University of Chicago Law School

Chicago Unbound

Articles

Scholarship

2021

Rethinking Nudge: An Information-Costs Theory of Default Rules

Oren Bar-Gill

Omri Ben-Shahar

Follow this and additional works at: https://chicagounbound.uchicago.edu/journal_articles

Part of the [Law Commons](#)

Recommended Citation

Oren Bar-Gill & Omri Ben-Shahar, "Rethinking Nudge: An Information-Costs Theory of Default Rules," 88 University of Chicago Law Review 531 (2021).

This Article is brought to you for free and open access by the Scholarship at Chicago Unbound. It has been accepted for inclusion in Articles by an authorized administrator of Chicago Unbound. For more information, please contact unbound@law.uchicago.edu.

Rethinking Nudge: An Information-Costs Theory of Default Rules

Oren Bar-Gill[†] and Omri Ben-Shahar^{††}

Policy makers and scholars—both lawyers and economists—have long pondered 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 donation, 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 nondefault 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 policy makers.

[†] William J. Friedman and Alicia Townsend Friedman Professor of Law and Economics, Harvard Law School.

^{††} Leo and Eileen Herzel Professor of Law, Kearney Director of the Coase-Sandor Institute for Law and Economics, The University of Chicago Law School. For helpful comments and suggestions, we thank Matthew Adler, Mireia Artigot i Golobardes, Ian Ayres, Lucian Bebchuk, Hanoah Dagan, John Donohue, Avihay Dorfman, Abigail Faust, Rosa Ferrer, Michael Frakes, Juan-José Ganuza, John Goldberg, Jacob Goldin, Fernando Gómez, Assaf Hamdani, Sharon Hanes, Alon Harel, Louis Kaplow, Kobi Kastiel, Roy Kreitner, Tamar Kricheli-Katz, Florencia Marotta-Wurgler, Alan Miller, A. Mitchell Polinsky, Ariel Porat, J. Mark Ramseyer, Barak Richman, Adriana Robertson, Steven Shavell, Henry Smith, Holger Spamann, Cass Sunstein, George Triantis, David Weisbach, and workshop participants at Bar-Ilan University, Chicago, Duke, Haifa University, Harvard, Stanford, Tel Aviv University, and Universitat Pompeu Fabra. Emily Feldstein and Haggai Porat provided outstanding research assistance.

INTRODUCTION	533
I. A THEORY OF INFORMATION COSTS AND DEFAULT RULES.....	544
A. A Simple Model.....	544
1. Framework of analysis.....	544
2. Perfect information.....	545
3. Imperfect information.....	545
4. Biased information.....	550
5. Beyond a binary default choice.....	552
6. Summary.....	553
B. Distributive Considerations	554
C. The Information-Costs Theory Versus Behavioral Accounts of Defaults	556
1. Present bias and procrastination.....	556
2. Optimism.....	558
3. Anchors and reference points.....	559
II. SOURCES OF INFORMATION AND BELIEFS	560
A. Deliberate Investment in Information	561
B. Endorsement Effects: The Informational Content of Defaults	563
C. Influencing Uninformed Beliefs.....	566
D. Forced Deliberation	567
E. Default Rules with Different Information Costs.....	569
F. Personalization	569
III. INFORMATION COSTS VERSUS MECHANICAL COSTS.....	571
A. Cost Engineering by Lawmakers.....	572
1. Increasing mechanical costs.....	572
2. Lowering information costs.....	572
B. Fighting Cost Engineering by Firms	573
1. Negative opt-out costs.....	573
2. Fighting costs with more costs.....	574
3. Regulating firm-engineered costs.....	576
IV. APPLICATIONS.....	576
A. Overdraft Protection.....	577
B. Privacy.....	580
C. Retirement Savings	584
D. Green Defaults.....	587
CONCLUSION	589
APPENDIX.....	592
A.I. FRAMEWORK	592
A.II. LOW DEFAULT.....	593
A.III. HIGH DEFAULT.....	596
A.IV. COMPARISON: LOW DEFAULT VERSUS HIGH DEFAULT	597
A.V. INACCURATE BELIEFS.....	600

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 transaction costs in disclaiming the defaults. If the population governed by a default rule is heterogeneous, continues this transaction-costs argument, it is better to 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

¹ See, e.g., RICHARD A. POSNER, *ECONOMIC ANALYSIS OF LAW* 99 (9th ed. 2014); ROBERT COOTER & THOMAS ULEN, *LAW AND ECONOMICS* 293 (6th ed. 2012); STEVEN SHAVELL, *FOUNDATIONS OF ECONOMIC ANALYSIS OF LAW* 301–04 (2004); Alan Schwartz, *The Default Rule Paradigm and the Limits of Contract Law*, 3 S. CAL. INTERDISC. L.J. 389, 399 (1993); Lucian Ayre Bebchuk & Steven Shavell, *Information and the Scope of Liability for Breach of Contract: The Rule of Hadley v. Baxendale*, 7 J.L. ECON. & ORG. 284, 289–92 (1991); Charles J. Goetz & Robert E. Scott, *The Mitigation Principle: Toward a General Theory of Contractual Obligation*, 69 VA. L. REV. 967, 971 (1983).

² See RICHARD H. THALER & CASS R. SUNSTEIN, *NUDGE: IMPROVING DECISIONS ABOUT HEALTH, WEALTH, AND HAPPINESS* 6–8 (2008); see also Cass R. Sunstein, *Deciding by Default*, 162 U. PA. L. REV. 1, 5 (2013) (“[D]efault rules . . . 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.”); *id.* at 9 (“Default rules may produce significantly better outcomes.”).

problems.³ From mortgage products⁴ to student loans,⁵ from data privacy protections⁶ to retirement savings⁷ and organ donation,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 would have 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. But the literature is surprisingly thin in characterizing these important opt-out costs. 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

³ See Michael S. Barr, Sendhil Mullainathan & Eldar Shafir, *A One-Size-Fits-All Solution*, N.Y. TIMES (Dec. 26, 2007), <https://perma.cc/6CEY-3E9C> (suggesting that default rules with consumer-friendly terms could solve a variety of contracting failures and could even prevent future mortgage crises). *But see* Jon M. Jachimowicz, Shannon Duncan, Elke U. Weber & Eric J. Johnson, *When and Why Defaults Influence Decisions: A Meta-Analysis of Default Effects*, 3 BEHAV. PUB. POL’Y 159, 177–78 (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 & ELДАР SHAFIR, AN OPT-OUT HOME MORTGAGE SYSTEM 22–25 (The Hamilton Project, Discussion Paper No. 2008-14, 2008).

⁵ See James C. Cox, Daniel Kreisman & Susan Dynarski, *Designed to Fail: Effects of the Default Option and Information Complexity on Student Loan Repayment* 6–7, 14–18 (Nat’l Bureau of Econ. Rsch., Working Paper No. 25258, 2018).

⁶ See 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, 37 [hereinafter GDPR] (“[T]he 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.”).

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, 1176 (2001) (“A substantial fraction of 401(k) participants hired under automatic enrollment exhibit default savings behavior, with a contribution rate and asset allocation corresponding to the automatic enrollment default.”). *But see* Ryan Bubb & Richard H. Pildes, *How Behavioral Economics Trims Its Sails and Why*, 127 HARV. L. REV. 1593, 1616–32 (2014).

⁸ See Eric J. Johnson & Daniel Goldstein, *Do Defaults Save Lives?*, 302 SCI. 1338, 1339 (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 arm’s length contracts are

contracts, opt-outs are usually cheap—predrafted 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 nontrivial social cost.¹⁰ Similarly, in noncontractual 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 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,

“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 (1990):

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 the 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 & Brigitte C. Madrian, *Behavioral Household Finance*, in 1 HANDBOOK OF BEHAVIORAL ECONOMICS—FOUNDATIONS AND APPLICATIONS 177, 231–32 (B. Douglas Bernheim et al. eds., 2018).

¹¹ *See, e.g.,* 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 (Dec. 22, 2016), <https://perma.cc/YS5P-WAUQ>; 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], JOURNAL OFFICIEL DE LA RÉPUBLIQUE FRANÇAISE [J.O.] [OFFICIAL GAZETTE OF FRANCE], Jan. 27, 2016, p. 98 (amending French law 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); *see also* Ryan Nakashima, *AP Exclusive: Google Tracks Your Movements, Like It or Not*, AP NEWS (Aug. 13, 2018), <https://perma.cc/9QHP-FKEC> (providing an example in which Google increased opt-out costs to make its privacy settings stickier).

¹² *See* Jamie Luguri & Lior Jacob Strahilevitz, *Shining a Light on Dark Patterns* 6–17 (Univ. of Chi., Pub. L. Working Paper No. 719, 2019), <https://perma.cc/92JE-Q9DD>.

and how it compares to the (sometimes obscure) nondefault 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 often 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, “[T]here is strong evidence that a lack of information on the part of choosers, including a lack of information about alternatives,

¹³ We began to lay the foundation for our information-costs theory in Oren Bar-Gill & 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, for example, 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. . . . [This] might be interpreted as a signal that the party suggesting the modification is more likely than previously thought to rely on his legal rights.

