Rethinking Nudge: An Information-Costs Theory of Default Rules

Oren Bar-Gill

Omri Ben-Shahar

HARVARD

RETHINKING NUDGE: AN INFORMATION-COSTS THEORY OF DEFAULT RULES

Oren Bar-Gill
Omri Ben-Shahar

Forthcoming in *University of Chicago Law Review*

Rethinking Nudge: An Information-Costs Theory of Default Rules

Oren Bar-Gill and Omri Ben-Shahar[*]

April 2020

Forthcoming in *University of Chicago Law Review*

Abstract

Policymakers and scholars – both lawyers and economists – have long been pondering the optimal design of default rules. From the classic works on "mimicking" defaults for contracts and corporations to the modern rush to set "sticky" default rules to promote policies as diverse as organ donations, retirement savings, consumer protection, and data privacy, the optimal design of default rules has featured as a central regulatory challenge. The key element driving the design is opt-out costs—how to minimize them, or alternatively how to raise them to make the default sticky. Much of the literature has focused on "mechanical" opt-out costs—the effort people incur to select a non-default alternative. This focus is too narrow. A more important factor affecting opt-out is information—the knowledge people must acquire to make informed opt-out decisions. But, unlike high mechanical costs, high information costs need not make defaults stickier; they may instead make the defaults "slippery." This counterintuitive claim is due to 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.

Table of Contents

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 of how to design optimal default rules has been the subject of enormous commentary. Traditional law and economics scholarship argued that default rules should mimic the will of the parties, because otherwise they would force people to waste transactions costs in disclaiming the defaults. If the population governed by a default rule is heterogeneous, continues this transaction-costs argument, better enact a "majoritarian" default so as to reduce the aggregate cost of opt-out.[1]

More recently, the behavioral economics literature has adopted a markedly different approach, asking not how to reduce opt-out costs but instead how to optimally exploit their presence. This Nudge-inspired literature suggests that the "stickiness" of default rules—the tendency of people not to override defaults because of high opt-out costs—is a blessed feature that could improve social welfare. Sticky default rules have been hailed as a major policy tool—an effective behaviorally-informed solution to the challenge of helping 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 many contracting failures and other social problems.[3] From mortgage products[4] to

---

[1] *See, e.g.*, ROBERT COOTER AND THOMAS ULEN, LAW AND ECONOMICS (6th ed., 2011); RICHARD A. POSNER, ECONOMIC ANALYSIS OF LAW (7th ed., 2007); STEVEN SHAVELL, FOUNDATION OF ECONOMIC ANALYSIS OF LAW (2004); 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 (1991); Charles Goetz and Robert Scott, *The Mitigation Principle: Toward A General Theory of Contractual Obligation*, 69 VA. L. REV. 967 (1983); Alan Schwartz, *The Default Rule Paradigm and the Limits of Contract Law*, 3 S. CAL. INTERDISC. L. J. 389 (1993).

[2] *See* RICHARD H. THALER & CASS R. SUNSTEIN, NUDGE: IMPROVING DECISIONS ABOUT HEALTH, WEALTH, AND HAPPINESS (2008); Cass R. Sunstein, *Deciding by Default*, 162 U. PA. L. REV. 1, 5, 9 (2013) ("[D]efault rule . . . count as prime "nudges," understood as interventions that maintain freedom of choice, that do not impose mandates or bans, but that nonetheless incline people's choices in a particular direction . . . . Default rules may produce significantly better outcomes.").

[3] *See* Michael S. Barr, Sendhil Mullainathan, and Eldar Shafir*, A One-Size-Fits-All Solution*, N.Y. TIMES (December 26, 2007), https://www.nytimes.com/2007/12/26/opinion/26barr.html. *But see* Jon M. Jachimowicz et al., *When and Why Defaults Influence Decisions: A Meta-Analysis of Default Effects*, 3 BEHAV. PUB. POL'Y 159 (2019) (adopting a more nuanced approach and explaining why some defaults are more effective, i.e., stickier, than others).

[4] *See* Michael S. Barr, Sendhil Mullainathan, and Eldar Shafir, *An Opt-Out Home Mortgage System* (The Hamilton Project, Brookings, Discussion Paper No. 2008-14, 2008).

student loans,[5] from privacy protections[6] to retirement savings[7] and organ donations,[8] proposals to produce better outcomes via sticky defaults abound.

Because the presence of opt-out costs is so profoundly critical for the design of default rules, you would think that a rich account has been developed to explain what exactly are these opt-out costs that the earlier tradition wants to minimize and the more recent approach wants to exploit. It hasn't. The typical view focuses on what we call "mechanical costs": the process of developing and drafting a tailored alternative to the default. In negotiated contracts, this process may be time consuming, and thus mechanical costs can be large.[9] In mass contracts, opt-outs are usually cheap—pre-drafted by the business (e.g., disclaiming a default warranty) and readily assented-to by the consumer (e.g., by clicking "I Agree"). Even so, when summed over a vast number of transactions they add up to a non-trivial social cost.[10] Similarly, in non-contractual settings—think organ donations—the per-person mechanical costs are not large, but they add-up quickly when millions of people opt-out. To make defaults stickier, lawmakers or businesses may try to increase the mechanical costs.[11] Lawmakers may require more disclosures, segregated signatures, and

---

[5] *See* James C. Cox et al., *Designed to Fail: Effects of the Default Option and Information Complexity on Student Loan Payment* (Nat'l Bureau of Econ. Research, Working Paper No. 25258, 2018).

[6] Regulation 2016/679, of the European Parliament and of the Council of 27 April 2016 on the Protection of Natural Persons with regard to the Processing of Personal Data and on the Free Movement of Such Data, and Repealing Directive 95/46/EC (General Data Protection Regulation), 2016 O.J. (L 119) 1, ch. II, art. 7, § 2 ("...the request for consent shall be presented in a manner which is clearly distinguishable from the other matters, in an intelligible and easily accessible form, using clear and plain language.") [hereinafter GDPR].

[7] *See* James J. Choi, David Laibson, Brigitte C. Madrian & Andrew Metrick, *For Better or for Worse: Default Effects and 401(k) Savings Behavior*, *in* PERSPECTIVES ON THE ECONOMICS OF AGING 81, 83 (David A. Wise ed., 2004) ("[A]utomatic enrollment has a large impact on contribution rates . . . 65–87 percent of new plan participants save at the default contribution rate."); Brigitte C. Madrian & Dennis F. Shea, *The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior*, 116 Q. J. ECON. 1149, 1150 (2001) ("[T]he default contribution rate and default investment allocation chosen by the company for automatic enrollment has a strong influence on the savings behavior."). For a critique, see Ryan Bubb & Richard H. Pildes, *How Behavioral Economics Trims its Sails and Why*, 127 HARV. L. REV. 1593, 1598–99 (2014).

[8] *See* Eric J. Johnson and Daniel Goldstein, *Do Defaults Save Lives?* 302 SCIENCE 1338 (2003).

[9] *See* Robert Cooter, *The Cost of Coase*, 11 J. LEGAL STUD. 1, 17 (1982) ("The obstacles to cooperation are portrayed as the cost of communicating, the time spent negotiating, the cost of enforcing agreements, etc. These obstacles can all be described as transaction costs of bargaining."). 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. *See, e.g.*, Avery Katz, *The Strategic Structure of Offer and Acceptance: Game Theory and the Law of Contract Formation*, 89 MICH. L. REV 215, 226 ("Strategic behavior costs, in contrast, are the losses suffered because bargainers have the incentive to maximize their individual gains rather than the total surplus from exchange. . . . Such actions may include selling or buying a lesser quantity . . . or extended haggling, which both takes up valuable time and delays enjoyment of the bargain.").

[10] At the individual level (namely, without aggregating across many individuals), time-inconsistent preferences can amplify the effects of small opt-out costs. *See* John Beshears, James J. Choi, David Laibson, and Brigitte C. Madrian, *Behavioral Household Finance*, *in* 1 HANDBOOK OF BEHAV. ECON., Vol. 1, Ch. 3, Sec. X (Douglas Bernheim, Stefano DellaVigna, and David Laibson, eds., 2018).

[11] *See, e.g.*, Loi 2016-41 du 26 janvier 2016 de modernization de notre système de santé [Law 2016-41 of January 26, 2016 on the modernization of our healthcare system], J.O. (French law designed to increase the number of organ donations in part by requiring explicit refusal—constituted of the submission of a form and proof of identity— to opt out of organ donation); Par Anne-Aël Durand, *Don D'organes: Que Change la Nouvelle Mesure Sur le Refus Explicite?* [*Organ Donation: What Does the New Measure on Explicit Refusal Change?*], LE MONDE (December 22, 2016), https://www.lemonde.fr/les-decodeurs/article/2016/12/22/don-

Electronic copy available at: https://ssrn.com/abstract=3582129

personalized interactions. And businesses—once they establish the default settings that apply within their platforms—may devise painstaking mechanics for opt-out, so that more consumers stick with the business's preferred settings.[12]

We think that mechanical costs alone provide a poor foundation for the theory of default rules, and we introduce an additional, arguably more important, factor affecting opt-out – information costs. A decision to opt-out is based on the parties' perceptions concerning the existence and relative value of the alternatives. People need information about the default, its value, and how it compares to the (sometimes obscure) non-default options.[13] Such information can be very costly to acquire. Indeed, information costs could greatly exceed the mechanical opt-out costs. The presence of information costs and the decisions made in their shadow produce a novel theory of default rules.

Consider the canonical example of retirement savings. Mechanical opt-out costs are relatively small—just another check-the-box selection during job enrollment. But information costs could be substantial. People have to project their lifetime income and evaluate their future consumption needs. An intensive session of financial counseling is required to make a good opt-out decision.[14] Likewise, in many other contexts—data sharing, overdraft protection, add-on rental-car insurance—mechanical costs are a mere "click," but the information needed to click smartly is complex and expensive to acquire.

Focusing on information costs as the primary impediment to optimal opt-out forces us to rethink the notion of stickiness. It is widely thought that high information costs—like high "mechanical" costs—prevent opt-out. For instance, according to Sunstein, "there is strong evidence that a lack of information on the part of choosers, including a lack of information

---

d-organe-quarante-ans-apres-la-loi-sur-le-consentement-presume-ou-en-est-on_5053129_4355770.html. *See also* Ryan Nakashima, *AP Exclusive: Google Tracks Your Movements*, AP NEWS (Aug. 13, 2018), https://apnews.com/f60bc112665b458cb6473d7ee9492932 (an example where Google increased opt-out costs to make its privacy settings stickier).

[12] *See* Jamie Luguri & Lior Strahilevitz, *Shining a Light on Dark Patterns* (U. of Chi., Public Law Working Paper No. 719, 2019), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3431205.

[13] 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). For a different notion of information costs as an impediment to opt out, focusing on asymmetric information, see, e.g., Lisa Bernstein, *Social Norms and Default Rules Analysis*, 3 S. CAL. INTERDISC. L.J. 59, 70 (1993) ("A party may be reluctant to suggest varying a particular default rule even if the 'direct transaction costs' are low and the variation would make both parties better off . . . [as this] might be interpreted as a signal that the party suggesting the modification is more likely than previously thought to rely on this legal rights . . . ."); Jason Scott Johnston, *Strategic Bargaining and the Economic Theory of Contract Default Rules*, 100 YALE L.J. 615, 617 (1990) ("… to bargain *around* this default, the promisor must convey information which is generally directly contrary to his strategic interest in bargaining *with* the default."); Kathryn E. Spier, *Incomplete Contracts and Signalling*, 23 RAND J. ECON. 432, 432–33 (1992) ("[A]n individual may refrain from including a particular clause is a contract in order to signal his type."); see generally Omri Ben-Shahar & John A. E. Pottow, *On the Stickiness of Default Rules*, 33 FLA. ST. U. L. REV. 651 (2006).

[14] *See* Joshua Blumenstock, Michael Callen, and Tarek Ghani, *Why Do Defaults Affect Behavior? Experimental Evidence from Afghanistan*, 108 AMER. ECON. REV. 2868, 2871 (2018) (finding substantial increases in opt-outs from the default contribution rate after employees receive a financial consultation); Beshears et al., *supra* note 10, at 231 ("[A] significant portion of the effort cost [of opting out] consists of figuring out the implications of alternative choices.").

about alternatives, helps to account for the power of defaults."[15] This is an alluring conjecture: why would people initiate any opt-out unless they have information driving them to do so? Without information, it is thought, people stick to the default. But there is another possibility. 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 due to high information costs people remain uninformed but nevertheless decide to affirmatively opt out, based on their perceptions about the relative value of the alternatives. 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. This suggests an important caveat: *stickiness is an artifact of high mechanical costs, not of high information costs.* Accordingly, if mechanical costs are low and information costs high, default rules would be less sticky and more slippery than otherwise hoped, vulnerable to uninformed opt-out.

Uninformed opt-out is indeed prevalent. Returning to the retirement savings example: Retirement savings defaults are perhaps the archetypal sticky defaults. And yet mounting evidence suggests that they are quite slippery, especially over time. The traditional zero contribution rate (or no enrollment) default triggered an opt out rate of approximately 60%. And a similar percentage of employees opted out from the predominant 3% contribution rate default.[16] Much of this wholesale opt out was uninformed. This is not to say that employees did not acquire *any* information; of course they did – some more, others less. But even the better-informed employees fell far short of the full information required for this complex, financial decision. Most employees had an uninformed, or partially informed, sense that a zero percent or 3% contribution is not right for them, but their lack of information did not stop them from opting out.

---

[15] *See* Cass R. Sunstein, *Deciding by Default*, 162 U. PA. L. REV. 1, 20 (2013). *See also* Jeffrey R. Brown, Anne M. Farrell, and Scott J. Weisbenner, *The Downside of Defaults* (Nat'l Bureau of Econ. Research, Unpublished Working Paper, 2011) ("…lack of adequate information about decision alternatives is a significant driver of the likelihood of default… information problems are especially important."); Cass R. Sunstein and Richard H. Thaler, *Libertarian Paternalism Is Not an Oxymoron*, 70 U. CHI. L. REV. 1159, 1197-98, 1201 (2003) ("[I]n many domains, people's preferences are labile and ill-formed, and hence starting points and default rules are likely to be quite sticky.").

