $G(\cdot)$. There is a threshold $\hat{x}$ (derived below), such that individuals with $x < \hat{x}$ invest and learn their type, while individuals with $x \geq \hat{x}$ remain uninformed. (This framework covers scenarios where some individuals initially know their type; in such scenarios the probability function would have a mass point at $x = 0$.) [23] We assume that uninformed individuals hold accurate beliefs about the share α and about the parameters *B* and *C*. The implications of inaccurate beliefs are explored in Section E below.

---

better on average. To get the perverse result that lower information costs reduce welfare, we need an overestimation of the share of Type 2 individuals and/or the benefit to Type 2 from High outcome – to get an overestimation of the benefit from acquiring information. And to maintain the assumption that inaccurate beliefs make Low default seem better, we also need an overestimation of the cost of High to Type 1 individuals.

[23] Of the $G(\hat{x})$ individuals who learn their type, $\alpha G(\hat{x})$ learn that they are type 2 and $(1-\alpha)G(\hat{x})$ learn that they are type 1. A share $1 - G(\hat{x})$ of individuals remain uninformed about their type and believe that with a probability α they are type 2 and with probability $1 - \alpha$ they are type 1. This group of uninformed individuals can be further divided into the $\alpha\big(1 - G(\hat{x})\big)$ type 1s and the $(1-\alpha)\big(1 - G(\hat{x})\big)$ type 2s. To summarize: There are four groups of individuals – Group 1, with a measure of $\alpha G(\hat{x})$ who know that they are type 2; Group 2 with measure $(1-\alpha)G(\hat{x})$ who know that they are type 1; Group 3 with measure $\alpha\big(1 - G(\hat{x})\big)$ who are type 2 but are uninformed about their type; and Group 4 with measure $(1-\alpha)\big(1 - G(\hat{x})\big)$ who are type 1 but are uninformed about their type.

The first question is whether an individual decides to become informed. Depending on this decision, we then have either informed or uninformed opt-out. Informed opt-out occurs, when (i) individuals who invest x and learn that they are type 2 decide to opt out of Low default (when $k < B$); or (ii) individuals who invest x and learn that they are type 1 decide to opt out of High default (when $k < C$). Uninformed opt-out occurs, when (i) the expected value of High is larger, i.e., $\alpha B - (1-\alpha)C > 0$, and uninformed individuals decide to opt-out of Low default (when $k < \alpha B - (1-\alpha)C$); or (ii) the expected value of Low is higher, i.e., $\alpha B - (1-\alpha)C < 0$, and uninformed individuals decide to opt-out of High default (when $k < (1-\alpha)C - \alpha B$). In our analysis, we assume, without loss of generality, that $\alpha B - (1-\alpha)C \geq 0$.[24]

B. Low Default

We study the two decisions faced by an individual: whether to become informed and whether to opt out. Consider an individual with (k, x). We map the information acquisition and opt-out decisions for different levels of opt-out costs, k, but then focus on the low opt-out cost scenario.

High opt-out costs. When $k \geq B$, the individual will not become informed, regardless of x. In this range, the mechanical opt-out costs prevent even informed opt-out, and thus there is no point in becoming informed. (And if there is no informed opt-out, there will be no uninformed opt-out: $k \geq B$ implies $k > \alpha B - (1-\alpha)C$.) To summarize: When $k \geq B$, the opt out rate is zero. In terms of welfare, when $k \geq B$, $W = 0$.

Intermediate opt-out costs. When $k \in (\alpha B - (1-\alpha)C, B)$, the mechanical opt-out costs are low enough to permit informed opt-out, but not uninformed opt-out. Specifically, an informed individual who learns that she is type 2 will opt out from Low default. If the individual becomes informed, her expected payoff is: $\alpha(B-k) + (1-\alpha)\cdot 0 - x = \alpha(B-k) - x$. If the individual remains uninformed, she will stick with Low default and earn a payoff of zero. Therefore, individuals will become informed iff $\alpha(B-k) - x > 0$, or $x < \alpha(B-k)$. To summarize, when $k \in (\alpha B - (1-\alpha)C, B)$, a share $G\big(\alpha(B-k)\big)$ of individuals will become informed and opt out with probability α; and a share $1 - G\big(\alpha(B-k)\big)$ will remain uninformed and stick with the Low default. For a given k, the opt out rate is: $\alpha G\big(\alpha(B-k)\big)$. In terms of welfare, for any $k \in (\alpha B - (1-\alpha)C, B)$, $W = \int_0^{\alpha(B-k)} (\alpha(B-k) - x) g(x) dx$.

Low opt-out costs. When $k < \alpha B - (1-\alpha)C$, the mechanical opt-out costs are low enough to permit both informed and uninformed opt-out. As with intermediate opt-out costs, an informed individual who learns that she is type 2 will opt out from Low default

[24] The case where $\alpha B - (1-\alpha)C < 0$, is captured by normalizing the High payoffs to be zero and redefining $\tilde{C} = B$ as the cost born by a share α under Low, and $\tilde{B} = C$ as the benefit enjoyed by a share 1-α under Low. The expected payoff in Low would then be: $(1-\alpha)\tilde{B} - \alpha\tilde{C} \geq 0$. We can further redefine: $\tilde{\alpha} = 1 - \alpha$, and get $\tilde{\alpha}\tilde{B} - (1-\tilde{\alpha})\tilde{C} \geq 0$.