See also, e.g., Omri Ben-Shahar & John A. E. Pottow, *On the Stickiness of Default Rules*, 33 FLA. ST. U. L. REV. 651, 653–54 (2006) (developing how opt-out proposals, and the information they convey, affect the stickiness of defaults); Kathryn E. Spier, *Incomplete Contracts and Signalling*, 23 RAND J. ECON. 432, 432 (1992) (“[A]n individual may refrain from including a particular clause in a contract in order to signal his type.”); Jason Scott Johnston, *Strategic Bargaining and the Economic Theory of Contract Default Rules*, 100 YALE L.J. 615, 617 (1990) (“[T]o bargain around this default, the promisor must convey information which is generally directly contrary to his strategic interest in bargaining with the default.” (emphasis in original)).

¹⁴ See Joshua Blumenstock, Michael Callen & Tarek Ghani, *Why Do Defaults Affect Behavior? Experimental Evidence from Afghanistan*, 108 AM. ECON. REV. 2868, 2890–94 (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.”).

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, rendering them 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%.¹⁶ And a similar percentage of employees opted to increase from the predominant 3% contribution rate default.¹⁷ Much of this wholesale opt-out was

¹⁵ Sunstein, *supra* note 2, at 20; *see also* Jeffrey R. Brown, Anne M. Farrell & Scott J. Weisbenner, The Downside of Defaults 3 (Sept. 16, 2011) (unpublished manuscript) (Nat’l Bureau of Econ. Rsch.) (“[L]ack of adequate information about decision alternatives is a significant driver of the likelihood of default. . . . [I]nformation problems are especially important.”); Cass R. Sunstein & Richard H. Thaler, *Libertarian Paternalism Is Not an Oxymoron*, 70 U. CHI. L. REV. 1159, 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*, for example, William E. Nessmith, Stephen P. Utkus & Jean A. Young, *Measuring the Effectiveness of Automatic Enrollment*, 31 VANGUARD CTR. FOR RET. RSCH. 1, 10 (2007) (predicting that voluntary enrollment would increase from 32% to 59% over the course of the first three years of employment); VANGUARD GRP., HOW AMERICA SAVES 35 (2019) (noting that 60% of employees whose retirement plans permit voluntary enrollment participate).

¹⁷ Nessmith et al., *supra* note 16, at 11 (“[A]fter 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. *See id.* On the predominant contribution rate default of 3%, *see id.* at 6 (“The median contribution rate in automatic enrollment designs is 2.9%.”); VANGUARD GRP., *supra* note 16, at 29 (noting the predominant 3% default). *See also infra*

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 their plan's default contribution rate was not right for them, but their lack of information did not stop them from opting out.¹⁸

Or take another example: overdrafts. A bank customer with no money in her checking account can still swipe her debit card and make a purchase. Banks call this "overdraft protection" and charge hefty fees for the service. In the past, banks automatically enrolled customers into their overdraft protection policies, requiring customers to opt out if they no longer wanted the service. Worried about harm to consumers from these fees, the Federal Reserve Board, in 2010, reversed and set a default of nonenrollment in overdraft protection. To receive overdraft protection, bank customers must now opt out of the regulatory default. And yet of the frequent overdrafters, the group that policy makers were most concerned about, 45% opted out.¹⁹ Here too the opt-out was largely uninformed.²⁰ 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 they are so often contracted around in the vendor's 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

Part IV.C; VANGUARD GRP., *supra* note 16, at 40 (noting that in automatic enrollment plans, which predominantly have a 3% default rate, the average contribution rate was 6.7%, as compared to a 7.1% contribution rate in voluntary-enrollment, i.e., zero-default, plans); *id.* at 35 (noting that the participation rate in voluntary-enrollment, i.e., zero-default, plans was 60%; and the participation rate in automatic enrollment plans was 91%).

¹⁸ See Nessmith et al., *supra* note 16, at 17 ("[E]mployees who quit an automatic enrollment arrangement often have both low levels of financial literacy and a mistrust of financial institutions.").

¹⁹ See CONSUMER FIN. PROT. BUREAU, CFPB STUDY OF OVERDRAFT PROGRAMS 30 (2013); see also *infra* Part IV.A.

²⁰ See ANDREA CAFLISCH, MICHAEL D. GRUBB, DARRAGH KELLY, JEROEN NIEBOER & MATTHEW OSBORNE, SENDING OUT AN SMS: THE IMPACT OF AUTOMATICALLY ENROLLING CONSUMERS INTO OVERDRAFT ALERTS 23–29 (Fin. Conduct Auth., Occasional Paper No. 36, 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).

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.²¹ 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. Recognition of 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.²² The standard account addresses scenarios of asymmetric information and advocates for default rules that force an informed contracting party to reveal 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 decision maker that needs to decide whether to opt out of a default.

²¹ For a discussion of applications, see *infra* Part IV.

²² For the standard accounts of penalty or information-forcing defaults, see Ian Ayres & Robert Gertner, *Filling Gaps in Incomplete Contracts: An Economic Theory of Default Rules*, 99 YALE L.J. 87, 95–107 (1989); Bebchuk & Shavell, *supra* note 1, at 287–92.

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 nondefault 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 counterpreferences 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's preferences given the majority's information*. In that sense, we still recommend majoritarian defaults, recognizing that the majority may act upon uninformed beliefs.²³

The value of this refined criterion is on full display when we consider the “intermediate” information-costs case, in which 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

²³ 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. Professors Richard Thaler and Cass Sunstein’s proposed rule of thumb for default choice “minimizes the number of opt-outs.” Richard H. Thaler and Cass R. Sunstein, *Libertarian Paternalism*, 93 AM. ECON. REV. 175, 178–79 (2003). For an article showing that this rule of thumb fares 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, 2822 (2015). But it does not distinguish between informed and uninformed opt-out.

information through the efforts of the default setter or of interested third parties (for example, 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 interests in mind.

A central payoff of the information-costs theory is to shed new light on, and suggest a reevaluation of, *Nudge*-type libertarian paternalistic ideas—in particular, the sticky-defaults paradigm. Regulation through defaults is premised on the expectation that less sophisticated people will stick with the default, while the more sophisticated are free to opt out.²⁴ 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.²⁵ This explanation helps bridge an uncomfortable gap between the academic appetite for sticky default rules and the reality of slippery defaults.²⁶ 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

²⁴ See Bubb & Pildes, *supra* note 7, at 1598–99 (“[T]he 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.”); Thaler & Sunstein, *supra* note 2, at 242 (“Most of the time, nudging helps those who need help while imposing minimal costs on those who do not.”); Cass R. Sunstein, *Boundedly Rational Borrowing*, 73 U. CHI. L. REV. 249, 257 (2006) (“[I]nterventions 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.”); Colin Camerer, Samuel Issacharoff, George Loewenstein, Ted O’Donoghue & Matthew Rabin, *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.”).

²⁵ 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 *infra* Part I.C.

²⁶ 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 190–94 (2014); Lauren E. Willis, *When Nudges Fail: Slippery Defaults*, 80 U. CHI. L. REV. 1155, 1210 (2013).

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.²⁷

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 choice²⁸—as contributing to the stickiness of defaults. But they mistakenly view information costs as simply another type of mechanical opt-out cost.²⁹ 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 antimanipulation law. Policy makers and commentators are increasingly concerned about ways in which

²⁷ 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) disproportionate representation of the default in the limited choice sets that individuals consider, and (7) cognitive dissonance. Beshears et al., *supra* note 10, at 231–33; *see also* B. Douglas Bernheim & Dmitry Taubinsky, *Behavioral Public Economics*, in 1 HANDBOOK OF BEHAVIORAL ECONOMICS, *supra* note 10, at 381, 465–67 (surveying the literature on sticky default rules in the context of retirement savings); Willis, *supra* note 26, at 1161–74.

²⁸ 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 [] this cost, together with present-biased preferences, creates procrastination.”).

²⁹ *See, e.g.*, 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.”).

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, nondefault option. Terms like "sludge" and "dark patterns" are sometimes used to describe such practices.³⁰ A different set of practices is designed to increase information costs, for example, 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 I develops our information-cost theory of defaults. Part II considers different sources of information and beliefs. Part III discusses the different regulatory implications of information costs and mechanical costs. Part IV 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 policy makers need to apply the information-costs theory in the design of default rules.

³⁰ The term "sludge" refers to frictions that make wise decision-making more difficult, whereas the term "dark patterns" refers to choice architecture that frustrates the ability of people to make their desired choices. See, e.g., Luguri & Strahilevitz, *supra* note 12, at 3; Cass R. Sunstein, *Sludge and Ordeals*, 68 DUKE L.J. 1843, 1850 (2019); see also Richard H. Thaler, *Nudge, Not Sludge*, 361 SCI. 431, 431 (2018).

I. 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.³¹ 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 of High versus 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.” Forty percent of the population are better off with High and are called “Type 2.” We use numbers to reflect the valuations assigned, as reflected in Table 1.

TABLE 1: SETTING UP THE EXAMPLE

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

To make the example simple, and without loss of generality, we assume that everyone assigns value $v = 0$ to Low. For Type 1,

³¹ A more formal and general derivation of the propositions is provided in the Appendix.