[16] On the opt-out rate from the no enrollment default, *see, e.g.*, William E. Nessmith, Stephen P. Utkus, & Jean A. Young, *Measuring the Effectiveness of Automatic Enrollment*, 31 Vanguard Center for Retirement Research, 1, 10 (2007) (showing that voluntary enrollment increased from 32% to 59% over the course of the first three years of employment); Vanguard, How America Saves (2019), p. 35 (the participation rate in voluntary enrollment, i.e., zero default, plans was 60%). On the opt-out rate from automatic enrollment plans, *see, e.g.,* Nessmith et al, *id*., p. 11 ("after 30 months, 57% of the employees hired under automatic enrollment [] have a rate higher than the default." [about half of the automatic enrollment plans featured automatic annual increases in the contribution rate; we are assuming that the "rate higher than the default" finding means higher than the annually-adjusted default]). On the predominant contribution rate default of 3%, *see id*. at 6 ("The median contribution rate in automatic enrollment designs is 2.9%."); Vanguard, HAS2019, p. 40 (noting the predominant 3% default). *See also* Part III.C. below.
[Vanguard, HAS2019, p. 40 (in automatic enrollment plans, with the predominant 3% default, the average contribution rate was 6.7%, as compared to a 7.1% contribution rate in voluntary enrollment, i.e., zero default, plans); p. 35 (the participation rate in voluntary enrollment, i.e., zero default, plans was 60%; the participation rate in automatic enrollment plans was 91%).

Or, take another example: overdrafts. In the past, a bank customer with no money in her checking account could still swipe her debit card and make a purchase. The banks called this "overdraft protection" and charged a hefty fee for the service. Worried about harm to consumers from these fees, the Federal Reserve Board, in 2010, set a no-overdraft-protection default. And yet of the frequent overdrafters, the group that policymakers were most concerned about, 45% opted out.[17] Here too the opt out was largely uninformed. Examples like this abound. Every sales contract contains a default (implied) warranty, but they are massively disclaimed in fine print. Every digital product is subject to privacy rules that govern by default, but so often contracted around in the vendors' terms of service. In all these scenarios, opt outs occur. They are by and large uninformed.

The first main theoretical contribution of this article is to recognize and analyze uninformed opt-outs. Defaults are everywhere, and opt-outs from these defaults are common. It is implausible to imagine that all opt-out decisions are or could ever be informed, given the staggering amounts of information that would be needed. Indeed, evidence suggests that many opt-out decisions are not only uninformed but in fact mistaken (or driven by cognitive biases), making people worse off. Consumers opt out of the default designed to protect them from high overdraft fees, without fully appreciating the financial consequences. Internet users opt out of the default designed to protect their privacy, without realizing how their personal information will be used. Employees opt out of the default contribution to their retirement plan, without understanding the effect on their retirement income. And so on.[18] Elsewhere, uninformed opt-out could also be fully rational. It is quite possible, for example, that people rationally allow firms to collect some personal information (thus opting out from the no-collection default), or to disclaim an implied warranty. Such opt-out behavior may be privately optimal, given the rational choice to remain uninformed, especially when information costs are high. Despite this prevalence of uninformed opt-out, the phenomenon has received surprisingly little, if any, attention. We identify the conditions for uninformed opt-out, and explain when it might be harmful and how it ought to affect the design of default rules.

The second theoretical contribution of the information costs theory is to highlight the effects of the default rule on information acquisition. The content of the default, we show, influences the incentives to acquire information. This, in turn, drives people's decisions to engage in informed or uninformed opt-out. Here, too, people's misinformation and irrational beliefs may distort their choice to acquire information and their resulting opt-out decisions. Recognizing these effects should play a role in the design of default rules. The information-forcing effect that we identify is different from the widely studied notion of penalty defaults.[19] The standard account addresses scenarios of asymmetric information, and advocates for default rules that force an informed contracting party to reveal

---

[17] *See* CFPB, *CFPB Study of Overdraft Programs*, at 29-30 (June 2013), https://files.consumerfinance.gov/f/201306_cfpb_whitepaper_overdraft-practices.pdf (last visited on Nov. 25, 2019). *See also* Part III.A. below.

[18] *See infra* Part III (where we discuss applications).

[19] For the standard accounts of penalty or information-forcing defaults, see Ian Ayres and Robert Gertner, *Filling Gaps in Incomplete Contracts: An Economic Theory of Default Rules* 99 YALE L.J. 87 (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 (1991).

information to another, uninformed contracting party. Our information-cost theory addresses a different problem. First, ours is not a theory of contract; it applies to a single decisionmaker that needs to decide whether or not to opt out of a default. Second, and related, our analysis does not assume asymmetry of information. It focuses on a single, uninformed party. The default rule does not force one party to reveal information to another; it induces uninformed parties to invest more in learning about themselves and how they would be affected by the default and non-default options.

The third theoretical contribution is a refinement of the standard majoritarian principle. The conventional approach assumes that mechanical costs are meaningful but not prohibitive, and thus to minimize these mechanical opt-out costs it advocates defaults that mimic the preferences of the majority. This approach is a good fit for environments with low information costs, where people with counter-preferences commit an informed opt-out. By contrast, when information costs are high, such selective opt-out does not occur, but the potential for uninformed opt-out suggests that the optimal default option is the one with the highest expected value (or perceived expected value), which minimizes the incidence and cost of uninformed opt-out. Both scenarios are unified under a general principle that prescribes a default that fits the *majority preferences given the majority's information*. In that sense, we still recommend majoritarian defaults, recognizing that the majority may act upon uninformed beliefs.[20]

The value of this refined criterion is at full display when we consider the "intermediate" information costs case, where people may choose to become informed under some default rules, but not others. In this region, sometimes a default rule that induces more information acquisition is superior, as it results in more tailored choices. But sometimes a default rule that does not induce any information acquisition is preferable. This may happen when the cost of information outweighs its value. Getting people to become informed, we show, is not necessarily better!

Information acquisition is a key ingredient in our model, and we recognize various ways in which it might occur. Apart from deliberate investment in information, people may acquire information through the efforts of the default setter or of interested third parties (e.g., when a firm "recommends" a particular choice). Importantly, people may glean information from the default option itself. This will happen, when they are unsure about their own information or think that the default-setter has better information. Here, the default-setter can influence opt-out and information-acquisition decisions. We should be especially worried about such "endorsement effects" when the default-setter does not have people's best-interest in mind.

---

[20] Our refined notion of majoritarian defaults minimizes opt-out costs – the costs of informed opt-out in the low-information-costs case and the costs of uninformed opt-out in the high-information-costs case. Thaler and Sunstein's proposed rule-of-thumb for default choice is 'minimize opt-out frequency.' *See* Richard H. Thaler and Cass R. Sunstein, *Libertarian Paternalism*, 93 AM. ECON. REV. 175 (2003). And Bernheim et al. show that this rule-of-thumb fairs well in certain cases. *See* B. Douglas Bernheim, Andrey Fradkin & Igor Popov, *The Welfare Economics of Default Options in 401(k) Plans*, 105 AM. ECON. REV. 2798 (2015). But they do not distinguish between informed and uninformed opt-out.

A central payoff of the information cost theory is to shed new light on, and suggest a reevaluation of, Nudge-type libertarian paternalistic ideas, and in particular the sticky-defaults paradigm. Regulation through defaults is premised on the expectation that less sophisticated people would stick with the default, while the more sophisticated are free to opt out.[21] But what counts as sophistication? If it is a proxy for low information costs, we show that libertarian paternalistic sorting could be defeated by uninformed opt out. Even unsophisticated people, with high information costs, may opt out of the default.[22] This explanation helps bridge an uncomfortable gap between the academic appetite for sticky default rules and the reality of slippery defaults.[23] It is the overlooked phenomenon of uninformed opt-out that accounts for the unintended slipperiness of so many default rules. Indeed, the information-costs theory provides a novel account of stickiness with surprising implications: we show that low information costs could *increase* stickiness, and that high information costs could *reduce* stickiness—a counterintuitive effect that prior accounts of stickiness did not recognize.[24]

The information costs theory complements conventional, behavioral accounts of default rules, but also qualifies them. Behavioralists have recognized a type of information costs—the "cognitive costs" associated with identifying optimal choice[25]—as contributing to the stickiness of defaults. But they mistakenly view information costs as simply another type

---

[21] *See* RICHARD H. THALER & CASS R. SUNSTEIN, NUDGE: IMPROVING DECISIONS ABOUT HEALTH, WEALTH, AND HAPPINESS, 42 (Yale Univ. 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."); 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.").

[22] Arguably, the libertarian paternalists would reject our attempt to link sophistication and information. They may argue that behavioral forces, not information costs, explain the stickiness of defaults. For further discussion about the relationship between information costs and behavioral forces – see Part VII below.

[23] Writers in the area of contract law have long recognized that stickiness is rare and opt-out is exceedingly common. *See, e.g.*, OMRI BEN-SHAHAR & CARL E. SCHNEIDER, MORE THAN YOU WANTED TO KNOW: THE FAILURE OF MANDATED DISCLOSURE 190-92 (2014); Lauren E. Willis, *When Nudges Fail: Slippery Defaults*, 80 CHI. L. REV. 1155 (2013).

[24] Existing accounts of stickiness include (1) effort costs, (2) an "endorsement effect" – trust in a benevolent default setter, (3) unawareness of the default and of the option to opt out, (4) an anchoring effect (the default as anchor), (5) loss aversion with the default as the reference point, (6) the default disproportionately features in the limited choice sets that individuals consider, and (7) cognitive dissonance. *See* B. Douglas Bernheim & Dmitry Taubinsky, *Behavioral Public Economics*, *in* 1 HANDBOOK OF BEHAV. ECON. – FOUNDATIONS AND APPLICATIONS 381 (B. Douglas Bernheim, Stefano DellaVigna & David Laibson eds., 2018) (surveying the literature on sticky default rules in the context of retirement savings); Beshears et al., *supra* note 10.

[25] Blumenstock et al., *supra* note 14, at 2871 ("[D]efault effects in savings persist because employees face significant cognitive costs associated with identifying their optimal contribution rate, and that this cost, together with present-biased preferences, creates procrastination.").

8

of mechanical opt-out costs.[26] While both information costs and mechanical costs affect the opt-out decision, they do so through different channels. In essence, people face two decisions: whether to acquire information, and whether to opt out and incur the mechanical costs. The behavioral literature flattens this two-stage process, collapsing the two types of costs into a single dimension. This mischaracterization leads the behavioral literature astray. It invokes myopia to explain stickiness. But, as we show, the main effect of myopia is to prevent people from acquiring information, not from opting out. The drive to find behavioral explanations for stickiness also obscures the nuanced role of misperceptions and how they might instead make the default slippery. We analyze these rich and underappreciated effects of misperception. Moreover, our information costs theory does not take any cognitive biases as fixed. It emphasizes that acquisition of information, which is affected by the default choice, can minimize the effects of misperception.

Finally, our analysis lays a novel and much needed foundation for a new consumer anti-manipulation law. Policymakers and commentators are increasingly concerned about ways in which firms manipulate consumers' choices. It has nevertheless proven challenging to distinguish such manipulations from the multitude of other permissible techniques used by firms to influence consumers' choices, such as most advertising campaigns, product shelf placement, and firms' sales recommendations. Our theoretical framework can help address this challenge. The distinction between mechanical costs and information costs imposes conceptual order on many potentially manipulative practices. Some of these practices are designed to increase mechanical costs, making it unnecessarily hard for consumers to keep the protective legal defaults, and all too easy to opt-into the firm's preferred, non-default option. Terms like "sludge" and "dark patterns" are sometimes used to describe such practices.[27] A different set of practices is designed to increase information costs, e.g., by peddling falsehoods or half-truths or by hiding important information in a haze of complex details. (Some of these practices are covered by false advertising and anti-deception laws; others are not.) For both of these categories – practices that increase mechanical costs and practices that increase information costs – our theory identifies those practices that reduce efficiency and harm consumers and should thus be prohibited as manipulation.

The remainder of the Article is organized as follows. Part II develops our information-cost theory of defaults. Part III considers different sources of information and beliefs. Part IV discusses the different regulatory implications of information costs and mechanical costs. Part V offers a more detailed analysis of several applications – policy domains where default rules play a key role: overdraft protection, privacy, retirement savings and "green" defaults. The Conclusion considers some of the assumptions underlying the analysis, specifically, what information policymakers need to apply the information-costs theory in the design of default rules.

---

[26] *See, e.g.*, Beshears et al., *supra* note 10 (suggesting that "a significant portion of the effort cost [of opting out] consists of figuring out the implications of alternative choices.").

[27] *See, e.g.*, Luguri & Strahilevitz, *supra* note 12, at 3; Cass R. Sunstein, *Sludges and Ordeals*, 68 DUKE L.J. 1843 (2019); Richard H. Thaler, *Nudge, Not Sludge*, 361 SCIENCE 431 (2018). The term "sludge" refers more narrowly to frictions, whereas the term "dark patterns" is broader.

II. A Theory of Information Costs and Default Rules

A. A Simple Model

We present here a theoretical illustration of the relationship between information costs and default rules. The analysis is presented through a numeric example.[28] Section 1 presents the framework of the analysis, and Section 2 demonstrates the conventional perfect information case, as a benchmark for the imperfect information analysis in Sections 3 and 4. Section 3 assumes that all uninformed parties have accurate beliefs about the distribution of types. Section 4 relaxes the accurate beliefs assumption.

