University of Chicago Law School

Chicago Unbound

Coase-Sandor Working Paper Series in Law and Economics Coase-Sandor Institute for Law and Economics

2016

Optimal Defaults in Consumer Markets

Oren Bar-Gill

Oren.BarGill@chicagounbound.edu

Omri Ben-Shahar

dangelolawlib+omribenshahar@gmail.com

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

Part of the **Law Commons**

Recommended Citation

Oren Bar-Gill & Omri Ben-Shahar, "Optimal Defaults in Consumer Markets," Coase-Sandor Working Paper Series in Law and Economics, No. 764 (2016).

This Working Paper is brought to you for free and open access by the Coase-Sandor Institute for Law and Economics at Chicago Unbound. It has been accepted for inclusion in Coase-Sandor Working Paper Series in Law and Economics by an authorized administrator of Chicago Unbound. For more information, please contact unbound@law.uchicago.edu.

CHICAGO

COASE-SANDOR INSTITUTE FOR LAW AND ECONOMICS WORKING PAPER NO. 764

COASE-SANDOR INSTITUTE
FOR LAW AND ECONOMICS
THE UNIVERSITY OF CHICAGO LAW SCHOOL

OPTIONAL DEFAULTS IN CONSUMER MARKETS

Oren Bar-Gill & Omri Ben-Shahar

THE LAW SCHOOL
THE UNIVERSITY OF CHICAGO

June 2016

ISSN 1936-5349 (print)
ISSN 1936-5357 (online)

HARVARD

JOHN M. OLIN CENTER FOR LAW, ECONOMICS, AND BUSINESS

OPTIMAL DEFAULTS IN CONSUMER MARKETS

Oren Bar-Gill
Omri Ben-Shahar

Discussion Paper No. 863

06/2016

Harvard Law School
Cambridge, MA 02138

This paper can be downloaded without charge from:

The Harvard John M. Olin Discussion Paper Series:
http://www.law.harvard.edu/programs/olin_center/

The Social Science Research Network Electronic Paper Collection:
<http://ssrn.com/abstract=2793128>

Optimal Defaults in Consumer Markets

Oren Bar-Gill and Omri Ben-Shahar^{*}

Abstract

The design of default provisions in consumer contracts involves an aspect that does not normally arise in other contexts. Unlike commercial parties, consumers have only limited information about the content of the default rule and how it fits with their preference. Inefficient default rules may not lead to opt outs when they deal with technical aspects consumers rarely experience and over which consumers' preferences are defined only crudely. This paper develops a model in which consumers are uninformed about their preferences, but can acquire costly information and then choose a contract term that best matches their preferences. The paper explores the optimal design of default rules in such environments, and how it differs from the existing conceptions of efficient default rule design.

^{*} Harvard Law School and University of Chicago Law School, respectively. For helpful comments and suggestions, we thank Barry Adler, Jennifer Arlen, Ian Ayres, Scott Baker, Lucian Bebchuk, Ryan Bubb, Emiliano Catan, Giuseppe Dari-Mattiacci, Kevin Davis, Barry Friedman, Scott Hemphill, Marcel Kahan, Lewis Kornhauser, Daryl Levinson, Florencia Marotta-Wurgler, Geoff Miller, Mark Ramseyer, Cathy Sharkey, Steve Shavell, workshop participants at NYU and conference participants at the JLS symposium on "Contracting for Privacy" at The University of Chicago Law School.

I. Introduction

A. Default Rules in Consumer Contracts

The design of default rules for contracts is one of the most thoroughly studied areas in the economics analysis of contract law (see, e.g., Goetz and Scott 1985, Shavell 2004, Ayres and Gertner 1989, Bebchuk and Shavell 1991). The fundamental insight on which the literature in this area rests is the saving of opt-out costs. Properly designed default rules save the parties the cost of drafting their own terms. We argue in this article that this transactions-cost framework fits poorly the area of consumer contracts, and we offer a novel extension to address this shortcoming.

The scenario motivating the standard transactions-costs account is the individually negotiated agreement. Parties must write down some specific terms – such as quantity, quality, price, and method of payment – but it would be costly for them to also write down all the contingent terms governing issues such as warranties, remedies, excuses, and dispute resolution. Default rules set by lawmakers to mimic what most parties would have chosen eliminate drafting costs and make more exchanges possible. If the default rules fail to mimic the preferences of some parties (with atypical preferences), opt out would occur. Similarly, default boilerplate terms pre-drafted by firms would be designed to maximize the joint surplus from the typical transaction and reduce redrafting costs.

But this theory becomes less plausible in consumer markets. Rather than two informed parties that together decide whether to opt-out of a default rule set by the lawmaker, one party, the business, can unilaterally opt-out. Is it even meaningful for the law to provide default rules given the reality in which businesses regularly and costlessly attach to each consumer transaction a long boilerplate which “deletes” the legally provided defaults, replacing them with pre-drafted comprehensive terms? (See Radin 2013, Kim 2013) The consumer enters the picture only after this first-stage opt-out by the business and could potentially opt-out of the boilerplate terms, e.g., by purchasing from a competing seller.

In the consumer contract setting, the design of defaults changes in two fundamental ways. First, the value of legally provided default rules is questionable given the quasi-legislative role that businesses have assumed, in privately drafting defaults. Accordingly, we should focus either on defaults drafted by businesses or on legal rules that force businesses to offer menus and defaults. Consumers can then opt-out of the default, which brings us to the second fundamental departure from standard theory: Consumers lack the information necessary to make optimal opt-out decisions.

One fundamental ingredient that consumers lack is precise knowledge of their preferences regarding the issues governed by default rules. These issues are often technical, complex, and numerous. Unlike salient product features (e.g., memory capacity of a laptop), those that deal with legal rights (e.g., data security or dispute resolution) are rarely invoked. Even if people intuitively know that they prefer pro-consumer terms (e.g., a broad warranty), they lack the information to make quality-price trade offs (e.g., to make the proper trade-off, the consumer would need to compare the expected value of a

better term with its effect on the price. It is a computation that requires information on values and probabilities that people rarely have, especially given the multitude of such concurrent issues. This difficult tradeoff might also be biased by firms' occasional incentive to "push" features that consumers don't really need (Baker and Siegelman 2013; Camerer et al 2013, at 1253-54).)

In addition, consumers will not be able to make efficient opt-out decisions if they do not understand the default provisions. These are notoriously complex and require much knowledge to interpret. What does it mean when the legally supplied default rule says, for example, that buyers are entitled to "consequential damages" in the event of breach? Or that buyers are entitled to goods that are "merchantable"? And conversely, what does it mean when the business-supplied default boilerplate disclaims such legal protections "to the maximal permissible extent"?

Contracting over privacy is particularly prone to these information problems. True, people may have an intuitive preference for data privacy and may worry about data security (Westin 1997, Pew Research Center 2013). But the strength of such preferences varies; there is not much evidence that consumers back their intuitive statements with any significant willingness to pay for anonymity or for greater protection (Hann et al 2007; Savage and Waldman 2013). Even if consumers know their preferences, they rarely know the default provision governing data collection, or spend the time to read the terms governing their transaction. Absent such knowledge, consumers cannot know if their "silence"—their failure to opt-out and seek a different arrangement—would hurt them.