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 nonzero 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. More people are better off with Low: 60% versus 40%. But higher expected welfare is obtained under High. The expected welfare is $0.6 * (-10) + 0.4 * 20 = 2$ 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: $W_{Low} = 0.6 * 0 + 0.4 * (20 - 1) = 7.6$. With a High default, Type 1 individuals opt out and average social welfare is: $W_{High} = 0.6 * (-1) + 0.4 * 20 = 7.4$.³² 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 $0.6 * (-10) + 0.4 * 20 = 2$. (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

³² We calculated welfare for a representative individual. To assess overall social welfare, we would need to multiply these values by the number of individuals who are subject to the default rule.

valuations, people prefer High.³³ 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 (which is less than 2, the expected value of High), 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.³⁴ Let us examine how people would behave, and the resulting social welfare, under the two defaults.

a) Low default. The first decision people face is whether to acquire information.³⁵ The value of information is:

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

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 also 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 nonselectively, uninformed. At the same time, high information costs reduce

³³ We assume that individuals are risk neutral and thus focus on expected values.

³⁴ 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 fare better with the nondefault 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 nondefault option.

³⁵ For a general exposition on the economics of information, see J. Hirshleifer & John G. Riley, *The Analytics of Uncertainty and Information—An Expository Survey*, 17 J. ECON. LITERATURE 1375, 1393–1414 (1979).

welfare, because they eliminate a “separating” outcome in which different people choose the outcome most suitable to them.

b) High default. Under a High default, the value of information is:

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

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: because 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.

c) Welfare comparison. First, notice that the value of information is greater with Low default ($I_{Low} > I_{High}$), which means that Low default leads to more acquisition of information. Why? Primarily because the expected value of 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:

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

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

With both defaults, informed parties get 8 (i.e., $0.4 * 20$) and uninformed parties get 2. The difference lies in the opt-out costs. With Low default, informed opt-out occurs 40% of the time (when information is acquired) and uninformed opt-out occurs 100% of the time (when information is not acquired). Information acquisition thus saves 0.6 in opt-out costs because it reduces opt-out from 100% to 40%. 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.³⁶

We proceed by distinguishing among three ranges of information costs.

(i) *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 0. 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.

(ii) *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.

(iii) *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

³⁶ 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.

individuals stick with the default; the opt-out rate is 0. 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 $0.4 * (20 - 1) = 7.6$ but requires costly investment in information acquisition. Social welfare then equals the difference between 7.6, the added value of the information, and the cost of acquiring the information. 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 (i.e., $7.6 - 2$); 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, the “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.³⁷ 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 in which 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 countermajoritarian 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

³⁷ For the standard accounts of penalty or information-forcing defaults, see Ayres & Gertner, *supra* note 22, at 95–107; Bebchuk & Shavell, *supra* note 1, at 287–92.

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 are accurate on average—they correctly anticipate the share of people who are Type 1 (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 they can do better.³⁸ 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%.³⁹ To be sure, there are many other ways in which beliefs could be inaccurate.⁴⁰ Here, we merely illustrate how inaccurate beliefs can change the effects of the different defaults and thus the optimal default choice.

(i) *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 $0.8 * (-10) + 0.2 * 20 = -4$.) 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). Regardless, 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.⁴¹

Here, inaccurate beliefs change the welfare ordering of the two defaults and thus alter our policy prescriptions. With

³⁸ 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, i.e., to stick with the default.

³⁹ 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, for example, when a business invests in manipulating consumer beliefs in order to maximize its profits.

⁴⁰ We examine the effects of other distortions more thoroughly in the Appendix.

⁴¹ 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.

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.

(ii) *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.

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

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

The perceived expected payoff if information is acquired decreases relative to the accurate-beliefs case (from $0.4 * (20-1)$ to $0.2 * (20-1)$). The perceived payoff if information is not acquired changes, relative to the accurate-beliefs case, from the uninformed opt-out payoff of 1 (i.e., $2 - 1$) to the no opt-out payoff, 0.

With High default, the perceived value of information is:

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

The perceived expected payoff if information is acquired decreases relative to the accurate-beliefs case (from $0.6 * (-1) + 0.4 * 20$ to $0.8 * (-1) + 0.2 * 20$). 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.)⁴²

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.⁴³

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 are not, either you have overdraft protection or do not. In other cases, there are multiple options: there are more than two possible privacy settings and

⁴² We also observe that the very definition of low versus intermediate versus high information costs changes, because the cutoff values of information are affected by the inaccurate beliefs.

⁴³ 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.

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, nondefault option. With multiple outcomes, the individual would need to compare the default option to several nondefault 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 retirement 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, among multiple discrete options, or along 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 transaction costs—mechanical costs and information costs—that are often treated interchangeably. The practical effects of these two types of transaction 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 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; 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. People 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 policy

maker could, under some conditions, 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 policy maker's distributive objective would be frustrated.

Return to our numerical example⁴⁴ and consider a policy maker who favors the Type 1 group and thus chooses Low default, best suited for Type 1, even though High default provides overall greater average value. For this policy to work, the uninformed Type 1 individuals must stick with the default. Otherwise, the favored group would just opt out to High, and the Low default would just impose extra opt-out costs on this group. The policy maker who intended to help Type 1 individuals would only end up hurting them.⁴⁵

For instance, banks used to default debit card holders into overdraft protection—an arrangement that may have been beneficial for a majority of card holders, but was harmful to the minority of poorer card holders who repeatedly overdrew their accounts and thus paid hefty fees (Type 1). Policy makers sought to help these poorer consumers by setting a no-overdraft-protection default (Low default). Had these consumers stuck with the new default, they would have avoided significant overdraft fees. But many of them did not stick with the new default. Prompted by the banks, they opted out of the new default and back into overdraft protection. This uninformed opt-out frustrated, at least in part, the policy makers' distributive objective.⁴⁶

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,

⁴⁴ See *supra* Part I.A.

⁴⁵ Others have noted that the possibility of opt-out limits the ability of default rules to achieve distributive goals, albeit without expressly distinguishing between informed and uninformed opt-out. See, e.g., Willis, *supra* note 26, at 1200–10.

⁴⁶ For a more detailed discussion of this example, see *infra* Part IV.A.

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.

The retirement savings context illustrates these distributive implications. Wealthier employees are likely to have lower information costs. They have financial advisers to consult with, and they can more easily access available information on optimal savings. For less affluent employees, information is less readily available. Policy makers can help these less-affluent and less-informed employees by choosing a default that maximizes the expected value of their retirement savings. For example, the informed majority of wealthier employees may prefer to invest their retirement savings in a higher-risk, higher-return fund, whereas the uninformed, less-affluent employees may prefer a more conservative investment. The information-costs theory would tell policy makers to choose the more conservative option as the default.

C. The Information-Costs Theory Versus Behavioral Accounts of Defaults

The theory of opt-out from default rules presented in Part I.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 responses 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.⁴⁷ Even if the nondefault 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.⁴⁸ 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 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, 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 nondefault option, procrastination results in uninformed opt-out and slippery defaults. Scholars such as Joshua Blumenstock, Michael Callen, and Tarek Ghani,⁴⁹ as well

⁴⁷ OREN BAR-GILL, SEDUCTION BY CONTRACT: LAW, ECONOMICS, AND PSYCHOLOGY IN CONSUMER MARKETS 21–22 (2012) (“Myopic consumers care more about the present and not enough about the future.”); Sunstein, *supra* note 24, at 252:

[S]ome borrowers procrastinate, ensuring that some bills are paid late.

...

Some borrowers are myopic, emphasizing the short term at the expense of the future.

⁴⁸ 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”).

⁴⁹ See *id.*

as John Beshears, James J. Choi, David Laibson, and Brigitte C. Madrian,⁵⁰ 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.)⁵¹

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 nondefault 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 nondefault 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

⁵⁰ See Beshears et al., *supra* note 10, at 231.

⁵¹ See Willis, *supra* note 26, at 1183 ("As regulators noted in promulgating the overdraft default, consumers are likely to assume overoptimistically they will not overdraw." (citing Electronic Fund Transfers, 74 Fed. Reg. 59,033, 59,044 (Nov. 17, 2009) (to be codified at 12 C.F.R. pt. 205))); see also TAL SHAROT, THE OPTIMISM BIAS: A TOUR OF THE IRRATIONALLY POSITIVE BRAIN 188–96 (2011); Christine Jolls, *Behavioral Economics Analysis of Redistributive Legal Rules*, 51 VAND. L. REV. 1653, 1659 (1998) ("An amazingly robust finding about human actors—and an important contributor to the phenomenon of risk underestimation—is that people are often unrealistically optimistic about the probability that bad things will happen to them.").

adherence to the default, especially in consumer markets where sellers profit from consumers' choice of nondefault 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.⁵² The information-costs theory helps policy makers identify the default rule that will induce more acquisition of information and thus minimize the effects of optimism.⁵³

3. Anchors and reference points.

A third set of behavioral explanations focuses 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 toward this number.⁵⁴ The default option may serve as such an anchor. 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.⁵⁵ If the default option becomes the reference point and relative costs of the nondefault option loom larger than

⁵² See BAR-GILL, *supra* note 47, at 9.

⁵³ The information-costs theory also shows how misperception, like optimism, can distort the decision whether to acquire information. See *supra* Part I.A.4; *infra* Part II.B.