and earn an expected payoff of $\alpha(B-k)-x$. If the individual remains uninformed, then she will opt-out from Low default and earn an expected payoff of $\alpha B-(1-\alpha)C-k$. Therefore, individuals will become informed iff $\alpha(B-k)-x>\alpha B-(1-\alpha)C-k$, or $x<(1-\alpha)(C+k)$. To summarize, when $k<\alpha B-(1-\alpha)C$, a share $G\big((1-\alpha)(C+k)\big)$ will become informed and opt out with probability α; and a share $1-G\big((1-\alpha)(C+k)\big)$ will remain uninformed and opt out. For a given k, the opt out rate is: $\alpha G\big((1-\alpha)(C+k)\big)+1-G\big((1-\alpha)(C+k)\big)$. In terms of welfare, for any $k\leq\alpha B-(1-\alpha)C$,

$$W=\int_0^{(1-\alpha)(C+k)}(\alpha(B-k)-x)g(x)dx+\Big(1-G\big((1-\alpha)(C+k)\big)\Big)(\alpha B-(1-\alpha)C-k)$$

Special Case: Perfect Information. We note that the perfect information case, where all individuals know their type without any need to invest in information acquisition, is a special case that is embedded in the preceding analysis. Specifically, with perfect information, we have $G(0)=1$. When opt-out costs are either intermediate or low, this implies an opt-out rate of α, and a welfare level of $W=\alpha(B-k)$. When opt-out costs are high, the opt-out rate is zero and welfare is zero, even with perfect information.

These results are summarized in the following lemma.

Lemma 1 (Low Default):

(a) For any $k\geq B$: The opt-out rate is zero and welfare is zero, with both perfect and imperfect information.

(b) For any $k\in(\alpha B-(1-\alpha)C,B)$: With perfect information, the opt-out rate is α and welfare is $W=\alpha(B-k)$; with imperfect information the opt-out rate is $\alpha G\big(\alpha(B-k)\big)<\alpha$ and welfare is $W=\int_0^{\alpha(B-k)}(\alpha(B-k)-x)g(x)dx$.

(c) For any $k\leq\alpha B-(1-\alpha)C$: With perfect information, the opt-out rate is α and welfare is $W=\alpha(B-k)$; with imperfect information the opt-out rate is $\alpha G\big((1-\alpha)(C+k)\big)+1-G\big((1-\alpha)(C+k)\big)>\alpha$ and welfare is

$$W=\int_0^{(1-\alpha)(C+k)}(\alpha(B-k)-x)g(x)dx+\Big(1-G\big((1-\alpha)(C+k)\big)\Big)(\alpha B-(1-\alpha)C-k)$$

The role of information costs. Ours is an information-costs theory. We thus focus on the role that information costs play in the analysis, specifically how the magnitude of information costs affects opt-out rates and welfare. We begin with the intermediate and high opt-out costs scenarios. In these scenarios (where $F(\alpha B-(1-\alpha)C)=0$), any opt-out will be informed. Therefore, a reduction in information costs, specifically when $G(x)$

is higher for all x (notion of first-order stochastic dominance), increases the opt-out rate and also increases welfare. This scenario captures the intuitive belief that high information costs create sticky defaults. And if we think of unsophisticated individuals as having high information costs, then we get the standard result that unsophisticated individuals always stick with the default, whereas sophisticated individuals opt-out when the default is not optimal for them.

The more interesting scenario is the low opt-out costs scenario. In this scenario (where $F(\alpha B - (1-\alpha)C) = 1$), we get both informed and uninformed opt-out. Specifically, individuals with high information costs will remain uninformed and opt out; and individuals with low information costs will opt-out only if they learn that the default is not optimal for them. A reduction in information costs, specifically when $G(x)$ is higher for all x (notion of first-order stochastic dominance), reduces the opt-out rate and increases welfare. We get the counterintuitive result that lower information costs increase stickiness. When information costs are high, few individuals become informed and, because opt-out costs are low (and $\alpha B - (1-\alpha)C > 0$) all the uninformed individuals opt out. When information costs are low, many individuals become informed and only a share α of them opt out.

Formally, for any $k \leq \alpha B - (1-\alpha)C$, the opt-out rate is: $\alpha G\big((1-\alpha)(C+k)\big) + 1 - G\big((1-\alpha)(C+k)\big) = 1 - (1-\alpha)G\big((1-\alpha)(C+k)\big)$. With lower information costs (i.e., when $G(x)$ is higher for all x [notion of first-order stochastic dominance]), the opt-out rate is lower and thus the default is more sticky. And, of course, lower information costs increase social welfare. Therefore, sticky defaults are associated with higher welfare. These results are summarized in the following proposition.

Proposition 1 (The Role of Information Costs):

(a) When $F(\alpha B - (1-\alpha)C) = 0$, lower information costs reduce stickiness and increase welfare.

(b) When $F(\alpha B - (1-\alpha)C) = 1$, lower information costs increase both stickiness and welfare.

C. High Default

With High default, there are only two possible ranges of opt-out costs. When $k \geq C$, there will be no informed opt-out, and thus no one will acquire information. With such high opt-out costs, the opt out rate is zero, and $W = \alpha B - (1-\alpha)C$.