In these environments of imperfect information, how should default rules be designed? In the standard model, which assumes that parties know their preferences, the design of optimal defaults is guided by the opt-out cost minimization principle. Our model cannot be based on this principle alone, because when consumers are initially uninformed there is a preliminary decision whether to invest some cost and become informed.

Because information is costly, some consumers may decide to remain uninformed. Then, their choices are based on expectations about "average preferences." We consider what happens when consumers form rational expectations – reflecting the true ex ante distribution of preferences among all consumers. Separately, we also consider the possibility of systematic bias in the formation of such estimates over average preferences. Either way, uninformed consumers have to decide whether to opt out of the default rule, and may do so if they conclude that it is ill-matched with their estimate of average preferences. We call this "uninformed opt out."

Other consumers may decide to become informed about their preferences so as to make a more accurate opt-out decision. Because "informed opt-out" is better tuned to serve a consumer's preferences than uninformed opt-out, there is a value to acquiring information, and so some consumers—those for whom the cost of information is low enough—will become informed.¹

¹ In some cases, consumers may acquire information over time, simply by using the default product with the default terms. The consumer can then opt out by switching to another product. In

In consumer markets, the concepts of “default” and “opt-out” have different meaning compared to the typical business-to-business context. As noted above, the relevant default is either a provision drafted by the business or an option, or menu of options, that the law forces the business to offer. Either way consumers must decide whether to accept the default or opt-out. This opt out can take several forms. First, a business may offer – voluntarily or by legal mandate – a menu of options, with one of the options pre-selected as the default, and allow the consumer to opt-out and select a different option. (The business may also offer a menu of options without pre-selecting any one of them.) For example, many websites and mobile apps default consumers into a pre-set level of privacy protection, but users can navigate through the settings menus and opt-out. Second, the business may offer multiple products, some of which are presented more prominently and effectively become the equivalent of the “pre-clicked” default. The consumer may select the default product, or opt out and choose another product. Finally, even if the business does not allow any opt outs from its preset default terms and features, consumers who choose to buy a similar product elsewhere under different terms are effectively opting out.

Because consumer opt-out can occur through several different procedures, the magnitude of opt-out costs varies – from very small, when the consumer can opt-out by clicking on a different item in an easily accessible menu of options, to quite large, when the only way to opt out is by searching for a new product with new terms offered by a different seller. As explained below, the magnitude of opt-out costs will critically affect the design of the optimal default.

B. Designing Default Rules in Consumer Markets

We evaluate each default rule by how well it matches consumers with their preferences, and by the costs consumers incur to improve this match. Under each possible default rule, a different group of consumers spends resources to become informed and a subset of those spends additional resources and opts out. The group of consumers who choose to remain uninformed may or may not opt out; for them, only average preferences are satisfied and a subgroup among them are ill-matched.

This framework allows us to look beyond the traditional majoritarian idea and identify two separate foundational principles. The first is the (majoritarian) principle of *opt-out cost minimization*, which prescribes a default rule that induces opt out by the fewest people and reduces the waste of contracting out. The second principle is *expected value maximization*. This principle tells us which rule is superior, assuming *no opt out*. In our model, both principles have strong appeal. The first principle—opt out cost minimization—is important because a default rule that is better matched with people’s preferences induces less investment not only in opting out, but also in acquisition of information by consumers. The second principle—expected value maximization—is important because some consumers remain uninformed and do not opt out even from a

such cases, the information acquisition cost is the cost of possibly using a suboptimal product during the learning period.

default that does not match their preferences, thus following this principle reduces the mismatch cost due to the stickiness of default rules.

The optimal default choice, then, depends on the relative weight of opt-out cost minimization versus expected-value-maximization. One way to understand this balance is as a tradeoff between the *distribution of preferences* versus the *intensity* of preferences. The distribution aspect helps design a default rule that would reduce the cost of transacting around it, whereas the intensity aspect helps design a rule that would mimic the intensity-weighted preferences.

The model draws out prescriptions that differ, in certain scenarios, from the standard model. To appreciate these differences, it is helpful to distinguish between two qualitative scenarios: high versus low opt-out costs.

When opt-out costs are high, opt out does not occur and people do not acquire information. In such cases, as in the standard model (which does not consider information costs), the most efficient arrangement is prescribed solely by the expected value maximization principle—the same criterion that applies to the design of a mandatory rule.

When opt-out costs are low, both informed and uninformed opt-outs occur. In the standard model, the optimal default would be determined solely by the opt-out cost minimization principle (because all mismatched consumers opt out). In our model, taking into account information costs, this principle has to be balanced against the expected value maximization principle. When information-acquisition costs are relatively high, fewer consumers become informed and uninformed opt out becomes more likely. The expected value maximization principle matches better average preferences and minimizes the incidences of uninformed opt out as well as the cost of acquiring information. Additional subtle incentives arise when opt-out costs are intermediate, such that only informed opt-out occurs. Here, too, some consumers remain uninformed and for them a default rule based on the expected value maximization principle would yield a superior result. As the cost of acquiring information rises, fewer consumers become informed and thus informed opt out becomes less likely. In such cases, the expected value maximization principle alone determines the optimal default rule.

Importantly, the choice of default rule endogenously affects the value of becoming informed about preferences, and a default rule for which the value of information is higher is more desirable because it leads to more informed and thus more efficient opt-out. Our model draws out the factors that affect the value of becoming informed. For example, we see two effects crossing. The first corresponds to the expected value maximization principle, operating “in reverse,” and can be thought of as an information-eliciting effect: the value of information under each default rule increases the more consumers have (uninformed) preferences that don’t match it. The second effect corresponds to the opt-out cost minimization principle: When opt-out costs are higher, the value of becoming informed diminishes, because information is less likely to lead to opt out.

The framework also allows us to explore additional factors. In particular, consumers' decisions whether to become informed may be distorted by misperceptions, such as overestimating the likelihood that they value a high-quality default (overlooking the price effect), or underestimating other high quality defaults such as greater privacy protection (overlooking non salient costs). We show how such misperceptions distort consumers' decisions to acquire information and thus reduce social welfare.

Our analysis has both a positive interpretation and a normative one. It can be read as a descriptive account of how a business will set default rules in its boilerplate consumer contract. And it can also be read as a normative account of how a lawmaker should design rules that force businesses to offer menus and defaults to their customers.

Our paper is structured as follows. Section II presents our basic insights through a numerical example. Section III presents the formal model, compares the different default rules, calculates welfare implications, and identifies the factors that determine which default is optimal. In Section IV, we briefly explore the implications of consumer misperception.