⁵⁴ The anchoring effect was first identified in Amos Tversky & Daniel Kahneman, *Judgment Under Uncertainty: Heuristics and Biases*, 185 SCI. 1124, 1128–30 (1974). In the retirement savings context, employees raise their 401(k) contribution rates if they are exposed to arbitrary high contribution examples in communications. See James J. Choi, Emily Haisley, Jennifer Kurkoski & Cade Massey, *Small Cues Change Savings Choices*, 142 J. ECON. BEHAV. 378, 392 (2017); 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. PUB. ECON. 80, 86–88 (2014); see also Bernheim et al., *supra* note 23, at 2818–19 (presenting an economic model which shows "anchoring . . . as the most plausible explanation for bunching at the default option"). This account is based on the literature review in Beshears et al., *supra* note 10, at 232.

⁵⁵ See Daniel Kahneman & Amos Tversky, *Prospect Theory: An Analysis of Decision Under Risk*, 47 ECONOMETRICA 263, 277–80 (1979) (developing Prospect Theory, which emphasizes the importance of reference points and identifies loss aversion).

relative gains from the nondefault option, then decision makers will tend to stick to the default.⁵⁶

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 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.⁵⁷

II. SOURCES OF INFORMATION AND BELIEFS

The theory presented in this Article identifies information and beliefs as key factors that affect opt-out decisions. Part I 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 surrounding the design of default rules. First, this Part takes the theoretical building blocks 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 toward accuracy or bias? Later, in Part III, we turn to the other feature of the model—the difference between mechanical and information costs—and explore regulatory implications arising from it.

⁵⁶ Moreover, Professors Daniel Kahneman and Amos 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 & Amos Tversky, *The Psychology of Preferences*, 246 SCI. AM. 160, 173 (1982). This would make the default option even stickier.

⁵⁷ See *supra* Part I.A; *infra* Part II.B.

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 nondefault 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.⁵⁸

Retirement savings investment 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 uninformed 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

⁵⁸ 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 all clients. With a 1% fee, therefore, individuals would have paid \$105 billion for investment advice. See NAT’L REG. SERVS. & INV. ADVISER ASS’N, 2019 EVOLUTION REVOLUTION: A PROFILE OF THE INVESTMENT ADVISER PROFESSION 8 (2019). The vast majority (95.5%) of fees paid to investment advisors are billed as a percentage of assets under management. *Id.* at 29.

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, which 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 and 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 Part I.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 increase, 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 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 nondefault options) and the information they use to shape these beliefs.

B. Endorsement Effects: The Informational Content of Defaults

Thus far we have 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 interests in mind or because they believe that the government mandated such a default option and the government has their best interests in mind—this would affect their decisions whether to collect information and whether to opt out.⁵⁹ The precise effects depend on the nature of the information that the employer has. Here, we assume that the employer

⁵⁹ 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 232 (“Individuals may believe that the default is a choice recommended by the default setter.”); Jachimowicz et al., *supra* note 3, at 172–73 (listing the endorsement effect as one of three main reasons why defaults are sticky); Craig R.M. McKenzie, Michael J. Liersch & Stacey R. Finkelstein, *Recommendations Implicit in Policy Defaults*, 17 PSYCH. SCI. 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 policy maker’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. In contractual settings, default rules can also provide information indirectly: the policy maker can set a default that is unfavorable to the sophisticated party, such that the act of opt-out (when this party contracts out of the default) conveys information to the less sophisticated counterparty. *See* J.H. Verkerke, *Legal Ignorance and Information-Forcing Rules*, 56 WM. & MARY L. REV. 899, 916 (2015) (“A surprisingly large number of common law and statutory rules . . . [are] designed to force a legally sophisticated party to inform unsophisticated parties about the prevailing legal standard [by opting out].”).

has information only about average values. The alternative assumption is considered in Part II.F below.

If employees hold accurate uninformed beliefs, the employer has no informational advantage and the default she 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).⁶⁰

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 in Part I.A.4, uninformed employees mistakenly thought that the prevalence of Type 1 was 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.⁶¹

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?⁶² Then an employer might set a default that does not maximize her employees' expected value. If employees are aware of this conflict of interest, little harm is done; the employees

⁶⁰ 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.

⁶¹ See, e.g., Beshears et al., *supra* note 10, at 232; McKenzie et al., *supra* note 59, at 418–19.

⁶² See Ryan Bubb & Patrick L. Warren, *An Equilibrium Theory of Retirement Plan Design*, 12 AM. ECON. J.: ECON. POLICY 22, 43 (2020) ("[I]f workers are subject to behavioral biases that affect retirement savings decisions, then employers have incentives to cater to rather than correct those biases."); Ryan Bubb, Patrick Corrigan & Patrick L. Warren, *A Behavioral Contract Theory Perspective on Retirement Savings*, 47 CONN. L. REV. 1317, 1338 (2015) ("[T]he labor market gives employers incentives to craft plan designs that cater to what biased workers perceive to be of value.").

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.⁶³ The concern is that employees would mistakenly attribute benevolence to a nonbenevolent 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 decision maker. This endorsement effect also depends crucially on the perceived alignment, or misalignment, of interests between the default setter and the decision maker. The greatest potential reduction in social welfare arises when decision makers 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, decision makers will prefer to rely on their own information. However, in some contexts in which 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.⁶⁴

⁶³ 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 nondefault 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.

⁶⁴ See, e.g., Kim Willsher, *France Introduces Opt-out Policy on Organ Donation*, GUARDIAN (Jan. 2, 2017), <https://perma.cc/8SC3-5KMC> (describing an official French campaign, which included a video advertisement discouraging people from opting out of organ donation).

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, policy makers 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, policy makers could actively warn consumers about the adverse consequences of choosing the nondefault 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, so-called 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 policy maker’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 forgoing this “recommended” option.⁶⁵ Indeed, when a firm

⁶⁵ See Willis, *supra* note 26, 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 a default is “recommended” or “advised.”

See also Lauren E. Willis, *Why Not Privacy by Default?*, 29 BERKELEY TECH. L.J. 61, 82–83 (2014) (“Opting out of a default also might be made more or less attractive through messages to the user . . . [by] warn[ing] the user that changing the setting could cause problems.”); *id.* at 95 (“One [firm] goes so far as to warn consumers that choosing not to be tracked will ‘spoil[] your experience of the website.’” (second alteration in original)); *id.* at 102 (“[F]inancial information defaults . . . [are] designed to trigger loss aversion and the endowment effect . . . [One example] 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*, ELEC. FRONTIER FOUND. (Dec. 9, 2009), <https://perma.cc/4GW7-2EVB>:

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

prefers a nondefault option, a contest between the policy maker and the firm over consumers' perceptions may ensue.⁶⁶ It is hard to imagine that lawmakers could win such tournaments;⁶⁷ 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 policy makers 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.⁶⁸ For example, an

“recommend”—preselect by default—the setting to share the content they post to Facebook.

⁶⁶ As with advertising generally, one could ask if policy makers and firms are targeting consumers' perceptions or trying to shape consumer preferences. *See generally, e.g.*, Christopher A. Summers, Robert W. Smith & Rebecca Walker Reczek, *An Audience of One: Behaviorally Targeted Ads as Implied Social Labels*, 43 J. CONSUMER RSCH. 156 (2016); Christina L. Brown & Aradhna Krishna, *The Skeptical Shopper: A Metacognitive Account for the Effects of Default Options on Choice*, 31 J. CONSUMER RSCH. 529 (2004); Peter Wright, *Marketplace Metacognition and Social Intelligence*, 28 J. CONSUMER RSCH. 677 (2002).

⁶⁷ *See, e.g.*, Willis, *supra* note 26, 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”).

⁶⁸ 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 218 (Till Grüne-Yanoff & Sven Ove Hansson eds., 2008) (“The cost of *Nudge* may be that we forego the chance to gain the virtue of self-command.”). *But see* N. Craig Smith, Daniel G. Goldstein & Eric J. Johnson, *Choice Without Awareness: Ethical and Policy Implications of Defaults*, 32 J. PUB. POL’Y & MKTG. 159, 163–64 (2013) (criticizing Bovens’s analysis for ignoring “the larger body of research (of which *Nudge* is a part) that questions the assumptions of rationality and active choice in many areas of human judgment”).

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.⁶⁹ 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 information toward an informed decision, which will lead to optimal sorting. It helps people prioritize their limited attention and information acquisition resources. As long as such a 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.⁷⁰

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.⁷¹

⁶⁹ See NAT'L SCI. & TECH. COUNCIL, EXEC. OFF. OF THE PRESIDENT, SOC. & BEHAV. SCI. TEAM, 2016 ANNUAL REPORT 40–41 (2016); Beshears et al., *supra* note 10, at 234; Gabriel D. Carroll, James J. Choi, David Laibson, Brigitte C. Madrian & Andrew Metrick, *Optimal Defaults and Active Decisions*, 124 Q.J. ECON. 1639, 1641–42 (2009).

⁷⁰ 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.

⁷¹ See Bernheim et al., *supra* note 23, at 2826–27 (noting that extreme defaults in combination with large penalties for passive choice can force active choice). Stop-and-think defaults are also related to penalty defaults. See Ayres & Gertner, *supra* note 22, at 91–93; Bechuk & Shavell, *supra* note 1, at 289–92.