When $k < C$, informed opt-out is possible. Specifically, an informed individual who learns that she is type 1 will opt out from High default. If the individual becomes informed, her expected payoff is: $\alpha B + (1-\alpha)\cdot(-k) - x = \alpha B - (1-\alpha)\cdot k - x$. With High default, there will be no uninformed opt-out, regardless of k (since $\alpha B - (1-\alpha)C > 0$). An individual who remains uninformed will stick with High default and earn a payoff of $\alpha B - (1-\alpha)C$. Therefore, individuals will become informed iff $\alpha B - (1-\alpha)\cdot k - x > \alpha B - (1-\alpha)C$, or $x < (1-\alpha)\cdot(C-k)$. To summarize, when $k < C$, a share $G\big((1-\alpha)\cdot$

$(C-k)\big)$ of individuals will become informed and opt out with probability $1-\alpha$; and a share $1-G\big((1-\alpha)\cdot(C-k)\big)$ will remain uninformed and stick with the High default. For a given k, the opt out rate is: $(1-\alpha)\cdot G\big((1-\alpha)\cdot(C-k)\big)$. In terms of welfare, for any $k<C$, $W=\int_0^{(1-\alpha)\cdot(C-k)}(\alpha B-(1-\alpha)\cdot k-x)g(x)dx+(1-G((1-\alpha)\cdot(C-k)))(\alpha B-(1-\alpha)C)$.

Special Case: Perfect Information. With perfect information, i.e., with $G(0)=1$, when $k<C$, the opt-out rate is $1-\alpha$, and the welfare level is: $W=\alpha B-(1-\alpha)\cdot k$. When $k\geq C$, the opt-out rate is zero and welfare equals $W=\alpha B-(1-\alpha)C$.

These results are summarized in the following lemma.

Lemma 2 (High Default):

(a) For any $k\geq C$: The opt-out rate is zero and welfare is $W=\alpha B-(1-\alpha)C$, with both perfect and imperfect information.

(b) For any $k<C$: With perfect information, the opt-out rate is $1-\alpha$ and welfare is $W=\alpha B-(1-\alpha)\cdot k$; with imperfect information the opt-out rate is $(1-\alpha)\cdot G\big((1-\alpha)\cdot(C-k)\big)$ and welfare is

$$W=\int_0^{(1-\alpha)\cdot(C-k)}(\alpha B-(1-\alpha)\cdot k-x)g(x)dx+(1-G((1-\alpha)\cdot(C-k)))(\alpha B-(1-\alpha)C)$$

The role of information costs. With High default, there is no possibility of uninformed opt out; only informed opt out is possible. Therefore, we obtain the standard result that lower information costs reduce stickiness and increase welfare.

D. Comparison: Low Default v. High Default

We can now compare the two defaults. We focus on the low opt-out costs scenario, to allow for both informed and uninformed opt-out. Specifically, we assume that $k\leq min(\alpha B-(1-\alpha)C,C)$. First, consider incentives for information acquisition. With Low default, information will be acquired when $x<(1-\alpha)(C+k)$. With High default, information will be acquired when $x<(1-\alpha)\cdot(C-k)$. We can state the following result.

Lemma 3 (Low Default v. High Default: Information Acquisition): Low default induces more information acquisition.

We note, however, that when information is costly to acquire, more information is not necessarily better.

We next compare the stickiness of the two defaults. With Low default, the opt-out rate is: $1-(1-\alpha)G\big((1-\alpha)(C+k)\big)$. With High default, the opt-out rate is: $(1-\alpha)\cdot$

$G\big((1-\alpha)\cdot(C-k)\big)$. We see that either rule can be stickier. But more can be said. Let $\Delta(k) = 1-(1-\alpha)\big[G\big((1-\alpha)(C+k)\big)+G\big((1-\alpha)\cdot(C-k)\big)\big]$ denote the difference between the two opt-out rates. When information costs are lower, $\Delta(k)$ is lower. When information costs are very low, i.e., when $G\big((1-\alpha)(C+k)\big)+G\big((1-\alpha)\cdot(C-k)\big) = 2$, the opt-out rate is higher with Low default if $\alpha > \frac{1}{2}$; and higher with High default if $\alpha < \frac{1}{2}$. When information costs are very high, i.e., when $G\big((1-\alpha)(C+k)\big)+G\big((1-\alpha)\cdot(C-k)\big) = 0$, the opt our rate is higher with Low default. These and other results are summarized in the following lemma.

Lemma 4 (Low Default v. High Default: Opt-out Rates):

(a) When information costs are lower, $\Delta(k)$ is lower. When information costs are high, the opt-out rate is higher with Low default. When information costs are low, the opt-out rate is higher with Low default if $\alpha > \frac{1}{2}$, and higher with High default if $\alpha < \frac{1}{2}$.

(b) When the share of type 1 individuals is higher, i.e., when $1-\alpha$ is larger, $\Delta(k)$ is lower. When $1-\alpha$ is small, the opt-out rate is higher with Low default. When $1-\alpha$ is large, the opt-out rate is higher with Low default if information costs are high, and higher with High default if information costs are low.

(c) When the cost that High imposes on type 1 individuals, C, is larger, $\Delta(k)$ is lower.

Finally, we turn to welfare levels. With Low default, welfare is:

$$W = \int_0^{(1-\alpha)(C+k)} (\alpha(B-k)-x)g(x)dx + \Big(1-G\big((1-\alpha)(C+k)\big)\Big)(\alpha B-(1-\alpha)C-k)$$

With High default, welfare is:

$$W = \int_0^{(1-\alpha)\cdot(C-k)} (\alpha B-(1-\alpha)\cdot k-x)g(x)dx + (1-G((1-\alpha)\cdot(C-k)))(\alpha B-(1-\alpha)C)$$

When information costs are high, the welfare comparison is determined by the difference: $[\alpha B-(1-\alpha)C-k]-[\alpha B-(1-\alpha)C] = -k$. Namely, welfare is higher with High default. When information costs are high, individuals do not acquire information. With High default, the uninformed individuals stick with the default (since $\alpha B-(1-\alpha)C > 0$). With Low default, the uninformed individuals engage in costly (uninformed opt out). Therefore, High default is more efficient. Stickiness – which, here, correlates with less need for costly uninformed opt-out – goes hand-in-hand with welfare outcomes. High default is both sticky and efficient.