1. Framework of Analysis

We consider a setting in which there are two possible arrangements, "High" and "Low". A typical issue addressed by default rules is the degree of protection to individuals, where High denotes more protection relative to Low—for example, a broader warranty, a higher pension savings rate, or greater privacy protection. 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 are better off with Low, and we call them "Type 1." 40% of the population are better off with 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, and without loss of generality, 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 assign a value of $v = 20$ to High. People can opt out of the default by incurring a "small" mechanical cost, 1. This cost is set to be non-zero but sufficiently low that any party would opt out from an unattractive default. In addition, people may invest some cost to acquire information about their type, if they don't already know it. One aspect of this example should be noted at the outset. Most people are better off with Low: 60% versus 40%. But higher expected welfare is obtained under High:

---

[28] A more formal and general derivation of the propositions is provided in the Appendix.

60%*(-10) + 40%*(20) = 2 under High versus 0 under Low. This duality will be critical to our analysis, which is aimed at identifying which rule is superior.

2. Perfect Information

In the perfect information benchmark, everyone knows their type. If the default rule is unattractive to them, they opt out. 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 opt out and average social welfare is: $W_{High} = 0.4 \cdot 20 + 0.6 \cdot (-1) = 7.4$.[29] 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.

3. Imperfect Information

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 $0.6 \cdot (-10) + 0.4 \cdot 20 = 2$. (The assumption that people know expected values but do not know their "type" is a common analytical representation of imperfect information. The term "type" should not be taken literally; it simply represents the information that people do not have.) Accordingly, based solely on such average valuations, people prefer High.[30] And if the default is set at Low, and people remain uninformed, 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" will 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.[31] Let us examine how people would behave, and the resulting social welfare, under the two defaults.

---

[29] We calculated welfare for a representative individual. To asses overall social welfare we would need to multiply these values by the number of individuals who are subject to the default rule.

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

[31] Note that when information is acquired, there will always be some level of informed opt-out, as long as the opt-out cost is sufficiently low (and some parties fair better with the non-default option). When information is not acquired, and the only possible opt-out is uninformed opt-out, there may be no opt-out at all, even with very low opt-out costs. Specifically, there will be no uninformed opt-out unless the expected (net) benefit of the default is smaller than the expected (net) benefit of the non-default option.

Low Default. The first decision people face is whether or not to acquire information.[32] The value of information is

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

The first component represents the payoff with perfect information: if an individual acquires information, she has a 40% chance of learning that she is Type 2; 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 of an uninformed individual: the individual will opt out – an uninformed opt-out – and earn 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.

High Default. Under a High default, the value of information is

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

The first component represents the payoff with perfect information: if an individual acquires information, she has a 60% chance of learning that she is Type 1; she will then opt out of the High default, incurring an opt-out cost of 1, and earning 0 instead of -10. The individual also 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.

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 thus 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:

---

[32] For a general exposition on the economics of information, see Jack Hirshleifer and John G. Riley, *The Analytics of Uncertainty and Information – An Expository Survey*, 17 J. ECON. LITERATURE 1375 (1979).

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

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

With both defaults, informed parties get 5 ($0.4 \cdot 20$) 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 (Low Default, in this example) 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.[33]

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 is 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. While the High outcome is optimal for only a minority of individuals, the High default is still majoritarian – it is what the majority of uninformed individuals, indeed all uninformed individuals, want.

(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 right match – 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.

---

[33] For expositional purposes, we compare the default rules without accounting for heterogeneity in information costs across people. When we allow for such heterogeneity in information costs, the optimal default will need to balance the different considerations listed above.

(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. With High default, individuals remain uninformed and thus forgo the optimal matching, but they avoid costly investment in information (as well as opt-out costs); social welfare equals the average value of 2. Therefore, Low default is more efficient when information costs are below 5.6 (= 7.6 – 2); otherwise, High default is more efficient. 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.

Note that, when information costs are at the lower end of the intermediate range, the theory prescribes Low default, even though High default maximizes expected value. Choosing a default option that is bad for the uninformed individual induces efficient information acquisition. In essence, the information costs theory identifies a new type of "penalty default" or "information-forcing default."

It is important to explain the difference between our notion of an information-forcing default and the familiar account of penalty defaults.[34] The standard account of penalty defaults addresses contracting scenarios with asymmetric information, and advocates default rules that would force an informed contracting party to reveal its private information or to otherwise act upon it. Our account, while applicable to contracting scenarios, is not focused on contracting. We study a single-party, decision-theoretic model, where an uninformed party decides whether to invest in information acquisition. If acquired, the information benefits the acquiring party, not some other party. As in the standard accounts of penalty defaults, our analysis justifies counter-majoritarian rules (given our refined notion of a majoritarian default). Indeed, when information costs are at the lower end of the intermediate range, our information-costs theory prescribes Low default, because it induces parties to acquire information, despite the fact that the majority of uninformed parties prefer High.

4. Biased Information

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

---

[34] For the standard accounts of penalty or information-forcing defaults, see Ayres and Gertner, *supra* note 19, and Bebchuk and Shavell, *supra* note 19.

can do better.[35] 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 people overestimate the likelihood that they are Type 1 to be 80%.[36] To be sure, there are many other ways in which beliefs could be inaccurate.[37] 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.[38]

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. Again, this is a majoritarian default – it gives the majority of individuals, indeed all individuals, what they want, given their inaccurate uninformed beliefs.

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.

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

---

[35] 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.

[36] 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.

[37] We examine the effects of other distortions more thoroughly in the Appendix.

[38] 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.

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.) [39]

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. [40]

---

[39] 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.

[40] 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 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.

5. Beyond a Binary 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 – either you are an organ donor or not, either you have overdraft protection or not. In other cases, there are multiple options: there are more than two possible privacy settings and thus more than two possible defaults. Even with organ donations, there can be more than two options – you may agree to donate some organs but not others. In yet other cases, the choice is continuous. In the retirement savings example, the choice of contribution level is a continuous choice.

Our framework can easily be adjusted to accommodate a multiplicity of possible outcomes and even a continuous outcome space. In the basic, binary model, an individual had to compare the default option to only one other, non-default option. With multiple outcomes, the individual would need to compare the default option to several non-default options. This comparison task would be even more challenging with a continuous outcome space. Consider the decision process of an employee who faces a specific contribution rate default, say 3%. The employee would need to calculate an expected payoff function, where the expected payoff is a function of the continuous contribution rate, and maximize this function. All of these comparisons – between two options, between multiple discrete options or between a continuum of options – are, at best, approximated by the individual. Such rough approximations are sufficient to support the insights generated by the theoretical model. Still, we acknowledge the added difficulty in assessing the optimal choice when the number of options is larger. We discuss in the Conclusion the impact of such informational burdens.

6. Summary

Information costs affect the design of optimal defaults. When information costs are high, people remain uninformed, but may still choose uninformed opt-out if they can increase their expected payoff. Thus, the default rule that is optimal when parties are informed—the one that tracks the informed preferences of the majority—has to be replaced with a rule that tracks the *uninformed* preferences of the majority or, put differently, the rule that maximizes expected payoffs. In both cases, the goal is to minimize the incidence of opt-out; but while the traditional, low-information-costs analysis focuses on informed opt-out, when information costs are high we should focus on uninformed opt-out.

Our analysis highlights the distinction between two types of transactions costs – mechanical costs and information costs – that are often treated interchangeably. The practical effects of these two types of transactions costs are drastically different. High mechanical costs can prevent opt out and lead to greater stickiness of defaults. High information costs, in contrast, do *not* prevent opt-out and thus do not render the default rule sticky. This distinction helps explain some of the contracting strategies deployed by firms

17

vis-à-vis consumers, and also helps guide lawmakers who seek to optimally intervene in these environments. We develop these lessons below.

Another insight from our model concerns the incentives of uninformed people to acquire information about the default and its relative value. Specifically, we saw that when information costs are positive but not prohibitive, people may choose to incur the information cost and make informed opt-out decisions. The choice of default affects this decision, and we identified an "intermediate" range of costs in which people acquire information under one default rule but not the other. Acquiring more information does not guarantee more opt-out, because informed people may decide to stick with a default that the uninformed reject. In this intermediate range, sometimes the default that leads to more information acquisition is superior, because it guarantees better matching; and other times the default that leads to less information acquisition is superior, because it saves information costs.

When people have biased prior beliefs about their type, uninformed opt-out still occurs and is more likely to result in a bad outcome—the one mistakenly expected to be better on average. If information costs are high, it is pointless for society to set a default that maximizes the true expected value, because uninformed people will act upon their biased beliefs and opt out. This is another implication of uninformed opt-out that prior literature ignored, and it sets an important practical limit on the social utility of paternalistic default rules: rather than being sticky, they merely impose opt-out costs. In addition, when information costs are not prohibitive, biased beliefs may distort people's decisions to acquire information. They might waste money acquiring information that they overvalue (high information costs can helpfully deter such waste – a counterintuitive result); and they might fail to acquire information that they undervalue. On the bright side, when people decide to acquire information, the pre-information biases become irrelevant.

B. Distributive Considerations

The possibility of uninformed opt-out forces us to rethink the role of distributive considerations in designing default rules. If most people are uninformed and stick with the default, the policymaker could successfully 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 (Section II.A.), 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.

18

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 tracks the preferences of the informed majority. 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.

C. The Information Cost Theory Versus Behavioral Accounts of Defaults

The theory of opt out from default rules presented in Part II.A focuses attention on the importance of information and information costs. It shows that costly information may not stop opt out from occurring, although it would affect the opt out choices people make. The possibility of uninformed opt out has the potential to explain how unsophisticated parties respond to defaults.

This is not the first attempt to find order in people's uninformed response to default rules. Different behavioral theories have been proposed. But while these behavioral theories have focused on explaining the stickiness of defaults, our information-costs theory explains why many defaults remain slippery. Our analysis is not intended to replace existing behavioral accounts of default rules. It does, however, qualify and shed new light on their conclusions.

1.   Present Bias and Procrastination

One of the most prominent accounts of default effects relies on the behavioral notions of myopia, present bias and procrastination. Even if the non-default option is better, the benefit from switching is in the future, whereas the cost of switching is in the present. Thus, the myopic or present-biased individual will defer the costly switching until the next period, and when the next period arrives to the period after that, and so forth. A standard pattern of procrastination emerges, and the individual ends up with the inefficient default.

The problem with this standard procrastination story is that it requires a switching cost that present biased individuals wish to postpone. The "mechanical" opt-out costs, in many cases, are too small to support the procrastination story. Information costs, on the other hand, can be much larger – definitely a cost that a present biased individual would want to postpone.[41] If information is costly, a present biased individual may remain uninformed, even when a rational individual would acquire the information. And if the optimal, uninformed choice is to stick with the default, this combination of information costs and present bias explain

---

[41] *See* Blumenstock et al., *supra* note 14, at 2871 (concluding that "default effects in savings persist because employees face significant cognitive costs associated with identifying their optimal contribution rate, and that this cost, together with present-biased preferences, creates procrastination").

the stickiness of defaults. To be sure, the information-costs theory can explain stickiness even without the behavioral add-on, but the interaction between the two accounts generates even more stickiness.

The procrastination story needs information costs. But adding information costs forces a revision of the standard procrastination story. This is a story of sticky defaults – defaults are sticky, because switching is costly and this cost is postponed, indefinitely. The information costs theory teaches that a present biased individual may postpone the costly act of acquiring information and thus remain uninformed. When uninformed individuals prefer the default option, procrastination produces sticky defaults. But when uninformed individuals prefer the non-default option, procrastination results in uninformed opt-out and slippery defaults. Scholars, such as Joshua Blumenstock, Michael Callen, and Tarek Ghani,[42] and John Beshears, James J. Choi, David Laibson, and Brigitte C. Madrian,[43] conclude that present bias, with information costs, generates stickiness, because they consider information costs as just another species of opt-out costs. Our information costs theory emphasizes the qualitative difference between information costs and "mechanical" opt-out costs. It is this difference that explains how present bias can result in slippery, not sticky, defaults.

2. Optimism

In some applications, optimism has been suggested as an explanation for the stickiness of defaults. Consider the overdraft protection example. Before the 2010 rule change, the 'overdraft protection' default was quite sticky. (This default allowed individuals to overdraw their checking account, and allowed the bank to charge a fee for each overcharge.) Arguably, this stickiness could have been attributed to people's optimism, specifically, their underestimation of the likelihood that they would incur multiple overdraft fees. (And after the rule change, the relative slipperiness of the 'no protection' default may be attributed to the same optimism.)[44]

The optimism account interacts with our information costs theory on two dimensions. First, optimism is an example of a misperception affecting the relative (net) benefit of the default and non-default options – the type of misperception incorporated into our theory. In this sense, the information costs theory generalizes the standard optimism story. But this is not an "innocuous" generalization; it is a generalization that highlights the limits of the standard account. Specifically, the optimism story assumes that the bias favors the default option. Our theory allows for misperceptions that favor either the default option or the non-

---

[42] *See id.*

[43] *See* Beshears et al., *supra* note 10.

[44] Willis (2013), *supra* note 23, at 1183 ("As regulators noted in promulgating the overdraft default, consumers are likely to assume overoptimistically they will not overdraw." Citing 74 Fed Reg at 59044.); *see also* Christine Jolls, *Behavioral Economics Analysis of Redistributive Legal Rules*, VAND. L. REV. 1653, 1659 (1998) ("An amazingly robust finding about human actors—and an important contributor to the phenomenon of risk estimation—is that people are often unrealistically optimistic about the probability that bad things will happen to them."); *see generally* TALI SHAROT, THE OPTIMISM BIAS, A TOUR OF THE IRRATIONALLY POSITIVE BRAIN (2011).

default option. This is an important practical expansion, because the legal default is often presented in a manner that triggers pessimism—an exaggerated likelihood of a negative event—to induce adherence to the default, especially in consumer markets where sellers profit from consumers' choice of non-default add-ons. For this reason, our stickiness result is more circumscribed.