II. Informal Analysis

Consider a market with two types of consumers. Some, labeled H, assign high value to privacy protections, and others, labeled L, assign low value. For simplicity, assume that there are only two possible levels of protection, High and Low. Assume that the value of High protection is \$100 for H consumers and \$0 for L consumers. The cost to sellers of providing High protection is \$30 and the cost of providing Low protection is 0. Assume also that H consumers constitute 40% of the market, and Ls the remaining 60%.

A contract stipulates two things: the level of protection and the price. The level of protection is initially determined by a legally supplied default. We consider two possible default rules, "L Default" and "H Default." This level can be changed by contract. Consumers can individually opt out and agree with the firm on a different level of protection, at a different price. We assume that sellers are operating in a competitive market, such that price exactly equals to expected cost. Accordingly, sellers charge a price of \$30 for High protection and \$0 for Low protection.

Our analysis differs from prior work on the following key assumption: Consumers do not know their individual preferences—whether they are H or L consumers—unless they spend some upfront cost to figure this out. Otherwise, they only know the distribution of types—namely, the 40% chance that they value High protection at \$100. Consumers choose to incur this information cost only if it is less than the value of the information (which we derive below).

Consumers – informed and uninformed – can opt out of the default, and pay an adjusted price. (As explained, such opt out occurs by choosing a different bundle from the same seller, or by switching to another firm.) We divide the analysis into three cases:

1. High opt-out costs

Assume that the transaction cost to opt-out is \$75. With such high opt-out costs, consumers will never opt-out. An H consumer facing an L Default would gain \$70 ($= \$100 - \30) from opting out; and an L consumer facing an H Default would gain \$30 from opting out. In both cases, the opt-out cost outweighs the benefit from opting out. Since there is no value in becoming informed, all consumers remain uninformed and stick with the default. Therefore, the default should be designed according to the expected value maximization principle. In our example, H Default is optimal, since it generates an expected value of \$10 ($= 40\% * \$100 - \30), as compared to L default that generates an expected value of zero.

2. Intermediate opt-out costs

Assume that the transaction cost to opt-out is \$15. Let us calculate the value of information under each default rule. Under L Default, uninformed consumers do not opt out because the expected net gain is less than the opt out cost. The expected gain from switching to High protection is \$40 (40% of \$100) and the price increase is \$30, for a net gain of \$10. But the opt-out cost of \$15 more than wipes out this net gain. Informed consumers behave differently: those who discover that they are H consumers opt out and enjoy a payoff increase of \$55 (\$100 minus \$30 price increase, minus \$15 opt-out cost); and those who discover that they are L consumers stick with the default rule and get a payoff of \$0. Thus, consumers expect a 40% chance of increasing their ex-post payoff from \$0 to \$55. The value of information is \$22 ($40\% \times \55).

Under H Default, uninformed consumers again do not opt out, because the expected net payoff from sticking with High protection is \$10 (40% of \$100 minus \$30 price), and there is no reason to spend the transaction cost of \$15 and shift to Low protection that would yield an expected payoff of \$0. Informed consumers again behave differently: those who discover that they are L consumers opt out and enjoy a net payoff increase of \$15 (they no longer have to pay the price of \$30 for the protection that they do not value, but they do incur an opt out cost of \$15); and those who discover that they are H consumers stick with the default rule. Thus, consumers expect a 60% chance of increasing their payoff by \$15. The value of information is \$9 ($60\% \times \15).

Notice that the value of information in this example is \$22 under L Default and only \$9 under H Default. The value of information is higher under L Default for two reasons: first, the informed opt-out that would occur under L Default is more valuable than the informed opt out that would occur under H Default, because the 40% of consumers who would shift from L to H would gain \$70 whereas the 60% of consumers who would shift from H to L would gain only \$30. This is a difference in expected value of \$10 ($40\% \times \$70 - 60\% \times \$30 = \10). Second, informed opt out is less frequent under L Default,

imposing the transaction cost of \$15 only 40% of the time, as compared to 60% under H Default. This is a difference in value of \$3.

We can now turn to compare the social value of the two default rules. Since the comparison depends on the incidence of opt out, which in turn depends on the incentive to acquire information, we must first specify the cost of information. Assume that information costs vary across consumers. Specifically, assume that half of the consumers (independent of their type) can become informed at a cost of \$8 and the other half at a cost of \$16.

Under L Default, all consumers acquire the information. Even if the cost is \$16, it is still less than the value of information, which we saw equals \$22. The expected total welfare is:

$$40\% \times (100 - 30 - 15) - (50\% \times 8 + 50\% \times 16) = \$10$$

It equals the expected payoff from perfectly tailored ex post protection, minus the cost of opt out for H consumers (who all opt out because they all become informed), minus the cost of information incurred by all consumers.

Under H Default, only consumers with information cost of \$8 acquire the information. Those with information cost of \$16 will not become informed, because the value of information under this rule is only \$9. The expected total welfare under H Default is:

$$50\% \times [40\% \times (100 - 30) + 60\% \times (-15) - 8] + 50\% \times [40\% \times 100 - 30] = \$10.5$$

Half of consumers acquire information at a cost of \$8, and among them 60% opt out after learning that they are L consumers. The other half remain uninformed and stick with the default High protection.

In this case, the total welfare is higher under H Default, even though the value of information is lower, fewer consumers acquire information, and fewer consumers end up matched with their privately optimal level of protection. The reason H Default does better here is that it reduces transactions costs. First, H consumers get their preferred outcome with less wasteful acquisition of information. Second, there is less incidence of costly opt out under H Default, because only half of the L types become informed and opt out. These two advantages – saving information costs and opt out costs – more than offset the downside of H Default, which leaves half of L consumers with an inefficiently high level of protection.

Effect of lower information costs. We just saw that H Default is superior, in part because it requires less wasteful acquisition of information. What happens when information costs are lower? Assume, as before, that information costs vary across consumers, and that some people can still spend only \$8 to become informed. But now assume that for those who have to spend the higher cost of information, the cost is not \$16 but rather \$12.

The expected welfare under H Default does not change, because consumers with high information costs remain uninformed, as before (the value of information to them is still only \$9), and so the decline in the high-cost of information does not affect them. Expected welfare remains \$10.5.

The expected welfare under L Default changes – it goes up. This is because consumers do spend the high cost of information, but now they only have to spend \$12, not \$16. Now expected welfare under this rule is:

$$40\% \times (100 - 30 - 15) - (50\% \times 8 + 50\% \times 12) = \$12$$

Total welfare is now higher under L Default. Here is why the comparison reversed: Like before, L Default continues to lead to better ex post levels of protection (all consumers get their preferred levels), and like before it continues to impose higher transactions costs—more opt out costs and more information acquisition costs. But now these higher costs are less burdensome given the assumption of lower information costs for half of the population.

3. Low opt-out costs

Assume that the transaction cost to opt-out is \$5. With such low opt-out costs, both informed and uninformed consumers opt out. Return to the case in which information cost is {8, 16}. When we reduce opt-out costs from \$15 to \$5, welfare under L Default rises to:

$$40\% \times (100 - 30 - 5) - (50\% \times 8 + 50\% \times 16) = \$14$$

(Note that, with the lower opt-out cost, uninformed opt-out becomes a viable option. But, in this example, informed opt-out is more attractive, even for consumers with high information costs. These consumers enjoy a payoff of $40\% \times (100 - 30 - 5) - 16 = \10 , if they become informed; whereas they would get $40\% \times 100 - 30 - 5 = \5 , if they opt-out without becoming informed.)