When information costs are low, the welfare comparison is determined by the difference: $[\alpha(B-k)-x]-[\alpha B-(1-\alpha)\cdot k-x]=(1-2\alpha)\cdot k$. Therefore, welfare is higher with Low default when $\alpha<\frac{1}{2}$, and welfare is higher with High default when $\alpha>\frac{1}{2}$. When information costs are sufficiently low to ensure informed opt out, the majoritarian principle determines the optimal default. The default that requires the least opt-out is more efficient. For this reason, stickiness – which, here, correlates with less need for costly informed opt-out – goes hand in hand with welfare outcomes: When $\alpha<\frac{1}{2}$, Low default is stickier and generates more welfare; and when $\alpha>\frac{1}{2}$, High default is stickier and generates more welfare.

When information cost are intermediate, such that information is acquired with Low default, but not with High default ($x\in\big((1-\alpha)\cdot(C-k),(1-\alpha)\cdot(C+k)\big)$, the welfare comparison is determined by the difference: $[\alpha(B-k)-x]-[\alpha B-(1-\alpha)C]=(1-\alpha)C-\alpha k-x$. Therefore, when $x<(1-\alpha)C-\alpha k$, Low default is more efficient; and when $x>(1-\alpha)C-\alpha k$, High default is more efficient. At the lower end of the intermediate information cost range, the benefit from information acquisition (and informed opt out) exceeds its cost and Low default is better. At the high end of the range, the cost of information outweighs its benefit and High default is better. Here, the opt-out rate with Low default is α, and with High default it is zero. The stickier default is more efficient, when information costs are higher; and the less sticky default is more efficient when information costs are lower.

These results are summarized in the following proposition.

Proposition 2 (Low Default v. High Default: Welfare Levels):

(a) With high information costs, High default is more efficient, and stickier.

(b) With low information costs, when $\alpha<\frac{1}{2}$, Low default is more efficient, and stickier; and when $\alpha>\frac{1}{2}$, High default is more efficient, and stickier.

(c) With intermediate information costs, at the low end of this range Low default is more efficient, and less sticky; at the high end of this range High default is more efficient, and stickier.

The results in parts (a) and (b) of Proposition 2 are not surprising. Part (a) states that, when parties remain uninformed, policymakers should follow the Expected-Value-Maximization (EVM) principle – to avoid costly uninformed opt out. And Part (b) instructs the policymaker to follow the majoritarian principle and thus avoid costly informed opt out, when parties are informed. In both Parts (a) and (b), stickiness goes hand-in-hand with efficiency, since stickiness implies fewer costly opt outs – informed or uninformed. Part (c) focuses on the differences in the incentives to acquire information under the two default rules. When information acquisition is costly (at the high end of the intermediate information costs range), High default, which keeps individuals uninformed and avoids (uninformed) opt out, is more efficient, and stickier. The link between stickiness and efficiency is maintained. But when information is less costly (at the low end of the

intermediate information costs range), Low default is the better rule – by inducing information acquisition, and informed opt out, Low default generates better matching between individuals and outcomes. Here, the slippery rule is more efficient.

E. Inaccurate Beliefs

The preceding analysis assumed that individuals, while (possibly) uninformed about their type, accurately assess the relative expected payoffs of the two outcomes, Low and High. Specifically, since the Low payoff was normalized to zero, the assumption was that the parties know the expected value of High: $\pi = \alpha B - (1-\alpha)C$. We now introduce the possibility of inaccurate beliefs and allow parties to hold beliefs $\hat{\pi} \neq \pi$ about the expected value of High. And since $\pi > 0$ (High is better on average), we will focus on inaccurate beliefs – about α, B or C – that result in $\hat{\pi} < 0$ (Low is perceived to be better on average). As we will see, the object of the inaccurate beliefs – α, B or C – affects the analysis, so we need to separately denote the perceived values of the three parameters: $\hat{\alpha}$, $\hat{B}$ and $\hat{C}$. (Inaccurate beliefs about k and x are also possible.)

We focus on the low opt-out costs scenario ($k < |\hat{\pi}|$), where the mechanical opt-out costs are low enough to permit both informed and uninformed opt-out. For informed parties, the analysis does not change. An informed individual who learns that she is type 2 will opt out from Low default and stick with High default. The inaccurate beliefs affect the decisions and behavior of uninformed parties. These parties who opted out of Low default and stuck with High default in the accurate beliefs analysis, now stick with Low default and opt out of High default.

Low default. An informed individual who learns that she is type 2 will opt out from Low default. The expected payoff of an individual who becomes informed is: $\alpha(B-k)-x$, and the perceived payoff is: $\hat{\alpha}(\hat{B}-k)-x$. An uninformed individual sticks with Low and earns a payoff of zero. Therefore, individuals will become informed iff $\hat{\alpha}(\hat{B}-k)-x > 0$, or $x < \hat{\alpha}(\hat{B}-k)$. A share $G\left(\hat{\alpha}(\hat{B}-k)\right)$ will become informed and opt out with probability α; and a share $1-G\left(\hat{\alpha}(\hat{B}-k)\right)$ will remain uninformed and stick with the Low default (opt-out rate of zero). These results are summarized in the following lemma.