More important, while the standard account assumes that optimism will always affect the individual's stick-or-switch decision, our information costs theory emphasizes the limits of optimism. Optimism, we show, is a function of information. It affects the uninformed, but not the informed. More generally, the more information you have, the less room there is for bias such as optimism to take hold.[45] The information costs theory helps policymakers identify the default rule that will induce more acquisition of information and thus minimize the effects of optimism.[46]

3. Anchors and Reference Points

A third set of behavioral explanations focus on defaults as anchors or reference points. Behavioral economists have shown that asking people to consider an arbitrary number (an "anchor") will bias subsequent judgments and estimations towards this number.[47] The default option may serve as an "anchor." A related, behavioral literature has emphasized the importance of reference points, against which decision makers evaluate gains and losses. This literature has also demonstrated that many individuals are loss averse, and experience losses more strongly than commensurate gains.[48] If the default option becomes the reference point and relative costs of the non-default option loom larger than relative gains from the non-default option, then decision-makers will tend to stick to the default.[49]

These behavioral explanations presume a rather sparse informational environment, where the default rule is the only possible anchor or reference point. Our information costs theory

---

[45] OREN BAR-GILL, SEDUCTION BY CONTRACT: LAW, ECONOMICS, AND PSYCHOLOGY IN CONSUMER MARKETS 9 (2012).

[46] The information-costs theory also shows how misperception, like optimism, can distort the decision whether to acquire information. *See supra* Secs. II.A.4 and IV.A.

[47] The anchoring effect was first identified by Tversky and Kahneman in Amos Tversky & Daniel Kahneman, *Judgment Under Uncertainty: Heuristics and Biases*, 185 SCIENCE 1124 (1974). In the retirement savings context, James J. Choi et al., *Small Cues Change Savings Choices*, 142 J. ECON. BEHAVIOR 379 (2017) and Gopi Shah Goda, Colleen Flaherty Manchester, & Aaron J. Sojourner, *What Will My Account Really Be Worth? Experimental Evidence on How Retirement Income Projections Affect Saving*," 119 J. PUBLIC ECON. 80 (2014) show that employees raise their 401(k) contribution rates if they are exposed to arbitrary high contribution examples in communications. *See also* B. Douglas Bernheim, Andrey Fredkin, & Igor Popov, *The Welfare Economics of Default Options in 401(k) Plans*, AM. ECON. REV. 2798 (2015). This account is based on the literature review in Beshears et al., *supra* note 10, at Ch. X.

[48] *See* Daniel Kahneman and Amos Tversky, *Prospect Theory: An Analysis of Decision Under Risk*, 47 ECONOMETRICA 263 (1979) (developing Prospect Theory that emphasizes the importance of reference points and identifies loss aversion).

[49] Moreover, Kahneman and Tversky argue that negative consequences of action (like affirmatively opting out of the default) are felt more strongly than negative consequences of inaction (like sticking with the default). *See* Daniel Kahneman and Amos Tversky, *The Psychology of Preferences*, 246 SCI. AMER. 160 (1982). This would make the default option even stickier.

Electronic copy available at: https://ssrn.com/abstract=3582129

highlights the possibility that individuals will acquire information before making the stick-or-switch decision. If individuals decide to acquire more information, then the salience of the default option would likely diminish. Allowing for such a richer informational environment, the anchoring or reference point explanations would predict a reduction in stickiness. In contrast, under the information-costs theory, more information can either increase or decrease stickiness.[50]

III. Sources of Information and Beliefs

The theory presented in this Article identifies information and beliefs as key factors that affect opt out decisions. Part II analyzed a framework in which people did not know their type, but were able at some cost to discover it. In the absence of specific information, we said that people were uninformed in the sense that they held onto prior beliefs—sometimes accurate on average but other times biased.

In the next two Parts we explore the relevance of these findings to several central questions surround the design of default rules. First, Part III takes the theoretical building block of information costs and beliefs and expands the lens to examine various factors that are captured by these concepts. What are the information costs that people can spend to become informed? What are the prior beliefs, and what drives them towards accuracy or bias? Later, in Part IV, we turn to the other feature of the model—the difference between mechanical and information costs—and explore regulatory implications arising from it.

A. Deliberate Investment in Information

Our basic model assumed that individuals deliberately weigh the cost of information acquisition against the benefit, or perceived benefit, of information. When the cost is small and the benefit (or perceived benefit) is large, individuals make a deliberate investment in information. Such acquisitions of information can take different forms.

Consider the retirement savings context. Individuals can invest more or less time researching the default and non-default options. They can read brochures or attend webinars sponsored by their employer. They can study third-party materials that explain the implications of higher versus lower contributions to a retirement account. And they sometimes discuss the question of optimal savings with colleagues, friends and family members. Individuals can also spend actual money on information. They can pay financial advisors who will explain the retirement savings options and help devise an optimal savings plan.[51]

---

[50] *See supra* Secs. II.A and IV.A.

[51] In 2019, almost 13,000 SEC-registered investment advisors managed $83.7 trillion for 43 million clients, including $10.5 trillion invested for individual clients who are 94.6% of the total, 43 million clients. With a 1% fee, individuals paid $105 billion for investment advice. See National Regulatory Services and

Electronic copy available at: https://ssrn.com/abstract=3582129

Retirement savings is a decision of great importance that people make one time (or once every few years), often with the benefit of some accessible advice. It pays to make a deliberate investment in information in this setting. There are other decisions that fall into this category. When buying property insurance, policyholders are offered a basic coverage default ("Low"), to which they can add supplemental coverages like a hurricane endorsement or coverage for damages from flooding ("High"). When purchasing big ticket items, buyers have to decide about adding extended warranties. In many transactions, buyers face a choice to extend the withdrawal period (for example, making an airline ticket refundable) by paying a premium. When these decisions arise, people are usually uniformed about the relative value of the add-ons, but have plenty of opportunity to acquire more information.

Some of the information that people deliberately acquire comes in the form of comparison shopping. Can an extended warranty be purchased more cheaply elsewhere? Other information can be teased out of the fine print. Under what conditions can the consumer withdraw from a "refundable" transaction? And some information comes from advice— for example, what coverages are critical to add to one's insurance policy.

Deliberate investment in information is probably more common when the stakes are large. In numerous other contexts, opt out decisions have to be made in day-to-day transactions, each of which contains a long list of opt outs. For example, opt out is common in the data privacy context. People can spend time studying the legal rules that provide default protections or firms' privacy policies that qualify these protections. But, realistically, most people will spend little (or no) time studying legal rules and reading privacy policies, leaving them with imprecise information. And, yet, when a particular issue (like data collection) comes up in multiple contexts, people may divert some attention to understanding the value of opt out. People might read newspaper articles about how firms use their data and reports about data breaches and their consequences. They may even consult with tech-savvy acquaintances. True, they are unlikely to have accurate information to evaluate the opt out options in each individual transaction, to support a truly informed decision. And even if they spend time and money trying to become informed, they might end up with inaccurate or even misleading information – not all the information available on the internet is properly vetted, and not all "expert advisers" are really expert advisers. Still, some consumers would immerse themselves more than others, acquire better information and make more informed decisions.

While the model in Section II.A identified the information cost as a parameter that characterizes each transaction, it is important to recognize that information costs interact across transactions. As the number and frequency of opt out decisions accumulate, people have to divide their attention across more contexts, and are thus able to acquire information only in those areas that really matter. Part of the information problem that people solve, then, is not only how much to invest in each case, but how to prioritize. Like the decision

---

Investment Adviser Association, 2019 Evolution Revolution: A Profile of the Investment Adviser Profession (2019) (available at: https://www.investmentadviser.org/publications/evolution-revolution). The vast majority (95.5%) of fees paid to investment advisors are billed as a percentage of assets under management.

to opt out, the decision to prioritize can be made with the benefit of some investment in "meta"-information, or it can be made uninformed. For example, having learned – from an employer or from a third party – about the importance of retirement savings, an individual will make an informed (or more informed) choice to prioritize the decision whether to opt out from the retirement savings default and to pay less attention to, say, an offer to buy an insurance add-on.

Even when deliberate investment in information is unrealistic, there are other ways to become informed. In the following discussion, we examine these alternative sources of information. We also explore people's beliefs (about the expected values of the default and non-default options) and the information they use to shape these beliefs.

B. "Endorsement Effects": The Informational Content of Defaults

Thus far we focused on deliberate investments in information acquisition. But an information theory of defaults must account for another channel through which decision-makers can become informed. People can glean information from the chosen default option itself. Consider a benevolent employer who sets a default retirement savings contribution rate that, according to the employer's expert opinion (formed after consulting with retirement savings professionals), maximizes the expected value for her employees. If employees believe that their employer has set the default option in this fashion—perhaps because they believe that the employer has their best interest in mind or because they believe that the government mandated such a default option and the government has their best interest in mind—this would affect their decisions, whether to collect information and whether to opt out.[52] The precise effects depend on the nature of the information that the employer has. Here, we assume that the employer has information only about average values. The alternative assumption is considered in Subsection 5 below.

If employees hold accurate uninformed beliefs, the employer has no informational advantage and the default it sets conveys no new information. The employer should choose the same default prescribed by our theory when there is no information conveyed by the default (which, in the case of high information costs, is the one that maximizes expected value).[53]

---

[52] This can be viewed as a formalization of the "endorsement effect" that is noted in the behavioral economics literature as a reason why people stick with the default. *See, e.g.*, Beshears et al., *supra* note 10, at Sec. X; Jachimowicz et al., *supra* note 3, (listing the endorsement effect as one of three main reasons why defaults are sticky); Craig R.M. McKenzie, Michael J. Liersch and Stacey R. Finkelstein, *Recommendations Implicit in Policy Defaults*, 17 PSYCHOL. SCIE. 414, 418 (2006) ("[P]olicymakers' choice of default leaks information regarding their beliefs or attitudes about the available options, and the public is sensitive to this information."). The policymaker's choice of default can also contain relevant information about what others are doing, about the social norm, which can affect an individual's payoffs.

[53] If, instead, the employer is known to choose the default option that is best for a majority of employees, rather than the one that maximizes the expected value to employees, the employees who remain uninformed will opt out.

By contrast, when employees have inaccurate uninformed beliefs, the employer's choice of default conveys information. Uninformed employees who recognize that the employer has better information about average values will update their beliefs after observing the employer's default choice. In our example above, uninformed employees mistakenly thought that the prevalence of Type 1 is 80%, when in fact it was 60%. With high information costs, the policy prescription was to "succumb" to the misperception and choose the default that maximizes *perceived* expected value, namely, Low default. This prescription may change once employees draw inferences from the employer's default choice. The employer could then choose High default, which maximizes actual (not perceived) expected value, and the employees, observing this default choice, would infer that the likelihood of being Type 1 is lower than 80% and stick with the default. Here, the presence of an endorsement effect contributes to the stickiness of the default.[54]

The preceding analysis assumed that the employer has her employees' best interests in mind and that the employees accurately perceive their employer's benevolence. But what if the employer's interests are not perfectly aligned with the employees' interests? [55] Then an employer might set a default that does not maximize her employees' expected value. If employees are aware of this conflict of interests, little harm is done; the employees would simply rely on their own imperfect information about average values (and not use the default option to update beliefs), as in our basic model.[56] The concern is that employees would mistakenly attribute benevolence to a non-benevolent employer. This could render the inefficient default sticky and prevent efficient uninformed opt-out.

We have thus far focused on the high-information-cost case, where parties do not acquire information. When information costs are lower, the endorsement effect can also influence the decision to acquire information. For example, uninformed employees who doubt the accuracy of their information about average values may decide to acquire more information. But if they get an informative signal from their employer, through the default choice, that alleviates uncertainty about average values, they may no longer feel the need to invest in information acquisition.

The informational content that default rules have depends on the perceived informational advantage that a default-setter enjoys in the eyes of the decisionmaker. This endorsement effect also depends crucially on the perceived alignment, or misalignment, of interests between the default-setter and the decisionmaker. The greatest potential reduction in social welfare arises when decisionmakers overestimate the informational advantage of the default-setter or mistakenly believe that the default-setter is looking after their interests. The importance of endorsement effects may vary across contexts and should not be exaggerated. In many contexts, trust in the default-setter will not be high or, simply,

---

[54] *See, e.g.*, Beshears et al., *supra* note 10, at Sec. X; McKenzie et al., *supra* note 52.

[55] *See* Ryan Bubb, Patrick Corrigan and Patrick L Warren, *A Behavioral Contract Theory Perspective on Retirement Savings*, 47 CONN. L. REV. 1317 (2015); Ryan Bubb and Patrick L. Warren, *An Equilibrium Theory of Retirement Plan Design* (work in progress).

[56] If the interests of the employer and the employee are in conflict, then when the employer chooses one default the employees may infer that the non-default option is better for them. Of course, anticipating such inferences, the employer may choose the default strategically. Such anti-endorsement effects are even more likely in the consumer context, where the interests of sellers and consumers are often in conflict.

decisionmakers will prefer to rely on their own information. However, in some contexts where the government is known to be setting the default, the endorsement effect could be significant, especially if it is further bolstered by public education campaigns that advertise the chosen default and explain its advantages.[57]

C. Influencing Uninformed Beliefs

The recognition that, because of high information costs, at least some people will remain uninformed emphasizes the importance of uninformed beliefs. Such beliefs affect the opt-out behavior of individuals that choose to remain uninformed, and distort the decision to acquire information. Accordingly, policymakers should pay more attention to uninformed beliefs and perhaps even seek to influence them.

How? We explained above that the legally prescribed default could at times shape uninformed beliefs. Even more aggressively, policymakers could actively warn consumers about the adverse consequences of choosing the non-default option. While we have our doubts about the ability of lawmakers to successfully educate people about the myriad of issues covered by default rules, it is possible that in select and particularly salient contexts such interventions would be desirable.