Welfare under H Default also rises, to:

$$50\% \times [40\% \times (100 - 30) + 60\% \times (-5) - 8] + 50\% \times [40\% \times 100 - 30] = \$13.5$$

Welfare increased under both rules when opt out costs went down, but the increase was larger under L Default because this rule led to more opt out (due solely to the fact that the value of information is higher under this rule), and with opt out being less costly, the overall performance of this rule is now superior.

Effect of higher information costs. Assume, as before, that information costs vary across consumers, and that some people can still spend only \$8 to become informed. But now assume that for those who have to spend the higher cost of information, the cost is not \$16 but rather \$24. Now, with L Default, consumers with high information costs will

choose uninformed opt-out (which gives them a payoff of $40\% \times 100 - 30 - 5 = \5 , as compared to a payoff of $40\% \times (100 - 30 - 5) - 24 = \2 if they become informed). Welfare under L Default becomes:

$$50\% \times [40\% \times (100 - 30 - 5) - 8] + 50\% \times [40\% \times 100 - 30 - 5] = \$11.5$$

Welfare under H Default remains: \$13.5. The high information costs eliminate L Default's information advantage and saddle L Default with additional transaction costs from uninformed opt-out. As a result, the optimal default flips back to H Default.

III. Formal Analysis

A. Framework

A certain product includes a binary quality dimension $q \in \{L, H\}$. Denote by c the cost to the seller of providing a product with $q = H$; the cost of providing $q = L$ is zero.² The seller operates in a competitive market and thus sets two prices – a price of zero for L quality and a price c for H quality.³

The benefit to consumers from a product with $q = L$ is zero. The benefit from a product with $q = H$ varies among consumers. For a share $\alpha \in [0, 1]$ of consumers (H types), the benefit is V ; and for a share $1 - \alpha$ of consumers (L types), the benefit from H quality is zero. We assume that $V > c$.

Initially, consumers do not know whether they are H types or L types. Consumers can invest x and learn their type. The investment x varies among consumers, according to cumulative distribution function $F(\cdot)$. There is a threshold $\hat{x}$ (derived below), such that consumers with $x < \hat{x}$ invest and learn their type, while consumers with $x \geq \hat{x}$ remain uninformed. (This framework covers scenarios where some consumers initially know their type; in such scenarios the probability function would have a mass point at $x = 0$. This framework covers the standard model, where all parties initially know their type, as a special case where $F(0) = 1$.)

Of the $F(\hat{x})$ consumers who learn their type, $\alpha F(\hat{x})$ learn that they are H types and $(1 - \alpha)F(\hat{x})$ learn that they are L types. A share $1 - F(\hat{x})$ of consumers remain uninformed about their type and believe that with a probability α they are H types and

² In some cases, higher quality can be provided at no additional cost. In such cases, $c = 0$ and the seller will charge the same price for $q = L$ and for $q = H$. For example, some products allow the consumer to opt out from the default privacy setting to a higher (or lower) privacy setting without any price adjustment.

³ The perfect competition assumption is in tension with at least one of our interpretations of opt-out costs – as the cost of switching from one seller to another.

with probability $1 - \alpha$ they are L types. This group of uninformed consumers can be further divided into the $\alpha(1 - F(\hat{x}))$ H types and the $(1 - \alpha)(1 - F(\hat{x}))$ L types.

To summarize: There are four groups of consumers – Group 1, with a measure of $\alpha F(\hat{x})$ who know that they are H types; Group 2 with measure $(1 - \alpha)F(\hat{x})$ who know that they are L types; Group 3 with measure $\alpha(1 - F(\hat{x}))$ who are H types but are uninformed about their type; and Group 4 with measure $(1 - \alpha)(1 - F(\hat{x}))$ who are L types but are uninformed about their type. We assume that the seller does not know the consumer's type and whether the consumer acquired information.

We consider two possible default rules: L Default, where the seller offers $q = L$ as the default; and H Default, where the seller offers $q = H$ as the default. Consumers can opt out of either default at a cost k (borne by the consumer), and pay an adjusted price reflecting the individually contracted quality level.

B. Designing Optimal Defaults

The first question is whether a consumer decides to become informed. As a function of this decision, we then have either informed or uninformed opt-out:

- Informed opt-out: Consumers who invest x and learn that they are H types (Group 1) would opt out of L Default, if $k < V - c$. And, consumers who invest x and learn that they are L types (Group 2) would opt out of H Default, if $k < c$.
- Uninformed opt-out: When $\alpha V > c$, uninformed consumers would opt-out of L Default if $k < \alpha V - c$. And, when $\alpha V < c$, uninformed consumers would opt-out of H Default if $k < c - \alpha V$.

We focus on the three cases where both defaults generate the same type of opt out: (1) both defaults lead to no opt-out (high k), (2) both defaults lead to only informed opt-out (intermediate k), and (3) both defaults lead to both informed and uninformed opt-out (low k).

1. High k : No opt-out

When opt-out costs are sufficiently high, there will be neither informed opt-out nor uninformed opt-out. The no opt-out scenario occurs when $k > \max(c, V - c)$. In this scenario, no consumer would become informed, because absent any potential opt out information has no value.

In the absence of opt-out, the optimal default rule is the one that maximizes expected value across all consumers. In particular, when $\alpha V - c < 0$, L Default is optimal; and when $\alpha V - c > 0$, H Default is optimal. The optimal default is determined solely by the expected value maximization principle.

In the high- k scenario, our model and the standard model result in an identical prescription. In the standard model, consumers know their preferences (without investing in information acquisition) but, since there is no opt-out, this informational difference is irrelevant. The expected value maximization principle dictates the optimal default rule.

2. Intermediate k : Only informed opt-out

In the intermediate- k scenario, opt-out costs allow for informed opt-out, but not uninformed opt-out. Formally, this scenario occurs, when opt-out costs satisfy: $\max(\alpha V - c, c - \alpha V) < k < \min(c, V - c)$. Note that the option to opt-out is more valuable for an informed consumer than for an uninformed consumer. That is why we have a range of (intermediate) opt-out costs, for which only informed opt-out occurs.

We analyze market outcomes and welfare levels with L Default (in subsection a) and with H Default (in subsection b). We then use these results to guide the choice of the optimal default rule (in subsection c).

(a) L Default

Informed consumers opt out of L Default if they are H type. Let us calculate the expected benefit from becoming informed. The uninformed consumer chooses $q = L$ and enjoys an expected benefit of zero. The informed consumer enjoys an expected benefit of $\alpha(V - c - k)$. The value of information is $I^L = \alpha V - c - k$. Therefore, consumers with $x < I^L$ invest x and learn their type, whereas consumers with $x \geq I^L$ remain uninformed. Social welfare is given by:

$$W^L = \int_0^{I^L} (I^L - x)f(x)dx$$

Define a general function $W(I) \equiv \int_0^I (I - x)f(x)dx$ and notice that $W(0) = 0$ and that $W(I)$ is increasing in I : $W'(I) = \int_0^I f(x)dx = F(I) > 0$. We can write: $W^L = W(I^L)$.