Lemma 1a (Low Default; Low Opt-out Costs; Inaccurate Beliefs): When $\hat{\pi} < 0$, for any $k \leq |\hat{\pi}|$: With perfect information, the opt-out rate is α and welfare is $W = \alpha(B-k)$; with imperfect information the opt-out rate is $\alpha G\left(\hat{\alpha}(\hat{B}-k)\right)$ and welfare is

$$W = \int_0^{\hat{\alpha}(\hat{B}-k)} (\alpha(B-k)-x)g(x)dx$$

Importantly, and counterintuitively, with inaccurate beliefs a policy aimed at reducing information costs might reduce efficiency. Specifically, when information costs are reduced from $\hat{\alpha}(\hat{B}-k)+\varepsilon$ to $\hat{\alpha}(\hat{B}-k)-\varepsilon$, welfare changes from zero to $\alpha(B-k)-$

x. When $\alpha(B-k) < \hat{\alpha}(\hat{B}-k)$ and $x \in \left(\alpha(B-k), \hat{\alpha}(\hat{B}-k)\right)$, the lower information costs reduce welfare. (The identified, perverse effect of lowering information costs requires $\hat{\alpha} > \alpha$ or $\hat{B} > B$. Our analysis focuses on inaccurate beliefs that imply $\hat{\pi} = \hat{\alpha}\hat{B} - (1-\hat{\alpha})\hat{C} < 0$, instead of the accurate $\pi = \alpha B - (1-\alpha)C > 0$. Within this constraint, it is still possible to get the perverse effect, if $\hat{C} > C$.) We summarize this result in the following corollary.

Corollary 1: With inaccurate beliefs, lower information costs might decrease welfare.

High default. An informed individual who learns that she is type 1 will opt out from High default. The expected payoff of an individual who becomes informed is: $\alpha B - (1-\alpha) \cdot k - x$, and the perceived payoff is: $\hat{\alpha}\hat{B} - (1-\hat{\alpha}) \cdot k - x$. An individual who remains uninformed will opt out to Low and earn a payoff of $-k$. Therefore, individuals will become informed iff $\hat{\alpha}\hat{B} - (1-\hat{\alpha}) \cdot k - x > -k$, or $x < \hat{\alpha}(\hat{B}+k)$. A share $G\left(\hat{\alpha}(\hat{B}+k)\right)$ of individuals will become informed and opt out with probability $1-\alpha$; and a share $1 - G\left(\hat{\alpha}(\hat{B}+k)\right)$ will remain uninformed and opt out with probability 100%. These results are summarized in the following lemma.

Lemma 2a (High Default; Low Opt-out Costs; Inaccurate Beliefs): When $\hat{\pi} < 0$, for any $k \leq |\hat{\pi}|$: With perfect information, the opt-out rate is $1-\alpha$ and welfare is $W = \alpha B - (1-\alpha) \cdot k$; with imperfect information the opt-out rate is $(1-\alpha) \cdot G\left(\hat{\alpha}(\hat{B}+k)\right) + 1 - G\left(\hat{\alpha}(\hat{B}+k)\right) = 1 - \alpha G\left(\hat{\alpha}(\hat{B}+k)\right)$ and welfare is

$$W = \int_0^{\hat{\alpha}(\hat{B}+k)} (\alpha B - (1-\alpha) \cdot k - x) g(x) dx - \left(1 - G\left(\hat{\alpha}(\hat{B}+k)\right)\right) \cdot k$$

As with Low default, here too lower information costs might reduce efficiency. Specifically, when information costs are reduced from $\hat{\alpha}(\hat{B}+k) + \varepsilon$ to $\hat{\alpha}(\hat{B}+k) - \varepsilon$, welfare changes from $-k$ to $\alpha B - (1-\alpha) \cdot k - x$. When $\alpha(B+k) < \hat{\alpha}(\hat{B}+k)$ and $x \in \left(\alpha(B+k), \hat{\alpha}(\hat{B}+k)\right)$, the lower information costs reduce welfare.

Comparison. Inaccurate beliefs alter the comparison between the two defaults. When information costs are high, welfare is greater with Low default, since it is Low default that now avoids the cost of uninformed opt out (albeit inefficient uninformed opt out). When information costs are low, information is acquired and thus beliefs, accurate or inaccurate, about average payoffs do not matter. (The assumption is that when an individual invests in information acquisition, she learns her type and obtains accurate information about all relevant parameters.)

When information costs are intermediate ($x \in \left(\hat{\alpha}(\hat{B}-k), \hat{\alpha}(\hat{B}+k)\right)$), we find that now High default generates stronger incentives to acquire information. The welfare comparison is determined by the difference: $[0] - [\alpha B - (1-\alpha)\cdot k - x] = -\alpha B + (1-\alpha)\cdot k + x$. Therefore, when $x < \alpha B - (1-\alpha)\cdot k$, High default is more efficient; and when $x > \alpha B - (1-\alpha)\cdot k$, Low default is more efficient. At the lower end of the intermediate information cost range, the benefit from information acquisition (and informed opt out) exceeds its cost and High default is better. At the high end of the range, the cost of information outweighs its benefit and Low default is better. Here, the opt-out rate with High default is $1-\alpha$, and with Low default it is zero. The stickier default is more efficient, when information costs are higher; and the less sticky default is more efficient when information costs are lower.

These results are summarized in the following proposition.

Proposition 3 (Low Default v. High Default: Welfare Levels; Inaccurate Beliefs):

(a) With high information costs, Low default is more efficient, and stickier.

(b) With low information costs, when $\alpha < \frac{1}{2}$, Low default is more efficient, and stickier; and when $\alpha > \frac{1}{2}$, High default is more efficient, and stickier.