E. Default Rules with Different Information Costs

The basic theory presented in Part I.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.⁷²

In the overdraft protection example, a probank default imposes higher information costs than an antibank default. If the default is overdraft protection and consumers are expected to stick with the probank 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 antibank default, then the bank would have a strong incentive to provide information to consumers about the benefits of the nondefault, overdraft protection option. The cost to consumers of acquiring at least some information—benefit information—is lower under the antibank 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 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 policy maker sets a single default rule for all relevant parties. But this need not be the case. When default setters have better information—individualized

⁷² Compare Ayres & Gertner, *supra* note 22, at 104, with Bebchuk & Shavell, *supra* note 1, at 289.

information—they can increase welfare by setting personalized defaults.⁷³ An employer will often have detailed personal information about 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 twenty-five-year-old employee is different from that of the sixty-five-year-old employee.⁷⁴

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.⁷⁵ Indeed, this strategy can be viewed as a means of reducing information costs. The policy maker 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.⁷⁶ 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.

⁷³ In Part II.B we discussed how the type of information that the policy maker 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 policy makers need to implement the information-costs theory). See *infra* Part III.

⁷⁴ 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), <https://perma.cc/7SLW-UQPA> (outlining the benefits and harms of target-date retirement funds). 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. 1417, 1425–31 (2014) (explaining the range of personalization available in designing majoritarian and minoritarian default rules); Sunstein, *supra* note 2, at 48–56 (exploring the types, reach, and feasibility of personalized default rules).

⁷⁵ 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.”

⁷⁶ OMRI BEN-SHAHAR & ARIEL PORAT, *DIFFERENT RULES FOR DIFFERENT PEOPLE: PERSONALIZED LAW IN THE ERA OF BIG DATA* (forthcoming 2021) (on file with authors) (exploring personalized rules derived with the use of algorithms sorting through Big Data).

The possibility of personalization interacts with the preceding discussion about the informational content of defaults.⁷⁷ 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 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 her own interests—interests that conflict with those of the employees. If employees recognize the conflict of interest, 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.

III. INFORMATION COSTS VERSUS MECHANICAL COSTS

The theory presented in Part I 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 “transaction 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 I 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 transaction 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 toward 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.

⁷⁷ See *supra* Part II.B.

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.⁷⁸

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, policy makers' 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 speed 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

⁷⁸ See, e.g., Michael S. Barr, Sendhil Mullainathan & Eldar Shafir, *The Case for Behaviorally Informed Regulation*, in *NEW PERSPECTIVES ON REGULATION* 25, 43 (D. Moss & J. Cisternino eds., 2009) ("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. . . . Deviation from the offer would require heightened disclosures and additional legal exposure for lenders in order to make the default sticky.").

that reducing information costs is desirable across the board, 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 even 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 nondefault 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 nondefault 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 policy makers' 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 policy maker'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.⁷⁹

2. Fighting costs with more costs.

Whether successful or not, the artificial mechanical costs imposed by such sludges are reason enough for policy makers 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.⁸⁰ But while justified, the practical value of such policies is questionable. The most typical tools policy makers use to increase mechanical

⁷⁹ See Luguri & Strahilevitz, *supra* note 12, at 27–29 (demonstrating, in experimental setting, the tendency of subjects to accept the seller's induced opt-out option).

⁸⁰ Cf. Ian Ayres, *Regulating Opt-Out: An Economic Theory of Altering Rules*, 121 YALE L.J. 2032, 2093 (2012) (arguing that in some cases policy makers should use "impeding altering rules"—with high opt-out costs—to reduce opt-out rates and noting that such "impeding altering rules" should be used when people overestimate the benefit from opting out).

opt-out costs are lengthy disclosures, educational prerequisites, segregated agreements, clause-by-clause signatures, and periodically renewed agreements.⁸¹ It is sometimes doubtful whether such hurdles succeed in increasing opt-out costs in a manner sufficient to render the defaults stickier, especially in the presence of a firm's sludges.⁸²

The privacy context illustrates this “battle of the mechanical costs.” Policy makers increased mechanical costs to strengthen the privacy-protective default. The European privacy regulation, the General Data Protection Regulation (GDPR),⁸³ requires more explicit consent to information collection and some European lawmakers have required renewed consent for every incidence of data collection.⁸⁴ Firms fought back, engineering their own mechanical costs. For example, Google increased opt-out costs to make its privacy settings stickier.⁸⁵ The battle will likely continue to rage on. At the end of the day, it seems doubtful that policy makers will succeed in creating a truly sticky default.

⁸¹ There are many examples in which regulators deliberately increase opt-out costs. In the privacy context, see *infra* Part IV.B. See also Law 2016-41 of January 26, 2016, art. 192 (Fr.) (requiring the submission of form and proof of identity in order to opt out of default organ donation); OFF. OF THE COMPTROLLER OF THE CURRENCY, OCC BULL. 2020-15, OVERDRAFT PROTECTION: OPT-IN REQUIREMENTS AND RELATED MARKETING ISSUES (2010) (separate disclosure and segregated assent are now necessary to enroll in overdraft protection); Law 2016-41 of January 26, 2016, art. 192 (Fr.) (requiring the submission of form and proof of identity in order to opt out of the default organ donation). There are also many examples of firms raising, or manufacturing, opt-out costs. See, e.g., Nakashima, *supra* note 11 (providing an example where Google increased opt-out costs to make its privacy settings stickier); see also Willis, *supra* note 26, at 1165 (“When the choice to opt out of a default is not made plain, people may perceive a default as unchangeable.”); *id.* at 1771 (“[F]irms actively work to increase the power of their defaults using . . . transaction barriers.”); *id.* (“[F]irms with automatically renewing subscriptions that consumers can sign up for in minutes online may require spending an hour on hold with customer service to cancel.”); *id.* (“Firms stymie consumers who might attempt to opt out, using fine print.”); Sumit Agarwal, Gene Amromin, Itzhak Ben-David, Souphala Chomsisengphet & Douglas D. Evanoff, *Do Financial Counseling Mandates Improve Mortgage Choice and Performance? Evidence from a Legislative Experiment* 32 (Fed. Rsrv. Bank of Chi., 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.”); Jeff Sovern, *Opting in, Opting out, or No Options at All: The Fight for Control of Personal Information*, 74 WASH. L. REV. 1033, 1083 (1999) (“[C]ompanies that offer opt-outs have an incentive to increase the transaction costs incurred by consumers who opt out.”); *id.* at 1089 (“[Some firms] provide[] subscribers with a lengthy, dull, and difficult-to-read statement of their rights and require[] subscribers wishing to opt out to communicate their intent in a separate writing.”).

⁸² See Willis, *supra* note 26, at 1224-25.

⁸³ 2016 O.J. (L 119) 1.

⁸⁴ See *infra* Part IV.B.

⁸⁵ See generally Nakashima, *supra* note 11.

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 is, in essence, 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.⁸⁶ 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 antimanipulation law.

IV. 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. Part IV.A considers the regulation of overdraft fees. Part IV.B considers the protection of privacy in digital information. Part IV.C returns to the canonical example of retirement savings. And Part IV.D discusses "green" defaults.

⁸⁶ 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 3 HANDBOOK OF INDUSTRIAL ORGANIZATION 1701, 1708–23 (Mark Armstrong & Robert Porter eds., 2007).

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.⁸⁷

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 Federal Reserve 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).⁸⁸ 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.⁸⁹

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 would knowingly and sparingly use it.⁹⁰ 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

⁸⁷ See, e.g., Peter Smith, *Report: FDIC Data Shows That Banks Collected \$11.45 Billion in Overdraft Fees in 2017*, CTR. FOR RESPONSIBLE LENDING (Aug. 7, 2018), <https://perma.cc/9C5G-FTFS>; Julia Chang, *Americans Paid \$34 Billion in Overdraft Fees Last Year. Here's How to Stop the Charges*, FORBES (Apr. 5, 2018), <https://perma.cc/7Y3Y-D9YU>; Jackie Wattles, *Americans Paid \$15 Billion in Overdraft Fees Last Year*, CFPB Says, CNN MONEY (Aug. 4, 2017), <https://perma.cc/V58Y-9EV7>.

⁸⁸ 12 C.F.R. § 205.17(b) (2012).

⁸⁹ See OFF. OF THE COMPTROLLER OF THE CURRENCY, *supra* note 81.

⁹⁰ 12 C.F.R. § 205.1(b) (2012) ("The primary objective of the act and this part is the protection of individual consumers engaging in electronic fund transfers.").

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.”⁹¹

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).⁹² If most consumers were indeed uninformed, but unbiased, then the purpose of the 2010 default switch

⁹¹ CONSUMER FIN. PROT. BUREAU, *A Closer Look: Overdraft and the Impact of Opting-in* (Jan. 19, 2017), <https://perma.cc/SB2Z-J36N>. 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 CONSUMER FIN. PROT. BUREAU, *CFPB Finds Small Debit Purchases Lead to Expensive Overdraft Charges* (July 31, 2014), <https://perma.cc/3TS5-7RM2> (“The study found that the majority of debit card overdraft fees are incurred on transactions of \$24 or less.”).