(b) H Default

Informed consumers opt out of H Default if they are L type. We first calculate the value of information. The uninformed consumer sticks to the default and gets $\alpha V - c$ (which can be either positive or negative). The informed consumer enjoys an expected benefit of $\alpha(V - c) - (1 - \alpha)k$. The value of information is $I^H = [\alpha(V - c) - (1 - \alpha)k] - [\alpha V - c] = (1 - \alpha)(c - k)$. Social welfare is given by:

$$\begin{aligned}
W^H &= \int_0^{I^H} [\alpha(V - c) - (1 - \alpha)k - x]f(x)dx + [1 - F(I^H)](\alpha V - c) = \\
&= \int_0^{I^H} [I^H - x]f(x)dx + (\alpha V - c)
\end{aligned}$$

Using the general function $W(I)$, we can write: $W^H = W(I^H) + (\alpha V - c)$.

(c) Comparing L and H Default

With L Default, social welfare is $W^L = W(I^L)$; and with H Default, social welfare is $W^H = W(I^H) + (\alpha V - c)$. To choose the optimal default, we must compare W^L and W^H . The comparison can be divided into two components:

- (1) Pre-information Welfare: Uninformed consumers stick with the default. With L Default, they get zero; with H Default they get $\alpha V - c$. When $\alpha V - c < 0$, L Default has a pre-information advantage. And when $\alpha V - c > 0$, H Default has a pre-information advantage.
- (2) Information-based Welfare: With L Default, the welfare generated by information acquisition and (possible) opt-out is $W(I^L)$; with H Default, the welfare is $W(I^H)$. When $I^L > I^H$, $W(I^L) > W(I^H)$ and L Default generates more information-based welfare. When $I^H > I^L$, $W(I^H) > W(I^L)$ and H Default generates more information-based welfare.

We must therefore compare the value of information with L Default, $I^L = \alpha(V - c - k) = \alpha(V - c) - \alpha k$, to the value of information with H Default, $I^H = [\alpha(V - c) - (1 - \alpha)k] - [\alpha V - c] = (1 - \alpha)c - (1 - \alpha)k$. We first look at the expected benefit from opt-out: $\alpha(V - c)$ is the benefit of opt-out from L Default, whereas $\alpha(V - c) + (c - \alpha V) = (1 - \alpha)c$ is the benefit of opt-out from H Default. When $\alpha V - c < 0$, the benefit of opt-out is larger by $c - \alpha V$ with H Default. When $\alpha V - c > 0$, the benefit of opt-out is smaller by $\alpha V - c$ with H Default. When $\alpha V - c = 0$, the benefit of opt-out is the same with both rules. We next look at the expected cost of opt out: αk is the cost of opt-out from L Default, and $(1 - \alpha)k$ is the cost of opt out from H Default. When $\alpha < \frac{1}{2}$, the cost of opt-out is larger by $(1 - 2\alpha)k$ with H Default. When $\alpha > \frac{1}{2}$, the cost of opt-out is smaller by $(2\alpha - 1)k$ with H Default. When $\alpha = \frac{1}{2}$, the cost of opt-out is the same with both rules. The overall comparison between I^L and I^H depends on both the relative benefits and the relative costs of opt-out, as detailed below.

To establish some basic intuition, we begin with the symmetric case where $\alpha = \frac{1}{2}$ and $\alpha V = c$. In this case, the benefit from opt-out and the cost of opt-out are the same with

both default rules; and there is no initial advantage or disadvantage to one rule ($\alpha V = c$). Therefore, the two defaults generate the same welfare level: $W^L = W^H$.

First, increase α (so that $\alpha > \frac{1}{2}$), while holding $\alpha V = c$. The benefit from opt-out is the same with both default rules, but the cost of opt-out is smaller with H Default. This means that the value of information is larger with H Default. And, since there is no pre-information advantage or disadvantage to one rule ($\alpha V = c$), H Default is the efficient rule. The results flip if we decrease α (so that $\alpha < \frac{1}{2}$), while holding $\alpha V = c$. As before, the benefit from opt-out is the same with both default rules, but now the cost of opt-out is smaller with L Default. This means that the value of information is larger with L Default. And, since there is no pre-information advantage or disadvantage to one rule ($\alpha V = c$), L Default is the efficient rule. To summarize, when $\alpha V = c$, the expected value maximization principle is neutral and the opt-out cost minimization principle solely determines the optimal default.

Next, we deviate from the symmetric case by moving away from $\alpha V = c$, while holding $\alpha = \frac{1}{2}$. The cost of opt-out is the same with both default rules (since $\alpha = \frac{1}{2}$). At $\alpha V = c$, the benefit from opt-out is also the same for both defaults. When we change the parameters to get $\alpha V < c$, namely when we reduce V , increase c or reduce α , W^H falls below W^L and L Default becomes the efficient rule. Conversely, when we change the parameters to get $\alpha V > c$, namely when we increase V , reduce c or increase α , W^H rises above W^L and H Default becomes the efficient rule. To summarize, when $\alpha = \frac{1}{2}$, the opt-out cost minimization principle is neutral and the expected value maximization principle solely determines the optimal default.

In the Appendix, we generalize the analysis allowing for simultaneous deviations from both $\alpha = \frac{1}{2}$ and $\alpha V = c$. The results of this analysis are summarized in the following Table:

	$\alpha V < c$	$\alpha V = c$	$\alpha V > c$
$\alpha < \frac{1}{2}$	L Default is better	L Default is better	L Default is better when α is smaller and αV is closer to c . H Default is better when α is closer to $\frac{1}{2}$ and $\alpha V - c$ is larger
$\alpha = \frac{1}{2}$	L Default is better	Both rules are equally efficient	H Default is better

$\alpha > \frac{1}{2}$	H Default is better when α is larger and αV is closer to c . L Default is better when α is closer to $\frac{1}{2}$ and $\alpha V - c$ is smaller	H Default is better	H Default is better
------------------------	---	---------------------	---------------------

Table 1: Identifying the Optimal Default Rule

Table 1 highlights the interaction between the two key forces that determine the optimal default rule: opt-out cost minimization and expected value maximization. In the middle of the table, both forces are neutral and thus the two rules are equally efficient. Moving up and down the middle column, we see the standard opt-out cost minimization principle at work – when $\alpha < \frac{1}{2}$, L Default is better; and when $\alpha > \frac{1}{2}$, H Default is better. Moving left and right on the middle row, we see the expected value maximization principle at work – when $\alpha V < c$, L Default is better; and when $\alpha V > c$, H Default is better. At the top-left corner of the table and at the bottom-right corner, the two principles push in the same direction: L Default is better at the top-left corner, when $\alpha < \frac{1}{2}$ and $\alpha V < c$; and H Default is better at the bottom-right corner, when $\alpha > \frac{1}{2}$ and $\alpha V > c$.