(c) With intermediate information costs, at the low end of this range High default is more efficient, and less sticky; at the high end of this range Low default is more efficient, and stickier.

When information costs are high, inaccurate beliefs flip the policy prescription – from High default to Low default. The driving force is, again, the uninformed opt out. Uninformed parties will inevitably end up with the inefficient outcome, Low. The best that the policymaker can do is avoid the cost of inefficient, uninformed opt outs. Inaccurate beliefs also flip the ordering of incentives to acquire information, such that, when information costs are intermediate, High default induces more information acquisition. The policymaker should thus prefer High default when information costs are at the low end of this range, and Low default when information costs are higher.

IV. Policy Implications

Having laid out our information-costs theory of default rules, we now explore the policy implications of the theoretical analysis. Section A considers policy interventions that increase (mechanical) opt-out costs. Section B considers policy interventions that reduce information costs. And Section C considers the distributional effects of different default rules, in light of the information-costs theory.

A. Increasing Opt-out Costs

As explained in the Introduction, there is a belief among both academics and policymakers that sticky defaults can be desirable. This belief has led to policy proposals that aim to increase (mechanical) opt-out costs, especially in contexts where firms try to lure consumers to easily, mindlessly, opt out. The hope fueling such proposals is that higher transaction costs would slow down opt outs and leave consumers with the default rules, which are usually more protective ("High").

When people have accurate beliefs, high mechanical opt out costs reduce welfare, and thus making defaults stickier makes no sense, as it hinders efficient informed and uninformed opt out. This, of course, is not surprising. And, indeed, the argument for sticky defaults rests on some account of inaccurate beliefs whereby people opt out without fully understanding the welfare implications of their decisions. We have seen that, when opt-out costs are low, inaccurate beliefs could lead people to agree to inefficient uninformed opt-out. In these settings, deliberately increasing opt out costs increases welfare. If people systematically underestimate the value of the default rule and mistakenly opt out to inferior alternatives, rules that make the opt out harder to complete eliminate some of the mistakes. Only people who strongly prefer the alternative outcome—and thus presumably are less likely to be mistaken—would be willing to incur the higher mechanical costs and opt out.

Our analysis thus provides a justification for policies that increase opt-out costs. The most typical tools employed to make defaults sticky are lengthy disclosures, educational pre-requisites, segregated agreements, clause-by-clause signatures, and periodically renewed agreements.[25] It is sometimes questionable whether such hurdles succeed in increasing opt out costs in a manner sufficient to render the defaults stickier. Pragmatically, it is hard to imagine, in the presence of a counterparty strongly motivated to induce smooth opt-out, that mechanical opt-out costs could be truly meaningful. Theoretically, however, the underlying goal of deliberately increasing mechanical costs is a correct response to the problem of welfare-reducing opt-out in the presence of inaccurate beliefs.

B. Increasing Information Costs

Our analysis is based on the premise that, more than mechanical opt-out costs, information costs are often the major impediment to efficient opt-out. It might therefore be tempting to think that reducing information costs is across the board desirable, and that any effort—either by lawmakers or by counterparties—to deliberately increase information costs is undesirable. Our analysis suggests, surprisingly, that when people have inaccurate beliefs about which default is better for them, lower information costs might reduce welfare. Specifically, when inaccurate beliefs result in an overestimation of the benefit from information, individuals will tend to invest excessively in information acquisition. High information costs limit the effects of this inefficient tendency; lowering the cost of information acquisition might exacerbate the problem.

[25] [RA: cite]

C. Distributive Considerations

The possibility of uninformed opt-out forces us to rethink the role of distributive considerations in designing default rules. In particular, when most people are uninformed and these uninformed parties stick with the default, the policymaker can shift resources to a preferred group by choosing a default that benefits that group, even if the default is less efficient overall. If, instead, uninformed members of the preferred group recognize the presence of an alternative arrangement that maximizes the expected value across all groups, and they do not recognize that the default is nevertheless better for them, then they would opt-out and the policymaker's distributive objective would be frustrated.

In our numerical example (Part II), a policymaker who prefers the Type 1 group may choose Low default, even though High default provides overall greater expected value. This policy, however, would work, only if the uninformed Type 1 individuals stick with the default. Otherwise, the preferred group would just opt out, to High; and the Low default would just impose extra opt out costs on this group. The policymaker who intended to help Type 1 would only end up hurting them.

While uninformed opt out might frustrate some distributive policies, the information-costs theory suggests other ways to achieve distributive goals. In particular, distributive concerns can influence default choice, when different people have different information costs. Choosing a default that induces only some people to acquire information may then be justified. We saw above that when information costs are large, the optimal default is the one that maximizes the expected value for uninformed individuals, and when information costs are small the optimal default is majoritarian. But what if poor people have greater information costs than the affluent? In that case, choosing a default based on the assumption that information costs are high could be justified.

V. Extensions

In this Part, we consider several extensions to the basic analysis. In Section A, we allow for information costs that depend on the chosen default. In Section B, we revisit the assumption that uninformed individuals know average values and thus can compare the default and non-default options. In Section C, we consider the (surprisingly common) case of negative opt-out costs. In Section D, we allow policymakers to choose from a continuous set of default rules (rather than a binary set). And in Section E we discuss the implications of our information-costs theory for forced choice regimes.