⁹² There is some evidence that overdrafting, when the cost (fee or interest rate) is so high, is harmful to consumers. See CAFLISCH et al., *supra* note 20, at 3–4, 23–29 (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 United Kingdom are smaller than overdraft fees in the United States, suggesting that the effect in the United States would be larger than the 21–25% figures. Compare Rupert Jones, *Overdrafts: Can You Cut the Cost of Yours?*, GUARDIAN (May 5, 2018), <https://perma.cc/6XNH-WQ6P> (describing U.K. overdraft fees as around seven pounds), with Chang, *supra* note 87 (describing the average U.S. overdraft fee as around thirty dollars). See also CAFLISCH et al., *supra* note 20, at 3 (citing OFF. OF FAIR TRADING, PERSONAL CURRENT ACCOUNTS IN THE UK: AN OFT MARKET STUDY 62 (2008) (“[E]vidence 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, TOM HOEHN, PHILP MARSDEN, JILL MAY & ED SMITH, COMPETITION & MKTS. AUTH., RETAIL BANKING MARKET INVESTIGATION: FINAL REPORT 536 (2016) (“[A] significant proportion of customer detriment experienced by overdraft users . . . may arise from lack of awareness and engagement with their [personal current accounts].”); PAUL ADAMS, MICHAEL D. GRUBB, DARRAGH KELLY, JEROEN NIEBOER & MATTHEW OSBORNE, TIME TO ACT: A FIELD EXPERIMENT ON OVERDRAFT ALERTS (Fin. Conduct Auth., Occasional Paper No. 40, 2018) (corroborating the earlier FCA study by Caflich et al.); STEFAN HUNT, DARRAGH KELLY & FABIAN GARAVITO, MESSAGE RECEIVED? THE IMPACT OF ANNUAL SUMMARIES, TEXT ALERTS AND MOBILE APPS ON CONSUMER BANKING BEHAVIOR 3 (Fin. Conduct Auth., Occasional Paper No. 10, 2015) (“[S]igning 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%.”).

would 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.⁹³ Why? It is possible that some of these consumers were engaging in informed opt-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.⁹⁴ 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 Federal Reserve hoped.

We highlight 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

⁹³ For frequent overdrafters, the group that policy makers 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 CONSUMER FIN. PROT. BUREAU, *supra* note 19, at 29–30.

⁹⁴ See Willis, *supra* note 26, at 1191–92 (emphasis in original):

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*."

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.⁹⁵ 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.

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.⁹⁶ And all users unknowingly "agree." Similarly, the European Union's GDPR and California's Consumer Privacy Act⁹⁷ establish a default that prohibits collection of some categories of personal information. 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 a long time, 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.⁹⁸ The mechanics of such opt-outs have been

⁹⁵ See, e.g., BAR-GILL, *supra* note 47, at 32–41.

⁹⁶ See 18 U.S.C. § 2511.

⁹⁷ Cal. Civ. Code §§ 1798.100–1798.199.100.

⁹⁸ See Willis, *supra* note 65, at 111 ("While a Don't-Track-Me setting would require firms to spend significant resources on maneuvering consumers out of the default, firms

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 versus low-privacy options became less relevant. Alternatively, people believed that the pay-off difference between high and 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, elevated the salience of data collection and its potential harms.⁹⁹ The enactment of the European data regulations—the GDPR—further heightened the public's awareness.¹⁰⁰ 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 nondefault options, and perhaps increasingly believing that the

determined to do so could be successful.”); Daniel J. Solove, *Introduction: Privacy Self-Management and the Consent Dilemma*, 126 HARV. L. REV. 1880, 1898–99 (2013) (“[M]any organizations 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.”); 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 on the site.”); 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.”).

⁹⁹ See, e.g., Andrew Perrin, *Americans Are Changing Their Relationship with Facebook*, PEW RSCH. CTR. (Sept. 5, 2018), <https://perma.cc/X92K-6RWG> (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); Julie Beck, *People Are Changing the Way They Use Social Media*, ATLANTIC (June 7, 2018), <https://perma.cc/T978-Q9VZ> (suggesting that breaches like Cambridge Analytica's led users to share less detailed information online); see also Kim Hart & Ina Fried, *Exclusive Poll: Facebook Favorability Plunges*, AXIOS (Mar. 26, 2018), <https://perma.cc/5JAZ-CHYM> (showing that Facebook's favorability had fallen from 33% to 5% in the wake of the revelations regarding Cambridge Analytica).

¹⁰⁰ 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).

no-information-collection regime is superior. By and large, these are 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.¹⁰¹ In some cases, lawmakers have required renewed consent for every incidence of data collection.¹⁰²

¹⁰¹ See GDPR, *supra* note 6, at 6:

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.

The GDPR creates a range of opt-out costs depending on the type of information collected. There are two main types of consent (they 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 (for example, using a site multiple times after agreeing to cookies clears the hurdle, or agreeing by submitting an email address). See *Explicit vs. Unambiguous Consent: What's the Difference?*, DATASTREAMS (Oct. 11, 2017), <https://perma.cc/7MFL-T57Q>.

On January 21, 2019, by force of the GDPR, the French National Data Protection Commission imposed a fine of €50 million 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 NAT'L DATA PROT. COMM'N, *The CNIL's Restricted Committee Imposes a Financial Penalty of 50 Million Euros Against Google LLC* (Jan. 21, 2019), <https://perma.cc/TEJ2-QR9P>.

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 OFF. OF THE ATT'Y GEN., CAL. DEPT OF JUST., CALIFORNIA CONSUMER PRIVACY ACT (CCPA) FACT SHEET (2019).

¹⁰² See, e.g., Délibération 2019-093 du 4 juillet 2019 portant adoption de lignes directrices relatives à l'application de l'article 82 de la loi 6 janvier 1978 modifiée aux opérations de lecture ou écriture dans le terminal d'un utilisateur (notamment aux cookies et autres traceurs) [Deliberation 2019-093 of July 4, 2019, Adopting Guidelines Relating to the Application of Article 82 of the Law of January 6, 1978, Amended to Reading or Writing Operations to a User's Computer Terminal (in Particular to Cookies and Other Tracers)], JOURNAL OFFICIEL DE LA RÉPUBLIQUE FRANÇAISE [J.O.] [OFFICIAL GAZETTE OF FRANCE], July 19, 2019, arts. 1–5 (providing the French government's guidelines for compliance with the GDPR and requiring renewed consent after thirteen months); see also Nat'l Data Prot. Comm'n, *Cookies and Other Tracking Devices: The CNIL Publishes New Guidelines* (July 23, 2019), <https://perma.cc/R8GE-635V>.

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.¹⁰³

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 toward 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 toward 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,”¹⁰⁴ 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 of 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 overreacting 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.¹⁰⁵ If this is the case, then the information-costs

¹⁰³ See GDPR, *supra* note 6, at 37 (“[T]he 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 §§ 1798.100–1798.199.100.

¹⁰⁴ See, e.g., Jack Schofield, *What Should I Do About All the GDPR Pop-ups on Websites?*, GUARDIAN (July 5, 2018), <https://perma.cc/6AXA-KP97>.

¹⁰⁵ A report from the Stigler Center advocates for “consumertarian defaults”—those preferred or expected by the majority of consumers—and high opt-out costs or opt-out procedures that would require firms to convince, rather than trick, consumers to opt out.

theory calls for additional efforts by policy makers to correct users' biased, uninformed beliefs. Such correction is unlikely to occur merely by enacting new default rules in the hope that people will be nudged toward their provisions.

Alternatively, policy makers could try to increase the number of informed users. Indeed, 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 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. If most uninformed opt-outs were driven by a failure to appreciate the benefit from the policy makers' protective default, the increased number of informed users would reduce opt-out rates. But if most uninformed users resisted data sharing and stuck with the default, because they overestimated the potential harm, then an increased number of informed users may actually result in *more* opt-out. Either way, reducing information costs improves the matching between different user types and the data protection regime that applies to them.

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,¹⁰⁶ and this evidence inspired a search for theoretical explanations. The behavioral economics literature viewed the evidence as consistent with decision makers' cognitive limitations. Based on that evidence and the behavioral interpretation, commentators called for using default rules as a solution to many social problems.¹⁰⁷

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

STIGLER COMM. ON DIGIT. PLATFORMS, STIGLER CTR. FOR THE STUDY OF THE ECON. & THE STATE, FINAL REPORT 234–37 (2019). 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?

¹⁰⁶ See, e.g., Beshears et al., *supra* note 10, at 230–33.

¹⁰⁷ See *supra* notes 3–8.

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 toward, or away from, a specific contribution level. 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.¹⁰⁸ 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.¹⁰⁹

Retirement savings defaults are somewhat sticky.¹¹⁰ 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.¹¹¹ A similar percentage of employees opted to increase, over time, from the common 3% contribution rate default.¹¹² Perhaps these were stop-and-think defaults

¹⁰⁸ See *supra* Part II.B.

¹⁰⁹ See Bubb & Warren, *supra* note 62, at 43–44.

¹¹⁰ See Madrian & Shea, *supra* note 7, at 1176 (“Automatic enrollment dramatically increases the average 401(k) participation rate.”).

¹¹¹ See, e.g., Nessmith et al., *supra* note 16, at 10 (predicting that voluntary enrollment increases from 32% to 59% over the course of the first three years of employment); VANGUARD GRP., *supra* note 16, at 35 (the participation rate in voluntary-enrollment, i.e., zero-default, plans was 60%).

¹¹² See, e.g., Nessmith et al., *supra* note 16, at 11 (“[A]fter 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 GRP., *supra* note 16, at 40 (noting the predominant 3% default).