At the same time, "educational" campaigns are attempted by other interested parties who hope to influence—or manipulate—people's uninformed beliefs so as to induce them to opt out of the policymaker's protective default. In the consumer context, firms work hard to exert such influence. They highlight the benefits of their preferred opt-out option, emphasizing some dimensions of its value to consumers, while magnifying the risks of foregoing this "recommended" option.[58] Indeed, when a firm prefers a non-default option,

---

[57] *See, e.g.*, Kim Willsher, *France Introduces Opt-Out Policy on Organ Donation*, GUARDIAN (Jan. 2, 2017, 8:21 AM), https://www.theguardian.com/society/2017/jan/02/france-organ-donation-law (describing an official French campaign, which included a video advertisement, discouraging people from opting out of organ donation).

[58] *See* Willis, *supra* note 23, at 1172–73 ("Firms exacerbate judgment and decision biases intentionally through framing devices. They advertise the benefits of the default, both to directly shape preferences and so that consumers will consider the benefits of the default before considering any alternatives … They trumpet the benefits and downplay the costs of the default. They explicitly tell consumers that the default is 'recommended' or 'advised'."); Lauren E. Willis, *Why Not Privacy by Default?*, 29 BERKELEY TECH. L.J. 61, 82–3, 95, 102 (2014) ("Opting out of a default also might be made more or less attractive through messages to the user … warn the user that changing the setting could cause problems"; "goes so far as to warn consumers that choosing not to be tracked will 'spoil your experience of the website'"; "financial information defaults [are] designed to trigger loss aversion and the endowment effect … warns that opting out will cause consumers to lose benefits they now have."); Kevin Bankston, *Facebook's New Privacy Changes: The Good, The Bad, and The Ugly*, ELECTRONIC FRONTIER FOUNDATION (Dec. 9, 2009), https://www.eff.org/deeplinks/2009/12/facebooks-new-privacy-changes-good-bad-and-ugly ("Although sold as a 'privacy' revamp, Facebook's new changes are obviously intended to get people to open up even more of their Facebook data to the public. The privacy 'transition tool' that guides users through the configuration will 'recommend' – preselect by default – the setting to share the content they post to Facebook."); Google Chrome, Privacy Settings ("allow websites to save and read cookie date **recommended**.").

a contest between the policymaker and the firm over consumers' perceptions may ensue.[59] It is hard to imagine that lawmakers could win such tournaments;[60] their best chance is to find novel ways to outlaw some of the firms' manipulative campaigns. While a general regulatory framework to disallow such manipulations does not presently exist, the law of deception could be stretched to deal with the worst cases.

D. Forced Deliberation

Our theoretical model assumed that any default prompts people to engage in the mental exertion of comparing its value to that of the opt-out option, if only in (uninformed) expected-value terms. But the number of issues that people encounter and that are governed by defaults is so vast that it casts doubt on this ideal of active evaluation and comparison. Indeed, many people likely stick to many defaults without thinking; they just decline to make an active decision. How could people know which issues are worth some deliberation? How could policymakers help by selectively identifying important issues and encourage deliberation, even uninformed deliberation, about these issues?

One such technique is enacting a "no-default regime." In it, people cannot remain passive, as the no-default requires active choice in order to complete the transaction.[61] For example, an employee must select a retirement plan or else the employment relationship cannot begin; or an applicant cannot apply for a driver's license without first choosing whether to be an organ donor.[62] Our analysis of uninformed opt-out makes it harder to justify such active choice structures. Even if they induce people to deliberate and choose, they usually lead to an uninformed opt-in, where people select what they perceive as the highest expected value option. Lawmakers could have chosen this option as the default, saving some mechanical costs of opt-in.

But active choice could be more subtly rationalized as a behaviorally-designed technique that forces people to think and acquire information when such acquisition is worthwhile. When forced to choose, people might be prompted to think harder and acquire more

---

[59] As with advertising generally, one could ask if policymakers and firms are targeting consumers' perceptions or trying to shape consumer preferences. *See generally, e.g.*, Christina L. Brown & Aradhna Krishna, *The Skeptical Shopper: A Metacognitive Account for the Effects of Default Options on Choice*, 31 J. CONSUMER RES. 529 (2004); Christopher A. Summers, Robert W. Smith, & Rebecca Walker Reczek, *An Audience of One: Behaviorally Targeted Ads as Implied Social Labels*, 43 J. CONSUMER RES. 156 (2016); Peter Wright, *Marketplace Metacognition and Social Intelligence*, 28 J. CONSUMER RES. 677 (2002).

[60] *See, e.g.*, Willis (2013), *supra* note 23, at 1184 (describing the impact of regulation that required consumers to opt into overdraft protection and noting that "[n]ot all banks energetically pursued overdraft revenue after the change in the law, but those that did have managed to achieve high opt-out rates . . . .").

[61] Such active choice regimes have been offered in response to a critique that sticky defaults are paternalistic. *See, e.g.*, Luc Bovens, *The Ethics of* Nudge, *in* PREFERENCE CHANGE: APPROACHES FROM PHILOSOPHY, ECONOMICS, AND PSYCHOLOGY (Till Grüne-Yanoff and Sven Ove Hansson, eds., 2008). *But see* Craig N. Smith, Daniel G. Goldstein, and Eric J. Johnson, *Choice Without Awareness: Ethical and Policy Implications of Defaults*, 32 J. PUBLIC POL'Y & MARKETING 159, 163–64 (2013).

[62] *See generally* NAT'L SCI. AND TECH. COUNCIL, EXEC. OFFICE OF THE PRESIDENT, SOCIAL AND BEHAVIORAL SCIENCES TEAM: 2016 ANNUAL REPORT (2016); Beshears et al. (2018), *supra* note 10; Gabriel D. Carroll et al., *Optimal Defaults and Active Decisions*, 124 Q. J. ECON. 1639 (2009).

information towards an informed decision, which will lead to optimal sorting. It helps people prioritize their limited attention and information acquisition resources. As long as such strategy is used sparingly, the increased mechanical costs would be outweighed by the superior ultimate outcomes. If used too often, it would lose its attention-alerting, information-inducing, effect.[63]

In addition, lawmakers could force people to notice and address an issue by enacting "stop-and-think" defaults. These are surprising or unexpected defaults that send some signal to people and force them to contemplate the issue. In the retirement savings context, a zero-contribution default may constitute such a stop-and-think default. It is clearly a suboptimal contribution rate. It does not directly provide information about the optimal choice, but it may encourage evaluation and information acquisition.[64]

E. Default Rules with Different Information Costs

The basic theory presented in Section II.A made the simplifying assumption that the distribution of information costs is independent of the chosen default. While this assumption is a plausible benchmark, in some applications information costs may depend on the default rule. Consider consumer contracts. The cost, to the consumer of becoming informed may be quite high unless the seller is motivated to provide information. If the seller's incentive to provide information is stronger under one default and weaker under the alternative default, then information costs vary with the chosen default.[65]

In the overdraft protection example, a pro-bank default imposes higher information costs than an anti-bank default. If the default is 'overdraft protection' and consumers are expected to stick with the pro-bank default, then the bank would not provide any relevant information and consumers would find it difficult to acquire information on their own – perhaps because the uninformed consumers are not even aware of the overdraft protection issue. In contrast, if the default is 'no overdraft protection' and consumers are expected to stick with the anti-bank default, then the bank would have a strong incentive to provide information to consumers about the benefits of the non-default, overdraft protection option. The cost to consumers of acquiring at least some information – benefit information – is lower under the anti-bank default. Other information – about the potential harm from overdraft protection – would remain costly to acquire.

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

---

[63] Active choice can also be attractive if we are concerned about a false endorsement effect. Namely, if people mistakenly trust an untrustworthy default setter and thus might stick to a harmful default, then an active choice regime that strips power from the default setter can be helpful.

[64] *Compare* Bernheim et al. (2015), *supra* note 20 (noting that extreme defaults can force active choice.) Stop-and-think defaults are also related to penalty defaults. *See* Ayres and Gertner, *supra* note 19, and Bebchuk and Shavell, *supra* note 19.

[65] *Compare* Ayres and Gertner, *supra* note 19, and Bebchuk and Shavell, *supra* note 19.

opt-out from a default that is less favorable to that seller. Also, as noted above, lower information costs might actually reduce welfare.

F. Personalization

We have thus far assumed that the policymaker sets a single default rule for all relevant parties. But this need not be the case. When default setters have better information – individualized information – they can increase welfare by setting personalized defaults.[66] An employer will often have detailed personal information about the employees, including age, income, number of dependents, education, and health. The employer can use this information to offer, as a default, personalized contribution rates or retirement funds (with personalized risk characteristics) to different employees. The optimal contribution and investment option for a 25-year-old employee is different from that of the 65-year-old employee.[67]

The default setter could affirmatively elicit such personal information from people. For example, an employer could ask (or require) the employees to complete a short survey and based on the survey response set the personalized, default contribution rate.[68] Indeed, this strategy can be viewed as a means of reducing information costs. The policymaker and the individual join forces, combine their information, to arrive at the optimal default.

Personalized rules in general, and personalized default rules in particular, require large amounts of information. At their purest form, they are derived from algorithmic analysis of Big Data. Personalization is thus a solution to information problems that underlie a heterogeneous society. It is a data-driven substitute for the solutions to the information problem that are developed in our model, in which parties either act uninformed or spend resources to acquire and analyze intuitive bits of information.

The possibility of personalization interacts with the preceding discussion about the informational content of defaults (in Subsection 2 above). That discussion assumed that the default setter knows only average values. What happens if the default setter knows individual values? If a benevolent employer knows the type of each individual employee and thus the optimal contribution rate for that employee, then the efficient outcome obtains

---

[66] *See also supra* Section 0. (on The Informational Content of Default), where we discussed how the type of information that the policymaker has – information on expected values of personalized information – affects the inferences that individuals will draw from the chosen default option. Personalization also relates to our Implementation discussion (where we consider the information that policymakers need to implement the information-costs theory). *See supra* Part 0.

[67] Some existing retirement defaults are age dependent. *See* Jill Cornfield, *The Good and the Bad in This Easy One-Stop-Shopping Retirement Fund*, CNBC (Oct. 16, 2018, 9:34 AM), https://www.cnbc.com/2018/10/16/target-date-funds-are-super-simple-but-is-that-all-you-need-to-know.html. This is a minimal type of personalization. *See, e.g.*, Ariel Porat & Lior Jacob Strahilevitz, *Personalizing Default Rules and Disclosure with Big Data*, 112 MICH. L. REV. 1147 (2014); Cass R. Sunstein, *Deciding by Default*, 162 U. PA. L. REV. 1, 48–55 (2013).

[68] The personalized default can be presented as a recommendation. For example, the employer could tell her employee: "Based on the information that you provided (or based on the information that we have on you), we think that a 7% contribution rate is optimal for you."

without any information acquisition by employees and without any opt out. Things change, of course, if the employer is not benevolent, but rather looks after its own interests – interests that conflict with those of the employees. If employees recognize the conflict of interests, they will not draw inferences from the personalized default option. But if employees mistakenly think that the employer is benevolent, they will not acquire information and will stick to the default – to their detriment.

IV. Information Costs versus Mechanical Costs

The theory presented in Part II identified the phenomenon of uninformed opt out and began to explore its implications. This phenomenon is the central implication of the distinction between the two types of opt out costs—mechanical and informational. While much of the literature considers the two as interchangeable components in the bin labeled 'transactions costs,' their implications turn out to be different. High mechanical costs prevent opt out, whereas high information costs can have the opposite effect and increase the incidence of opt out. We saw in Part II that this difference has important implications for the design of default rules.

In this Part we begin to examine more broadly the different regulatory implications of the two types of transactions costs. The analysis examines two ways in which these costs could be regulated. First, lawmakers may attempt to engineer the magnitude of these costs, so as to achieve socially optimal outcomes. These efforts, we show, depend on the type of costs involved. Second, firms dealing with consumers may also try to manipulate the costs their counterparties incur when making opt out choices, but this time with the goal of steering people towards the firms' preferred choices. We examine what lawmakers can do to counteract these efforts, and how these measures depend on the type of cost involved.

A. Cost-Engineering by Lawmakers

1. Increasing Mechanical Costs

Increasing mechanical opt-out costs could make opt out more difficult and the default stickier. The motivation for such a policy is a reality where individuals too often agree to opt-out of protective and socially desirable defaults. Making these defaults stickier via mandates that increase mechanical costs (such as requiring more meticulous contract formation routines) can thus be welfare enhancing. Increasing mechanical opt-out costs is especially attractive when firms try to lure consumers to disclaim important protections.[69]

---

[69] *See, e.g.*, Barr, Mullainathan, and Shafir, *supra* note 4, at 8 ("Given the strong market pressures to deviate from the default offer, we would need to require more than a simple opt-out to make the default sticky enough . . . [t]hus, we propose that deviation from the offer would require heightened disclosures and additional legal exposure for lenders in order to make the default sticky.").

Our analysis suggests that mechanical opt-out costs should be increased less often than commonly intuited. We saw that when people have accurate beliefs, high mechanical opt-out costs reduce welfare, and thus increasing such costs and making defaults stickier makes no sense; it hinders efficient informed and uninformed opt-out. Only when people have inaccurate beliefs that would lead them to poorly-judged opt-outs should lawmakers create mechanical obstacles. This, of course, is not surprising. Indeed, policymakers' motivations for such interventions sometimes rest on the notion that people opt out mistakenly. For example, if people systematically underestimate the value of the default rule and agree to opt out into inferior alternatives (being prompted by firms in these directions), building road bumps against such hurried and harmful opt out would be good. 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 to opt out. In this context, sludge can be welfare-enhancing.

2. Lowering 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.

This is not to say that lower information costs are generally, or even commonly, bad. Indeed, it will generally be advisable to reduce information costs. Counterintuitively, even the potential downside of low information costs noted above can be mitigated by lower information costs. Low information costs can be harmful only when parties overestimate the value of information. Such overestimation, and indeed any misperception, will be mitigated when individuals are better informed. Thus, lower information costs reduce the misperception that makes low information costs potentially harmful. This argument is not circular. Information is not all-or-nothing; people can acquire less information or more information. Accordingly, lower information costs can induce acquisition of the first batch of information, and this information will limit any misperception that may otherwise have led to inefficient acquisition of the second batch of information.

