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

HARVARD

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

MANIPULATION BY MISLAID PRIORITIES

Oren Bar-Gill
Omri Ben-Shahar

Discussion Paper No. 1063

05/2021

Harvard Law School
Cambridge, MA 02138

Manipulation By Mislaid Priorities

Oren Bar-Gill and Omri Ben-Shahar*

Abstract

This paper lays a foundation for a new theory of manipulation, based on the misprioritization of (truthful) information. Since consumers review only a subset of all available information, firms can harm consumers by prioritizing information that maximizes firms' profits but has a smaller impact on the utility that consumers stand to gain from the purchase. Moreover, the distortions due to misprioritized information can arise not only from firms' boastful disclosures, but also from the warnings and disclosures mandated by lawmakers. The paper identifies the product and market characteristics that determine the optimal prioritization of information and, correspondingly, the incidence of harm when the wrong information is prioritized for disclosure—either voluntarily by sellers or by legal mandate. It provides a framework for optimal legal intervention.

Keywords: Manipualtion, Asymmetric information, Consumer markets, Disclosure.
JEL Codes: D11, D18, D82, D83, D91, K20, L15,

* Bar-Gill is the William J. Friedman and Alicia Townsend Friedman Professor of Law and Economics at Harvard Law School; Ben-Shahar is the Leo and Eileen Herzel Professor of Law and Kearney Director of the Coase-Sandor Institute for Law and Economics at the University of Chicago Law School. For valuable comments and conversations, we thank Lucian Bebchuk, Oliver Board, Tom Brennan, Glenn Cohen, David Gilo, John Goldberg, Moshe Halbertal, Howell Jackson, Louis Kaplow, Yoram Margalioth, Florencia Marotta-Wurgler, Holger Spamann, Kathy Spier, Rebecca Tushnet, Mark Tushnet, and workshop participants at Harvard, Tel-Aviv University, New York University and the University of Chicago. Emily Feldstein, Haggai Porat, and Sahar Segal provided excellent research assistance.

1. Introduction

Manipulation has many shades.[1] While it is a staple of almost every human interaction, it is particularly prevalent in consumer markets, where firms promote their business by diverting people's attention, shifting their preferences, influencing their beliefs and otherwise affecting their choices. The concern over manipulative tactics has stimulated growing attention by policymakers and commentators, seeking to better understand the scope of the problem and to establish boundaries for permissible tactics (see, e.g., Akerlof and Shiller 2015; Hanson and Kysar 1999).

The law has little problem prohibiting lying. The challenge posed by manipulation becomes tougher in the absence of lies. To address this challenge, much of the literature focuses on manipulation that takes advantage of people's irrationality and seeks to subvert their rational capacities. Under these accounts, firms exploit people's mistakes and misjudgments, or inflame them.[2]

In this paper, we identify and study a general source of manipulation that does not fall within the traditional characterization. It is a form of manipulation that does not require preying on incorrect assumptions and inferences that people make. This category of manipulation relies solely on how the truthful information that sellers present is *prioritized*. Indeed, we define manipulation as prioritizing information in a way that conflicts with its recipients' interests.

Consider the following example. A food product has various quality dimensions, each of which can be regarded by byers as "good" or "bad." These dimensions can include nutritional aspects like sugar or fat content, health claims, or production qualities ("organic," "GMO"). When designing the product's packaging, the seller could refer to all of the aspects equally, for example in a list that resembles the familiar Nutrition Data disclosure. Alternatively, the seller may highlight only the good dimensions ("low fat" or "non-GMO") and either remain silent about the bad dimensions, present them in a less salient manner or otherwise make it more costly for buyers to learn the bad information.[3] If buyers have unlimited capacity to review information, the way in which the seller prioritizes the ingredients would not matter. Buyers would view both the good and bad quality attributes and become perfectly informed. But buyers rarely become perfectly informed. When buyers review only some disclosures, how sellers prioritize their disclosures determines which information buyers get.

A key feature in the environments that exhibit 'manipulation by mislaid priorities' is the *complexity* of products. Most products have multiple dimensions of quality, and it is often impractical to effectively communicate to buyers the quality of each dimension. Even if sellers disclose information about all quality dimensions, buyers will review and digest only some of the information. Rational buyers will stop reviewing information when the cost of additional review outweighs the expected benefit from the additional information (Stigler 1961). And imperfectly

[1] Large literatures in philosophy, psychology, and other disciplines attempt to answer the basic question 'what is manipulation.' See, e.g., Coons and Weber (2014); Wertenbroch (2016).

[2] See Gorin (2014, 90); Sunstein (2016) (a statement or action is manipulative "if it does not sufficiently engage or appeal to people's capacity for reflective and deliberative choice."). See also Wilkinson (2013); Hanson and Kysar (1999); Calo (2014).

[3] See Mantikas v. Kellogg Co., 910 F.3d 633 (2d Cir. 2018).

rational buyers will focus only on the more salient information (see, e.g., Taylor and Thompson 1982; Milosavljevic, et al. 2012; Bordalo, Gennaioli and Shleifer 2013). Anticipating buyers' review patterns, sellers could choose what true information to prioritize.

The "priority" of information, as we define the term, derives from any aspect of how information is presented that increases the likelihood of buyers reviewing this particular information ahead of, other, de-prioritized information. In a rational-choice framework, information is prioritized when the relative cost of acquiring this information is reduced. Allowing for imperfect rationality, information can also be prioritized or de-prioritized by altering its salience in terms of mode, place, time or context of presentation. Sellers control the relative costs of acquiring information about different product dimensions and they control the relative salience of different bits of information.

Manipulation by mislaid priorities occurs even when buyers recognize that they are only partially informed or that sellers selectively prioritize some information, and draw rational inferences about what they don't know. Buyers' inferences will be imperfect, causing them to make purchases they end up regretting or, due to the prudence they rationally exert, refrain from making purchases they would value.[4]

The distortion due to misprioritized information can arise not only from sellers' strategic disclosures, but also from the warnings and disclosures mandated by lawmakers. This might catch our readers by surprise. You might think that consumers' misinformation is uniquely due to how firms self-servingly dim some disclosures while highlighting others. When governments mandate warnings, so goes the thought, people will not be misled about a product's value, firms' manipulation would be countered, and the problem of mislaid priorities would be solved. We show that this intuition is valid only sometimes, and entirely misguided otherwise. To understand why, let us take a step back and explain the primary contribution of the article—when is it that buyers, who are able to review only a subset of the disclosed dimensions, are harmed by poorly-prioritized information.

Buyers' interest is to reserve their attention to the product dimensions that have the greatest impact on their welfare. For some products, it is critical for buyers to know that there are some *low-quality* dimensions, because the presence of such dimensions undermines the value of the product as a whole, even in the presence of other, high-quality attributes. For other products, however, buyers' interest is reversed: it is critical for them to know that there are some *high-quality* dimensions, and it matters less to them that these high-quality aspects are accompanied by other, low-quality dimensions.

The first category of products—where low-quality dimensions matter more—includes products with dimension of quality that we loosely characterize as *complements.* Buyers' value from such

[4] Prior literature has recognized that asymmetric information is rampant in consumer markets, despite important, but only partially successful, attempts to inform consumers through disclosure, reputational mechanism, private or state-sponsored certifications, ratings systems and other market and non-market mechanisms. For theory papers that explain the prevalence of asymmetric information in consumer markets—see, e.g., Milgrom (1981), Daughety and Reinganum (2008) and Sun (2011). For empirical papers that document and measure information deficits—see, e.g., Alba and Hutchinson (2000), Mocan (2007), Lacko and Pappalardo (2010) and Bakos, Marotta-Wurgler and Trossen (2014). This paper does not add to that literature; we take the prevalence of some asymmetric information as given. Our contribution is to show how suboptimal prioritization exacerbates the harm from asymmetric information.

products is significantly reduced if *either* of the critical dimensions is low. For example, when people purchase a vacation package in a resort, several dimensions are critical. If their suite is shabby, if the food is bad, if the service is poor, if the pool is out of order—each of these dimensions can single-handedly destroy much of the benefit from the vacation. Before they purchase a vacation package, it is important for consumers to know that none of the critical dimensions are low quality. Or, to take another example, when purchasing a printer, buyers might care about several dimensions: durability, the cost of replacement ink, and the quality of the printer's wireless connection. Again, if any of these dimensions is of poor quality, the value of the device is significantly diminished, even if other dimensions are high quality. For such products, buyers want information about the low-quality aspects to be prioritized, so that they could focus their limited attention on these dimensions. Sellers, however, prefer to de-prioritize the information about low-quality dimensions and highlight instead the high-quality ones.

In this *complements* scenario, sellers manipulate buyers by prioritizing the high-quality dimensions, which have a smaller impact on buyers' utility. Sellers lead buyers to review less important information, thus draining their capacity to review more important information. Notice that sellers are not misleading buyers, nor falsely implying that other dimensions are high. They are "only" choosing to prioritize some information – information that maximizes their profits – knowing that buyers would have preferred to receive other information. Legal intervention could reduce the harm to buyers. Since the problem is due to the manipulative prioritization of information, the solution is straightforward: require sellers to prioritize the disclosure of the low-quality dimensions. We call this a *Warning* regime and show that it prescribes disclosures based on a criterion that differs from typical product labeling laws.

We then turn to examine a second category of products, for which the *Warning* regime can do more harm than good. In this second category, the dimensions of quality are not complements but instead are (loosely characterized as) *substitutes*. If one of the product dimensions is high quality, it matters little to buyers that other dimensions are low. Accordingly, it is more important for buyers to turn their attention to the presence of high-quality dimensions. For example, a visit to a theme park or a restaurant could be greatly satisfying even if some of the rides or menu items are low-quality, as long as others are high. Or, a college course could be successful if *either* the professor or the teaching assistant is effective, they don't both have to be high quality. Because buyers' interest is to be told about the high-quality aspects, the problem of manipulation due to sellers' mislaid priorities does not arise. The interests of buyers and sellers coincide: buyers wish to learn about high-quality dimensions, and sellers wish to prioritize the disclosure of information about high-quality dimensions.

In the *substitutes* scenario, no legal intervention is needed. Caveat Emptor—the regime that allows sellers to make any disclosures they want, or none—guarantees optimal dissemination of information. In fact, distortions arise if the law steps in to mandate that sellers prioritize the low-quality dimensions. If the *Warning* regime, devised to address manipulation in the *complements* scenario, is extended to the *substitutes* scenario—if, that is, the law requires prominent warnings of the low-quality dimensions even when buyers prefer to direct their attention to high-quality dimensions—buyers' ability to attend to the most critical information and to make optimal purchase decisions would be frustrated.

We show that the distortions arising from sellers' mislaid disclosure priorities in the complements case mirror the distortions arising from lawmakers' mislaid disclosure mandates in the substitutes case. We are accustomed to thinking of sellers' selective disclosure of high-quality dimensions as manipulative and harmful, and of lawmakers' mandates that low dimensions be warned against as protective and helpful. Our analysis shows that this instinctive view is valid only for the first category of products—when it is more important for buyers to know about low-quality dimensions. When, instead, it is more important for buyers to know about high-quality dimensions, sellers' unregulated disclosures are in fact helpful and well-prioritized, whereas legal mandates to warn buyers about low-quality dimensions are harmful. It would be awkward to call these mandated warnings 'manipulation,' but they do in fact cause exactly the same type of harm as manipulation by sellers' selective (voluntary) disclosures. Unlike sellers, lawmakers issuing disclosure mandates are not *strategically hoping* to manipulate. But like sellers' selective presentation of information, warnings can prioritize disclosed information in a manner that harms consumers by diverting attention from more important aspects of the product.

Our analysis of manipulation by mislaid disclosure priorities informs the optimal design of legal interventions. First, as we already said, warnings—which require sellers to prominently disclose the low-quality dimensions of their products—should be used for some products, but not all. They are helpful when product dimensions are complements, yet harmful when the dimensions are substitutes. Second, our analysis reveals a subtle and surprising effect: the value of warnings depends also on the product's price. A *Warning* regime is more valuable for high-priced products, where buyers are more likely to experience post-purchase regret due to the presence of undisclosed low-quality dimensions. Caveat Emptor is more valuable for low-priced products, where buyers' regret is primarily due to missing out on a valuable purchase, and this in turn could be reduced by allowing sellers to focus buyers' attention on the high-quality dimensions. Third, we show that a legal regime of *Full Disclosure*, which requires unprioritized disclosure of all quality dimensions, is unambiguously inferior to either *Warning* or *Caveat Emptor*. Unlike the *Warning* regime, it fails to fully alert buyers to the existence of low-quality dimensions when such priority is necessary. And unlike *Caveat Emptor*, it fails to effectively communicate the existence of high-quality dimensions when such information is most valuable. Put differently, we formally establish one reason why a legal mandate to disclose *more* information could backfire.[5]

Our baseline analysis assumes that buyers draw rational inferences from their partial information, but it can be extended to allow for biased inferences.[6] In the *Caveat Emptor* regime, the key question is what inferences buyers draw when sellers disclose high quality dimensions. Rational buyers recognize the payoff-relevance of undisclosed dimensions, and this cabins their estimate of the product's value. Imperfectly rational buyers place insufficient weight on the undisclosed dimensions and thus overestimate the product's value. Accordingly, the need for legal intervention

[5] The intuition that too much information could be harmful is typically grounded in a behavioral account of cognitive overload. See, e.g., Ben-Shahar and Schneider (Ch.6, 2014). This intuition has been recognized by courts. See, e.g., Ford Motor Credit Co. v. Milhollin, 444 U.S. 555, 568 (1980) ("*Meaningful* disclosure does not mean *more* disclosure. Rather, it describes a balance between 'competing considerations of complete disclosure… and the need to avoid… [information overload]'…. And striking the appropriate balance is an empirical process that entails investigation into consumer psychology….") (brackets and emphasis in the original; citations omitted).

[6] The behavioral economics literature has studied such irrational inferences, specifically the failure to appreciate the importance of undisclosed information. See, e.g., Enke (2020) and Jehiel (2018). For a recent survey paper on errors in statistical reasoning – see Benjamin (2019).

is greater when consumers are imperfectly rational. In markets where buyers are more sophisticated, it is less critical to warn people about low quality dimensions. A similar comparison applies in the *Warning* regime: Rational buyers exposed to mandated warnings will optimally account for the undisclosed, potentially high-quality dimensions, whereas imperfectly rational buyers will place excessive weight on the warning and underestimate the product's value. The cost to consumers of misprioritized information mandated by law is lower for rational consumers, in the same manner that the cost of manipulation by sellers under *Caveat Emptor* was lower for rational consumers.

At the most abstract level, our analysis begins with the recognition that buyers will review only a subset of information about the product, and that disclosure regulation affects the subset of information that buyers will review—removing some, good information about the product and adding other, bad information about the product. To help buyers, regulators must ensure that the added information is more important to buyers than the information it replaces. Our analysis helps regulators identify which information is more important, primarily through two factors: the distinction between complements versus substitutes in quality dimensions, and the surplus that buyers expect to secure.

This paper contributes to a growing literature on manipulation in consumer markets. Whereas much of the literature focuses on manipulation as the exploitation of irrationality and bias and a failure to engage "people's capacity for reflective and deliberative choice" (Sunstein 2016), we focus on manipulation by mislaid priority and we show that it harms consumers even if they are aware of their imperfect information and make deliberative choices.[7] That people have a limited capacity to review information is considered in the literature on bounded rationality and satisficing (see, e.g., Simon 1956), in the literature on information overload (see, e.g., Jacoby 1984; Eppler and Mengis 2004; Persson 2018), and in the behavioral industrial organization literature (see, e.g., Spiegler 2016; Spiegler 2006; Gabaix and Laibson 2006; Heidhues, Kőszegi and Murooka 2016).[8] These various inquiries set up the starting point for our analysis, which considers the interaction between the reviewed and non-reviewed product dimensions—complements vs. substitutes—and how this interaction dictates the optimal disclosure regime.