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 increases the ability of policy makers and employers to help employees by choosing the right default. It also places a heightened burden on the default setter to choose the right default, since we cannot count on opt-outs to avoid the consequences of ill-chosen defaults. In choosing the optimal default, policy makers and employers should focus on the interest of the subgroup of employees who stick with the default and maximize the expected value for this group as long as the information and opt-out costs of other employees are not too high.

The relative stickiness of the retirement savings defaults has inspired support for default rules as a consumer protection technique in many other contexts. Such uncritical borrowing from the retirement savings context 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.¹¹³ The endorsement effect may thus be weaker. And in consumer markets, even if the lawmaker's default is proconsumer and should benefit from an endorsement effect, sellers try hard to shift uninformed beliefs away from the default and toward the seller's preferred option.¹¹⁴ Overall, the stickiness observed in the retirement savings context is probably not representative.

¹¹³ 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 PROC. NAT'L ACAD. OF SCI. 15201, 15204 (2012). See generally Kenworthy 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).

¹¹⁴ See Barr et al., *supra* note 78, at 43 (contrasting the employment and consumer contexts and arguing that market forces make defaults slippery, without emphasizing the information effects).

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.¹¹⁵ While there are some success stories, a recent meta-analysis suggests a relatively small effect of defaults in the environmental context.¹¹⁶ 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 2013 field experiment conducted by Jorge Araña and Carmelo León, subjects were attendees of different academic conferences and conventions across disciplines.¹¹⁷ 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, a switch to the green default increased the participation rate (in the carbon offset program) from 62% to 81%.¹¹⁸ The results of this study are most informative, when compared to a very similar 2012 study by Magnus Hennlock and Professors Åsa Löfgren, Peter Martinsson, and Thomas Sterner. The only

¹¹⁵ Simon Hedlin & Cass R. Sunstein, *Does Active Choosing Promote Green Energy Use? Experimental Evidence*, 43 *ECOLOGY L.Q.* 107 (2016) (describing results of experimental tests to promote selection of green electricity); Johan Egebark & Mathias Ekström, *Can Indifference Make the World Greener?*, 76 *J. ENVTL. ECONS. & MGMT.* 1, (2016) (same, for green printing); Felix Ebeling & Sebastian Lotz, *Domestic Uptake of Green Energy Promoted by Opt-out Tariffs*, 5 *NATURE CLIMATE CHANGE* 868 (2015) (green electricity); Zachary Brown, Nick Johnstone, Ivan Haščič, Laura Vong & Francis Barascud, *Testing the Effect of Defaults on the Thermostat Settings of OECD Employees*, 39 *ENERGY ECON.* 128 (2013) (green thermostat use); Isaac Dinner, Daniel G. Goldstein, Eric J. Johnson & Kaiya Liu, *Partitioning Default Effects: Why People Choose Not to Choose*, 17 *J. EXPERIMENTAL PSYCH.* 332 (2011) (green lightbulbs); Daniel Pichert & Konstantinos V. Katsikopoulos, *Green Defaults: Information Presentation and Pro-Environmental Behaviour*, 28 *J. ENVTL. PSYCH.* 63 (2008) (green electricity).

¹¹⁶ See Jachimowicz et al., *supra* note 3, at 176.

¹¹⁷ See Jorge E. Araña & Carmelo J. León, *Can Defaults Save the Climate? Evidence From a Field Experiment on Carbon Offsetting Programs*, 54 *ENVTL. & RES. ECON.* 613, 615–17 (2013).

¹¹⁸ *Id.* at 619.

difference was that, in the Löfgren study, subjects were attendees of a single academic conference—on environmental economics.¹¹⁹ In this study, the default choice had no statistically significant effect. (The participation rate, for a €10 carbon offset payment, was approximately 50%.)¹²⁰

The information-costs 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 policy using a green default, the strategy is more likely to succeed when the decision makers are uninformed.

This observation, however, cannot account for the limited success of green defaults in all contexts. Indeed, it would seem that, in many contexts, decision makers are uninformed. The problem, in these other contexts is likely unrelated to the question of information. Consider, again, the examples where green defaults have been used—to encourage purchasing electricity from clean, renewable sources, using energy-efficient lightbulbs, enduring a lower temperature on the office thermostat, and utilizing double-sided printing. In these applications, decision makers are asked to bear higher costs—monetary costs, or costs in time and convenience—to support the environmental goal. If these extra costs exceed the cost of opt-out, the default will not stick. To increase effectiveness, policymakers could increase the mechanical opt-out costs, or try to affect people’s preferences. Merely changing the content of the default, however, would not suffice.

¹¹⁹ See Åsa Löfgren, Peter Martinsson, Magnus Hennlock & Thomas Sterner, *Are Experienced People Affected by a Pre-Set Default Option—Results from a Field Experiment*, 63 J. ENVTL. ECON. & MGMT. 66, 67 (2012).

¹²⁰ *Id.* at 68.

CONCLUSION

This Article develops a new information-costs theory of default rules and uses this theory not only to help policy makers 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 nondefault 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 versus nondefault 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 possess. 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. In order to do so, lawmakers 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 policy maker, 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 policy maker needs to know the value of information, I_{Low} or I_{High} in our example. To calculate the value of information, the policy maker needs the same information that an uninformed individual has. (Recall that the uninformed individual calculates the value of information and thus decides whether to become informed.) And, when individuals hold inaccurate uninformed beliefs, the policy maker needs to know the perceived value of information. The policy maker 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 policy makers need to know whether they are trusted. If people trust the policy maker, then the policy maker can use the choice of default to inform people. And, as we have seen, this allows the policy maker 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 overly optimistic.

We therefore urge lawmakers to exercise more caution before relying on the stickiness of defaults.

APPENDIX

The Appendix generalizes and extends the numerical example of Part I, using a formal model. In Part A.I, we present our framework of analysis. In Part A.II, we analyze outcomes and welfare with Low default. In Part A.III, we analyze outcomes and welfare with High default. In Part A.IV, we compare the two defaults and provide guidance to policy makers 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 Part A.V.

A.I. FRAMEWORK

Consider a binary choice between two options that we will call Low and High. We normalize the net benefit from Low to 0. The benefits and costs generated by High differ across individuals. Specifically, a share $a \in [0, 1]$ of individuals enjoy a net benefit $B > 0$, whereas the remaining $1 - a$ 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 $x = 0$.)¹²¹ We assume that uninformed individuals hold accurate beliefs about the share a and about the parameters B and C . The implications of inaccurate beliefs are explored in Part A.V 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., $aB - (1 - a)C > 0$, and uninformed individuals decide to opt out of Low default (when $k < aB - (1 - a)C$); or (ii) the expected value of Low is higher, i.e., $aB - (1 - a)C < 0$, and uninformed individuals decide to opt out of High default (when $k < (1 - a)C - aB$). In our analysis, we assume, without loss of generality, that $aB - (1 - a)C \geq 0$.¹²²

A.II. 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 $k > aB - (1 - a)C$.) To

¹²¹ Of the $G(\hat{x})$ individuals who learn their type, $aG(\hat{x})$ learn that they are Type 2 and $(1 - a)G(\hat{x})$ learn that they are Type 1. A share $a(1 - G(\hat{x}))$ of individuals remain uninformed about their type and believe that with a probability a they are Type 2 and with probability $1 - a$ they are Type 1. This group of uninformed individuals can be further divided into the $a(1 - G(\hat{x}))$ Type 1s and the $(1 - a)(1 - G(\hat{x}))$ Type 2s. To summarize: There are four groups of individuals: Group 1, with a measure of $aG(\hat{x})$ who know that they are Type 2; Group 2 with measure $(1 - a)G(\hat{x})$ who know that they are Type 1; Group 3 with measure $a(1 - G(\hat{x}))$ who are Type 2 but are uninformed about their type; and Group 4 with measure $(1 - a)(1 - G(\hat{x}))$ who are Type 1 but are uninformed about their type.

¹²² The case where $aB - (1 - a)C < 0$, is captured by normalizing the High payoffs to be zero and redefining $\tilde{C} = B$ as the cost borne by a share a under Low, and $\tilde{B} = C$ as the benefit enjoyed by a share $1 - a$ under Low. The expected payoff in Low would then be: $(1 - a)\tilde{B} - a\tilde{C} \geq 0$. We can further redefine: $\tilde{a} = 1 - a$, and get $\tilde{a}\tilde{B} - (1 - \tilde{a})\tilde{C} \geq 0$.

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 $k \in (aB - (1-a)C, B)$, 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: $a(B-k) + (1-a) \cdot 0 - x = a(B-k) - x$. If the individual remains uninformed, she will stick with Low default and earn a payoff of zero. Therefore, individuals will become informed if and only if $a(B-k) - x > 0$, or $x < a(B-k)$. To summarize, when $k \in (aB - (1-a)C, B)$, a share $G(a(B-k))$ of individuals will become informed and opt out with probability a ; and a share $1 - G(a(B-k))$ will remain uninformed and stick with the Low default. For a given k , the opt-out rate is: $aG(a(B-k))$. In terms of welfare, for any $k \in (aB - (1-a)C, B)$,

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

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

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

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 a , and a welfare level of $W = a(B - k)$. 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 $k \in (aB - (1 - a)C, B)$: With perfect information, the opt-out rate is a and welfare is $W = a(B - k)$; with imperfect information the opt-out rate is $aG(a(B - k)) < a$ and welfare is $W = \int_0^{a(B - k)} (a(B - k) - x)g(x)dx$.