A. Rule-Dependent Information Costs

We thus far assumed that the distribution of information costs is independent of the chosen default. While this assumption strikes us as a plausible benchmark, we note that, in some applications, information costs may depend on the default rule. Consider consumer contracts. The cost, to the consumer, of becoming informed on her own may be quite high. On the other hand, if the seller provides information, then it could become much easier, and cheaper, for the consumer to become informed. If the seller's incentive to provide information is stronger with one default, and weaker with another, then we get rule-dependent information costs.[26]

The policy implications of this observation are not obvious. It may seem that a rule that induces lower information costs is the better rule. But there is a real risk that the low-cost information will be biased, especially when it is provided by a seller who is trying to induce opt-out from a default that is less favorable to that seller. Also, as noted above, lower information costs might actually reduce welfare.

B. Do People Really Know Average Values?

Our analysis assumes that people know, or think they know, the average payoff under each default. This seemingly strong assumption should not be taken literally. All we need to assume is that, before any information is acquired, people form some estimate (beliefs) about the net benefit – the average value – of the default vs. non-default options. This estimate need not be accurate and, indeed, our analysis allows for inaccurate beliefs.

A harder question concerns the assumption that people know average values, but do not know their own "type." Intuitively, people may know more about themselves than about the abstract average value across a population of heterogeneous individuals. But this, too, is consistent with our analysis, if we reinterpret "type" as the unknown idiosyncratic element over which people, even within the same subgroup, vary.

To illustrate, consider the warranty example, where High and Low represent the duration and scope of the warranty. Some people—call them the "careless"—generally benefit from a High warranty, even when it is reflected in a higher product price. Other people—the "careful"—benefit less from costly warranties and generally prefer Low. Within each group, however, there is variation. Some of the careful may be highly risk-averse, enough to make the High warranty a better fit; and some of the careless may have risk preferences that make a Low warranty more suitable (or may simply be unable to afford the High warranty).

---

[26] Compare: Ian Ayres and Robert Gertner, "Filling Gaps in Incomplete Contracts: An Economic Theory of Default Rules," 99 YALE L.J. 87-130 (1989); Lucian Arye Bebchuk and Steven Shavell, "Information and the Scope of Liability for Breach of Contract: The Rule of "Hadley v. Baxendale"," 7 J.L. ECON. & ORG. 284-312 (1991).

In this setting, people form beliefs as to whether they are careless or careful (beliefs that could be accurate or inaccurate). Even if they have accurate beliefs about the subgroup to which they belong, they need additional information in every specific context to know their ad-hoc type: the actual value they assign to the warranty in that context. Our analysis thus applies: Once we zoom in on the subgroup, we can assume that without information the members of the subgroup form beliefs about the average value within the subgroup, but they need information about the actual valuation (i.e., about their type) to make the optimal decision.

The model we sketched thus fits such information environments, as long as lawmakers could set separate defaults for each subgroup – one default for the careless and another for the careful. Sometimes this is possible: the careless could be groups of younger and less educated people, whereas the careful could be older, more affluent, or people that have other forms of insurance (that could substitute for a broad warranty). Different default rules could be triggered by different observable characteristics. But other times, it is not possible to segregate people and tailor default rules. When lawmakers are constrained to set a common default rule across subgroups, this would add another layer of complexity to the analysis. In essence, each person would have to estimate the (net) expected benefit that the common default generates for his or her group. We would then need to separately examine information-acquisition and opt-out decisions in each group.

C. Negative Opt-out Costs

Assume that the policymaker sets a default rule that hurts the seller's bottom line, e.g., a no-warranty default. The seller may then prompt the consumer to opt-out of the default – "click here if you want to add a warranty." And, if the consumer declines to "click here," the seller might prompt her again, and again, and again. In this scenario, opting out of the default is relatively easy – the consumer just needs to "click here" once, whereas sticking to the default becomes quite costly – the consumer is repeatedly interrupted by the seller's relentless warranty offers and must, time and again, close the pop-up window. This scenario characterizes many consumer transactions, where sellers seek to push various add-on features, like store extended warranties, other warranties/insurance, service plans, in-app purchases and recurring purchases.

What is special, and interesting, in such cases is that, contrary to our basic assumption, opting out is easier than sticking to the default. If the benchmark cost is the cost of sticking to the default, then we have negative opt-out costs. In essence, sellers are unilaterally changing the default – to "yes warranty" – and forcing consumers to incur real costs, if they want the policymaker's default. A similar phenomenon occurs, when sellers unilaterally change the default using unread, fine-print contracts.

Our information-costs theory sheds new light on the normative assessment of such unilateral default changes. Information costs might render the "seller's default" stickier, which provides further cause for concern. But, with uninformed opt-out, the seller's attempt to impose a term with negative expected value (to consumers) may be unsuccessful;

the seller's default may be quite slippery. Still, the cost of such uninformed opt-out may be reason enough for policymakers to intervene and strengthen the original default rule against unilateral changes by sellers. And, of course, if sellers are able to impose opt-out costs that are sufficiently high to prevent uninformed opt-out (from the seller's default), then there is further justification for policy intervention. Perhaps even more troubling: Uninformed consumers might not realize that the add-on is welfare-reducing (on average) and stick to the seller's default.

D. From Binary to Continuous Default Choice

For simplicity, our analysis assumed only two outcomes – Low and High, and thus only two defaults – Low default and High default. In some cases, this binary-choice assumption is realistic; think organ donations. In other cases, the choice is actually continuous. Indeed, in our leading retirement savings example, the choice of contribution level is a continuous choice. Our framework can easily be adjusted to accommodate a continuous outcome space. The hardest question is how to think about the decision process of an employee who faces a 3% contribution rate default. In the basic, binary model, the employee compared expected payoffs in Low and High outcomes. In a continuous model, the employee would need to calculate an expected payoff function, where the expected payoff is a function of the continuous contribution rate.