The article is organized as follows. Section 2 develops a simple model of transacting over a complex product under asymmetric information. It characterizes the problem of manipulation by

[7] The economics literature has examined the main building blocks of our model—asymmetric information between buyers and sellers and limits on sellers' ability to overcome the information asymmetry with disclosure. But this literature commonly assumes a single quality dimension with costly disclosure on even this one dimension. (For surveys – see Milgrom 2008; and Dranove and Jin 2010.) A few papers have considered more than one quality dimension: Hotz and Xiao (2013) study a model with two quality dimensions, but one of them captures heterogeneity in consumer preferences. A separate literature developed multi-dimensional product differentiation models, but assumed that buyers know both quality dimensions and thus cannot be manipulated, focusing instead on sellers' quality and price decisions (see, e.g., Vandenbosch and Weinberg 1995; Barigozzi and Ma 2018).

[8] Spiegler (2006) assumes that consumers randomly observe a single dimension of the product. Gabaix and Laibson (2006) assume that some consumers observe only one of the product's dimensions and are non-Bayesian. Heidhues, Kőszegi and Murooka (2016) assume that consumers do not observe the add-on price when all firms coordinate to shroud it. Spiegler (2016) reviews the behavioral industrial organization literature. Persson (2018) studies a model where a decisionmaker (buyer) rationally allocates scarce attention among the many, mostly irrelevant information cues that an expert (seller) sends in attempt to strategically induce information overload (Persson then argues that regulation mandating the disclosure of a relevant information cue would not increase welfare).

mislaid priorities, and compares the positive and normative implications of three legal regimes—Caveat Emptor, Warning, and Full Disclosure. It also introduces the principle of information priority as the underlying regulatory criterion. Section 3 then extends the basic analysis in multiple directions. It considers scenarios in which some product dimensions are more important than others, showing that such ranking makes legal intervention potentially more effective, but also potentially more destructive. It also considers how the potential manipulation of buyers affects sellers' incentives to invest in product quality and to make products more complex. Section 4 offers concluding remarks about the rationality assumption, additional extensions, doctrinal applications and institutional design.

2. Model

We begin with a simple benchmark model of trade over a product that has two dimensions, under conditions of asymmetric information. This framework allows us to examine how information is revealed and how purchase decisions are made, in light of sellers' strategic disclosure and in light of the legal interventions designed to guides such disclosures.[9]

2.1 Informal Summary

To simplify matters, we consider a product with only two dimensions, where each dimension can be of either high or low quality. When buyers have perfect information, they would be willing to pay a high price when both dimensions are high quality; a lower price if only one dimension is high quality; and if neither dimension is high-quality, buyers would not want to purchase the product. The problem is that buyers don't have perfect information, and thus cannot always distinguish between products with two, one or zero high-quality dimensions. This asymmetric-information problem might lead buyers to pay too much for low-quality products or to inefficiently refrain from purchasing high-quality products. Disclosure – when made voluntarily by sellers or mandated by policymakers – can mitigate the asymmetric-information problem. But when buyers are unable to review or digest all relevant information, selective disclosure can be manipulative. We assume that buyers are able to review only one dimension. Specifically, if only one dimension is disclosed (or presented more prominently), buyers will review this disclosure. When both dimensions of quality are disclosed in equal prominence, buyers will review one of the two, determined randomly.

Consider first a *Caveat Emptor* regime, where sellers can freely decide what information to disclose, and what information not to disclose (as long as any voluntary disclosure is truthful). In this regime, if the product has two high-quality dimensions, the seller will disclose information on one or both dimensions (but buyers will only view one); if the product has one high-quality dimension, the seller will disclose information on that dimension; and if the product has two low-quality dimensions, the seller will remain silent. This means that buyers can identify, with certainty, products with two low-quality dimensions and avoid purchasing such products.

[9] A more general model would have a state space, a negotiation protocol including a message spaces and an outcome space. The key element is the payoff-relevant uncertainty that is created, or exacerbated, by Seller's strategic disclosure.

However, buyers cannot distinguish products with one versus two high-quality dimensions – in both cases buyers observe a (truthful) disclosure touting the product's high-quality on a single dimension. Because buyers can review information on only one quality dimension, observing a disclosure about one high-quality dimension leaves buyers uncertain about the quality of the other, unobserved dimension.

Buyers' inability to distinguish between products with one versus two high-quality dimensions is not a problem when the price is low, specifically, when the price is below the value of a product with one high-quality dimension. Buyers always make an optimal decision to buy such products. The problem arises when the price exceeds the value of a product with one high-quality dimension, but is lower than the value of a product with two high-quality dimensions. If buyers take the chance and purchase the product, they might pay too much if the product ends up having only one high-quality dimension. And if they act prudently, buyers might inefficiently fail to purchase a product with two high-quality dimensions. Depending on the product's price, each of these distortions could occur any time buyers observe a high-quality disclosure from the seller.

Next, consider a *Warning* regime, where the law forces sellers to disclose and prioritize information about low-quality dimensions. For example, a seller of a food product that has both bad and good qualities (high-sugar, low-fat) must prioritize a warning about the dimension with the bad quality. If the product has two low-quality dimensions (high-sugar *and* high-fat), the seller will disclose with equal priority information on both dimensions and buyers will randomly view one of them. Only if the product has two high-quality dimensions, will the seller be free to disclose and prioritize information on a high-quality dimension. This means that when no warning is viewed, buyers can safely infer that the product has two high-quality dimensions. But buyers cannot distinguish between products with one versus two low-quality dimensions—in both cases they observe a low-quality warning. Because buyers can review information on only one quality dimension, observing a warning leaves buyers uncertain about the quality of the other, unobserved dimension.

Buyers' inability to distinguish between products with one versus two low-quality dimensions is not a problem when the price is high, specifically, when the price is above the value of a product with one low-quality dimension. Buyers always make an optimal decision not to buy such products. The problem arises when the price is below the value of a product with one low-quality dimension. Buyers will pay too much for products with two low-quality dimensions, because of the (unrealized) possibility that the unobserved dimension is high-quality. And buyers will inefficiently fail to purchase a product with one high-quality dimension, when the high price outweighs the (unrealized) risk that the unobserved dimension is low-quality.

There is a third regime that should also be considered – a *Full Disclosure* regime that forces sellers to disclose information on both product dimensions, regardless of quality. Full Disclosure turns out to be inferior to the previous two regimes. Caveat Emptor allows buyers to identify, with certainty, products with two low-quality dimensions – when buyers face a silent seller. There are no silent sellers under Full Disclosure. Warning allows buyers to identify, with certainty, products with two high-quality dimensions – when buyers observe a disclosure about a high-quality dimension. Under Full Disclosure, buyers who observe such a disclosure remain uncertain about the other quality dimension. Similarly, under Full Disclosure, buyers who observe a disclosure

about a low-quality dimension remain uncertain about the other quality dimension. Full Disclosure does not allow buyers to identify any product category with certainty.

As we see, Caveat Emptor enjoys an advantage when the price of the product, relative to its value, is low, whereas Warning enjoys an advantage when the price is high. Because the distortions under Caveat Emptor are due to buyers' inability to distinguish between products with one versus two high-quality dimensions, the uncertainty is irrelevant when the price is low enough to make either purchase worthwhile. Conversely, because the distortions under Warning are due to buyers' inability to distinguish between products with one versus zero high-quality dimensions, the uncertainty is irrelevant when the price is high enough to assure that neither is purchased.

This relative ranking of Caveat Emptor and Warning suggest the typical cases where each regime is superior. Warning is the better rule when the two quality dimensions are *complements*, such that it is important for buyers to have high-quality on both dimensions and even one low-quality dimension significantly lowers buyers' payoffs. When the two dimensions are complements it is important for buyers to distinguish products with two high-quality dimensions from all other products, hence the superiority of Warning. The implication is that consumer products and services for which quality dimensions are complements would be better served by the Warning regime. A warranty, for example is valuable if it both covers a broad range of defects and provides coverage for a long period of time. There is little value in long duration if it does not cover the important defects, and there is little value in an all-inclusive warranty if it expires before defects manifest. A home improvement is valuable only if both the materials used and the workmanship are of high quality. And a cruise ship vacation is luxurious only if the private suite and the public amenities are both high-quality. These are scenarios in which even one low-quality dimension can destroy the value, and the Warning regime's ability to focus buyers' attention on such deficiencies accounts for its superiority.

The Caveat Emptor regime is the better rule when the two quality dimensions are *substitutes.* Here, it is critical for buyers to know that there is at least one high-quality dimension, namely, to separate and distinguish products with two low-quality dimensions. Having a second high-quality dimension is of relatively less value. The fact that Caveat Emptor makes sellers with one and two high-quality dimensions indistinguishable matters less to buyers who care primarily about securing at least one high-quality dimension. The substitutes case captures products with dimensions that overlap, or serve a similar function. A college course could be valuable if *either* the professor or the teaching assistant is effective, they don't both have to be high quality. A visit to a theme park or a restaurant can be valuable if some of the rides or menu items are high-quality, they don't all have to be. And used products are valuable if either their working condition or warranty are high quality. In general, a product that has more features than a buyer has capacity to enjoy could be highly valuable even if not all features are high-quality. Caveat Emptor's ability to focus buyers' attention on the high-quality dimensions accounts for its superiority.

Our analysis suggests a general principle that identifies the optimal legal regime—the Principle of Informational Priority (PIP). It is a conditional disclosure regime: sellers must warn about low-quality dimensions when such information is critical to buyers; otherwise, when it is more important for buyers to know about a high-quality dimension, sellers are under no obligation to disclose. PIP requires the Warning regime when a single high-quality provides minimal value (the

complements case), and when product prices exceed this minimal value. Otherwise, when high-quality on only one dimension provides significant value (the substitutes case) or product prices are below this value, PIP rejects a warning requirement and supports Caveat Emptor.

We can now return to the motivating concern about manipulation. Caveat Emptor allows sellers to prioritize information about high-quality dimensions. This strategic prioritization is detrimental to buyers, when it is more important for them to learn whether there is a single low-quality dimension (complements with higher prices). Sellers' mislaid priority in disclosure is manipulative. But a problem that exactly mirrors manipulation occurs also when policymakers intervene and force sellers to issue warnings. The Warning regime prioritizes information about low-quality dimensions, which could be detrimental to buyers who are more interested in learning whether there is a single high-quality dimension (substitutes with lower prices).

2.2 Framework of Analysis

Consider a product with two quality dimensions, d_1 and d_2, each of which can be either high or low ("H" or "L"), with equal likelihood, and independently distributed among Sellers. Accordingly, there are four equally likely types of Sellers – HH, HL, LH, and LL.[10] Sellers and Buyers are randomly matched. Buyers each purchase at most one unit of the good.

Information is asymmetric. Sellers know the quality of their product, while Buyers initially know only the distribution of Seller types. Sellers can disclose information whether a dimension is H or L, and may only do so in a truthful manner.[11] We make the crucial assumption that Buyers have limited capacity to digest information: they can review information about only one dimension (that is disclosed to them). If Seller discloses information about both dimensions, Buyer reviews information about one, randomly selected dimension.[12] This assumption is motivated by the typical cases in which products have a large number of quality dimensions that require costly effort to review, and thus rational Buyers who invest optimally in information choose to review only a subset of the dimensions. Buyers form rational expectations about the remaining, unreviewed dimensions. In the two-dimension model, this structure is captured by the assumption that Buyers review only one dimension—namely, that the product is sufficiently complex so that Buyers can overcome some, but not all, asymmetric information.[13]

The value of the product to risk-neutral Buyers depends on its mix of qualities, and is denoted by $u(d_1, d_2)$. We set the LL utility to equal 0 (same as the no-purchase option), and the HH utility to 1. In the basic model, we focus on the symmetric case, where the HL and LH utilities are identical, and equal to $x \in [0,1]$. This utility parameter x captures the value of a product with only a single H-dimension, and it will play a critical role in the analysis. The utility levels are summarized in Table 1 below.

[10] More precisely, with a sufficiently large number of firms, 25% is the expected mean share of firms in each category.
[11] We assume that the message space is limited to H or L. This assumption is discussed further in Section 4.5 below.
[12] The assumption is that Buyer does not know whether Seller made a disclosure on a non-reviewed dimension. Otherwise, rational inferences would eliminate the asymmetric information problem.
[13] For other studies that take a similar approach, see, e.g., Spiegler (2006) (assuming, as in Osborne and Rubinstein (1998), that consumers randomly sample pricing information on a single dimension of the product); Persson (2018) (assuming that consumers rationally allocate their 'attention budget').

	$d_2 = H$	$d_2 = L$
$d_1 = H$	$u(H,H) = 1$	$u(H,L) = x$
$d_1 = L$	$u(L,H) = x$	$u(L,L) = 0$

Table 1: Quality Combinations and Utility Levels

Let $\hat{u}$ denote Buyers' expected utility, given Sellers' disclosure. On the Sellers' side, we assume in the basic model that the per-unit cost is c, regardless of quality. We also assume that the market is competitive, such that price is equal to cost: $p = c$. If there is a sale, then the Seller's profit is zero, and the Buyer's expected surplus (the consumer surplus, "CS"), is $CS = \hat{u} - c$. In this setup, the entire surplus is captured by Buyers, and thus the consumer surplus is equal to overall welfare.[14]

The assumption of quality-independent per-unit cost fits scenarios in which quality is the result of a prior investment (a fixed cost). Higher investment affects the quality of the product, but not the per-unit production cost (c). Recognizing that this assumption is limiting, in Section 3 we relax the perfect competition assumption and allow the per-unit cost to vary with endogenously-determined quality.[15]

We compare three legal regimes. The first regime involves no legal restriction on which dimensions Sellers disclose (as long as they don't lie), and we refer to it as the Caveat Emptor (CE) regime. In particular, Sellers may, if they so wish, disclose only H-quality dimensions.

The second regime requires that Sellers warn Buyers about L-quality dimensions, if such exist. In this Warning (W) regime, the disclosure of L-quality must be prioritized, and it cannot be overshadowed by voluntary disclosures of H-quality features. If HL/LH Sellers choose to disclose the H dimension, they must prioritize the L-quality warning, which means that Buyers will be able to effectively review only the L-quality dimension. For LL Sellers, the Warning regime requires the disclosure of both L dimensions, even though Buyers will randomly review only one.

The third regime, which we call Full Disclosure (FD), requires that Sellers disclose equally-prioritized information on both dimensions, regardless of their quality, even though Buyers will be able to review only one. In practice, FD could require the disclosure of L-quality dimensions, without requiring the prioritization of these L disclosures (in contrast to the Warning regime).

[14] Technically, we need to assume that Seller gets a positive but arbitrarily small fraction epsilon of the overall surplus. Otherwise, Seller might not adopt the disclosure strategy that we derive below, and there would be multiple equilibria.

[15] If we retain the exogenous quality assumption, then it is straightforward to relax the assumption of quality-independent costs and assume that the cost of producing a product is higher when more of its dimensions are high quality. This extension retains the assumption of a uniform price, which exceeds the cost of an HH product; if price was allowed to depend on the product's quality, then buyers would infer quality from price and the underlying asymmetric information problem would go away. (In Section 3.2, we assume that sellers have enough bargaining power to set prices that extract the full expected surplus. Since the expected surplus is a function of sellers' disclosures, and of buyers' inferences from these disclosures, the price varies with quality only to the extent that buyers can distinguish between different quality combinations.) The move to quality-dependent costs does not change our analysis or results, as long as quality remains exogenous. A detailed analysis of this extension is provided in the Appendix.

Since Sellers will voluntarily disclose H-quality dimensions, the only additional requirement is to forbid the prioritization of such H-disclosures. So understood, Full Disclosure is probably the most common regime.

To evaluate these three regimes, we begin by deriving a benchmark: the maximum consumer surplus that can be attained under conditions of perfect information. We then proceed to analyze how each of the three regimes performs vis-à-vis this benchmark, and to rank their performance.

2.3 Benchmark: Perfect Information

With perfect information, Buyers know with certainty which type of Seller they face. When matched with an HH Seller, Buyers' utility is $u = \hat{u} = 1$, and they purchase the product if $c < 1$. When matched with an HL/LH Seller, Buyers' utility is $u = \hat{u} = x$, and they purchase the product if $c < x$. When matched with an LL Seller, Buyers' utility is 0 and they do not purchase the product. Accounting for the distribution of Seller types, the overall consumer surplus is:

$$CS = \begin{cases} \frac{1}{4} \cdot (1-c) + \frac{1}{2} \cdot (x-c) & , \quad c < x \\ \frac{1}{4} \cdot (1-c) & , \quad c \in [x, 1) \\ 0 & , \quad c \geq 1 \end{cases}$$

2.4 Caveat Emptor

Under this regime, Sellers need not disclose the L-dimensions and may (truthfully) disclose only H-dimensions. Rational Buyers recognize Sellers' strategic disclosure and make Bayesian inferences regarding the undisclosed dimension.