B. Fighting Cost-Engineering by Firms

1. Negative Opt-out Costs

Firms often manipulate mechanical opt out costs to get consumers to forgo a protective default and opt-out into the firm's preferred non-default option. The firms make it easy to

opt out and mechanically painstaking to stick with the default. This is attractive to firms that are hoping to profit by selling add-ons and other non-default features that consumers would otherwise decline, or by avoiding socially valuable but costly (to the firm) consumer protections. Rather than allowing the consumer to simply say "no" to the opt-out, firms require a complicated ritual. Note that, in this second scenario, the mechanical costs of opting out of the default are reduced, not increased. The increase is in the mechanical cost of sticking to the default. The result is *negative* opt-out costs.

While policymaker's motivations to inflate mechanical costs could be desirable (when used in the right circumstances), the same cannot be said about firms' motivations. As explained, firms engineer mechanical costs to induce, rather than prevent, opt out. These are situations in which consumers want to stick with the default but firms make it artificially hard to do so. For example, consumers want to buy a standard product but firms prompt them to select the (more profitable) premium version, and nudge them to do so again and again. In some cases, firms give people only two options—"Yes" or "Not Now"—denying people the preferred choice of "No" in the hope that eventually the not-now choosers will surrender or inadvertently say "Yes." These commercially motivated nudges, or rather "sludges" or "dark patterns," increase the mechanical costs of adopting the default. Contrary to our model, where opt out was mechanically costlier than the default, sludges make opt out effortless while the preservation of the default becomes unnecessarily cumbersome. In essence, firms are automatically changing or "unclicking" the policymaker's default. They must elicit consumers' consent to these reversals, and consumers—even uninformed— might be unobliging. Even when uninformed, the consumers may regard these changes as carrying negative expected value, resist them and, at some cost, keep the legal default. As these mechanical costs increase, consumers' resistance dissipates.[70]

2. Fighting Costs with More Costs

Whether successful or not, the artificial mechanical costs imposed by such sludges are reason enough for policymakers to intervene. They could do so by strengthening the original default rule against unilateral changes by firms. This strategy results in a tug of war between good nudges and bad sludges: To combat the "bad" mechanical costs of sticking with the default, policymakers could mandate "good" mechanical costs for any opt out. But while justified, the practical value of such policies is questionable. The most typical tools policymakers use to increase mechanical opt-out costs are lengthy disclosures, educational pre-requisites, segregated agreements, clause-by-clause signatures, and periodically renewed agreements.[71] It is sometimes doubtful whether such hurdles succeed

---

[70] *See* Luguri and Strahilevitz, *supra* note 12 (demonstrating, in experimental setting, the tendency of subjects to accept the seller's induced opt out option).

[71] There are many examples where regulators deliberately increase opt-out costs. The GDPR requires more explicit consent to information collection and some European lawmakers have required renewed consent for every incidence of data collection. *See infra* Sec. III.B. *See also* OFFICE OF THE COMPTROLLER OF THE CURRENCY, DEP'T OF THE TREASURY, OVERDRAFT PROTECTION: OPT-IN REQUIREMENTS AND RELATED MARKETING ISSUES (2010) (separate disclosure and segregated assent are now necessary to enroll in overdraft protection); Loi 2016-41, *supra* note 11 (requiring the submission of form and proof of identity in order to opt out of the default). There are also many examples of firms raising, or manufacturing, opt-out costs. *See,*

in increasing opt-out costs in a manner sufficient to render the defaults stickier, especially in the presence of a firm's sludges.[72]

3. Regulating Firm-Engineered Costs

Rather than fight costs with more costs, lawmakers can regulate the costs, or sludges, that firms try to impose. The problem with this strategy is that it faces a line drawing challenge: how to distinguish sludges from the multitude of other techniques used by firms to influence consumers' choices. All of advertising, in essence, is a campaign to affect people's choices, to ignite behavioral and cognitive mechanisms, and to deluge consumers with invitations to opt in. Our theoretical framework provides a clue on how to design pinpointed intervention. One difference between legal advertising and ought-to-be-illegal sludges tracks the difference between mechanical and information costs. Much advertising operates on the information dimension and does not create mechanical impediments.[73] Because the information it carries could be productive, its regulation is and should be governed by false advertising and anti-deception laws. These laws have the dexterity to prohibit practices that increase, rather than reduce, information costs. Sludges, by contrast, are not informational; they operate within the dimension of mechanical costs. Because the increased mechanical costs they inflict are never productive, they should be prohibited. Indeed, regulating them could be a cornerstone of a new consumer anti-manipulation law.

---

e.g., Willis, *supra* note 23, at 1165, 1171 (2013) ("When the choice to opt out of a default is not made plain, people may perceive a default as unchangeable"; "[f]irms actively work to increase the power of their defaults using […] transaction barriers"; "firms with automatically renewing subscriptions that consumers can sign up for in minutes online may require spending an hour on hold with costumer service to cancel"; "[f]irms stymie consumers who might attempt to opt out, using fine print…"); Jeff Sovern, *Opting In, Opting Out, or No Options At All: The Fight for Control of Personal Information*, 74 WASH. L. REV. 1033, 1083, 1089 (1999) ("[C]ompanies that offer opt-outs have an incentive to increase the transaction costs incurred by consumers who opt out"; "provides subscribers with a lengthy, dull, and difficult-to-read statement of their rights and requires subscribers wishing to opt out to communicate their intent in a separate writing."); Nakashima *supra* note 11 (an example where Google increased opt-out costs to make its privacy settings stickier); Sumit Agarwal et al., *Do Financial Counseling Mandates Improve Mortgage Choice and performance? Evidence from a Legislative Experiment* 32 (Federal Reserve Bank of Chicago, Working Paper No. 2009-07, 2009) ("Those who were required to attend counseling … tended to not walk away from the original offer following counseling and reapply … which would have required another counseling session").

[72] *Compare* Ian Ayres, *Regulating Opt-Out: An Economic Theory of Altering Rules*, 121 YALE L.J. 2032, 2093 (2012) (arguing that in some cases policymakers should use "impeding altering rules", i.e. with high opt-out costs, to reduce opt-out rates; Ayres notes that such "impeding altering rules" should be used when people overestimate the benefit from opting out.)

[73] While a large subset of advertising is informational, we recognize that another large subset of advertising seeks to shape preferences or stir emotions. *See* Kyle Bagwell, *The Economic Analysis of Advertising*, in Mark Armstrong and Robert Porter (eds.), HANDBOOK OF INDUSTRIAL ORGANIZATION, Vol. 3, Ch. 28, pp. 1701-1844 (2007).

Electronic copy available at: https://ssrn.com/abstract=3582129

V. Applications

As noted in the Introduction, default rules are increasingly being used across diverse policy domains. Lawmakers are increasingly aware that default rules can be readily disclaimed, and are thus working to design stickier defaults. High hopes for better social outcomes have been hanging on this technique. In this Part, we consider, in some detail, several important applications of regulation by default rules, and highlight the different ways in which our information-costs theory informs these applications. Section A considers the regulation of overdraft fees. Section B considers the protection of privacy in digital information. Section C returns to the canonical example of retirement savings. And Section D discusses "green" defaults.

A. Overdraft Protection

Debit card holders who do not have sufficient funds in their checking account are able to complete debit transactions by borrowing from the bank. To do so, they have to enroll in the bank's overdraft protection plan. Until 2010, the law allowed banks to automatically enroll their checking account customers; overdraft protection was the default. This policy came under scrutiny, because card holders were charged high fees any time they borrowed via overdraft withdrawal, and banks were collecting many billions of dollars, mostly from low-income customers.[74]

Recognizing that many people are either able to receive short term credit more cheaply elsewhere, or have learned to regret the costly overdraft fees, in 2010, the Federal Reserve reformed the law. Seeking to reduce the prevalence of overdraft transactions, the Fed reversed the default. The previous overdraft protection default (auto-enrollment unless consumers opt out) was replaced by a no-overdraft default (and express opt-in was required for overdraft protection).[75] The new default was intended to be sticky, requiring more mechanical effort to disclaim it. Separate disclosure and segregated assent were now necessary to make enrollment effective.[76]

The purpose of the new default was to prevent unsophisticated consumers from incurring high overdraft fees, in the hope that only those who truly needed this exceptional measure

---

[74] *See, e.g.*, Peter Smith, *Report: FDIC Data Shows that Banks Collected $11.45 Billion in Overdraft Fees in 2017*, CENTER FOR RESPONSIBLE LENDING (Aug. 7, 2018), https://www.responsiblelending.org/media/report-fdic-data-shows-banks-collected-1145-billion-overdraft-fees-2017 (reporting on regarding only the ten largest banks in the country); Julie Chang, Americans Paid *$34 Billion in Overdraft Fees Last Year. Here's How to Stop the Charges*, FORBES (April 5, 2018, 9:00 AM), https://www.forbes.com/sites/learnvest/2018/04/05/americans-paid-34-billion-in-overdraft-fees-last-year-heres-how-to-stop-the-charges/#7981b7703ce9; Jackie Wattles, *Americans Paid $15 Billion in Overdraft Fees Last Year CFPB Says*, CNN (Aug. 4, 2017, 6:29 PM), https://money.cnn.com/2017/08/04/pf/overdraft-fees-cfpb/index.html.

[75] 12 CFR § 205.17(c).

[76] OFFICE OF THE COMPTROLLER OF THE CURRENCY, DEP'T OF THE TREASURY, OVERDRAFT PROTECTION: OPT-IN REQUIREMENTS AND RELATED MARKETING ISSUES (2010).

would knowingly and sparingly use it. This seemed like a perfect environment to use sticky defaults, which protect vulnerable consumers from high overdraft fees, while allowing those consumers who truly need a different regime to opt out.

Our analysis questions the ability of the new opt-in default to achieve its stated goal. As in our model, it is useful to think of banking customers as consisting of two types: (1) a majority who would overdraft rarely and thus gain a small benefit from overdraft protection; and (2) a minority who would overdraft frequently and suffer a large loss from the multiple, high overdraft fees. Indeed, evidence shows that the high overdraft fees, between $30-$35 for each overdraft, were incurred by a minority of consumers. As explained by the Consumer Financial Protection Bureau (CFPB): "In a given year, only 30% of consumers overdraw their checking account. The 8% of consumers who overdraft more than 10 times per year pay 74% of overdraft fees. These consumers are charged $380 in overdraft fees on average annually." [77]

If consumers are uninformed about their type, how would they assess the expected value of overdraft protection? A consumer with unbiased uninformed beliefs would likely prefer no overdraft protection – the small benefit enjoyed by Type 1s is outweighed by the large loss incurred by Type 2s (even if the chance of being Type 1 is larger).[78] If most consumers were indeed uninformed, but unbiased, then the purpose of the 2010 default switch would

---

[77] CONSUMER FINANCIAL PROTECTION BUREAU (CFPB), *A Closer Look: Overdraft and the Impact of Opting-In*, January 19, 2017 (https://files.consumerfinance.gov/f/documents/201701_cfpb_Overdraft-and-Impact-of-Opting-In.pdf) (last visited on Nov. 25, 2019). The significant loss that Type 2 consumers incur is evident when the high overdraft fees are compared to the small average charge that triggers this fee. *See* CFPB, *CFPB Finds Small Debit Purchases Lead to Expensive Overdraft Charges*, July 31, 2014 (https://www.consumerfinance.gov/about-us/newsroom/cfpb-finds-small-debit-purchases-lead-to-expensive-overdraft-charges/) (last visited Nov. 25, 2019) ("The study found that the majority of debit card overdraft fees are incurred on transactions of $24 or less.")

[78] There is some evidence that overdrafting, when the cost (fee or interest rate) is so high, is harmful to consumers. *See* Paul Adams et al., *Time to Act: A Field Experiment on Overdraft Alerts* (Financial Conduct Authority Occasional Paper 40, 2018) (finding that overdraft alerts reduce overdrafting by 21-25%, suggesting that many consumers, when they are made aware of the overdraft decision, choose not to overdraft). Moreover, the cost of "unarranged overdrafts" in the UK are smaller than overdraft fees in the US, suggesting that the effect in the US would be larger than the 21-25% figures. *Compare* Rupert Jones, *Overdrafts: Can You Cut the Cost of Yours?*, GUARDIAN (May 5, 2018), https://www.theguardian.com/money/2018/may/05/cut-cost-overdraft-uk-banks (describing UK overdraft fees as around seven pounds) with Chang, *supra* note 74, (2018) (describing the average US overdraft fee as around thirty dollars). *See also* OFFICE OF FAIR TRADING, *Personal Current Accounts in the UK: An OFT Market Study*, at 62 (2008) ("evidence from several recent market investigations suggests that some of these incidental charges could have been avoided if consumers had been aware of their balance."); Alasdair Smith et al., *Retail Banking Market Investigation*, COMPETITION AND MARKETS AUTHORITY, at 536 (2016) ("This suggests that a significant proportion of customer detriment experience by overdraft users . . . may arise from lack of awareness and engagement with their [personal current accounts]."); Andrea Caflisch et al., *Sending out an SMS: The impact of automatically enrolling consumers into overdraft alerts*, (Financial Conduct Authority, Occasional Paper 36, 2018) (finding that "[a]utomatic enrolment into unpaid item alerts reduces charges by 21-24%" and that "[a]utomatic enrolment into unarranged overdraft alerts reduces charges by 25%"); Stefan Hunt et al., *Message Received? The Impact of Annual Summaries, Text Alerts and Mobile Apps on Consumer Banking Behavior* (Financial Conduct Authority, Occasional Paper 10, 2015) ("signing up to text alerts or mobile banking apps reduces the amount of unarranged overdraft charges incurred by 5% to 8%, and signing up to both services has an additional effect, resulting in a total reduction of 24%.").

have been to save the costs of uninformed opt-out. Or, if these costs were high, to save consumers from a harmful arrangement – overdraft protection – imposed by banks.