The most interesting cases are at the top-right corner of the table and at the bottom-left corner, where the two principles push in opposite directions. At the top-right corner, $\alpha < \frac{1}{2}$ makes L Default more attractive, whereas $\alpha V > c$ makes H Default more attractive. The optimal default is determined by the relative strength of the two considerations, namely, how far are we from $\alpha = \frac{1}{2}$ and how far are we from $\alpha V = c$.

Conversely, at the bottom-left corner, $\alpha > \frac{1}{2}$ makes H Default more attractive, whereas $\alpha V < c$ makes L Default more attractive. Again, the optimal default is determined by the relative strength of the two considerations.

The intermediate k scenario highlights the differences between our model and standard model. In the standard model, the optimal default is determined solely by the opt-out cost minimization principle. In that model, consumers know their type in advance and so, when k allows for informed opt out, any mismatched consumer will opt-out. Opt-out costs are minimized by choosing L Default when $\alpha < \frac{1}{2}$, and H Default when $\alpha > \frac{1}{2}$. In our model, consumers are initially uninformed; and some remain uninformed and incur a mismatch cost. A default rule based on the expected value maximization principle protects these uninformed consumers. Of course, some consumers do become informed and opt-out, and so opt-out cost minimization is still an important consideration. Our model shows how the two principles should be balanced to maximize social welfare.

3. Low k : Both informed and uninformed opt-out

In the low- k scenario, opt-out costs allow for both informed and uninformed opt-out. Formally, this scenario occurs, when opt-out costs satisfy: $k < \max(\alpha V - c, c - \alpha V)$. (There are two, mutually exclusive cases: (i) $\alpha V - c > 0$, in which we assume $k < \alpha V - c$, and (ii) $\alpha V - c < 0$, in which we assume $k < c - \alpha V$.) We analyze market outcomes and welfare levels with L Default (in subsection a) and with H Default (in subsection b). We then use these results to guide the choice of the optimal default rule (in subsection c).

(a) L Default

When $\alpha V > c$ and $k < \alpha V - c$, uninformed consumers will opt out of L Default. To find the social welfare level in this scenario, we start by calculating the value of information. The uninformed consumer opts out and gets $\alpha V - c - k$. The informed consumer gets $\alpha(V - c - k)$, as before. The value of information is $I^{L1} = \alpha(V - c - k) - (\alpha V - c - k) = (1 - \alpha)(c + k)$. Intuitively, information is valuable for L types who stick with L Default and save opt-out costs (k) and get a lower price (lower by c). Social welfare is given by:

$$\begin{aligned} W^{L1} &= \int_0^{I^{L1}} (\alpha(V - c - k) - x)f(x)dx + [1 - F(I^{L1})](\alpha V - c - k) \\ &= \int_0^{I^{L1}} (I^{L1} - x)f(x)dx + (\alpha V - c - k) \end{aligned}$$

Using the general function $W(I)$, we can write: $W^{L1} = W(I^{L1}) + (\alpha V - c - k)$.

(b) H Default

When $\alpha V < c$ and $k < c - \alpha V$, uninformed consumers will opt out of H Default. The uninformed consumer opts out and incurs a cost of k . The informed consumer enjoys an expected benefit of $\alpha(V - c) - (1 - \alpha)k$. The value of information is $I^{H1} = [\alpha(V - c) - (1 - \alpha)k] - [-k] = \alpha[V - c + k]$. Intuitively, information is valuable for H types who stick with H Default and save opt-out costs (k) and gain $V - c$. Social welfare is given by:

$$\begin{aligned} W^{H1} &= \int_0^{I^{H1}} [\alpha(V - c) - (1 - \alpha)k - x]f(x)dx - [1 - F(I^{H1})]k \\ &= \int_0^{I^{H1}} [I^{H1} - x]f(x)dx - k \end{aligned}$$

Using the general function $W(I)$, we can write: $W^{H1} = W(I^{H1}) - k$.

(c) Comparing L and H Default

When $\alpha V > c$, opt out by uninformed consumers occurs only with L Default. With H Default, uninformed consumers stick with the default. To identify the optimal rule, we must therefore compare: $W^{L1} = W(I^{L1}) + (\alpha V - c - k)$ and $W^H = W(I^H) + (\alpha V - c)$. H Default has the pre-information advantage, since $\alpha V - c > \alpha V - c - k$. Moving on to information-based welfare, we compare $I^{L1} = \alpha(V - c - k) - (\alpha V - c - k) = (1 - \alpha)(c + k)$ to $I^H = [\alpha(V - c) - (1 - \alpha)k] - [\alpha V - c] = (1 - \alpha)(c - k)$. L Default generates more information-based welfare, since $I^{L1} > I^H$. Specifically, $I^{L1} = I^H + (1 - \alpha)(2k)$. Intuitively, L Default is more costly for uninformed consumers, as it induces costly opt-out. By learning their type, consumers can potentially avoid these opt-out costs (specifically, if they learn that they are low type).

When $k = 0$, H Default's pre-information advantage disappears and L Default's added information-based welfare also disappears. Hence, the two rules are equally efficient. Formally, when $k = 0$, we have $I^{L1} = I^H = (1 - \alpha)c$ and $W^{L1} = W^H = W((1 - \alpha)c + \alpha V - c)$. What happens when k is larger? Let $\Delta W \equiv W^H - W^{L1} = W(I^H) - W(I^{L1}) + k$ and take the derivative of ΔW w.r.t. k :

$$\begin{aligned} \frac{d\Delta W}{dk} &= W'(I^H) \frac{dI^H}{dk} - W'(I^{L1}) \frac{dI^{L1}}{dk} + 1 = -F(I^H)(1 - \alpha) - F(I^{L1})(1 - \alpha) + 1 \\ &= 1 - (1 - \alpha)(F(I^H) + F(I^{L1})) \end{aligned}$$

When $\alpha \geq \frac{1}{2}$, we have $\frac{d\Delta W}{dk} > 0$ (since $F(I^H) + F(I^{L1}) \leq 2$) and so H Default is better.

When $\alpha < \frac{1}{2}$, we may get $\frac{d\Delta W}{dk} < 0$, which means that L Default may be better.

Parallel analysis applies when $\alpha V < c$, such that opt out by uninformed consumers occurs only with H Default. With L Default, uninformed consumers stick with the default. To identify the optimal rule, we must therefore compare: $W^L = W(I^L)$ and $W^{H1} = W(I^{H1}) - k$. L Default has the pre-information advantage. Moving on to information-based welfare, we compare $I^L = \alpha(V - c - k)$ to $I^{H1} = [\alpha(V - c) - (1 - \alpha)k] - k = \alpha V - c + k$. H Default generates more information-based welfare, since $I^{H1} > I^L$. Specifically, $I^{H1} = I^L + \alpha(2k)$. Intuitively, H Default is more costly for uninformed consumers, as it induces costly opt-out. By learning their type, consumers can potentially avoid these opt-out costs (specifically, if they learn that they are high type).