Sellers' disclosure strategy is the following: An HH Seller randomly discloses one of the H-dimensions. (It may disclose both H-dimensions, but Buyer will randomly review one of them.) HL/LH Sellers disclose the H-dimension. And an LL Seller remains silent. The Buyer would draw the following inferences: If Seller remains silent, Buyer infers that Seller is LL-type and does not purchase the product. If Seller discloses an H-dimension, Buyer knows that either the other dimension is L (an ex ante likelihood of $\frac{1}{4}$), or the other dimension is H and Seller randomly chose this specific dimension for disclosure (an ex ante likelihood of $\frac{1}{4} \times \frac{1}{2}$). Conditional on viewing a H-disclosure, the likelihood of the other dimension being H is $\frac{1}{3}$, and of being L is $\frac{2}{3}$. Buyer's expected utility is: $\hat{u} = \frac{1}{3} \cdot 1 + \frac{2}{3}x$. When $c < \frac{1}{3} + \frac{2}{3}x$, Buyer purchases the product and consumer surplus is $\frac{1}{3} + \frac{2}{3}x - c = \frac{1}{3}(1-c) + \frac{2}{3}(x-c)$.[16] If the price is higher, Buyer does not purchase.

[16] When c is between x and $\frac{1}{3} + \frac{2}{3} \cdot x$, Buyer would have the incentive to spend costly effort to review the other dimension. In general, Buyer can spend a cost k to acquire full information and, thus, in this region, to avoid the

Ex ante, there is a probability of $\frac{1}{4}$ that Buyer meets a silent Seller, whom Buyer infers to be LL-type (with consumer surplus of 0), and a probability of $\frac{3}{4}$ that Buyer meets a Seller who discloses H, in which case the expected surplus is $\frac{3}{4}\left(\frac{1}{3}(1-c)+\frac{2}{3}(x-c)\right)=\frac{1}{4}(1-c)+\frac{1}{2}(x-c)$. Thus, the consumer surplus under Caveat Emptor is:

$$CS=\begin{cases}\frac{1}{4}(1-c)+\frac{1}{2}(x-c) & , \quad c<\frac{1}{3}+\frac{2}{3}x \\ 0 & , \quad c\geq\frac{1}{3}+\frac{2}{3}x\end{cases}$$

Comparison to the Perfect Information Benchmark. At the low range of costs, when $c<x$, Buyers should optimally purchase the good, as long as they are facing either an HH or an HL/LH Seller. Caveat Emptor allows Buyers to identify LL Sellers with certainty (when Seller is silent and makes no disclosure). Buyers make the optimal purchase decision in this range of low prices because their only remaining uncertainty—whether the product has one or two H-dimensions—does not matter.

At the high end, when $c\in(x,1)$, things don't look so good under Caveat Emptor. In the perfect information benchmark, Buyers distinguish between HH Sellers and HL/LH Sellers and make a purchase only from the HH Seller. In the Caveat Emptor regime, Buyers cannot distinguish between HH and HL/LH Sellers and cannot make such selective purchase decisions. As a result, Caveat Emptor distorts Buyers' decisions. Buyers make *excessive* purchase when they receive a disclosure of H and $c\in[x,\frac{1}{3}+\frac{2}{3}x)$. At this price range, Buyers purchase the product despite the remaining uncertainty. If the other dimension is L, Buyers end up with negative ex-post surplus. Buyers are also at risk of making *insufficient* purchases, when they receive a disclosure of H and $c\in[\frac{1}{3}+\frac{2}{3}x,1)$. At this price range, Buyers do not purchase the product, and lose a positive surplus when the other dimension is H.[17]

Figure 1 below depicts these distortions. The black line represents the maximal, perfect-information surplus and the red line represents the surplus under Caveat Emptor. When $c<x$, the two lines merge and there is no distortion. When $c\in[x,\frac{1}{3}+\frac{2}{3}x)$, there is positive consumer surplus under Caveat Emptor (the red line is above 0), but less than the maximal surplus. When $c\in[\frac{1}{3}+\frac{2}{3}x,1)$, there is no purchase under Caveat Emptor and surplus is zero, whereas perfect information allows for valuable purchases from HH Sellers.

---

inefficient purchase from an HL or LH Seller. Buyer would make this expenditure if and only if $k<\frac{2}{3}\cdot(c-x)$. Our Buyer-learns-one-dimension assumption is equivalent to an assumption that $k>\frac{2}{3}\cdot(c-x)$.

[17] If $c\in[x,\frac{1}{3}+\frac{2}{3}x)$, CS is smaller by $\frac{1}{4}\cdot(1-c)-\frac{3}{4}\left(\frac{1}{3}+\frac{2}{3}x-c\right)=\frac{1}{2}(c-x)$, as compared to the perfect information benchmark. If $c\in[\frac{1}{3}+\frac{2}{3}x,1)$, CS is smaller by $\frac{1}{4}(1-c)$, as compared to the perfect information benchmark. Note that, since $x<1$, we have $x<\frac{1}{3}+\frac{2}{3}x$, and both distortions occur.

This analysis suggests a key insight: Caveat Emptor is distortive when $c \in (x, 1)$. As the value of *x* decreases—namely, as the importance of having two high dimensions is greater—Caveat Emptor draws further from the perfect information benchmark. Informally, a low *x* makes the difference between one versus two H-dimensions greater, and it becomes all the more important for Buyers to know that both dimensions are H, namely, that there is at least one L-dimension. Caveat Emptor does not deliver this information to Buyers.

2.5 Warning

When it is important for Buyers to recognize the presence of at least one L-dimension, a straightforward legal intervention would require that Sellers prioritize this information. Under a "Warning" regime, LH/HL Sellers have to disclose the L-dimension and prioritize it over their voluntary disclosure of the H dimension.[18] LL Sellers have to disclose both L-dimensions, and Buyers will randomly review one. And HH Sellers may voluntarily disclose one or both H-dimensions (and if both are disclosed, Buyer will review one of them randomly.) We show that the Warning regime prevents the manipulation created by Caveat Emptor, but introduces a different, unintended, distortion.

If Buyers observe an H-dimension, they infer that the Seller is HH. They purchase the product if the price is below 1, and enjoy a consumer surplus of: $CS = 1 - c$. If Buyers observe a warning about an L-dimension, they know that either the other dimension is H (an ex ante likelihood of $\frac{1}{4}$), or the other dimension is L, and Buyer randomly chose this specific dimension for review (an ex ante likelihood of $\frac{1}{4} \times \frac{1}{2}$). Thus, conditional on viewing a warning of L, the likelihood of the other dimension being L is $\frac{1}{3}$ (in which case the value of the product is 0), and the likelihood of the other dimension being H is $\frac{2}{3}$ (in which case the value of the product is *x*). The expected utility from purchasing a good in this case is: $\hat{u} = \frac{2}{3}x$. When $c < \frac{2}{3}x$, Buyer purchases the good and enjoys a consumer surplus of: $CS = \frac{2}{3}x - c = \frac{2}{3}(x - c) - \frac{1}{3}c$. If cost, and price, are higher, Buyer will not purchase the good.

Ex ante, there is a probability of $\frac{1}{4}$ that Buyer meets an HH Seller who discloses H; the expected surplus is: $\frac{1}{4}(1 - c)$. And there is a probability of $\frac{3}{4}$ that Buyer meets a Seller who discloses L; the expected surplus is: $\frac{3}{4}\left(\frac{2}{3}(x - c) - \frac{1}{3}c\right) = \frac{1}{2}(x - c) - \frac{1}{4}c$. Consumer surplus is:

[18] Generally, a Warning regime could allow sellers to also disclose, if they so choose, any H dimension, but it requires that such voluntary High disclosures not compete for attention with the L warning. Since Buyers have limited capacity to review, they will only observe the L dimension, and any additional H disclosure would be futile for Sellers. For simplicity, then, we assume that under a Warning regime only HH sellers will disclose an H dimension.

$$CS = \begin{cases} \frac{1}{4} \cdot (1-c) + \frac{1}{2} \cdot (x-c) - \frac{1}{4} \cdot c & , \quad c < \frac{2}{3}x \\ \frac{1}{4} \cdot (1-c) & , \quad c \in \left[\frac{2}{3}x, 1\right) \\ 0 & , \quad c \geq 1 \end{cases}$$

Comparison to the Perfect Information Benchmark. At the higher range of prices, when $c \in [x, 1)$, Buyers should purchase the good only when facing an HH Seller. The Warning regime allows Buyers to identify HH Sellers with certainty—when a Seller does not issue an L warning. When a warning is observed, Buyers are uncertain whether the product has one or two L-dimensions, but at such high prices they should not buy either way, and thus their uncertainty does not affect the optimal decision.

At the low end, when $c < x$, Buyers should purchase only from HL/LH Sellers, but Buyers who observe an L warning cannot distinguish HL/LH Sellers from LL Sellers. As a result, Buyers make *excessive* purchases, when $c < \frac{2}{3}x$. At such a low price, Buyers purchase the product and, if it turns out that the other dimension is also L, Buyers ends up with a negative ex-post surplus. Buyers also make *insufficient* purchases, when $c \in \left[\frac{2}{3}x, x\right)$. Recognizing the risk that the other dimension might (also) be L, Buyers do not buy at this price, forgoing the value of a profitable purchase when the other dimension turns out to be H.[19]

Figure 1 below depicts these distortions. The black line represents the maximal, perfect-information surplus and the blue line represents the surplus in the Warning regime. When $c \geq x$, the two lines merge and there is no distortion. When $c < x$, there is still positive consumer surplus in the Warning regime, but less than the maximal surplus.

This analysis suggests a key insight (the mirror image of the insight about Caveat Emptor): The Warning regime is distortive when $c < x$. As the value of x decreases—namely, as the difference between LL and HL/LH decreases—Warning comes closer to the perfect information benchmark, delivering a more efficient outcome. Informally, a lower x makes the difference between one versus zero H-dimensions smaller. In this scenario, it is all the more important for Buyers to distinguish the HH product from all other products. Warning delivers this information to Buyers.

2.6 Full Disclosure

We now consider a different mandated disclosure regime: Full Disclosure (FD). It requires that Sellers disclose the quality of both dimensions without prioritizing. Buyers, limited in their ability to review, randomly view one of the two disclosed dimensions. It captures many of the common disclosure laws, which allow sellers to make quality claims but also require "conspicuous" disclaimers and alerts regarding low quality.

[19] If $c < \frac{2}{3}x$, CS is smaller by $\frac{1}{4}c$, as compared to the perfect information benchmark. If $c \in \left[\frac{2}{3}x, x\right)$, CS is smaller by $\frac{1}{2}(x-c)$, as compared to the perfect information benchmark.

Regardless of the dimension they review, Buyers know that the other dimension is either H or L, with equal likelihood. Thus, if they view an H-dimension, they expect the product to be either HH or HL with an expected utility of $\hat{u} = \frac{1}{2}(1 + x)$. When the price is below this value, Buyers make a purchase and enjoy a payoff of $CS = \frac{1}{2}(1 + x) - c = \frac{1}{2}(1 - c) + \frac{1}{2}(x - c)$. If Buyers view an L-dimension, they expect the product to be either LL of LH with an expected utility of $\hat{u} = \frac{1}{2}x$. They purchase when the price is below this value with a payoff: $CS = \frac{1}{2}x - c = \frac{1}{2}(x - c) - \frac{1}{2}c$. Accounting for the distribution of Seller types (and for the equal likelihood that Buyers will review any of the two product dimensions), the expected consumer surplus under Full Disclosure is:

$$CS = \begin{cases} \frac{1}{4}(1-c) + \frac{1}{2}(x-c) - \frac{1}{4}c & , \quad c < \frac{1}{2}x \\ \frac{1}{4}(1-c) + \frac{1}{4}(x-c) & , \quad c \in \left[\frac{1}{2}x, \frac{1}{2}(1+x)\right) \\ 0 & , \quad c \geq \frac{1}{2}(1+x) \end{cases}$$

Comparison to the Perfect Information Benchmark. Regardless of the observed dimension, Buyers remain uncertain about the Seller's type. Therefore, they end up making either excessive purchases (when the undisclosed dimension is L) or insufficient purchases (when the undisclosed dimension is H). Figure 1 below depicts these distortions. The green line represents the surplus in the FD regime, and it is always below the perfect-information black line.

2.7 Comparing the Three Regimes

No regime achieves the perfect-information outcome. In each regime, some uncertainty about the product's quality remains, resulting in both excessive and insufficient purchases. But the magnitude of these distortions, and the resulting consumer surplus, depend on the regime, as we now show.[20]

Under the Caveat Emptor regime, distortions occur when $c > x$. At these high prices, Buyers need to know if a product shown to have at least one H dimension has an additional H dimension. When they observe one H, Buyers make a purchase only when the price is below the $\frac{2}{3}x + \frac{1}{3}$ cutoff. They might lose potential surplus when, due to their prudence, they decide not to purchase; or they might incur actual loss when they do purchase. Importantly, under Caveat Emptor there is no distortion when $c < x$. At such low prices, Buyer's know for sure (when they observe an H dimension) that the product is worth purchasing.

Under the Warning regime, in contrast, distortions occur when $c < x$. At these low prices, Buyers need to know if a product warned to have at least one L dimension has an additional L dimension.

[20] Table A1 in the Appendix collects the consumer surplus values under CE, W, FD and the perfect-information benchmark. The information is presented graphically in Figure 1.

When they observe L, Buyers make a purchase only when the price is below the $\frac{2}{3}x$ cutoff. Again, the decision to purchase and the decision not to purchase could turn out to be ex-post inefficient. In contrast with Caveat Emptor, under the Warning regime there is no distortion when $c > x$. At such high prices, it suffices to observe a warning about a single L dimension to make the efficient decision not to purchase.

Unlike the Caveat Emptor and Warning regimes, the Full Disclosure regime does not allow Buyers to identify any Seller with certainty. When the disclosure observed is L, Buyers purchase if the price is below $\frac{1}{2}x$. And when the disclosure observed is H, Buyers purchase if the price is below $\frac{1}{2}(1+x)$. Either decision can result in ex-post inefficiency. For some parameter values, Full Disclosure yields higher consumer surplus relative to either the Caveat Emptor or Warning regimes. This can be seen in Figure 1, in the range of c values for which the FD payoff line lies above either the CE or the W lines. It represents the subtle fact that Full Disclosure causes a milder version of the distortions that occur under Caveat Emptor or Warning. For example, Full Disclosure dominates Caveat Emptor in the $c \in [x, \frac{1}{3}+\frac{2}{3}x)$ range, where both regimes lead to inefficient purchases from HL/LH Sellers when Buyer observes H; under Caveat Emptor, Buyers facing HL/LH Sellers always observe H (because of these Sellers' strategic disclosure), whereas under FD Buyers facing these Sellers observe H only when randomly viewing the H-dimension. As a result, the likelihood of observing H, and thus making an inefficient purchase, is lower under FD. (Full Disclosure also dominates Caveat Emptor in the $c \in \left[\frac{1}{3}+\frac{2}{3}x, \frac{1}{2}(1+x)\right)$ range, where FD still leads to some inefficient purchases from HL or LH Sellers and CE leads to a costlier failure to purchase from HH Sellers.)[21]

While Full Disclosure dominates Caveat Emptor or Warning in certain cases, the policy implications of this "local" superiority are more subtle. First, Full Disclosure is *never* the optimal regime—it is always inferior to either Caveat Emptor or Warning. As can be seen in Figure 1, when $c < x$ Caveat Emptor is the most efficient rule, achieving the first-best. And, when $c > x$, Warning is the most efficient rule, achieving the first-best. There are no c values for which Full Disclosure is optimal. These policy implications are depicted graphically is Figure 2 below, which depicts all combinations of (c, x) and the corresponding optimal regime. Moreover, if the law is constrained to choose one regime for all combinations of (c, x), there is no distribution of parameters for which FD dominates both CE and W. This result follows from the symmetry between FD's "local" advantage vis-à-vis CE and its disadvantage vis-à-vis W, and vice versa.[22]

[21] Full Disclosure dominates Warning in the $c \in \left[\frac{2}{3}x, x\right)$ range, where both regimes lead to an inefficient failure to purchases from HL or LH Sellers following a Low disclosure; under Warning Buyers facing HL/LH Sellers always observe Low, whereas under FD Buyers facing these Sellers observe Low only if they happen to review the Low dimension. (Full disclosure also dominates Warning in the $c \in \left[\frac{1}{2}x, \frac{2}{3}x\right)$ range, where FD still leads to some inefficient failures to purchase from HL or LH Sellers and W leads to a more costly purchase from LL Sellers.)