This account is challenged by evidence about the aftermath of the 2010 reform. With unbiased beliefs, we would have expected minimal opt-out from the post-2010, no-overdraft-protection default. And yet many consumers opted out. In particular, 45% of the frequent overdrafters opted out.[79] Why? It is possible that some of these consumers were engaging in informed opted out, having learned from past experience that overdraft protection is beneficial, despite the high fees. But it is also possible, indeed likely, that the observed opt-out was largely uninformed and, moreover, based on false, uninformed beliefs – that overdraft protection is a good deal.

These inaccurate beliefs were sustained by banks' marketing efforts. The "overdraft protection" label itself suggests a benefit, and the arrangement is promoted as a "free" perk that allows the customer "to enjoy peace of mind"—namely, the option to make debit purchases even with a zero balance.[80] Banks highlight the upside (avoiding declined transactions), not the downside (high fees). And uninformed consumers, including those who would ultimately incur multiple overdraft fees, opt out of the no-overdraft default, often to their detriment. The new default is not as sticky as the Fed hoped.

We bring this example because a more general lesson can be learned from it. It is difficult to change outcomes for consumers without addressing the uninformed opt-out phenomenon, especially when it is fueled by inaccurate beliefs. Most attempts by regulators to make a default sticky focus on the wrong method: making the mechanical costs of opt-out higher. In the overdraft regulation, this increased cost amounted to an additional disclosure-and-signature. These attempts fail because even with costlier mechanics, opt-out remains easy; especially when the firm on the other side is motivated to make it so.

Lawmakers could, instead, try to affect people's uninformed beliefs, so that uninformed consumers would learn to prefer the socially targeted outcome. But lawmakers could educate people only on so many issues, and, besides, their attempts to influence uninformed beliefs would need to overcome the industry's own marketing campaigns. Lawmakers could also try to lower people's information costs to help them acquire information about their individual type, or force firms to disclose such type-specific information.[81] Here, too, we might worry that the corrective policy would disproportionately affect the more educated consumers and might not prevent the irrational uninformed opt-out by others.

---

[79] For frequent overdrafters, the group that policymakers were most concerned about, the opt-out rate is 45%. (The overall opt-out rate is 16%, and for new accounts it is 22%.) *See* CFPB, *CFPB Study of Overdraft Programs*, at 29-30 (June 2013), https://files.consumerfinance.gov/f/201306_cfpb_whitepaper_overdraft-practices.pdf (last visited on Nov. 25, 2019).

[80] *See* Willis (2013), *supra* note 23, at 1191-92 ("In their communications with consumers, banks refer to opting out of the policy default as 'opting in' to a bank's "overdraft service'. Thus, opting out of the default is framed as gaining a service rather than losing an endowed reference position"; "In their marketing, banks explicitly invoked loss aversion to encourage opting out with copy such as 'Don't lose your ATM and Debit Card Overdraft Protection' and 'STAY PROTECTED with [] ATM and Debit Card Overdraft Coverage.'").

[81] *See, e.g.*, BAR-GILL, *supra* note 45.

B. Privacy

The basic default rule in many jurisdictions does not allow firms and digital platforms to collect, use, and share the large quantities of personal information that many companies rely on. People must consent to any opt-out from that default, and such consent is often solicited through unread fine print. For example, under federal wiretap laws, Google may not scan the text of its users' email messages, unless the users agree.[82] And all users unknowingly "agree." Similarly, the EU's General Data Protection Regulation (GDPR) and California's Consumer Privacy Act establish a default that prohibits collection of some categories of personal information. And, again, most users opt out of this default without an informed understanding of the implications of such opt out. How to design and police such consent-based information collection has been the subject of much debate in privacy law and of ongoing legal reform. While a possible regulatory approach would be to outlaw the collection or sharing of some personal data, the far more common technique is to redesign the default rules and the process of opt-out.

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.[83] The mechanics of such opt-outs have been designed by firms to be so easy and cheap that opt-out is achieved smoothly. This was uninformed opt-out at the extreme. In our model, we regarded uninformed opt-out as a deliberate choice based on imperfect information. In the privacy context, opt-out was uninformed in the sense that people were not even aware that they are opting out. In this context, beliefs about the relative (expected) payoffs of the high-privacy vs. low-privacy options became less relevant. Alternatively, people believed that the payoff difference between high- vs. low-privacy does not merit the attention to firms' opting-out practices.

Then things began to change. Facebook's lax data-sharing practices were revealed after the Cambridge Analytica fiasco. This, together with a series of massive data security breaches,

---

[82] 25 U.S.C. 2511.

[83] *See* Willis (2014), *supra* note 71 ("While a Don't-Track-Me setting would require firms to spend significant resources on maneuvering consumers out of the default, firms determined to do so could be successful"); Allyson W. Haynes, *Online Privacy Policies: Contracting Away Control over Personal Information*, 111 PENN ST. L. REV. 587, 617 (2007) ("[P]rivacy policies are often presented in terms of browsewrap. Users are deemed to have agreed to them simply by being on the website or by disclosing information."); Mark A. Lemley, *Terms of Use*, 91 MINN. L. REV. 459, 460 (2006) ("[A]n increasing number of courts have enforced "browsewrap" licenses, in which the user does not see the contract at all but in which the license terms provide that using a Web site constitutes an agreement"); Paul M. Schwartz, *Privacy Inalienability and the Regulation of Spyware*, 20 BERKELEY TECH. L.J. 1269, 1274 (2005) ("[M]any data-processing institutions are likely to be good at obtaining consent on their terms regardless of whether the default requires consumers to authorize or preclude information-sharing."); Daniel J. Solove, *Introduction: Privacy Self-Management and the Consent Dilemma*, 126 HARV. L. REV. 1880, 1898–99 (2013) ("[O]rganizations will have the sophistication and motivation to find ways to generate high opt-in rates … a requirement of affirmative consent for most new uses of data will likely lead to more buttons to click and more forms to sign, but not to more meaningful privacy protection.").

elevated the salience of data collection and its potential harms.[84] The enactment of the European data regulations—the GDPR—further heightened the public's awareness.[85] People's beliefs have shifted to some degree, with more people noticing that a choice is being made, forming more deliberate beliefs about the default and non-default options, and perhaps increasingly believing that the no-information-collection regime is superior. By and large, these are is still uninformed beliefs, because people need more information to know how they are personally impacted.

Recent legal reforms have begun to address the process of opt-out. Lawmakers have taken actions to increase mechanical opt-out costs, and also to reduce information costs. Mechanical costs were increased by requiring more explicit consent to information collection, in the hope of making it more difficult for firms to induce opt-out.[86] In some

---

[84] *See, e.g.*, Julie Beck, *People are Changing the Way They Use Social Media*, THE ATLANTIC (June 7, 2018), https://www.theatlantic.com/technology/archive/2018/06/did-cambridge-analytica-actually-change-facebook-users-behavior/562154/ (suggesting that breaches like Cambridge Analytica led users to share less detailed information online); Andrew Perrin, *Americans are Changing Their Relationship with Facebook*, PEW RESEARCH CENTER (Sept. 5, 2018), https://www.pewresearch.org/fact-tank/2018/09/05/americans-are-changing-their-relationship-with-facebook/ (finding, in a survey months after the Cambridge Analytica story broke, that 54% of Facebook users had adjusted their privacy settings in the past year); *see also* Kim Hart & Ina Fried, *Exclusive Poll: Facebook Favorability Plunges*, AXIOS (Mar. 26, 2018), https://www.axios.com/exclusive-poll-facebook-favorability-plunges-1522057235-b1fa31db-e646-4413-a273-95d3387da4f2.html (showing that Facebook's favorability had fallen 28 points in the wake of the revelations regarding Cambridge Analytica).

[85] *See, e.g.*, European Commission Press Release IP/19/2956, Data Protection Regulation One Year on: 73% of Europeans Have Heard of At Least One of Their Rights (June 13, 2019).

[86] *See* GDPR, recital 32 ("Consent should be given by a clear affirmative act establishing a freely given, specific, informed and unambiguous indication of the data subject's agreement to the processing of personal data relating to him or her, such as by a written statement, including by electronic means, or an oral statement … Silence, pre-ticked boxes or inactivity should not therefore constitute consent ... When the processing has multiple purposes, consent should be given for all of them."). GDPR creates a range of opt-out costs depending on the type of information collected. There are two main types of consent (which can be further varied by each EU member state): "explicit" consent (as defined in article 9), which applies to sensitive information (health, sexual orientation, politics, etc.); and "unambiguous" (implied) consent (as defined in articles 4 and 7), which applies to all other information. Explicit consent imposes higher opt-out costs, because the user must be given notice of the purpose and type of information collected and the user must explicitly assent to them. Unambiguous consent imposes lower opt-out costs (e.g., using a site multiple times after agreeing to cookies clears the hurdle, or agreeing by submitting an email address, etc.). *See Explicit vs. Unambiguous Consent*, DATA STREAMS.IO (Oct. 11, 2017), https://www.datastreams.io/explicit-vs-unambiguous-consent-whats-the-difference/.
[On January 21st, 2019, by force of the GDPR, the French National Data Protection Commission imposed a fine of 50 million euros on Google LLC, due to several breaches, one of them being the unlawful acquisition of consent to the processing of personal data for personalized advertisement. *See* NATIONAL DATA PROTECTION COMMISSION, *The CNIL's Restricted Committee Imposes a Financial Penalty of 50 Million Euros Against Google LLC* (Jan. 21, 2019), https://www.cnil.fr/en/cnils-restricted-committee-imposes-financial-penalty-50-million-euros-against-google-llc.]
The new California Consumer Privacy Act, recognizing the reality of wholesale opt-out by Privacy Policy, makes it easier for consumers to opt back into the no collection default by requiring that firms add a "Do Not Sell My Personal Information" link to their websites. *See* OFFICE OF THE ATT'Y GENERAL, CAL. DEP'T OF JUST., CALIFORNIA CONSUMER PRIVACY ACT (CCPA) FACT SHEET (2019).

cases, lawmakers have required renewed consent for every incidence of data collection.[87] Information costs were reduced through mandates requiring explanations in simple language and easy-to-use privacy tools, so that people interested in making an informed choice could do so.[88]

Our analysis sheds light on these reforms. Again, it is useful to think of users as consisting of two types. High-harm types are more sensitive to data privacy concerns than low-harm types. Some consumers acquire information and, if they discover that they are high-harm types, make informed decisions to preserve the legal default (which is often not easy, when opt-out costs are negative and a choice to maintain the default entails deliberate and careful rejection of repeat invitations to opt out). But how do the uninformed consumers behave? It is possible that the public anger towards some data platforms persuaded a fraction of the consumer body to change its uninformed behavior and to stick with, rather than opt out of, the privacy-protective default. This could be interpreted as a shift towards more accurate uninformed beliefs (or, rather, to less accurate uninformed beliefs, if the privacy costs are relatively small). Many, however, continue to take the path of least resistance charted by firms and opt out of the privacy-protective default. They incur slightly higher mechanical costs, which create some "annoyance,"[89] but not enough annoyance to make the default sticky. (Indeed, sticking with the default, when firms repeatedly invite opt-out can be more annoying.)

Without a better empirical sense which default maximizes expected welfare, it is hard to interpret which outcome is desirable. Uninformed opt-out may be privately optimal if private harms from data collection are small and private benefits large. If that is the case, those who stick with the default while uninformed are over-reacting to the public outcry. But it is also possible that uninformed opt-out is welfare-reducing, and users agreeing to it are misjudging the harms that surrendering their data would ultimately cause.[90]

Unlike the overdraft regulation, which only required additional mechanical rituals to opt out, privacy laws are trying to reduce uninformed opt-out by also targeting information

---

[87] *See, e.g.*, Commission Nationale de l'Informatique et des Libertés, Délibération no. 2019-03 (July 4, 2019) (providing the French government's guidelines for compliance with the GDPR and requiring renewed consent after 13 months).

[88] *See* GDPR, ch. II, art. 7, § 2 (" . . . the request for consent shall be presented in a manner which is clearly distinguishable from the other matters, in an intelligible and easily accessible form, using clear and plain language."). Following the enactment of the GDPR, the California State Legislature passed a similar bill to enhance privacy protection, that includes mandates aimed at facilitating consumers' understanding of contractual terms regarding the collection and usage of information by firms. *See* California Consumer Privacy Act of 2018, Cal. Civ. Code § 1798.100 (effective January 1, 2020).

[89] *See, e.g.*, Jack Schofield, *What Should I Do About All The GDPR Pop-ups on Websites?*, Guardian (July 5, 2018, 11:01 AM), https://www.theguardian.com/technology/askjack/2018/jul/05/what-should-i-do-about-all-the-gdpr-pop-ups-on-websites.

[90] The Stigler Center Report advocates for "consumertarian defaults" and high opt-out costs or opt-out procedures that would require firms to convince, rather than trick, consumers to opt out. It is hard to believe that firms will provide individualized, type information. Rather firms will try to influence uninformed beliefs. If accurate uninformed beliefs support the consumertarian default, there is a real concern that firms will promote inaccurate beliefs and induce inefficient, uninformed opt-out. If accurate uninformed beliefs do not support the consumertarian default, then is the default really consumertarian?

costs. If lawmakers succeed in reducing information costs, some uninformed action will be replaced with informed choice. Ideally, people will self-select according to their type. It is possible, however, as our model showed, that some of the uninformed resistance to data sharing would subside and as a result there will be *more* opt-out.