When $k = 0$, L Default's pre-information advantage disappears and H Default's added information-based welfare also disappears. Hence, the two rules are equally efficient. Formally, when $k = 0$, we have $I^L = I^{H1} = \alpha(V - c)$ and $W^L = W^{H1} = W(\alpha(V - c))$. What happens when k is larger? Let $\Delta W \equiv W^L - W^{H1} = W(I^L) - W(I^{H1}) + k$ and take the derivative of ΔW w.r.t. k :

$$\begin{aligned}\frac{d\Delta W}{dk} &= W'(I^L) \frac{dI^L}{dk} - W'(I^{H1}) \frac{dI^{H1}}{dk} + 1 = -F(I^L)\alpha - F(I^{H1})\alpha + 1 \\ &= 1 - \alpha(F(I^L) + F(I^{H1}))\end{aligned}$$

When $\alpha \leq \frac{1}{2}$, we have $\frac{d\Delta W}{dk} > 0$ (since $F(I^L) + F(I^{H1}) \leq 2$) and so L Default is better.

When $\alpha > \frac{1}{2}$, we may get $\frac{d\Delta W}{dk} < 0$, which means that H Default may be better.

Like the intermediate k scenario (Sec. B.2), the low k scenario also highlights the differences between our model and standard model. In the standard model, the optimal default is determined solely by the opt-out cost minimization principle. In our model, the expected value maximization principle must be balanced against the opt-out cost minimization principle. First consider the $\alpha V > c$ case. When $\alpha < \frac{1}{2}$, the opt-out cost minimization principle suggest that L Default has an advantage. But, in our model, there is an additional concern about costly uninformed opt-out – a concern that is addresses by the expected value maximization principle. When this concern about uninformed opt-out is sufficiently large, our model prescribes H Default, whereas the standard model prescribes L Default. When $\alpha > \frac{1}{2}$, the opt-out cost minimization principle suggest that H Default has an advantage. In our model, there is an additional concern about costly uninformed opt-out. But here the expected value maximization principle joins the opt-out cost minimization principle in recommending H Default.

Next consider the $\alpha V < c$ case. When $\alpha > \frac{1}{2}$, the opt-out cost minimization principle suggest that H Default has an advantage. But, in our model, there is an additional concern about costly uninformed opt-out – a concern that is addresses by the expected value maximization principle. When this concern about uninformed opt-out is sufficiently large, our model prescribes L Default, whereas the standard model prescribes H Default. When $\alpha < \frac{1}{2}$, the opt-out cost minimization principle suggest that L Default has an advantage. In our model, there is an additional concern about costly uninformed opt-out. But here the expected value maximization principle joins the opt-out cost minimization principle in recommending L Default, and so our model and the standard model offer the same prescription.

IV. Consumer Misperceptions

We have thus far assumed that consumers hold rational expectations. While initially uninformed about their type, consumers hold accurate beliefs about the relevant parameters: V , c , α and k (and also about the distribution function $F(\cdot)$). We now relax this assumption and explore the implications of consumer misperceptions.

We focus on the intermediate k case, where informed consumers, and only informed consumers, opt out. In this case, consumers decide whether to become informed and only then some of the informed consumers opt out. Misperception might distort this decision.

Specifically, the decision whether to become informed will now be determined by the *perceived* value of information, rather than the actual value of information. With L Default, the perceived value of information is $\hat{I}^L = \hat{\alpha}(\hat{V} - \hat{c} - \hat{k})$, where $\hat{\alpha}$, $\hat{V}$, $\hat{c}$ and $\hat{k}$ represent the perceived values of the relevant parameters. Similarly, with H Default, the perceived value of information is $\hat{I}^H = (1 - \hat{\alpha})(\hat{c} - \hat{k})$.

In terms of social welfare, with L Default, social welfare is:

$$W^L = \int_0^{\hat{I}^L} (I^L - x)f(x)dx$$

Define a general function $W(I, \hat{I}) \equiv \int_0^{\hat{I}} (I - x)f(x)dx$ and notice that misperception can lead to excessive investment (in becoming informed) when $\hat{I} > I$ and to insufficient investment when $\hat{I} < I$. We can thus write $W^L = W(I^L, \hat{I}^L)$.

With H Default, social welfare is $W^H = W(I^H, \hat{I}^H) + (\alpha V - c)$.

To choose the optimal default, we must compare W^L and W^H . The comparison can be divided into two components: (1) Pre-information welfare, and (2) Information-based welfare. Pre-information welfare is not affected by the misperception. Information-based welfare, however, is distorted by the misperception.

We focus on misperception about the likelihood that the consumer benefits from high quality, α . With such misperception, the perceived value of information with L Default is $\hat{I}^L = \hat{\alpha}(V - c - k)$, and the perceived value of information with H Default is $\hat{I}^H = [\hat{\alpha}(V - c) - (1 - \hat{\alpha})k] - [\hat{\alpha}V - c] = (1 - \hat{\alpha})(c - k)$. When consumers underestimate the probability of being high-type, i.e., when $\hat{\alpha} < \alpha$, they underestimate the value of information with L Default and overestimate the value of information with H Default. This means that we get excessive investment in information with H Default and insufficient investment with L Default. And the opposite holds when consumers overestimate the probability of being high-type, i.e., when $\hat{\alpha} > \alpha$.

These results contrast with what we would expect in a model with misperception but without investment in information. The no investment model generates clear predictions. For example, if most consumers benefit from high quality (α is large) but mistakenly think they don't ($\hat{\alpha}$ is small), then L Default is better since it avoids many costly opt-outs. (Recall that we are focusing on the intermediate k case where informed, including mistakenly informed, consumers opt out.) In our model, with investment in information, the results are more nuanced. Specifically, H Default can be better: with misperception, H Default increases the (mis)perceived value of information, resulting in an increased level of investment in information; if the benefit from more investigation – corrected misperceptions and more efficient opt out decisions – outweighs the cost of excessive investigation, then H Default would be the better rule.

V. Conclusion

We argued in this paper that contracting around defaults in consumer contracts is a different process, and requires a different model, than the traditional account. Consumers are not informed about the default terms and their value, and thus an essential first step is the decision to acquire information.

In this environment, default rules affect the incentive to acquire information. When information is more likely to alter the opt-out decision, its value increases and more consumers decide to become informed. These informed consumers enjoy a more accurate match between the contract terms and their own preferences. Default rules may be designed to induce such investment in information. Or, they may be designed to prevent such investment, so as to save the transaction costs of information acquisition and of informed opt out. We showed how these design choices are affected by various parameters.⁴

Some features of consumer markets were missing from our analysis and can be added to the model. For example, we assumed that opt out costs are exogenously determined. But the cost to opt out may be influenced by the law. If consumers are rational, it would be optimal for the law to push opt out costs lower, so that more efficient opt out and information acquisition would ensue. But if consumers act upon misperceptions, it is often thought that increasing opt out costs (making defaults stickier) is a good policy, to prevent undesirable opt outs. Our model casts some doubt on this pro-stickiness view. If consumers can invest in information that would correct their misperceptions and lead to more efficient contracts, there is good reason to keep opt-out costs low. Another possible extension would consider psychological opt-out costs (inertia, procrastination, status quo bias). It is not clear how to account for such costs when calculating social welfare, but it is clear that their presence can affect the incentive to acquire information.