[22] The result is formally proved in the Appendix.

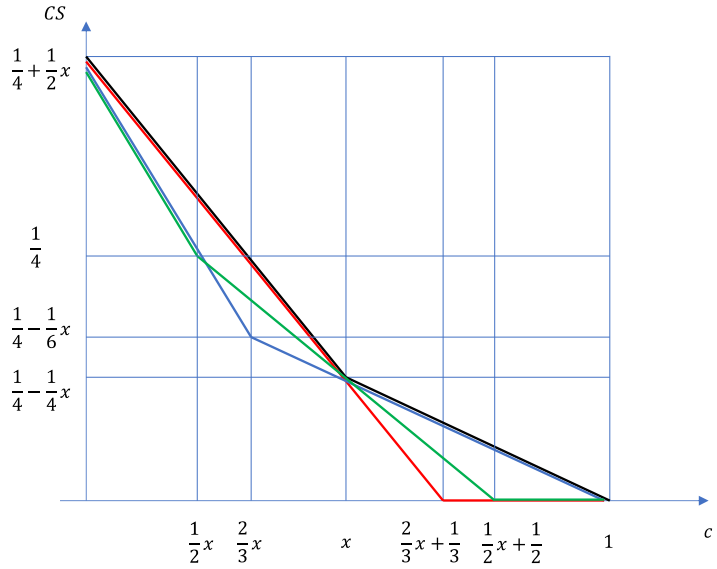

Figure 1

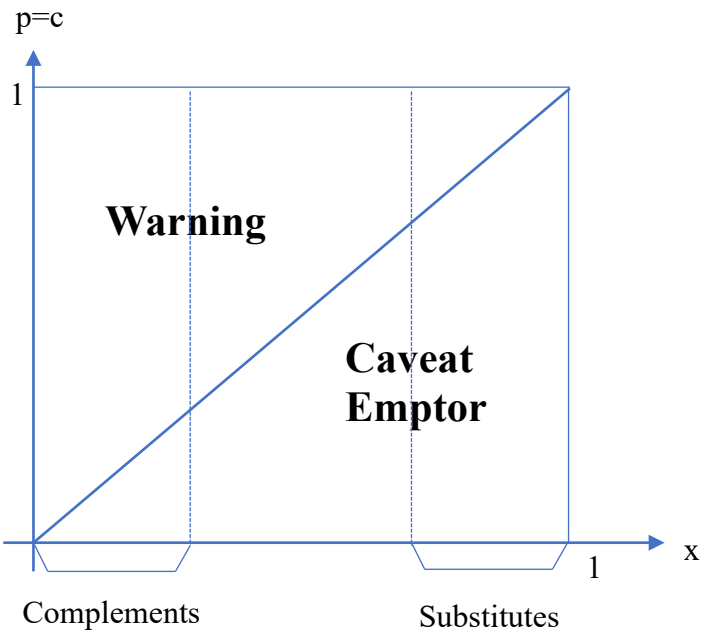

Figure 2

2.8 The Principle of Informational Priority (PIP)

From the preceding analysis, we can derive a general principle that identifies the optimal legal regime—the Principle of Informational Priority (PIP). It prescribes a conditional disclosure regime: "issue a warning if and only if $x < c$." When this condition is satisfied, Buyers' priority is to know about an L dimension. Otherwise, when $x > c$, no warning is necessary because Buyers' priority to identify H dimensions is fully addressed by Sellers' voluntary disclosure.

The PIP condition requires courts or lawmakers to verify the c and x parameters. Several heuristics could help approximate this conditional mandate. First, when x is low, the environment is more likely to satisfy this warning condition. This is the case where quality dimensions can be loosely thought of as *complements*—where Buyers need to know that none are L. Conversely, when x is high, PIP does not require a warning. This is the case where quality dimensions are *substitutes*—where Buyers need to know that at least one dimension is H. Sellers have the incentive to supply this information.

Second, when c is high, it is more likely that a warning would be necessary. Here we are dealing with high-cost, high-price, products, for which it is less likely that a single H-dimension will give Buyer enough utility, thus a warning is critical. Conversely, when c is low, a product with even one H-dimension is more likely to generate enough utility to justify its cost, and a warning could be harmful by deprioritizing Sellers' disclosure of H.

These practical criteria are, of course, crude approximations for the underlying condition "issue a warning if and only if $x < c$." If, say, both x and c are high, these simple criteria offer conflicting prescriptions. This is the case of *expensive substitutes*—costly products where it suffices to have one H-dimension to gain significant value, but a second H-dimension might be needed to justify the high price. (High end restaurants?) Or, if both x and c are low, again the simple criteria are pointing in opposite directions. This is the case of *cheap complements*—low cost products where a single H-dimension does not yield much value, but perhaps enough to justify the low price. (discount vacation packages?)

3. Extensions

In this Section, we consider several extensions in which key assumptions of the model are relaxed, exploring the generality of the analysis and identifying additional refinements. We focus on three aspects. First, we introduce asymmetry in the quality dimensions, allowing some to contribute more than others to the overall value. We show that such ranking across dimensions can make legal intervention more effective, but also potentially more harmful. Second, we examine sellers' incentive to invest in quality, showing how the prospect of ex post manipulation distorts this incentive. Third, we examine sellers' incentives to make products more complex, and thus more manipulable. As in Part II, we begin each of the three extensions with an informal summary of the analysis, followed by a more complete derivation of the claims.

3.1 Asymmetric Quality Dimensions

3.1.1 Informal Summary

In the basic model, we considered a product with two symmetrically valued dimensions, which meant that, if there was only one high dimension, buyers didn't care which one it was. We now relax this assumption and consider asymmetric dimensions, where one ("primary") dimension is more important than the other ("secondary") dimension. If the product is to have only one high quality dimension, buyers prefer that it would be the primary dimension.

The asymmetry across dimensions offers a key new element in the analysis: a focal priority in the disclosure of information. The primary dimension is more important and ought to be prioritized both by sellers and by the law's disclosure mandates. Recognizing the focal priority of the primary dimension, buyers would be make inferences about this dimension even when it is not disclosed.

Consider *Caveat Emptor*. Buyers know that sellers will prioritize disclosure of the primary dimension. Accordingly, if they observe a seller disclosing high quality in the secondary dimension, buyers infer that the primary dimension must be low quality—otherwise it would have been prioritized for disclosure. The inherent ranking of dimensions allows buyers to infer more information than they could in the symmetric-dimensions scenario, and as a result they will never have any uncertainty about the primary dimension. If it is high quality, sellers will disclose it. Any other disclosure (or silence) by sellers means that the primary dimension is low-quality. The only remaining uncertainty under Caveat Emptor surrounds the secondary dimension, and it occurs when the primary dimension is (disclosed to be) high-quality. Some distortion remains, but its incidence and magnitude decline as the difference – in terms of utility impact – between the primary and secondary dimensions increases.

The presence of a primary dimension also improves the warning regime. In the symmetric case, buyers who viewed a warning on a certain dimension were unable to infer anything about the other dimension. Now, the law can mandate a priority: if both dimensions are low-quality, sellers have to disclose only the primary dimension, and are prohibited from issuing a warning regarding the secondary dimension, so as to allow buyers to direct their limited attention to the more important information. (Or, more realistically, the law would require that the primary dimension warning be more conspicuous.) Under this *Priority Warning* regime, buyer will always know the quality of the primary dimension. Either it is disclosed to be low, or, if they observe a warning on the secondary dimension, buyers infer that the primary dimension is high. The only remaining uncertainty under Priority Warning surrounds the secondary dimension, and it occurs when the primary dimension is (warned to be) low-quality. Some distortion remains, but again its incidence and magnitude decline as the difference – in terms of utility impact – between the primary and secondary dimensions increases.

The analysis of the asymmetric case yields several insights. First, the problem of manipulation by mislaid priorities, and the welfare loss it creates, is mitigated, because buyers are able to infer more information. Second, the problem of manipulation does not go away, requiring a thoughtful regulatory intervention. One aspect of the optimal intervention, derived in the basic model, still holds: Caveat Emptor is optimal when prices are high, and a Warning regime is optimal when

prices are low. But the complements/substitutes heuristic needs to be refined, and as the relative importance of the two dimensions diverges, it becomes easier to choose between the Caveat Emptor and Warning regimes. When a product has a "super-dimension" that accounts for much of the value, Caveat Emptor is superior regardless of the secondary dimension's value contribution, because it provides buyers with perfect information about this super dimension. Conversely, when the product has a "slack dimension" that does not produce much value, Priority Warning is superior regardless of the primary dimension's value contribution because, here too, buyers perfectly infer the quality of the primary dimension, and the remaining uncertainty is about the slack dimension which has minimal impact on purchasing decisions. Third, as the payoff-gap between the primary and secondary dimensions increases, the choice between Caveat Emptor and Priority Warning becomes less important, because the large payoff gap creates a wide range of prices for which both regimes succeed in solving buyers' imperfect information problem.[23] In the asymmetric dimensions case, the *principle of information priority* should be used both to guide the choice between the Caveat Emptor and Priority Warning regimes and to select the prioritized warning in the (Priority) Warning regime.

3.1.2 Revised Formal Framework

We now demonstrate these claims within a more formal setting. Our baseline model, in Section 2, assumed that HL and LH products are payoff-identical, i.e., $u(1,0) = u(0,1) = x$, where $x \in [0,1]$. To relax this symmetry assumption, let $u(1,0) = x_1$ and $u(0,1) = x_2$, where $x_1, x_2 \in [0,1]$. Let $x_1 > x_2$, namely, D1 is the more important quality dimension.

Perfect information benchmark. If Buyers know with certainty which type of Seller they face, they purchase from HH Sellers whenever $c < 1$, from HL Sellers whenever $c < x_1$, from LH Sellers whenever $c < x_2$, and they do not make a purchase when facing LL Sellers.

3.1.3 Caveat Emptor

In the symmetric-dimensions model, an HH Seller randomly chose to disclose one of the two quality dimensions. In the asymmetric case, an HH Seller discloses the more important dimension, D1. Buyers therefore know with certainty when they are facing an LH Seller—if Seller discloses H on D2. And, as in the baseline model, Buyers also know with certainty when they are facing an LL Seller—if Seller is silent. The only remaining uncertainty is due to Buyers' inability to distinguish between HH and HL Sellers, who both disclose H on D1. In this case, Buyers estimate the expected value at $\frac{1}{2}(1 + x_1)$, and make the purchase if the price is lower. This uncertainty accounts for some loss of surplus: Buyer inefficiently purchases the product from HL Sellers when $c \in [x_1, \frac{1}{2} \cdot (1 + x_1))$, for a negative expected payoff of $\frac{1}{4}(x_1 - c)$, and inefficiently fails to purchase the product from HH Sellers when $c \in [\frac{1}{2} \cdot (1 + x_1), 1)$, forgoing a positive expected payoff of $\frac{1}{4}(1 - c)$. Both distortions occur when $c > x_1$.

[23] Under Caveat Emptor, as long as the product price is below the value of the primary dimension buyers make the optimal decision to purchase. Under Priority Warning, as long as the product price is above the value of the secondary dimension buyers make the optimal decision not to purchase. Thus, under both regimes buyers make optimal decisions when the price is between the value contribution of the primary and the secondary dimension.

3.1.4 Priority Warning

In the asymmetric-dimensions extension, a legal mandate can do better than require the disclosure, with priority, of any or all L dimensions. It can now require that Sellers disclose the more important L dimension. In the model, LL Sellers have to disclose only D1, and will be prohibited from disclosing the Low on D2. We show that the Priority Warning regime is superior to other disclosure regimes, and compare Priority Warning to Caveat Emptor.

Under Priority Warning, Buyers continue to infer the presence of an HH Seller (when there is no warning), and now they can also infer the presence of an HL Seller (when the warning they observe is on D2). This additional inference accounts for Priority Warning's superiority over the (regular) Warning regime. But Buyers are still unable to distinguish between LL and LH Sellers, who both disclose L on D1. In this case, Buyers estimate the expected value at $\frac{1}{2}x_2$, and make the purchase if the price is lower. This uncertainty accounts for some loss of surplus: Buyer inefficiently purchases from LL Sellers, for a negative expected payoff of $-\frac{1}{4}c$; and inefficiently fails to purchase from LH Sellers when $c \in \left[\frac{1}{2}x_2\,, x_2\right)$, forgoing a positive expected payoff of $\frac{1}{4}(x_2 - c)$. Both distortions occur when $c < x_2$.

3.1.5 Caveat Emptor v. Priority Warning

As in the basic model, Caveat Emptor and Priority Warning share important parallels. They both prioritize D1, enabling Buyers to make additional inferences. And while both regimes perform better, as compared to the baseline, symmetric-dimensions case, each regime retains its relative advantages (and disadvantages) vis-à-vis the other. Caveat Emptor enjoys an advantage when the product's cost, and price, are low, and Priority Warning enjoys an advantage when the cost, and price, are high. Intuitively, the distortions under Caveat Emptor are due to Buyer's inability to distinguish an HH Seller from an HL Seller, which is irrelevant when the price is low enough to justify a purchase from either Seller. Thus, distortions under Caveat Emptor occur only when $c > x_1$. By contrast, the distortions under Priority Warning are due to Buyer's inability to distinguish an LL Seller from an LH Seller, which is irrelevant when the price is high enough to deter a purchase from either Seller. These distortions therefore occur only when $c < x_2$.

Accordingly, when a product has a "super dimension" that accounts for much of the value (x_1 close to 1), Caveat Emptor is superior regardless of the value of the other, lesser, dimension. And when a product has a "slack dimension" that does not produce much value (x_2 close to 0), Priority Warning is superior regardless of the value of the more important dimension. Moreover, in the entire intermediate range, in which $c \in [x_2, x_1]$, the choice of regime is irrelevant, because no distortion occurs under either regime. As the degree of asymmetry increases, namely, as the interval of no-distortion $[x_2, x_1]$ grows, both Caveat Emptor and Priority Warning approach the perfect information benchmark.[24] See Figure 3 below.

[24] At the extreme, where $x_1 = 1$ and $x_2 = 0$, the product has only one payoff-relevant dimension, and the problem of asymmetric information is fully resolved, regardless of the legal regime.

In all, the Principle of Information Priority provides guidance to policy makers in two ways. First, it identifies which information the warning ought to prioritize – information about the dimension with the higher "x". Second, it determines the range of costs, and prices, for which the warning is needed: "warning only if $c > x_1$." (PIP would work equally well if the criterion is expanded, "warning only if $c > x_2$," but expanding the warning requirement is not necessary to achieve the perfect information benchmark.) If a conditional rule is infeasible, PIP guides the choice between Caveat Emptor and Priority Warning, by fine-tuning the complements v. substitutes heuristic: We still have a complements case, where both x_1 and x_2 are relatively low, and Priority Warning is the better rule. And we still have a substitutes case, where both x_1 and x_2 are relatively high, and Caveat Emptor is the better rule. But we also have additional cases. There is a "super dimension" case, where x_1 is high and x_2 is intermediate, and Caveat Emptor is the better rule. And there is the "slack dimension" case, where x_2 is low and x_1 is intermediate, and Priority Warning is the better rule. Finally, there are products with both a "super dimension" (high x_1) and a "slack dimension" (low x_2); for these products, Caveat Emptor and Priority Warning are roughly equivalent.