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

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

The role of information costs. Ours is an information-costs theory. We thus focus on the role that information costs play in the analysis, specifically how the magnitude of information costs affects opt-out rates and welfare. We begin with the intermediate and high opt-out costs scenarios. In these scenarios (where $F(aB - (1 - a)C) = 0$), 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 $F(aB - (1 - a)C) = 1$), 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 $aB - (1-a)C > 0$), all the uninformed individuals opt out. When information costs are low, many individuals become informed and only a share a of them opt out.

Formally, for any $k \leq aB - (1-a)C$, the opt-out rate is: $aG((1-a)(C+k)) + 1 - G((1-a)(C+k)) = 1 - (1-a)G((1-a)(C+k))$. 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 $F(aB - (1-a)C) = 0$, lower information costs reduce stickiness and increase welfare.
- (b) When $F(aB - (1-a)C) = 1$, lower information costs increase both stickiness and welfare.

A.III. 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 $W = aB - (1-a)C$.

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: $aB + (1-a) * (-k) - x = aB - (1-a) * k - x$. With High default, there will be no uninformed opt-out, regardless of k (since $aB - (1-a)C > 0$). An individual who remains uninformed will stick with High default and earn a payoff of $aB - (1-a)C$. Therefore, individuals will become informed if and only if $aB - (1-a) * k - x > aB - (1-a)C$, or $x < (1-a) * (C-k)$. To summarize, when $k < C$, a share $G((1-a) * (C-k))$ of individuals will

become informed and opt out with probability $1 - \alpha$; and a share $1 - G((1 - \alpha) * (C - k))$ will remain uninformed and stick with the High default. For a given k , the opt-out rate is: $(1 - \alpha) * G((1 - \alpha) * (C - k))$. In terms of welfare, for any $k < C$,

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

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: $W = aB - (1 - \alpha) * k$. When $k \geq C$, the opt-out rate is zero and welfare equals $W = aB - (1 - \alpha)C$.

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 $W = aB - (1 - \alpha)C$, with both perfect and imperfect information.

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

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

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.

A.IV. COMPARISON: LOW DEFAULT VERSUS 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 $k \leq \min(aB - (1 - \alpha)C, C)$. 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 $x < (1 - \alpha) * (C - k)$. We can state the following result.

LEMMA 3: INFORMATION ACQUISITION FOR LOW DEFAULT
VERSUS HIGH DEFAULT

Low default induces more information acquisition than High default.

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: $1 - (1 - a)G((1 - a)(C + k))$. With High default, the opt-out rate is: $(1 - a) * G((1 - a) * (C - k))$. We see that either rule can be stickier. But more can be said. Let $\Delta(k) = 1 - (1 - a)[G((1 - a)(C + k)) + G((1 - a) * (C - k))]$ 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 $G((1 - a)(C + k)) + G((1 - a) * (C - k)) = 2$, the opt-out rate is higher with Low default if $a > 1/2$, and higher with High default if $a < 1/2$. When information costs are very high, i.e., when $G((1 - a)(C + k)) + G((1 - a) * (C - k)) = 0$, the opt-out rate is higher with Low default. These and other results are summarized in the following lemma.

LEMMA 4: OPT-OUT RATES FOR LOW DEFAULT VERSUS HIGH
DEFAULT

(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 $a > 1/2$, and higher with High default if $a < 1/2$.

(b) When the share of Type 1 individuals is higher, i.e., when $1 - a$ is larger, $\Delta(k)$ is lower. When $1 - a$ is small, the opt-out rate is higher with Low default. When $1 - a$ 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:

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

With High default, welfare is:

$$W = \int_0^{(1-a) * (C-k)} (aB - (1-a) * k - x)g(x)dx \\ + (1 - G((1-a) * (C-k)))(aB - (1-a)C)$$

When information costs are high (i.e., above $(1-a)(C+k)$), the welfare comparison is determined by the difference: $[aB - (1-a)C - k] - [aB - (1-a)C] = -k$. 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 $aB - (1-a)C > 0$). 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 $(1-a)(C-k)$), the welfare comparison is determined by the difference: $[a(B-k) - x] - [aB - (1-a) * k - x] = (1-2a) * k$. Therefore, welfare is higher with Low default when $a < 1/2$, and welfare is higher with High default when $a > 1/2$. 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 $a < 1/2$, Low default is stickier and generates more welfare; when $a > 1/2$, 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, $x \in ((1-a) * (C-k), (1-a) * (C+k))$, the welfare comparison is determined by the difference: $[a(B-k) - x] - [aB - (1-a)C] = (1-a)C - ak - x$. Therefore, when $x < (1-a)C - ak$, Low default is more efficient; and when $x > (1-a)C - ak$, 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 0. 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: WELFARE LEVELS FOR LOW DEFAULT VERSUS
HIGH DEFAULT

- (a) With high information costs, High default is more efficient, and stickier.
- (b) With low information costs, when $\alpha < 1/2$, Low default is more efficient, and stickier; and when $\alpha > 1/2$, 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 Proposition 2(a) and 2(b) are not surprising. Proposition 2(a) states that, when parties remain uninformed, policy makers should prefer the default rule that maximizes expected value and thus tracks the preferences of the uninformed majority—to avoid costly uninformed opt-out. And Proposition 2(b) instructs the policy maker to follow the preferences of the informed majority and thus avoid costly informed opt-out, when parties are informed. In both Propositions 2(a) and 2(b), stickiness goes hand in hand with efficiency, since stickiness implies fewer costly opt-outs—informed or uninformed. Proposition 2(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.

A.V. 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 0, 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 if and only if $\hat{\alpha}(\hat{B} - k) - x > 0$, or $x < \hat{\alpha}(\hat{B} - k)$. A share $G(\hat{\alpha}(\hat{B} - k))$ will become informed and opt out with probability α ; and a share $1 - G(\hat{\alpha}(\hat{B} - k))$ will remain uninformed and stick with the Low default (opt-out rate of 0). These results are summarized in the following lemma.

LEMMA 1A: LOW DEFAULT; LOW OPT-OUT COSTS; INACCURATE BELIEFS

When $\hat{\pi} < 0$, for any $k \leq |\hat{\pi}|$: With perfect information, the opt-out rate is α and welfare is $W = \alpha(B - k)$; with imperfect information the opt-out rate is $\alpha G(\hat{\alpha}(\hat{B} - k))$ 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 0 to $\alpha(B - k) - x$. When $\alpha(B - k) < \hat{\alpha}(\hat{B} - k)$ and $x \in (\alpha(B - k), \hat{\alpha}(\hat{B} - k))$, the lower

information costs reduce welfare. (The identified perverse effect of lowering information costs requires $\hat{a} > a$ or $\hat{B} > B$. Our analysis focuses on inaccurate beliefs that imply $\hat{\pi} = \hat{a}\hat{B} - (1 - \hat{a})\hat{C} < 0$, instead of the accurate $\pi = aB - (1 - a)C > 0$. 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: 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: $aB - (1 - a) * k - x$, and the perceived payoff is: $\hat{a}\hat{B} - (1 - \hat{a}) * k - x$. An individual who remains uninformed will opt out to Low and earn a payoff of $-k$. Therefore, individuals will become informed if and only if $\hat{a}\hat{B} - (1 - \hat{a}) * k - x > -k$, or $x < \hat{a}(\hat{B} + k)$. A share $G(\hat{a}(\hat{B} + k))$ of individuals will become informed and opt out with probability $1 - a$; and a share $1 - G(\hat{a}(\hat{B} + k))$ 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 $1 - a$ and welfare is $W = aB - (1 - a) * k$; with imperfect information the opt-out rate is $(1 - a) * G(\hat{a}(\hat{B} + k)) + 1 - G(\hat{a}(\hat{B} + k)) = 1 - aG(\hat{a}(\hat{B} + k))$ and welfare is

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

As with Low default, here too lower information costs might reduce efficiency. Specifically, when information costs are reduced from $\hat{a}(\hat{B} + k) + \varepsilon$ to $\hat{a}(\hat{B} + k) - \varepsilon$, welfare changes from $-k$ to $aB - (1 - a) * k - x$. When $a(B + k) < \hat{a}(\hat{B} + k)$ and $x \in (a(B + k), \hat{a}(\hat{B} + k))$, 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, $x \in (\hat{a}(\hat{B} - k), \hat{a}(\hat{B} + k))$, we find that now High default generates stronger incentives to acquire information. The welfare comparison is determined by the difference: $[0] - [aB - (1 - a) * k - x] = -aB + (1 - a) * k + x$. Therefore, when $x < aB - (1 - a) * k$, High default is more efficient; and when $x > aB - (1 - a) * k$, 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 $1 - a$, and with Low default it is 0. 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: WELFARE LEVELS WITH INACCURATE BELIEFS
FOR LOW DEFAULT VERSUS HIGH DEFAULT**

- (a) With high information costs, Low default is more efficient, and stickier.
- (b) With low information costs, when $a < 1/2$, Low default is more efficient, and stickier; and when $a > 1/2$, 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 policy maker 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 policy maker should thus prefer High default when information costs are at the low end of this range, and Low default when information costs are higher.