E. Forced Choice

In response to a critique that selecting a default can be paternalistic (given the stickiness of defaults), some have proposed an alternative regime – forced choice. In a forced choice regime, the individual is required to select among several options, but no option is pre-selected as a default. For example, an employee cannot start to work, before she selects a retirement plan; or an individual cannot get a driver's license, before she chooses whether or not to be an organ donor.

Our framework does not support forced choice regimes. As long as people are uninformed, they will select the highest expected value option, and the only effect of a forced-choice regime is to require more mechanical costs—those that could have been saved if the highest expected value option was preselected as the default. It is hard to imagine that the need to make an affirmative selection would reduce in any manner the cost of information. It is possible, however, that the argument for forced choice is based on a behavioral model of decision making – perhaps when forced to choose, people would feel more compelled to acquire information and make informed choices.

VI. Applications

As noted in the Introduction, default rules are increasingly being used across diverse policy domains. In this Part, we consider, in some detail, several important applications, and highlight the different ways in which our information-costs theory can inform policymakers. Section A considers privacy and data security. Section B considers overdraft protection. [More applications TBA]

A. Privacy

The basic legal default in many jurisdictions does not allow firms to collect people's personal information (in many contexts). People must consent to any opt-out from that default. How to design and police such consent-based information collection has been the subject of much debate and ongoing legal reform.

For long, the legal default of no-information-collection has been routinely subject to uninformed opt out. Firms interested in collecting personal information could easily guide consumers to opt out.[27] The mechanics of such opt-outs are designed by firms to be so simple and cheap that opt out is achieved smoothly. This was uninformed opt out at the extreme – people were not even aware that they were opting out. Put differently, the beliefs people had about the implications of data collection did not allow them to distinguish between the payoffs in the default and non-default options.

Then things began to change. A series of massive data breaches, in which the personal information of tens or hundreds of millions of individuals was accessed for socially harmful purposes, shocked our collective consciousness. What was the effect of these massive data breaches? First, people's beliefs changed to some degree, and many now believe that the no-information-collection default is better than the opt-out alternative. This is still an uninformed belief, because people need more information to know how they, individually, are impacted.[28] Second, lawmakers have increased opt-out costs, making it more difficult for firms to induce mindless consent to information collection.[29] Third, lawmakers have been trying to provide for more informed decisionmaking by consumers. Specifically, they have enacted mandates to make it easier for people to read and learn the terms of firms' privacy policies.[30] Using our terminology, lawmakers want to reduce information costs.

Some consumers examine what is actually going on and make informed decisions to opt out (or to stick). Other consumers, perhaps prompted by the public outcry or by misperceptions, stick with the high-privacy default while uninformed. As the information

[27] [RA: Collect references on how easy it is to opt out: conspicuous browsewrap in *Specht*, etc.]
[28] [RA: Document attempts by firms, and also policymakers(?), to influence people's uninformed beliefs. [Ask: are these attempts to help people identify their own true preferences, or rather to shape these preferences?]
[29] [RA: Cite GDPR requirements. Look at recent $50M fine imposed on Google by French regulators.]
[30] [RA: Cite]

costs become lower, we have less of the latter and perhaps more of the former, and it is even possible that there will be *more* opt out. It is clear that the goal of the GDPR is to regulate the incidence and ease of such opt outs. And it is interesting, theoretically, that as the GDPR reduces information costs, the unintended result might be *more* opt out and reduced privacy.

B. Overdraft Protection

Until 2010, debit card holders could swipe their card and charge their checking account, even if they did not have sufficient funds in their checking account. Of course, card holders were charged a fee for this "overdraft protection." In 2010, an Amendment to Regulation E went into force – an amendment that established a no-overdraft-protection default. Now card holders, if they want overdraft protection, must affirmatively opt into the service.[31]

The purpose of the new default is to prevent unsophisticated consumers from incurring high overdraft fees. Our analysis questions the ability of the default to achieve its stated goal. Assume, for example, that overdraft protection is beneficial for the average consumer who overdrafts rarely. Uninformed consumers, including those who would incur multiple overdraft fees, might opt out to their detriment. The default would not be as sticky as the policymaker hoped. Indeed, there is evidence of substantial opt out from the new Regulation E default.[32] [33]

A more general lesson can be learned from the overdraft example: It is difficult to change outcomes for consumers without addressing the uninformed opt-out phenomenon. Most attempts by regulators to make a new default sticky focus on the wrong method: making the mechanical costs of opt out higher. These attempts fail because it is fundamentally very easy to opt out, especially when the firm on the other side is motivated to make such opt-outs smooth. So what can lawmakers do? They could try to affect people's uninformed beliefs, so that the many uninformed consumers would prefer the socially targeted outcome. Or they could try to lower people's information costs to make them less devoted to the firm's manipulated choice. But neither of these options is without problems. The lawmaker's attempt to influence uninformed beliefs will be countered by the industry's own marketing campaigns.[34] And, as we have seen, policies that seek to reduce the cost of information to individual consumers might backfire.

[31] [RA: Cite to CFR.]
[32] [RA: Cite evidence from Willis and Sarin papers.]
[33] [Why was the 'no overdraft protection' default less sticky than the 'no overlimit protection" default for credit cards, which was adopted more or less at the same time? One possible answer: Less uninformed opt-out in the credit card context.]
[34] [RA: Cite: overdraft "protection" campaigns]

Conclusion

[TBA]