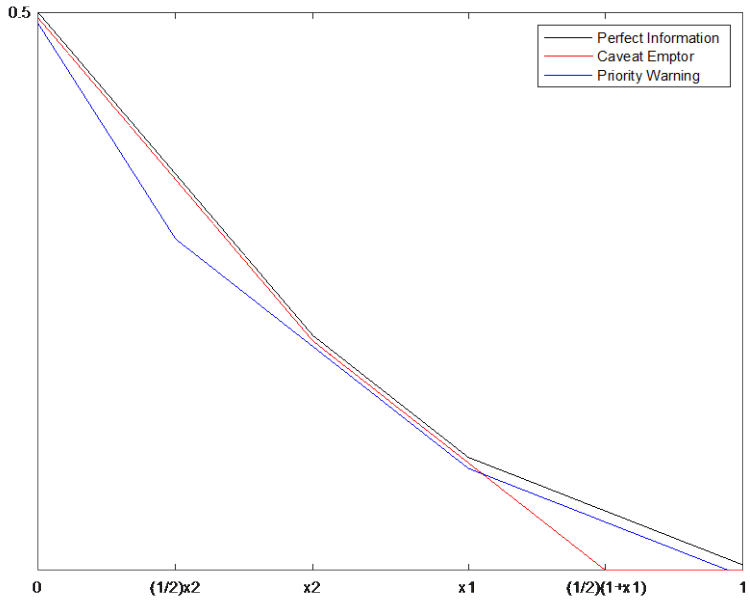

Figure 3

3.2 Investment in Quality

3.2.1 Informal Summary

The analysis in Section 2 made the simplifying assumptions that product quality is exogenously given, and that the share of products with different quality combinations is fixed and predetermined. It allowed us to focus on mislaid priorities as a source of ex post allocative inefficiency. But it is also unrealistic and limiting. If sellers are able to invest in quality, we now ask whether the potential for buyer manipulation affects their investment decisions, and whether

it does so in a harmful way. Does the allocative inefficiency due to manipulation lead to productive distortions?

We allow sellers to make investments that determine the distribution of quality levels. Sellers will make such investments only if they can capture at least some of the benefit from their investment, and we assume, for simplicity, that they have enough bargaining power to set prices that expropriate buyers' entire expected value from the product. The introduction of endogenous quality significantly changes the analysis.

Under *Caveat Emptor*, the key new insight is that sellers would never produce products with two high-quality dimensions. Because buyers cannot distinguish between products with one versus two high-quality dimensions, sellers with two high-quality dimensions will not be able to distinguish themselves and charge a higher price that could justify the cost of a second high-quality dimension. Recognizing this, buyers who observe a disclosure of a high-quality dimension would correctly infer that the other dimension is low quality. This unravelling of information guarantees that there is no manipulation in equilibrium. There is no allocative distortion, but there is welfare loss whenever it is socially desirable to produce products with two high-quality dimensions. Buyers are now hurt by being sold lower quality products.

Once again, a symmetric distortion occurs in the *Warning* regime. Because buyers are unable to distinguish between products with one versus two low-quality dimensions, sellers will not be able to charge more for a product with one high-quality dimension, and will not produce it. Thus, if buyers observe a warning, they will correctly conclude that the product has two low-quality dimensions. Sellers will therefore produce only two kinds of products: all-high, or all-low. Manipulation of buyers does not happen, but a loss of surplus occurs whenever it is socially desirable to produce products with one high-quality dimension. Some buyers will have to pay extra for an all-high product, when they would have preferred a cheaper product with one high-quality dimension.

This analysis reveals a new source of productive inefficiency due to the problem of mislaid priorities, and it reinforces the relative strength of each regime, as identified in Section 2. When the product's dimensions are complements, buyers are not interested in a single high-quality dimension. In their search for a product with two high-quality dimensions, buyers would be better off under a Warning regime that allows them to perfectly identify such products and thus guarantees that sellers will produce them. In contrast, when the product's dimensions are substitutes, buyers are particularly interested in a product with a single high dimension. Caveat Emptor incentivizes the production of such products.

3.2.2 Revised Formal Framework

We employ the setup from Section 2, but allow Sellers to choose the quality of their product. We denote the costs of the four types of products as c_{LL}, c_{LH}, c_{HL} and c_{HH}, and the shares of the products in the market as ρ_{LL}, ρ_{LH}, ρ_{HL} and ρ_{HH}. Producing higher quality products is assumed to be costlier, and so $c_{LL} < c_{LH} = c_{HL} < c_{HH}$. We normalize c_{LL} to zero, and let $c_{LH} = c_{HL} \equiv c_{1H}$. We focus on symmetric equilibria where $\rho_{LH} = \rho_{HL} \equiv \rho_{1H}$.

In the baseline model of Section 2, we assumed perfect competition, such that Sellers' profits were always zero. The competition assumption precludes investment in quality. In this extension, which focuses on incentives to invest, we adopt the opposite assumption: Sellers have all the bargaining power, such that $p = \hat{u}$.[25] Since consumer surplus is zero ($\hat{u} - p = 0$), we focus on overall welfare, $W = \hat{u} - c$, as our normative criterion.

Sellers' disclosure strategies are as in the baseline model. But Buyer's inferences must be adjusted to account for the unequal distribution among the different types of sellers. Under Caveat Emptor, if Seller remains silent, Buyer infers that Seller is an LL type. And if Seller discloses H, Buyer infers that the probability of facing an HH Seller is $\frac{\frac{1}{2}\rho_{HH}}{\frac{1}{2}\rho_{HH}+\rho_{1H}}$, and that the probability of facing an HL or LH Seller is $\frac{\rho_{1H}}{\frac{1}{2}\rho_{HH}+\rho_{1H}}$. Under the Warning regime, if Seller discloses H, Buyer infers that Seller is an HH type. And if Seller discloses L, Buyer infers that the probability of facing an LL Seller is $\frac{\frac{1}{2}\rho_{LL}}{\frac{1}{2}\rho_{LL}+\rho_{1H}}$ and that the probability of facing an HL or LH Seller is $\frac{\rho_{1H}}{\frac{1}{2}\rho_{LL}+\rho_{1H}}$.

Perfect Information. In the perfect information benchmark, Buyers identify Seller's quality with certainty and are thus willing to pay a price that is equal to the actual utility that they gain from the product. Therefore, HH Sellers enjoy a profit of $1 - c_{HH}$, HL and LH Sellers enjoy a profit of $x - c_{1H}$, and LL Sellers get nothing. When $max(1 - c_{HH}, x - c_{1H}, 0) = 1 - c_{HH}$, such that HH products create the largest net benefit, Sellers invest in high quality, on both dimensions, and as a result all sellers are HH. When $max(1 - c_{HH}, x - c_{1H}, 0) = x - c_{1H}$, such that HL and LH products create the largest net benefit, Sellers invest in high quality on only one dimension, and all sellers are HL or LH Sellers. And when $max(1 - c_{HH}, x - c_{1H}, 0) = 0$, Sellers do not invest in quality and all are LL.

Caveat Emptor. Under Caveat Emptor, the consumer cannot distinguish between HH and HL/LH Sellers. Therefore, Sellers will have no incentive to ensure H on both dimensions, and $\rho_{HH} = 0$.[26] Thus, when Buyers observe H, they realize it must be from an HL/LH Seller, and would be willing to pay x. Therefore, HL and LH Sellers enjoy a profit of: $\pi_{HL} = \pi_{LH} = x - c_{1H}$. If $x > c_{1H}$, the profit enjoyed by Seller with one High dimension is positive, and all Sellers will be either HL or LH Sellers. If $x < c_{1H}$, all Sellers will be LL Sellers. When $max(1 - c_{HH}, x - c_{1H}, 0) = 1 - c_{HH}$, Caveat Emptor fails to produce the efficient outcome. Caveat Emptor produces the efficient outcome only when $max(1 - c_{HH}, x - c_{1H}, 0) = x - c_{1H}$ or $max(1 - c_{HH}, x - c_{1H}, 0) = 0$.

---

[25] The analysis is qualitatively similar, when Sellers get some fraction of the overall surplus (not all of the surplus). [Verify]

[26] If Seller remains silent, Buyer's expected utility is: $\hat{u} = 0$. If Seller discloses High, Buyer's expected utility is: $\hat{u} = \frac{\frac{1}{2}\rho_{HH}}{\frac{1}{2}\rho_{HH}+\rho_{1H}} \cdot 1 + \frac{\rho_{1H}}{\frac{1}{2}\rho_{HH}+\rho_{1H}} \cdot x = \frac{\frac{1}{2}\rho_{HH}+\rho_{1H} \cdot x}{\frac{1}{2}\rho_{HH}+\rho_{1H}}$. An HH Seller discloses High (on either D1 or D2, each with probability $\frac{1}{2}$) and gets: $\pi_{HH} = \frac{\frac{1}{2}\rho_{HH}+\rho_{1H} \cdot x}{\frac{1}{2}\rho_{HH}+\rho_{1H}} - c_{HH}$. HL and LH Sellers disclose the one High dimension and get: $\pi_{HL} = \pi_{LH} = \frac{\frac{1}{2}\rho_{HH}+\rho_{1H} \cdot x}{\frac{1}{2}\rho_{HH}+\rho_{1H}} - c_{1H}$. And an LL Seller remains silent, and gets: $\pi_{LL} = 0$. Since $c_{HH} > c_{1H}$, we have: $\pi_{HH} < \pi_{HL} = \pi_{LH}$; and this implies: $\rho_{HH} = 0$.

Note that there is no manipulation in equilibrium. It is the potential for manipulation that destroys the incentives of Sellers to invest in high quality on both dimensions.[27] Of course, in the real world, outside of the confines of our stylistic model, the imperfect information problem will result in ex ante distortions in Sellers' quality investments *and* in ex post manipulation.

Warning. Under the Warning regime, the consumer cannot distinguish between LL and HL/LH Sellers. Therefore, when HH is not feasible, Sellers will choose LL; there is no incentive to ensure High on only one dimension. We thus have: $\rho_{1H} = 0$.[28] Therefore, when Buyers observe L they realize it must be from an LL Seller, and would be willing to pay zero. If $1 > c_{HH}$, the profit enjoyed by a Seller with two High dimensions is positive, and all Sellers will be HH Sellers. If $1 < c_{HH}$, all Sellers will be LL Sellers. The Warning regime produces the efficient outcome, when $max(1 - c_{HH}, x - c_{1H}, 0) = 1 - c_{HH}$ *and* when $max(1 - c_{HH}, x - c_{1H}, 0) = 0$. The Warning regime fails to produce the efficient outcome, when $max(1 - c_{HH}, x - c_{1H}, 0) = x - c_{1H}$.

Comparison. Under Caveat Emptor, Sellers will not invest in HH products. If HH products create the largest net benefit, then Caveat Emptor is inefficient; otherwise, if either HL, LH or LL products create the largest net benefit, then Caveat Emptor is efficient. In the Warning regime, Sellers will not invest HL/LH products. If such one-H products create the largest net benefit, then Warning is inefficient. Otherwise, Warning is efficient. We see that the two regimes diverge in scenarios where either HH is optimal or HL/LH are optimal. If it is cost-effective to have high quality on both dimensions (which is more likely in the complements case), then Warning is the better rule. If it is cost-effective to have high quality on only one dimensions (the substitutes case), then Caveat Emptor is the better rule.

3.3 Multiple, Endogenous Dimensions

3.3.1 Informal Summary

The baseline model considered a 2-dimensional product. It captured the idea that products have more quality dimensions than buyers can effectively review. We now consider the possibility of added complexity, which increases the number of dimensions buyers are unable to review. We first ask a static question: If buyers are able to review a smaller fraction of the attributes, are they more vulnerable to manipulation? Is legal intervention more critical? We then explore a dynamic aspect: Do sellers have an incentive to increase product complexity? We show that some sellers can profit from spurious complexity which allows them to extract more value through manipulation. We study the effects of the different legal regimes on these incentives.

---

[27] This is an illustration of the general unravelling phenomenon. Compare: Akerlof (1970); Shapiro (1982); Beales, Craswell and Salop (1981, 510) ("poor information about the quality … the marketplace responds by channeling competition toward more easily observable product attributes").

[28] If Seller discloses High, Buyer's expected utility is: $\hat{u} = 1$. If Seller discloses Low, Buyer's expected utility is: $\hat{u} = \frac{\rho_{1H} \cdot x}{\frac{1}{2}\rho_{LL} + \rho_{1H}}$. An HH Seller discloses High (on either D1 or D2, each with probability $\frac{1}{2}$) and gets: $\pi_{HH} = 1 - c_{HH}$. HL and LH Sellers disclose the one Low dimension and get: $\pi_{HL} = \pi_{LH} = \frac{\rho_{1H} \cdot x}{\frac{1}{2}\rho_{LL} + \rho_{1H}} - c_{1H}$. And an LL Seller discloses one of the Low dimensions (each with probability $\frac{1}{2}$) and gets: $\pi_{LL} = \frac{\rho_{1H} \cdot x}{\frac{1}{2}\rho_{LL} + \rho_{1H}} - c_{LL}$. Since $c_{1H} > c_{LL}$, we have: $\pi_{LL} > \pi_{HL} = \pi_{LH}$; and this implies: $\rho_{1H} = 0$.

The static question. What happens when the number of dimensions increases, while buyers' ability to review information remains unchanged? Unsurprisingly, as buyers lack information about more attributes, things get worse. With more dimensions, the information content of the disclosures goes down and buyers make less efficient purchasing decisions. The aggravated information problem affects both the Caveat Emptor and Warning regimes, but the relative advantages and disadvantages of the two regimes remain as in Section 2: Caveat Emptor has an advantage when the dimensions are substitutes and the price is lower, and the Warning regime has an advantage when the dimensions are complements and the price is higher. Still, we note the increased concern about manipulation, when the number of dimensions is higher. Under Caveat Emptor, in the complements case (with a high price), the social cost of manipulation – through strategic disclosure of high dimensions – is greater, and thus the benefit from legal intervention, through a Warning regime, is greater. On the other hand, in the substitutes case (with a low price), the social cost of misprioritized mandated warnings also increases with the number of dimensions.

The dynamic question. We next consider the incentives of sellers to add quality dimensions. The analysis demonstrates that low-quality sellers may want to increase the number of dimensions, in attempt to produce a new dimension which they can then manipulatively showcase. Specifically, under Caveat Emptor, if the new dimension is of high quality, the seller will disclose the new dimension and thus pool with sellers that offer high quality on more dimensions. Thus, adding welfare-decreasing dimensions for the purpose of extracting surplus from buyers should itself be considered a type of manipulation. In the Warning regime, we demonstrate the opposite problem: high-quality sellers will be reluctant to add welfare-increasing dimensions. The reason is that, if they end-up with low-quality on the new dimension, they will have to prioritize information about this dimension (rather than about the old, high-quality dimension) and thus pool with sellers that offer low quality on more dimensions, thus losing the ability to distinguish themselves and charge higher prices.

Extending the formal analysis to answer the static question is straightforward and we thus relegate this extension to the Appendix. We next present a formal analysis of the dynamic question.

3.3.2 Sellers' Incentives to Increase the Number of Dimensions

Do Sellers have an incentive to increase the number of dimensions? To study Sellers' incentives, we have to assume that in equilibrium Sellers make non-zero profit. For simplicity, assume that Sellers have all the bargaining power, i.e., they set a price $p = \hat{u}$ and enjoy the full surplus. We start with a one-dimensional product and explore Sellers' incentives to move to a two-dimensional product (as characterized in Section 2).

With one dimension, there are two types of Sellers: H Sellers (50%) and L Sellers (50%). In the 1-dimensional case, we have perfect information, since the buyers can absorb 1-dimensional disclosures. There is no difference between the Caveat Emptor and Warning regimes. Letting u denote the value of High on the single dimension, we see that H Sellers get $u - c$ if $c < u$, and L sellers get zero.

When sellers can add a new dimension, the market is characterized by the following distribution of sellers $(\rho_H, \rho_L, \rho_{HH}, \rho_{HL}, \rho_{LH}, \rho_{LL})$: A share ρ_H and ρ_L of H and L Sellers, respectively, who did not add a new dimension; a share ρ_{HH} of HH Sellers, a share ρ_{HL} of HL Sellers, a share ρ_{LH} of LH Sellers and a share ρ_{LL} of LL Sellers. We assume that Buyer observes the number of dimensions.[29] Therefore, Buyer can identify H Sellers and L Sellers with certainty. A product with one High dimension (H, HL or LH) creates a value of u to Buyers and costs Sellers c to produce. A product with two High dimension (HH) creates a value of $2u$ to Buyers and costs Sellers $2c$ to produce. A product with no High dimension (LL) creates a value of zero to Buyers and costs zero to produce.