As noted in the Introduction, defaults in consumer markets can be set either by the lawmaker or by the seller. Accordingly, our analysis of optimal defaults should inform policymakers seeking to maximize social welfare, but also firms that seek to attract consumers in a competitive market.

Finally, we note that the basic tradeoff between the opt-out cost minimization principle and the expected value maximization principle applies more broadly. It is not limited to consumer contracts and it applies also in the standard framework where parties know their preferences at the outset. We develop the general implications of this fundamental tradeoff in other work.

⁴ Since the relevant parameters can take different values in different consumer markets, the optimal default can also change from market to market. We acknowledge the empirical challenge of measuring the parameter values across different markets.

References

Acquisti, Alessandro, Leslie K. John, and George Loewenstein (2013), *What is Privacy Worth?* 42 J. LEGAL STUD. 249-74.

Ayres, Ian, and Robert Gertner (1989), *Filling Gaps in Incomplete Contracts: An Economic Theory of Default Rules*, 99 YALE L.J. 87-130.

Baker, Tom, and Peter Siegelman (2013), *You Want Insurance With That? Using Behavioral Economics to Protect Consumers from Add-On Insurance Products*, 20 CONN. INS. L. J. 1-60.

Bebchuk, Lucian Arye, and Steven Shavell (1991), *Information and the Scope of Liability for Breach of Contract: The Rule of “Hadley v. Baxendale”*, 7 J.L. ECON. & ORG. 284-312.

Camerer, Colin, Samuel Issacharoff, George Loewenstein, Ted O’Donoghue, and Matthew Rabin (2003), *Regulation for Conservatives: Behavioral Economics and the Case for “Asymmetric Paternalism”*, 151 U. PA. L. REV. 1211-54.

Goetz, Charles J., and Robert E. Scott (1985), *The Limits of Expanded Choice: An Analysis of the Interactions Between Express and Implied Contract Terms*, 73 CAL. L. REV. 261-322.

Hann, Il-Horn, Kai-Lung Hui, Sang-Yong Tom Lee, and Ivan P.L. Png (2007), *Overcoming Online Information Concerns: An Information-Processing Theory Approach*. J. MANAGEMENT INFORMATION SYSTEMS 24 (2), 13-42.

Kim, Nancy S. (2013), *WRAP CONTRACTS: FOUNDATIONS AND RAMIFICATIONS* (Oxford University Press, New York, NY).

Pew Research Center (2013), *Anonymity, Privacy, and Security Online* (<http://www.pewinternet.org/2013/09/05/anonymity-privacy-and-security-online/>).

Radin, Margaret Jane (2013), *BOILERPLATE: THE FINE PRINT, VANISHING RIGHTS, AND THE RULE OF LAW* (Princeton University Press, Princeton, NJ).

Savage, Scott, and Donald M. Waldman (2013), *The Value of Online Privacy* (Working Paper)

Shavell, Steven (2004), *FOUNDATIONS OF ECONOMIC ANALYSIS OF LAW* (The Belknap Press of Harvard University Press, Cambridge, MA).

Westin, Alan F. (1997), *“Whatever Works”--The American Public’s Attitudes Toward Regulation and Self-Regulation on Consumer Privacy Issues*, in *PRIVACY AND SELF-REGULATION IN THE INFORMATION AGE* ch.1, §F (<http://www.ntia.doc.gov/page/chapter-1-theory-markets-and-privacy>).

Appendix: Simultaneous deviations from both $\alpha = \frac{1}{2}$ and $\alpha V = c$

We repeat the exercise, from Section III.B.2, of starting at $\alpha V = c$ and changing the parameter values, but without assuming $\alpha = \frac{1}{2}$. We begin by establishing a few results. Let $\Delta W \equiv W^H - W^L = W(I^H) - W(I^L) + (\alpha V - c)$ and take the derivative of ΔW w.r.t. the relevant parameters:

$$\frac{d\Delta W}{dV} = W'(I^H) \frac{dI^H}{dV} - W'(I^L) \frac{dI^L}{dV} + \alpha = -F(I^L)\alpha + \alpha = (1 - F(I^L))\alpha > 0$$

$$\frac{d\Delta W}{dc} = W'(I^H) \frac{dI^H}{dc} - W'(I^L) \frac{dI^L}{dc} - 1 = F(I^H)(1 - \alpha) + F(I^L)\alpha - 1 < 0$$

$$\frac{d\Delta W}{d\alpha} = W'(I^H) \frac{dI^H}{d\alpha} - W'(I^L) \frac{dI^L}{d\alpha} + V = -F(I^H)(c - k) - F(I^L)(V - c - k) + V$$

Since $k < \min(c, V - c)$, we can write:

$$\begin{aligned} \frac{d\Delta W}{d\alpha} &= -F(I^H)(c - k) - F(I^L)(V - c - k) + V > -(c - k) - (V - c - k) + V \\ &= 2k > 0 \end{aligned}$$

At $\alpha V = c$, the benefit from opt-out is the same for both defaults. When $\alpha > \frac{1}{2}$, the cost of opt-out is smaller with H Default. This means that the value of information is larger with H Default: $I^H > I^L$, which implies $W(I^H) > W(I^L)$. We thus have:

$$\Delta W(\alpha V = c) = W(I^H) - W(I^L) > 0$$

Welfare is higher with H Default. When we change the parameter values, to get $\alpha V > c$, namely, when we increase V , decrease c or increase α , this reinforces the advantage of H Default (since $\frac{d\Delta W}{dV} > 0$, $\frac{d\Delta W}{dc} < 0$, and $\frac{d\Delta W}{d\alpha} > 0$, as shown above). When we change the parameter values, to get $\alpha V < c$, namely, when we reduce V , increase c or reduce α , this reduces the advantage of H Default.

When $\alpha < \frac{1}{2}$, the cost of opt-out is smaller with L Default. This means that the value of information is larger with L Default: $I^L > I^H$, which implies $W(I^L) > W(I^H)$. We thus have:

$$\Delta W(\alpha V = c) = W(I^H) - W(I^L) < 0$$

Welfare is higher with L Default. When we change the parameter values, to get $\alpha V < c$, namely, when we reduce V , increase c or reduce α , this reinforces the advantage of L Default (since $\frac{d\Delta W}{dV} > 0$, $\frac{d\Delta W}{dc} < 0$, and $\frac{d\Delta W}{d\alpha} > 0$, as shown above). When we change the

parameter values, to get $\alpha V > c$, namely, when we increase V , decrease c or increase α , this reduces the advantage of L Default.