C. Retirement Savings

The retirement savings defaults have featured as a canonical example for the power of default rules to change behavior. A large empirical literature demonstrated that the auto-enrollment defaults stick,[91] and this evidence inspired a search for theoretical explanations. The behavioral economics literature viewed the evidence as consistent with decisionmakers' cognitive limitations. Based on that evidence and the behavioral interpretation, commentators called for using default rules as a solution to many social problems.[92]

Our information-costs theory sheds new light on emerging evidence of opt-out behavior in the retirement savings context. We offer a different framework to understand the empirically observed stickiness. If the mechanical costs of opt-out are low, what explains this stickiness? Why do we not observe uninformed opt-out from some defaults? This area of contracting differs from many others by the role that the informed party—here, the employer—plays. First, unlike many other contexts, the employer does not have a strong interest in the content of the default and does not actively propel employees towards, or away from, a specific arrangement. The firm-induced uninformed opt-out that takes place in the overdraft and privacy contexts does not happen here. Second, retirement contribution defaults may be sticky if employees attribute an informational signal to the default presented to them—an endorsement effect. Many employees trust that their employer is looking after their best interests and adjust their uninformed beliefs based on the default that the employer chooses. While much of the literature regards this endorsement effect as a socially desirable information inference, recent work notes the potential conflicts of interests between employers and employees and suggests that the employer's default could bias employees' uninformed beliefs and lead to inefficient decisions to stick with the default.[93]

Retirement savings defaults are somewhat sticky.[94] But even in this archetypal example of sticky defaults, there is mounting evidence of wholesale opt out—even in this context some defaults are quite slippery, especially over time. The traditional zero contribution rate (or no enrollment) default may have stuck for the short term, but over time it was relatively slippery – approximately 60% of employees opted out.[95] A similar percentage of

---

[91] *See, e.g.*, Beshears et al., *supra* note 10.

[92] *See* sources cited *supra* notes 3-4.

[93] *See* Ryan Bubb & Patrick L. Warren, *An Equilibrium Theory of Retirement Plan Design*, AM. ECON. J.: ECON. POL'Y (forthcoming).

[94] Brigitte C. Madrian & Dennis F. Shea, *The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior*, 116 Q. J. ECON. 1149 (2001).

[95] *See, e.g.*, William E. Nessmith, Stephen P. Utkus, & Jean A. Young, *Measuring the Effectiveness of Automatic Enrollment*, 31 Vanguard Center for Retirement Research, 1, 10 (2007) (showing that voluntary

employees opted out, over time, from the common 3% contribution rate default.[96] Perhaps these were stop-and-think defaults encouraging some employees to acquire information. For most employees, however, the opt out was likely uninformed (or largely uninformed). Employees just had a sense that 0% or 3% is too low and opted for something higher.

But even if retirement-savings defaults are not as sticky as commonly believed, they have still proven more resilient to uninformed opt-out, as compared to many other defaults. This relative stickiness has inspired support for default rules as a consumer protection technique in many other contexts. Such uncritical borrowing from the retirement-savings contexts is perilous. In other contexts, the informational structure is significantly different. It is not clear whether the lawmaker's default contains the same informational content as the employer's default; lawmakers may be looking out for other groups and can be motivated by political pressure and popular sentiment.[97] The endorsement effect may thus be weaker. And in consumer markets, even if the lawmaker's default is pro-consumer and should benefit from an endorsement effect, sellers try hard to shift uninformed beliefs away from the default and towards the seller's preferred option.[98] Overall, the stickiness observed in the retirement savings context is probably not representative.

D. "Green" Defaults

Default rules have been extensively used to encourage environmentally-friendly ("green") outcomes, such as purchasing electricity from clean, renewable sources, using energy-efficient light bulbs, enduring a lower temperature on the office thermostat, and utilizing double-sided printing.[99] While there are some success stories, a recent meta-analysis

---

enrollment increased from 32% to 59% over the course of the first three years of employment); Vanguard, *How America Saves* (2019), p. 35 (the participation rate in voluntary enrollment, i.e., zero default, plans was 60%).

[96] *See, e.g.,* Nessmith et al, *id*., p. 11 ("after 30 months, 57% of the employees hired under automatic enrollment [] have a rate higher than the default." [about half of the automatic enrollment plans featured automatic annual increases in the contribution rate; we are assuming that the "rate higher than the default" finding means higher than the annually-adjusted default]). On the predominant contribution rate default of 3%, *see id*. at 6 ("The median contribution rate in automatic enrollment designs is 2.9%."); Vanguard, *id*., at 40 (noting the predominant 3% default).

[97] This is not to say that the lawmaker's default can never have an informational effect. In the organ donations context, it has been argued that the lawmaker's default contains information about social norms. *See, e.g.*, Shai Davidai, Thomas Gilovich, & Lee D. Ross, *The Meaning of Default Options for Potential Organ Donors*, 109 PROCEEDINGS OF THE NATIONAL ACADEMY OF SCIENCES (PNAS) 15201 (2012); *see generally* Kenworthey Bilz & Janice Nadler, *Law Moral Attitudes, and Behavioral Change*, *in* THE OXFORD HANDBOOK OF BEHAVIORAL ECONOMICS AND THE LAW 241 (Eyal Zamir & Doron Teichman eds., 2014).

[98] *See* Michael S. Barr, Sendhil Mullainathan, & Eldar Shafir, *The Case for Behaviorally Informed Regulation*, 25 NEW PERSPECTIVES ON REG. 41 (2009) (contrasting the employment and consumer context and arguing that market forces make defaults slippery, without emphasizing the information effects).

[99] *See* Zachary Brown et al., *Testing the Effect of Defaults on the Thermostat Settings of OECD Employees*, 39 ENERGY ECON. 128 (2013) (green thermostat); Isaac Dinner et al., *Partitioning Default Effects: Why People Choose Not to Choose*, 17 J. EXPERIMENTAL PSYCHOL. 332 (2011) (green lightbulbs); Felix Ebeling & Sebastian Lotz, *Domestic Uptake of Green Energy Promoted by Opt-Out Tariffs*, 5 NATURE CLIMATE CHANGE 868 (2015) (green electricity); Johan Egebark & Mathias Ekström, *Can Indifference Make the World Greener?*, 76 J. ENVTL. ECON. AND MGMT. 1 (2016) (green printing); Simon Hedlin & Cass R. Sunstein, *Does Active Choosing Promote Green Energy Use? Experimental Evidence*, 43 ECOLOGY LQ 107 (2016)

suggests a relatively small effect of defaults in the environmental context.[100] The information-cost theory sheds light on these green defaults and helps predict when default rules are more or less likely to affect outcomes.

Two studies are particularly instructive, and particularly susceptible to an information-costs analysis. In a field experiment conducted by Araña and León (2013), subjects were attendees of different academic conference and conventions across disciplines.[101] These subjects were asked whether they would like to pay to offset the carbon emissions caused by their travel to the conference. For some subjects the default option was to pay the carbon offset and they had to actively decline to avoid payment. Other subjects faced the opposite default, and had to actively choose the carbon offset payment. The effect of the default choice was significant, at least when the carbon offset payment was relatively modest. For example, when the carbon offset payment was 10 EUR, a switch to the green default increased the participation rate (in the carbon offset program) from 62% to 81%.[102] The results of this study are most informative, when compared to a very similar study by Löfgren et al. (2012). The only difference was that, in Löfgren et al. (2012), subjects were attendees of a single academic conference – on environmental economics.[103] In this study, the default choice had no statistically significant effect. (The participation rate, for a 10 EUR carbon offset payment, was approximately 50%.)

The information-cost theory can help reconcile these studies. The theory tells us that an informed party is less likely to be affected by the chosen default. If I know the environmental costs of travel to the conference, then I would participate in the carbon offset program, or not, depending on how much I care about the environment. The default would not have a large effect. In contrast, if I am uninformed and thus unsure about the environmental cost of travel, the default would have a stronger effect: I might glean information from the default (an endorsement effect). Or, I might simply stick to the default, because my weaker "average" preferences – for or against the carbon offset – are insufficient to overcome the opt-out cost. If the goal is to promote an environmental goal using a green default, the strategy is more likely to succeed when the decision makers are uninformed.

(green electricity); Daniel Pichert & Konstantinos V. Katsikopoulos, *Green Defaults: Information Presentation and Pro-Environmental Behaviour*, 28 J. ENVTL. PSYCHOL. 63 (2008) (green electricity).

[100] Jon M. Jachimowicz et al., *When and Why Defaults Influence Decisions: A Meta-Analysis of Default Effects*, 3 BEHAVIORAL PUBLIC POLICY 159, 162 (2019).

[101] Jorge E. Araña & Carmelo J. León, *Can Defaults Save the Climate? Evidence From a Field Experiment on Carbon Offsetting Programs*, 54 ENVTL. AND RESOURCE ECON. 613, 616–17 (2013).

[102] *Id*., at 619.

[103] Åsa Löfgren et al., *Are Experienced People Affected by a Pre-Set Default Option—Results From a Field Experiment*, 63 J. ENVTL. ECON. AND MGMT. 66, 66 (2012).

VI. Conclusion

This Article develops a new information-costs theory of default rules, and uses this theory to help policymakers choose the best default option, but also to evaluate the limits of regulation-by-default. The prescriptions derived from the information-costs theory rely on various assumptions, and we cannot end without discussing the realism of these assumptions and the applicability of the theory.

The theory assumes that people have less than full information about the value of the default rules, and thus behave on the basis of their uninformed expectations regarding the average values of the default and non-default options. This raises two related questions. First, is it realistic to assume that people make decisions based on average values? And second, do lawmakers have the information necessary to effectively use our information-costs theory?

We recognize that people often lack information about average payoffs under each potential default. Indeed, the assumption that people know average values should not be taken literally. All we need to assume is that, before any information is acquired, people form some estimate 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 separate assumption underlying our analysis applies to information that lawmakers have. In order to set the optimal default, lawmakers need various types of information, primarily regarding the preferences of groups of people, but also regarding the information costs that people have. Consider first the case where information costs are clearly low enough that most people will choose to become informed. This is the scenario assumed by most traditional accounts of default choice. In this scenario, lawmakers have to set the default that most people would prefer – a majoritarian default. For this, they have to know which option is favored by a majority of people.

By contrast, in the case where information costs are sufficiently high, lawmakers have to set a default with the highest expected value. For this, they need information about expected values (the same information that uninformed people have). This becomes more complicated when individuals hold inaccurate uninformed beliefs, because lawmakers now have to identify the default with the highest *perceived* expected value, namely, they need to have some sense of the direction and magnitude of people's misperceptions. Things become significantly easier, if people derive information from the content of the default. Then, lawmakers face a lesser informational burden, as they may choose the option with the highest expected value, knowing that many people will use the chosen default to correct their misperception.

The hard case, in terms of informational demands on the policymaker, is the case where information costs are intermediate. Here, information will be acquired under one default rule, but not another; or by some individuals, but not others. To assess whether people will

acquire information given a specific default rule, the policymaker needs to know the value of information, I_{Low} or I_{High} in our example. To calculate the value of information, the policymaker needs the same information that an uninformed individual has. (Recall, the uninformed individual calculates the value of information and thus decides whether to become informed.) And, when individuals hold inaccurate uninformed beliefs, the policymaker needs to know the *perceived* value of information. The policymaker also needs to know the distribution of information costs in the population. Or, at least, she needs to know for how many people the cost of becoming informed is smaller than the value of information and for how many the cost of becoming informed is larger than the value of information.

Finally, if people draw inferences from the content of the chosen default (endorsement effect), then policymakers need to know whether or not they are trusted. If people trust the policymaker, then the policymaker can use the choice of default to inform people. And, as we have seen, this allows the policymaker to achieve higher welfare levels, correcting misperceptions rather than accommodating them.

In some cases, lawmakers will have the kind of information that our model requires them to know in order to design optimal defaults. In other cases, they won't. More generally, there are valuable insights from the model that could inform regulatory design without the need to rely on complex information. Our key insight—that uninformed opt-out makes defaults less sticky than otherwise assumed—should help lawmakers avoid regulatory failures. We showed that lawmakers rush to endorse regulation-by-default, in the hope that good outcomes would ensue when these new defaults stick. The most important information that lawmakers need to have is that these hopes are over-optimistic. We therefore urge lawmakers to exercise more caution before relying on the stickiness of defaults.

44

Appendix

[Note to the Editor: This Appendix can be published in print with the main text or as an online Appendix. Or it can be removed altogether. We leave it up to you.]

The Appendix generalizes and extends 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 will analyze choices and welfare for different opt-out cost levels, k. A more general model would assume that k is distributed across contracting pairs according to $F(\cdot)$ and derive expected welfare levels based on this distribution. Since our focus is on information costs and not on opt-out costs, this more general framework is not needed for our purposes.)

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 $G(\cdot)$ and the density function $g(\cdot)$. (The distribution of information costs, x, is the same for both types.) 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 $x = 0$.) [104] We assume that uninformed individuals hold accurate beliefs about the share α and

---

[104] 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 $\alpha(1 - G(\hat{x}))$ type 1s and the $(1 - \alpha)(1 - G(\hat{x}))$ 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

about the parameters B and C. The implications of inaccurate beliefs are explored in Section E below.

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 $\alpha B - (1 - \alpha)C \geq 0$.[105]

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, for any $k \geq B$, $W = 0$.

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

$\alpha(1 - G(\hat{x}))$ who are type 2 but are uninformed about their type; and Group 4 with measure $(1 - \alpha)(1 - G(\hat{x}))$ who are type 1 but are uninformed about their type.

[105] 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 $\tilde{\alpha}\tilde{B} - (1 - \tilde{\alpha})\tilde{C} \geq 0$.

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. If the individual becomes informed, her expected payoff is: $\alpha(B - k) - x$. 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 $k \leq \alpha B - (1 - \alpha)C$:

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

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

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

47

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:

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

50

When information costs are high (i.e., above $(1 - \alpha)(C + k)$), 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 (i.e., below $(1 - \alpha)(C - k)$), 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 costs are intermediate, such that information is acquired with Low default, but not with High default ($x \in \left((1 - \alpha) \cdot (C - k), (1 - \alpha) \cdot (C + k)\right)$, 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 prefer the default rule that maximizes

51

expected value and thus tracks the preferences of the uninformed majority – to avoid costly uninformed opt-out. And Part (b) instructs the policymaker to follow the preferences of the informed majority 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.

52

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.