How should we think about the move from one dimension to two dimensions? Assume that: (i) an H Seller (in a one-dimensional world) will become either HH or HL (each with probability $\frac{1}{2}$). And (ii) an L Seller (in a one-dimensional world) will become either LH or LL (each with probability $\frac{1}{2}$).[30] From a social perspective, L Sellers and H Sellers should add a dimension, when $c < u$ and when the investment in adding a dimension does not exceed the expected social benefit $\frac{1}{2} \cdot (u - c)$.[31]

Caveat Emptor. Under Caveat Emptor, if Seller remains silent, Buyers will know that they are facing an L Seller or an LL Seller, identifying each with certainty. If a 2-dimensions Seller discloses High, Buyer will know that the probability of facing an HH Seller is $\frac{\rho_{HH}}{\rho_{HH}+\rho_{HL}+\rho_{LH}}$, that the probability of facing an HL Seller is $\frac{\rho_{HL}}{\rho_{HH}+\rho_{HL}+\rho_{LH}}$, and that the probability of facing an LH Seller is $\frac{\rho_{LH}}{\rho_{HH}+\rho_{HL}+\rho_{LH}}$. Let $u_H^{CE} = \frac{\rho_{HH}}{\rho_{HH}+\rho_{HL}+\rho_{LH}} \cdot 2u + \frac{\rho_{HL}+\rho_{LH}}{\rho_{HH}+\rho_{HL}+\rho_{LH}} \cdot u$ denote te expected utility enjoyed by a Buyer who observes High (from a 2-dimensions Seller). The expected profit of HL and LH Sellers is: $\pi_{1H}^{CE} = u_H^{CE} - c$ if $c < u_H^{CE}$. The expected profit of HH Sellers is: $\pi_{HH}^{CE} = u_H^{CE} - 2c$ if $c < \frac{1}{2}u_H^{CE}$.

If an L Seller adds a new dimension, then: With probability $\frac{1}{2}$, he will become LL, disclose Low and get zero; and with probability $\frac{1}{2}$, he will become LH, disclose High and get π_{1H}^{CE}, if $c < u_H^{CE}$. Without adding a new dimension, L Seller gets zero. If $c < u_H^{CE}$, the L Seller will add a dimension if the investment does not exceed $\frac{1}{2}(u_H^{CE} - c)$. Since $u_H^{CE} > u$, the L Seller will have a socially excessive incentive to add a dimension.

---

[29] This is consistent with the assumption in Section 2 that Buyers know the number of dimensions (which was always two in Section 2). The alternative assumption is that Buyer does not observe the number of dimensions. This would allow pooling between 1-dimension Sellers and 2-dimensions Sellers.

[30] We can easily generalize from $\frac{1}{2}$ to a general probability of getting H on the new dimension.

[31] Looking at the Buyers who are matched with an (ex ante) L Seller: When $c < u$, the resulting transactions generate welfare of $\frac{1}{2} \cdot (u - c)$ if a new dimension is added, rather than zero if Seller sticks with one dimension. When $c > u$, the welfare with and without a new dimension is zero. Looking at the Buyers who are matched with an (ex ante) H Seller: When $c < u$, the resulting transactions generate welfare of $\frac{1}{2} \cdot (2u - 2c) + \frac{1}{2} \cdot (u - c)$, rather than $u - c$ without the new dimension. In both cases, when $c < u$, the social benefit from adding a dimension is $\frac{1}{2} \cdot (u - c)$; and when $c > u$, the social benefit is zero. This normative benchmark requires that the welfare-increasing, 2-dimensions transactions (when L Sellers become LH Sellers and H Sellers become HH Sellers) take place.

If an H Seller adds a dimension, then: With probability $\frac{1}{2}$, he will become HH and get π_{HH}^{CE}, if $c < \frac{1}{2}u_H^{CE}$; and with probability $\frac{1}{2}$, he will become HL and get π_{1H}^{CE}, if $c < u_H^{CE}$. Without adding a new dimension, H Seller gets $u - c$, if $c < u$. If $c < \frac{1}{2}u_H^{CE} < u$, the private benefit from adding a dimension is: $\frac{1}{2}\pi_{HH}^{CE} + \frac{1}{2}\pi_{1H}^{CE} - (u - c) = u_H^{CE} - u - \frac{1}{2}c$. Comparing to the social optimum, we find that the H Seller will have either excessive or insufficient incentives to add a dimension, depending on the c values and on the distribution of Sellers in the market.[32] If $c \in \left(\frac{1}{2}u_H^{CE}, u\right)$, the private benefit from adding a dimension is: $\frac{1}{2}\pi_{1H}^{CE} - (u - c) = \frac{1}{2}(u_H^{CE} - c) - (u - c)$. Once again, we find that the H Seller will have either excessive or insufficient incentives to add a dimension, depending on the c values and on the distribution of Sellers in the market.[33] If $c \in (u, u_H^{CE})$, the private benefit from adding a dimension is: $\frac{1}{2}\pi_{1H}^{CE} = \frac{1}{2}(u_H^{CE} - c)$, and the Seller will have an excessive incentive to add a dimension *iff* $c \in (u, u_H^{CE})$.

Warning. Under the Warning regime, if Seller discloses High, Buyer will know that she is facing an H Seller or an HH Seller, identifying each with certainty. An H Seller's expected profit will be: $u - c$; and an HH Seller's expected profit will be: $2u - 2c$. If a 2-dimensions Seller discloses Low, Buyer will know that the probability of facing an LL Seller is $\frac{\rho_{LL}}{\rho_{LL}+\rho_{HL}+\rho_{LH}}$, that the probability of facing an HL Seller is $\frac{\rho_{HL}}{\rho_{LL}+\rho_{HL}+\rho_{LH}}$, and that the probability of facing an LH Seller is $\frac{\rho_{LH}}{\rho_{LL}+\rho_{HL}+\rho_{LH}}$. Let $u_L^W = \frac{\rho_{HL}+\rho_{LH}}{\rho_{LL}+\rho_{HL}+\rho_{LH}} \cdot u$ denote te expected utility enjoyed by a Buyer who observes Low (from a 2-dimensions Seller). The expected profit of HL and LH Sellers is: $\pi_{1L}^W = u_L^W - c$ if $c < u_L^W$. The expected profit of LL Sellers is: $\pi_{LL}^W = u_L^W$.

If an L Seller adds a new dimension, then: With probability $\frac{1}{2}$, he will become LL and get π_{LL}^W; and with probability $\frac{1}{2}$, he will become LH and get π_{1L}^W if $c < u_L^W$. Without adding a new dimension, L Seller gets zero. If $c < u_L^W$, the private benefit from adding a dimension is: $\frac{1}{2}\pi_{LL}^W + \frac{1}{2}\pi_{1L}^W = u_L^W - \frac{1}{2}c$. Comparing to the social optimum, we find that the L Seller will have either excessive or

---

[32] If $c < \frac{1}{2}u_H^{CE} < u$, the Seller will add a dimension if $c < min\left(\frac{1}{2}u_H^{CE}, 2(u_H^{CE} - u)\right)$ and if the investment does not exceed $u_H^{CE} - u - \frac{1}{2}c$. The H Seller will inefficiently fail to add a dimension when $u_H^{CE} < \frac{4}{3}u$ and $c \in \left(2(u_H^{CE} - u), \frac{1}{2}u_H^{CE}\right)$, regardless of the required investment. Otherwise, the Seller will add a dimension if the investment does not exceed $u_H^{CE} - u - \frac{1}{2}c$. Thus, the H Seller will have an insufficient incentive to add a dimension when $u_H^{CE} < \frac{3}{2}u$; and an excessive incentive when $u_H^{CE} > \frac{3}{2}u$.

[33] If $c \in \left(\frac{1}{2}u_H^{CE}, u\right)$, the Seller will add a dimension if $c \in \left(max\left(\frac{1}{2}u_H^{CE}, 2u - u_H^{CE}\right), u\right)$ and if the investment does not exceed $\frac{1}{2}(u_H^{CE} - c) - (u - c)$. The H Seller will inefficiently fail to add a dimension when $u_H^{CE} < \frac{4}{3}u$ and $c \in \left(\frac{1}{2}u_H^{CE}, 2u - u_H^{CE}\right)$, regardless of the required investment. Otherwise, the Seller will add a dimension if the investment does not exceed $\frac{1}{2}(u_H^{CE} - c) - (u - c)$. Thus, the H Seller will have an insufficient incentive to add a dimension when $u_H^{CE} < 3u - 2c$; and an excessive incentive when $u_H^{CE} > 3u - 2c$.

insufficient incentives to add a dimension, depending on the distribution of Sellers in the market.[34] If $c > u_L^W$, the private benefit from adding a dimension is: $\frac{1}{2}\pi_{LL}^W = \frac{1}{2}u_L^W$. Here, any investment in adding a dimension is inefficient, since an L Seller who becomes an LH Seller will exit the market; only an L Seller who becomes an LL Seller will transact and profit from pooling with HL and LL Sellers.

If an H Seller adds a dimension, then: With probability $\frac{1}{2}$, he will become HH and get $2u - 2c$; and with probability $\frac{1}{2}$ he will become HL and get π_{1L}^W if $c < u_L^W$. Without adding a new dimension, H Seller will get $u - c$. If $c < u_L^W$, the private benefit from adding a dimension is: $\frac{1}{2} \cdot (2u - 2c) + \frac{1}{2} \cdot \pi_{1L}^W - (u - c) = \frac{1}{2} \cdot \pi_{1L}^W = \frac{1}{2} \cdot (u_L^W - c)$. Since $u_L^W < u$, the H Seller will have a socially insufficient incentive to add a dimension. If $c \in (u_L^W, u)$, the H Seller will inefficiently fail to add a dimension.[35]

Comparison. Under Caveat Emptor, L Sellers will have socially excessive incentives to add a dimension; and H Sellers will have either excessive or insufficient incentives to add a dimension, depending on the *c* values and on the distribution of Sellers in the market. In the Warning regime, H Sellers will have socially insufficient incentives to add a dimension; and L Sellers will have either excessive or insufficient incentives to add a dimension, depending on the *c* values and on the distribution of Sellers in the market.

4. Concluding Remarks

We conclude by offering a few remarks on the relationship between rationality and manipulation, and on the implications of the model for the design of optimal legal intervention.

4.1 Rationality and Manipulation

In much of the academic literature, manipulation is defined as the exploitation of irrationality. We depart from this conventional definition and identify an alternative mode of manipulation—one based on mislaid priorities. Priority matters when consumers review only a subset of all relevant information about product attributes. Perfectly rational consumers will often choose to remain imperfectly informed in this way. Boundedly rational consumers will also often remain imperfectly informed. The underlying premise of the 'manipulation by mislaid priorities' theory—that consumers review only a subset of all relevant information—allows for both perfect rationality and bounded rationality foundations. But even when we allowed for a bounded rationality foundation, our analysis assumed an arguably high degree of consumer sophistication. We posited that consumers are aware of their limited information, that they know how sellers strategically

[34] Specificlly, the L Seller will have an excessive incentive to add a dimension when $u_L^W > \frac{1}{2}u$; and an insufficient incentive when $u_L^W < \frac{1}{2}u$.

[35] If $c \in (u_L^W, u)$, the private benefit from adding a dimension is: $\frac{1}{2} \cdot (2u - 2c) - (u - c) = 0$.

decide what information to prioritize for disclosure, and that they draw rational inferences about undisclosed, or de-prioritized, dimensions.

Our analytical framework can be extended to allow for irrational, biased inferences, which compound the problem of mislaid priorities. In the Caveat Emptor regime, imperfectly rational buyers place insufficient weight on the undisclosed dimensions and thus overestimate the product's value. This overestimation exacerbates the excessive purchases distortion and attenuates the insufficient purchases distortion. But these countervailing effects do not offset each other. The overestimation disrupts the balance achieved by rational inference. Imperfect rationality increases the harm from manipulation. Accordingly, the need for legal intervention is greater when consumers are imperfectly rational. In markets where buyers are more sophisticated, it is less critical to warn people about low quality dimensions; warnings can be used more sparingly—targeting only the most important low-quality features of the product.

But these conclusions about the relationship between rationality and the proper scope of regulation are incomplete. The aggravation of the distortion that occurs with irrational inferences applies also to mandated low-quality warnings in the Warning regime. Rational buyers will optimally account for the undisclosed, potentially high-quality dimensions, whereas imperfectly rational buyers will place excessive weight on the warning and underestimate the product's value. The cost, to consumers, of "manipulation" by (poorly designed) regulation is lower for rational consumers, in the same manner that the cost of manipulation by sellers under Caveat Emptor was lower for rational consumers. These results push back against the prior conclusion that buyer rationality justifies less regulation. We are thus left with an indeterminate relationship between buyer sophistication and the scope of optimal regulation. The first-order consideration, for regulators, should be the relative importance of high- versus low-quality information—the substitutes versus complements distinction—rather than buyers' sophistication.

4.2 Doctrinal Applications

The analysis in this paper compared the performance of several legal regimes, identifying when warnings are needed, how to design them so as to provide the most useful (rather than the maximal, or most negative) information, and when they can do more harm than good. We now add a few more remarks on specific aspects of legal design that are informed by the theoretical framework.

Harmful Warnings. Disclosures, we showed, can be harmful. This observation is not new.[36] But it received a novel foundation in our analysis of the Full Disclosure and Warning regimes. Moreover, prior accounts state that a new disclosure mandate might obscure other, more important information, but say little about when one piece of information is more important than another. Our analysis precisely identifies the circumstances when a low-quality warning is more or less important than voluntarily disclosed high-quality information.

[36] See Ben-Shahar and Schneider (2014, 169-182).

To illustrate, consider the privacy disclosures typically appearing as prominent upfront warnings upon entry into websites.[37] They are prioritized, requiring attention to this matter as a precondition to accessing the service, thus delaying information about any other quality aspect. For some people privacy is a complement to any other quality dimension, such that the warnings result in superior browsing choices. For other people, privacy is a substitute to other quality dimensions (e.g., the unbiased, comprehensive news coverage that the website provides). For these internet users, the prioritization of the privacy warning is harmful; it might lead them to avoid the service altogether, when the high-quality dimensions more than offset the data privacy downside. Further, prioritizing data-privacy disclosure harms users for whom data privacy is what we called a slack dimension—one that has a small impact on the overall payoff. Forced to pay attention to this dimension, users might not be informed of other, more important, low-quality aspects.

Disclaimers. Another lesson from the analysis concerns the impact of disclaimers—mandated disclosures that accompany and qualify sellers' claims of high quality.[38] But what if buyers cannot attend to both the claims and the disclaimers? Our analysis provides a clue as to which of the two—the claim or the disclaimer—should have priority. When the quality dimensions are substitutes, the disclaimer may (as it often does) take the back seat. But when the dimensions are complements, it is more critical for buyers to know the information in the disclaimer, and it will be effective only if it gains priority over sellers' praiseful claims.

Such prioritization is radical. Should an advertiser be required to make the low-quality aspects *more prominent* than any other information it seeks to promote in the ad? So much advertising and front-of-the-package labeling is loaded with information about quality attributes with positive but small impact, while so much information about critical and negative aspects is relegated to less salient disclaimers. This priority is harmful, we showed, when the attributes are complements. For such products, it is not enough to require (as the law presently does) that the disclaimers come in prominent format. The priority has to be reversed.

It may be impractical, and perhaps even unconstitutional, for the law to require that sellers advertise primarily their product's bad dimensions. But it is entirely possible for advertising law to recognize that non-prioritized disclaimers and warnings have little effect on buyers and should thus have no legal consequences. If sellers decide to deprioritize warnings that are critical to buyers, which could significantly affect buyers' assessment of the product's value, these sellers should be treated as if they did not issue the warning at all. Disclaimers that currently satisfy conspicuousness standards and help sellers avoid liability, but fall short of the more demanding prioritization standard, should be held by the law as ineffective and unfit to provide the liability shield. In such a legal environment, sellers might reprioritize their disclosures.[39]

---

[37] The use of such entry-blocking prominent warnings is mandated by the combined directives of the GDPR (General Data Protection Regulation, Article 13) and the ePrivacy Directive (Directive 2009/136/EC of the European Parliament and of the Council); see also California digital privacy law (California Consumer Privacy Act of 2018).

[38] The Uniform Commercial Code, for example, mandates that warranty disclaimers in the sale of goods be conspicuous (UCC § 2-316(2)), and the FTC requires that advertisement disclaimers be "clear and conspicuous." (Federal Trade Commission, Enforcement Policy Statement on Deceptively Formatted Advertisements, p. 3).

[39] To some extent, advertising law is sensitive to questions of prioritization or relative prominence—of the disclaimer vs. the claim of high quality. Disclaimers are considered "in the context of the ad." See, e.g., Rent A Car System, Inc. v. Hertz Corp., 782 F.2d 381 (2d Cir. 1986); Fink v. Time Warner Cable, 714 F.3d 739, 742 (2d Cir. 2013); Mantikas v. Kellogg Co., 910 F.3d 633, 636 (2d Cir. 2018). Also, when sellers try to distract the consumer when presenting the

4.3 Complex Prices

The analysis in this paper assumed that products have quality dimensions that are numerous and complex, while the product's price is simple and one-dimensional. Indeed, the criteria for determining which disclosure regime is superior, Caveat Emptor versus Warning, depended critically on the product's price. But prices, like other attributes, could be complex, with various elements of the product separately and contingently priced. Loans, medical services, and auto repair are often priced along multiple dimensions. How would the analysis change with such compound complexity?

Sometimes, prices alone are complex, while other attributes of quality are relatively simple. Broadband cable service has one important quality dimension (bandwidth) but complex pricing. A medical or dental provider may have known expertise and quality, but again the pricing for the service is notoriously complex. In these scenarios, the analysis is flipped, and the information regime applies to the disclosure of prices. A price dimension would be regarded as high-quality when the price reflects an excellent bargain; it would be regarded as low-quality when the price includes a significant markup. Price attributes can be complements—when the presence of some high-price dimensions significantly destroys the net value to buyers. And price attributes can be substitutes—when it suffices that some price dimensions are low to secure a significant net value to buyers. Finally, the choice between the different price-disclosure regimes depends, in part, on the quality of the product.

Other times, both prices and quality attributes are complex and surrounded with uncertainty. For some products, price and quality dimensions both have to represent high value for buyers to derive much net surplus. For others, it suffices that either quality or price dimensions are high-valued. Thus, we can adapt the framework of the model to examine the interaction between the two categories – quality and price. It is normally the case that they are complements—low-value on either quality or price destroys much of the net surplus to buyers—and then priority disclosure should be given to the dimension that carries more negative information. But there could be cases in which quality and price are substitutes, where much surplus is gained either from an expensive high-quality product or from a cheap low-quality product. In these cases, sellers should be allowed to prioritize the attractive part of the deal.[40]

disclaimer, the disclaimer will not shield the seller from liability. See, e.g., Kraft, Inc., 114 F.T.C. 40 (1991), aff'd, 970 F.2d 311 (7th Cir. 1992); Federal Trade Commission (2013, 19) (available at https://www.ftc.gov/sites/default/files/attachments/press-releases/ftc-staff-revises-online-advertising-disclosure-guidelines/130312dotcomdisclosures.pdf; see also Andrews (2011, 152). But, while advertising law considers questions of prioritization or relative prominence, at least to some extent, the key distinction between complements and substitutes, which determines the need for prioritized disclaimers, is absent from current doctrine.

[40] Formally, we can adjust the Section 2 framework such that one of the two dimensions is a price dimension (and the other remains a quality dimension). We would then posit an exogenous (reservation) utility threshold, above which the consumer buys the product.

4.4 The Use of "Scores" to Aggregate Information

Complex information about product quality or prices has bred a unique form of simplified disclosure—the "score." When the multiple dimensions can each be quantified along some measurement metric and then aggregated, a one-dimensional score can provide critical information without having to prioritize some dimensions over others. The use of scores to simplify the cost of credit is a staple of Truth-in-Lending laws, which mandate a single APR score that sums up different financing costs associated with a consumer loan. Scores are also used to aggregate quality dimensions, as in the A/B/C rating system for restaurants' hygiene practices.

When scores are able to cut through complexity and effectively aggregate multiple dimensions of price or quality, they work better than the partial information provided by prioritized disclosures. Scores seem particularly suitable when the various dimensions they seek to aggregate are *quantitative* and easy to add up.[41] But such aggregation is not always possible. The value of a score is limited, when it is not clear how to combine and weigh the different components that the score seeks to aggregate. This is the case, for example, when a loan or cable service entail different, contingent fees that are incurred by only some buyers some of the time.

Our analysis helps identify another problem with scores, not previously noted. A critical feature of the model was the recognition that the marginal value of any dimension—quality or price—depends on the value of other dimensions. For example, an expensive feature of a product that could be easily substituted by a less expensive feature should be weighed in the aggregate score differently than a similarly expensive dimension that operates as a necessary complement. The score given to a cable subscription plan that charges a high price for a specific premium channel should depend on the cost of, and substitutability with, other channels.

Thus, while scores are hailed for their ease of *use* by buyers—how they save buyers the need to review cumulative or prioritized information—their *production* can be quite challenging. In creating a useful score, the priority criteria derived from our model must be recognized and applied. When the dimensions aggregated into the score are substitutes, the score should overweigh the high-quality ones; and when the dimensions are complements, added weight should be given to the low-quality ones.

4.5 Market Solutions

Buyers are prone to manipulation because of the information asymmetry in consumer markets. Scores are one possible method for reducing information asymmetry. Other market solutions include trademarks and reputational mechanisms, ratings systems, and quality certification. As with scores, these solutions do not eliminate the information asymmetry problem. Indeed, a large literature has convincingly demonstrated the prevalence of asymmetric information in consumer

[41] Additional examples of price disclosures using scores include mandating disclosure of total payments imposed by a consumer lease (12 C.F.R. § 213.4(c),(e)); disclosure of total life cycle cost of products that use electricity (see Deutsch (2010)); disclosure of closing costs in real-estate transactions (12 U.S.C. § 2604); Google's price rating ('$'–'$$$$') of restaurants. See also (Bar-Gill, 2012).

markets even when such mitigating mechanisms exist.[42] Our goal is not to offer additional proof of information asymmetry. Rather, taking the prevalence of asymmetric information as our starting point, our contribution is to show how suboptimal prioritization of information for disclosure – voluntary or mandated – exacerbates the harm from asymmetric information.

Nevertheless, there are several market solutions that more specifically engage with different elements of our model and thus merit further discussion. First, the model limited the message space, allowing sellers to disclose either High or Low. A broader message space could sometimes solve the information asymmetry problem. For example, under Caveat Emptor, an HH Seller could separate itself by communicating "no-Low" or "nothing to warn about" (rather than disclosing High)—a statement that the HL/LH Sellers cannot make. But this solution becomes less feasible as the number of quality dimensions increases. If there are, say, 10 dimensions, only a seller with 10-High would be able to separate itself. To affect fuller separation, sellers with a large number of High dimensions would need more complex characterizations, which would be costly for buyers to review.

A related solution invokes product warranties. Rather than informing buyers of specific dimensions, sellers could warrant overall quality, offering buyers full redress if any dimension is Low. Indeed, warranties are widely regarded as a solution to the asymmetric information problem (Grossman 1981). Our analysis suggests that the impact of warranties is more limited when products are multi-dimensional because it is more difficult to know what the warranty covers. Indeed, we suggested that a warranty is itself a quality dimension, and often one of the more complicated dimensions. A warranty can itself be multi-dimensional, varying in its scope, duration, and other details.

Another market response to manipulation would originate with buyers, who might adapt to sellers' mislaid priorities by strategically selecting the disclosures they review. Buyers for whom quality dimensions are substitutes, who care more about high-quality attributes, could ignore mandated warnings and focus on sellers' voluntary disclosures. For example, internet users could ignore GDPR-mandated warnings. Conversely, buyers who care more about low-quality attributes could ignore sellers' (voluntary) disclosures and focus on the mandated warnings. For example, consumers could look past sellers' claims on packaged foods and focus instead on the nutrition table. And yet it is not always practical to ignore information. First, it may be necessary to know what the claim is before making sense of the disclaimer. Second, it is often unclear before reading what the nature of the disclosure is and whether it is a mandated warning or a seller's claim, especially in "real" environments with many dimensions. (This reality motivated our assumption that Buyer randomly chooses a dimension for review without knowing if it is High or Low, namely, that Buyer has to invest in reading and understanding – to figure out the type and content of the disclosure.) Third, sellers could defeat buyers' attempts to filter disclosures just by their "look." For example, sellers could dress up their unsolicited mailings to look official and trick buyers into reading them. Or, dietary supplements could be labeled in a manner that makes it hard to distinguish between the mandated and voluntary disclosures. Fourth, it is impractical for consumers to mute TV ads and listen only to the disclaimer.

[42] See *supra* note 4.

Finally, our model allows only for disclosure by a seller. In some cases, disclosure can originate from other sources. For example, if an HL Seller discloses High, then a competitor can advertise the Low aspects of the first seller's product (see, e.g., Gilo and Porat 2010). We note that, in these cases, the outcome would similar to the outcome under our Full Disclosure regime, which is generally inferior to the outcome under the other regimes.

4.6 Tiers of Priority

We have shown that the distinction between complements and substitutes helps to identify the information—good versus bad—that should be prioritized for disclosure. In our baseline model, the two quality dimensions could be either complements or substitutes. With more than two dimensions, some could be substitutes and others complements. A restaurant service offers substitutes (the different items on the menu) as well as complements (food, service quality, ambience, food safety).

Our analytical approach allows us to examine buyers' priority problems in tiers. First, a product's attributes can be divided into subsets or *clusters* of attributes; and, defining each cluster as a "dimension," we can ask whether these clusters are substitutes or complements. If only some clusters can be prioritized for disclosure, the substitutes v. complements classification would determine how to choose among them and which disclosure regime should govern. Second, zooming in, we further examine the attributes within each cluster as "dimensions" and again ask whether they are substitutes or complements. The answer to this question would determine how to prioritize disclosure within the cluster.

Consider a product with two clusters, each containing two attributes that are either H or L. Assume that buyers can review only one of the four attributes. The optimal warning regime would depend in a more subtle way on the substitutability and complementarity of the tiers. In one scenario, the clusters are complements but within each cluster the attribute are substitutes. Here, buyers' interest in evaluating each cluster is to know if it has at least one H attribute. With each cluster thus represented by its best attribute, buyers' interest is to receive a warning of an L-attribute only if one of the clusters contains both L attributes. By contrast, in another scenario the clusters are substitutes but their inner attributes are complements. Now, each cluster is valued primarily by its low quality. Here buyers' interest is to receive a warning only there is at least one L-attribute in each of the clusters. If one of the clusters is HH, the seller should be exempt from warning and be allowed under Caveat Emptor to highlight one of the H-attributes of this cluster.

A related, tier-like question arises when buyers' preferences reflect a hybrid, complements-substitutes characterization. Specifically, the most important thing for buyers to know is whether the quality of any dimension falls below a minimum threshold – a complements-like reason for prioritizing Low disclosures. Then, if quality on all dimensions meets the minimum threshold, buyers want to learn about the highest-quality dimension – a substitutes-like reason for prioritizing High disclosures. When buyers have such hybrid preferences, the optimal regime would require a warning when an attribute falls below the minimum quality threshold and otherwise give sellers free rein to disclose high quality attributes.

4.7 Buyer Heterogeneity

The model assumed that buyers are homogeneous. But, of course, consumer markets are populated with heterogeneous buyers who vary in how they value products and in their capacity to review information about products. To some, the various dimensions are complements, to others – substitutes. Some buyers will rationally decide to review information about fewer attributes, while others will examine more attributes and resolve more of the uncertainty. And when the dimensions are asymmetric (as in Section 3.1), buyers differ on which dimension matters more. Short of personalized disclosures, how would the information priority principle apply in such an environment?

This is a standard problem in protective regulation, but it is sharpened in the present context because a warning that helps some buyers is not only useless to others; it may in fact be harmful. To minimize this harm, policymakers can take one of two approaches: They could either segment the market and regulate separately for the different buyer types; or they could identify the subgroup of buyers most in need of protection and follow a one-size disclosure regime tailored to the interests of this group.

Thus, if a product is HL—with one high dimension and one low dimension—a segmentation approach would address the low-quality warning only to buyers for whom the dimensions are complements. If such selective disclosure is impossible, the regulatory choice could take one of several approaches. First, policymakers could follow a majoritarian approach and choose the legal regime that protects the larger group. This could apply both to the choice between Caveat Emptor and Warning, and to the choice about which dimension to prioritize within a Warning regime. Second, policymakers could follow a minoritarian approach, if it is clear that the minority group has more intense preferences, stands to lose more, or is otherwise more in need of protection. Third, policymakers could follow the approach we developed elsewhere (Bar-Gill and Ben-Shahar 2016, 2021), which accounts also for the differences in information cost. Information could be prioritized based on the need of buyers that typically acquire less information. Other buyers can make up for the less-than-ideal disclosure by exerting more effort to review other dimensions.[43]

4.8 Implementing the Principle of Information Priority

Our analysis suggests that the optimal anti-manipulation regime varies by market and by product, based on two primary factors. First, is it more important for buyers to know about bad or about good quality dimensions? This factor was captured by the substitutes versus complements distinction. Second, does the price of the product leave buyers with significant surplus? When prices are in the low range, we saw, buyers are more likely to value products even if they have fewer high-quality dimensions, and Caveat Emptor is better at exposing these high-quality dimensions. And when prices are high, buyers' priority is to know about low-quality dimensions, and here a Warning regime is better.

[43] Policymakers should prioritize information based on exogenous information acquisition costs, not on buyers' endogenous decisions whether to acquire information—decisions that could be made strategically to influence the disclosure regime.

Because products differ along these factors, sector-specific regulatory expertise would be needed. A general consumer protection agency in charge of anti-manipulation regulation in all markets is less likely to identify the critical trade-offs between product dimensions compared to specialized agencies familiar with specific markets and products. Unlike the FTC, which regulates many different sectors, agencies like the FDA (for food and dietary supplements), CFPB (for financial products), or state insurance commissioners (for insurance policies) would be better positioned to identify the proper priorities for the disclosure rules in their domain and issue specific commands product-by-product.

Another interesting aspect of regulatory design is the competence of courts in implementing the information priority principle. Because so much manipulation occurs in advertising, and so much of the law of advertising is managed through Lanham Act and consumer-initiated litigation, rather than through ex ante mandates, a key question is whether courts can identify the factors necessary for optimal disclosure policy. On one hand, courts are already accustomed to examining how partial information in ads and labeling creates misleading impressions. They conduct fact-intensive inquiries that can identify buyers' priorities. On the other hand, courts are often asked to address whether a single component of the transaction had to be highlighted, without regard to how it might crowd out other, more important, information, which is not part of the litigation. A court might thus require that a prominent disclaimer qualify a misleading claim, but both the claim and the disclaimer should ideally yield in priority to other issues. The tendency of litigation to scrutinize one issue at a time, and to build sequential mandates issue-by-issue, could defeat the design of a priority regime.

References

Akerlof, George A. 1970. "The Market for "Lemons": Quality Uncertainty and the Market Mechanism." *The Quarterly Journal of Economics* 84 (3): 488-500.

Akerlof, George A., and Robert J. Shiller. 2015. *Phishing for Phools: The Economics of Manipulation and Deception.* Princeton University Press.

Alba, Joseph W., and J. Wesley Hutchinson. 2000. "Knowledge Calibration: What Consumers Know and What They Think They Know." *Journal of Consumer Research* 27(2): 123–156.

Andrews, J. Craig. 2011. "Warnings and Disclosures." Chap. 15 in *Communicating Risks and Benefits: An Evidence-Based User's Guide*, edited by Baruch Fischhoff, Noel T. Brewer and Julie S. Downs, 149-162. Food and Drug Administration (FDA).

Bakos, Yannis, Florencia Marotta-Wurgler, and David R. Trossen. 2014. "Does Anyone Read the Fine Print? Consumer Attention to Standard-Form Contracts." *Journal of Legal Studies* 43(1): 1-35.

Bar-Gill, Oren. 2012. *Seduction by Contract: Law, Economics and Psychology in Consumer Markets.* Oxford University Press.

Bar-Gill, Oren, and Omri Ben-Shahar. 2021. "Rethinking *Nudge*: An Information-Costs Theory of Default Rules." *University of Chicago Law Review* 88: 531-604.

Bar-Gill, Oren, and Omri Ben-Shahar. 2016. "Optimal Defaults in Consumer Markets." *Journal of Legal Studies* S2: S137-S161.

Barigozzi, Francesca, and Albert Ma. 2018. "Product differentiation with multiple qualities." *International Journal of Industrial Organization* 61: 380-412.

Beales, Howard, Richard Craswell, and Steven C. Salop. 1981. "The Efficient Regulation of Consumer Information." *The Journal of Law & Economics* 24: 491-539.

Ben-Shahar, Omri, and Carl E. Schneider. 2014. *More than you wanted to know: The Failure of Mandated Disclosure.* Princeton University Press.

Benjamin, Daniel J. 2019. "Errors in Probabilistic Reasoning and Judgmental Biases." In Handbook of Behavioral Economics: Applications and Foundations 1, Vol. 2, pp. 69–186 (Elsevier).

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2013. "Salience and Consumer Choice." *Journal of Political Economy* 121 (5): 803-843.

Calo, Ryan. 2014. "Digital Market Manipulation." *George Washington Law Review* 82 (4): 995-1051.

Coons, Christian, and Michael Weber, eds. 2014. *Manipulation: Theory and Practice.* New York: Oxford University Press.

Daughety, Andrew. F., and Jennifer F. Reinganum. 2008. "Communicating Quality: A Unified Model of Disclosure and Signalling." *RAND Journal of Economics* 39(4): 973-989.

Deutsch, Matthias. 2010. "Life Cycle Cost Disclosure, Consumer Behavior, and Business Implications." *Journal of Industrial Ecology* 14 (1): 103-120.

Dranove, David, and Ginger Zhe Jin. 2010. "Quality Disclosure and Certification: Theory and Practice." *Journal of Economic Literature* 48 (4): 935-63.

Enke, Benjamin. 2020. "What You See Is ALL There Is." *Quarterly Journal of Economics* 135(3): 1363-1398.

Eppler, Martin, and Jeanne Mengis. 2004. "The Concept of Information Overload: A Review of Literature from Organization Science, Accounting, Marketing, MIS, and Related Disciplines." *The Information Society* 20: 325-344.

Federal Trade Commission. 2013. ".com Disclosures: How to Make Effective Disclosures in Digital Advertising."

Gabaix, Xavier, and David Laibson. 2006. "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets." *The Quarterly Journal of Economics* 121 (2): 505-540.

Gilo, David, and Ariel Porat. 2010. "Viewing Unconscionability Through a Market Lense." *William & Mary Law Review* 52(1): 133-195.

Gorin, Moti. 2014. "Towards a Theory of Interpersonal Manipulation." In *Manipulation: Theory and Practice*, edited by Christian Coons and Michael Weber, 73-97. New York: Oxford University Press.

Grossman, Sanford J. 1981. "The Informational Role of Warranties and Private Disclosure about Product Quality." *Journal of Law and Economics* 24(3): 461-483.

Hanson, Jon D., and Douglas A. Kysar. 1999. "Taking Behavioralism Seriously: Some Evidence of Market Manipulation." *Harvard Law Review* 112 (7): 1420-1572.

Häubl, Gerald, and Valerie Trifts. 2000. "Consumer Decision Making in Online Shopping Environments: The Effects of Interactive Decision Aids." *Marketing Science* 19 (1): 4-21.

Heidhues, Paul, Botond Kőszegi, and Takeshi Murooka. 2016. "Exploitative Innovation." *American Economic Journal: Microeconomics* 8 (1): 1-23.

Hotz, V. Joseph, and Mo Xiao. 2013. "Strategic Information Disclosure: The Case of Multi-Attribute Products with Heterogeneous Consumers." *Economic Inquiry* 51 (1): 865-881.

Jacoby, Jacob. 1984. "Perspectives on Information Overload." *Journal of Consumer Research* 10 (4): 432-435.

Jehiel, Philippe. 2018. "Investment Strategy and Selection Bias: An Equilibrium Perspective on Overoptimism." *American Economic Review* 108: 1582–1597.

Lacko, James M., and Janis K. Pappalrdo. 2010. "The Failure and Promise of Mandated Consumer Mortgage Disclosures: Evidence from Qualitative Interviews and a Controlled Experiment with Mortgage Borrowers." *American Economic Review: Papers & Proceedings* 100: 516-521.

Mocan, Naci. 2007. "Can Consumers Detect Lemons? An Empirical Analysis of Information Asymmetry in the Market for Childcare." *Journal of Population Economics* 20: 743-780.

Milgrom, Paul. 2008. "What the Seller Won't Tell You: Persuasion and Disclosure in Markets." *Journal of Economic Perspectives* 22 (2): 115-131.

Milgrom, Paul. 1981. "Good News and Bad News: Representation Theorems and Applications." *Bell Journal of Economics* 12(2): 380-391.

Milosavljevic, Milica, Vidhya Navalpakkam, Christof Koch, and Antonio Rangel. 2012. "Relative visual saliency differences induce sizable bias in consumer choice." *Journal of Consumer Psychology* 22: 67-74.

Osborne, Martin J., and Ariel Rubinstein. 1998. "Games with Procedurally Rational Players." *American Economic Review* 88 (4): 834-847.

Persson, Petra. 2018. "Attention Manipulation and Information Overload." *Behavioural Public Policy* 2 (1): 78-106.

Shapiro, Carl. 1982. "Consumer Information, Product Quality, and Seller Reputation." *Bell Journal of Economics* 13 (1): 20-35.

Simon, Herbert A. 1956. "Rational Choice and the Structure of the Environment." *Psychological Review* 63 (2): 129-138.

Spiegler, Ran. 2016. "Choice Complexity and Market Competition." *Annual Review of Economics* 8: 1-25.

Spiegler, Ran. 2006. "Competition Over Agents with Boundedly Rational Expectations." *Theoretical Economics* 1: 207-231.

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

Sun, Monic. 2011. "Disclosing Multiple Product Attributes." *Journal of Economics & Management Strategy* 20(1): 195-224.

Sunstein, Cass R. 2016. "Fifty Shades of Manipulation." *Journal of Marketing Behavior* 1 (3-4): 213-244.

Taylor, Shelley, and Suzanne Thompson. 1982. "Stalking the Elusive Vividness Effect. *Psychological Review* 89(2): 155–81.

Vandenbosch, Mark B., and Charles B. Weinberg. 1995. "Product and Price Competition in a Two-Dimensional Vertical Differentiation Model." *Marketing Science* 14 (2): 224-249.

Wertenbroch, Klaus, ed. 2016. *Journal of Marketing Behavior* 1 (3-4): 209-384.

Wilkinson, Tess M. 2013. "Nudging and Manipulation." *Political Studies* 61: 341-355.

Appendix

The Appendix contains: (i) the proof of the result from Sec. 2.6 that Full Disclosure is always dominated by either the Caveat Emptor or Warning regime; (ii) Table A1, referenced in Sec. 2, which collects the consumer surplus values under CE, W, FD and the perfect-information benchmark; (iii) a formal analysis of the quality-dependent costs extension, referenced in Sec. 2; (iv) a formal analysis of the static question from Section 3.3.

(i) Full Disclosure is Never Globally Optimal

We now show that if the law is constrained to choose one regime for all combinations of (*c, x*), there is no distribution of parameters for which FD dominates both CE and W.

To see this, assume that FD generates a larger surplus than CE:

$$\int_{0}^{\frac{1}{2}x}\left[-\frac{1}{4}c\right]f(c)+\int_{\frac{1}{2}x}^{\frac{2}{3}x}\left[-\frac{1}{4}(x-c)\right]f(c)+\int_{\frac{2}{3}x}^{\frac{2}{3}x+\frac{1}{3}}\left[-\frac{1}{4}(x-c)\right]f(c)+\int_{\frac{2}{3}x+\frac{1}{3}}^{\frac{1}{2}(1+x)}\left[\frac{1}{4}(1-c)+\frac{1}{4}(x-c)\right]f(c)>0$$

This implies that:

$$\int_{0}^{\frac{1}{2}x}\left[-\frac{1}{4}c\right]f(c)+\int_{\frac{1}{2}x}^{\frac{2}{3}x}\left[-\frac{1}{4}(x-c)\right]f(c)+\int_{\frac{2}{3}x+\frac{1}{3}}^{\frac{1}{2}(1+x)}\left[\frac{1}{4}(1-c)+\frac{1}{4}(x-c)\right]f(c)>\int_{\frac{2}{3}x}^{\frac{2}{3}x+\frac{1}{3}}\frac{1}{4}[(x-c)]f(c)$$

which we call the "FD-CE Inequality".

Now consider the difference between the surplus under FD and the surplus under W:

$$\int_{\frac{1}{2}x}^{\frac{2}{3}x}\left[\frac{1}{4}c-\frac{1}{4}(x-c)\right]f(c)+\int_{\frac{2}{3}x}^{\frac{2}{3}x+\frac{1}{3}}\left[\frac{1}{4}(x-c)\right]f(c)+\int_{\frac{2}{3}x+\frac{1}{3}}^{\frac{1}{2}(1+x)}\left[\frac{1}{4}(x-c)\right]f(c)+\int_{\frac{1}{2}(1+x)}^{1}\left[-\frac{1}{4}(1-c)\right]f(c)$$

Substituting the FD-CE Inequality, we know that this difference is smaller than:

$$\int_{0}^{\frac{1}{2}x}\left[-\frac{1}{4}c\right]f(c)+\int_{\frac{1}{2}x}^{\frac{2}{3}x}\left[\frac{1}{4}c-\frac{1}{2}(x-c)\right]f(c)+\int_{\frac{2}{3}x+\frac{1}{3}}^{\frac{1}{2}(1+x)}\left[\frac{1}{4}(1-c)+\frac{1}{2}(x-c)\right]f(c)+\int_{\frac{1}{2}(1+x)}^{1}\left[-\frac{1}{4}(1-c)\right]f(c)$$

Since the expressions in each of the four integrals is negative, we know that, for any distribution $f(c)$, if FD generates a larger surplus than CE, then it necessarily generates a smaller surplus than W.

(ii) Table A1: Consumer Surplus Values

	Consumer Surplus			
	Perfect Information	Caveat Emptor	Full Disclosure	Warning
$c < \frac{1}{2}x$	$\frac{1}{4}(1-c)+\frac{1}{2}(x-c)$	$\frac{1}{4}(1-c)+\frac{1}{2}(x-c)$	$\frac{1}{4}(1-c)+\frac{1}{2}(x-c)-\frac{1}{4}c$	$\frac{1}{4}(1-c)+\frac{1}{2}(x-c)-\frac{1}{4}c$
$c \in \left[\frac{1}{2}x, \frac{2}{3}x\right)$	$\frac{1}{4}(1-c)+\frac{1}{2}(x-c)$	$\frac{1}{4}(1-c)+\frac{1}{2}(x-c)$	$\frac{1}{4}(1-c)+\frac{1}{4}(x-c)$	$\frac{1}{4}(1-c)+\frac{1}{2}(x-c)-\frac{1}{4}c$
$c \in \left[\frac{2}{3}x, x\right)$	$\frac{1}{4}(1-c)+\frac{1}{2}(x-c)$	$\frac{1}{4}(1-c)+\frac{1}{2}(x-c)$	$\frac{1}{4}(1-c)+\frac{1}{4}(x-c)$	$\frac{1}{4}(1-c)$
$c \in \left[x, \frac{2}{3}x+\frac{1}{3}\right)$	$\frac{1}{4}(1-c)$	$\frac{1}{4}(1-c)+\frac{1}{2}(x-c)$	$\frac{1}{4}(1-c)+\frac{1}{4}(x-c)$	$\frac{1}{4}(1-c)$
$c \in \left[\frac{2}{3}x+\frac{1}{3}, \frac{1}{2}\cdot(1+x)\right)$	$\frac{1}{4}(1-c)$	0	$\frac{1}{4}(1-c)+\frac{1}{4}(x-c)$	$\frac{1}{4}(1-c)$
$c \in \left[\frac{1}{2}\cdot(1+x), 1\right)$	$\frac{1}{4}(1-c)$	0	0	$\frac{1}{4}(1-c)$
$c \geq 1$	0	0	0	0

(iii) Quality-Dependent Costs

The analysis, from a consumer surplus perspective, is identical to the analysis in the basic model, after replacing the parameter c from the basic model with the parameter p which is above-cost in the extension.

The analysis, from a social welfare perspective, requires some adjustment. Table A2 collects the social welfare values under CE, W, FD and the perfect-information benchmark. The information is presented graphically in Figure A1. To summarize:

- When $p < \frac{1}{2}x$: All three regimes induce the perfect-information outcome.
- When $p \in \left[\frac{1}{2}x, \frac{2}{3}x\right)$: CE and Warning induce the perfect-information outcome. FD is less efficient.
- When $p \in \left[\frac{2}{3}x, x\right)$: CE induces the perfect-information outcome. FD is less efficient. And Warning is the least efficient.
- When $p \in \left[x, \frac{2}{3}x + \frac{1}{3}\right)$: Warning induces the perfect-information outcome. If $x < c$, then CE and FD are worse than the perfect information outcome, where FD dominates CE. If $x > c$, then CE and FD are *better* than the perfect information outcome, where CE dominates FD.
- When $p \in \left[\frac{2}{3}x + \frac{1}{3}, \frac{1}{2}x + \frac{1}{2}\right)$: Warning induces the perfect-information outcome. If $x < c$, then CE and FD are worse than the perfect information outcome, where FD dominates CE. If $x > c$, then FD is *better* than the perfect-information outcome, while CE is worse than the perfect-information benchmark.
- $p \in \left[\frac{1}{2}x + \frac{1}{2}, 1\right)$: Warning induces the perfect-information outcome. CE and FD are (equally) less efficient.

	Social Welfare			
	Perfect Information	Caveat Emptor	Full Disclosure	Warning
$p < \frac{1}{2}x$	$\frac{1}{4}(1-2c)+\frac{1}{2}(x-c)$	$\frac{1}{4}(1-2c)+\frac{1}{2}(x-c)$	$\frac{1}{4}(1-2c)+\frac{1}{2}(x-c)$	$\frac{1}{4}(1-2c)+\frac{1}{2}(x-c)$
$p \in \left[\frac{1}{2}x, \frac{2}{3}x\right)$	$\frac{1}{4}(1-2c)+\frac{1}{2}(x-c)$	$\frac{1}{4}(1-2c)+\frac{1}{2}(x-c)$	$\frac{1}{4}(1-2c)+\frac{1}{4}(x-c)$	$\frac{1}{4}(1-2c)+\frac{1}{2}(x-c)$
$p \in \left[\frac{2}{3}x, x\right)$	$\frac{1}{4}(1-2c)+\frac{1}{2}(x-c)$	$\frac{1}{4}(1-2c)+\frac{1}{2}(x-c)$	$\frac{1}{4}(1-2c)+\frac{1}{4}(x-c)$	$\frac{1}{4}(1-2c)$
$p \in \left[x, \frac{2}{3}x+\frac{1}{3}\right)$	$\frac{1}{4}(1-2c)$	$\frac{1}{4}(1-2c)+\frac{1}{2}(x-c)$	$\frac{1}{4}(1-2c)+\frac{1}{4}(x-c)$	$\frac{1}{4}(1-2c)$
$p \in \left[\frac{2}{3}x+\frac{1}{3}, \frac{1}{2}\cdot(1+x)\right)$	$\frac{1}{4}(1-2c)$	0	$\frac{1}{4}(1-2c)+\frac{1}{4}(x-c)$	$\frac{1}{4}(1-2c)$
$p \in \left[\frac{1}{2}\cdot(1+x), 1\right)$	$\frac{1}{4}(1-2c)$	0	0	$\frac{1}{4}(1-2c)$
$p \geq 1$	0	0	0	0

Table A2: Social Welfare Values

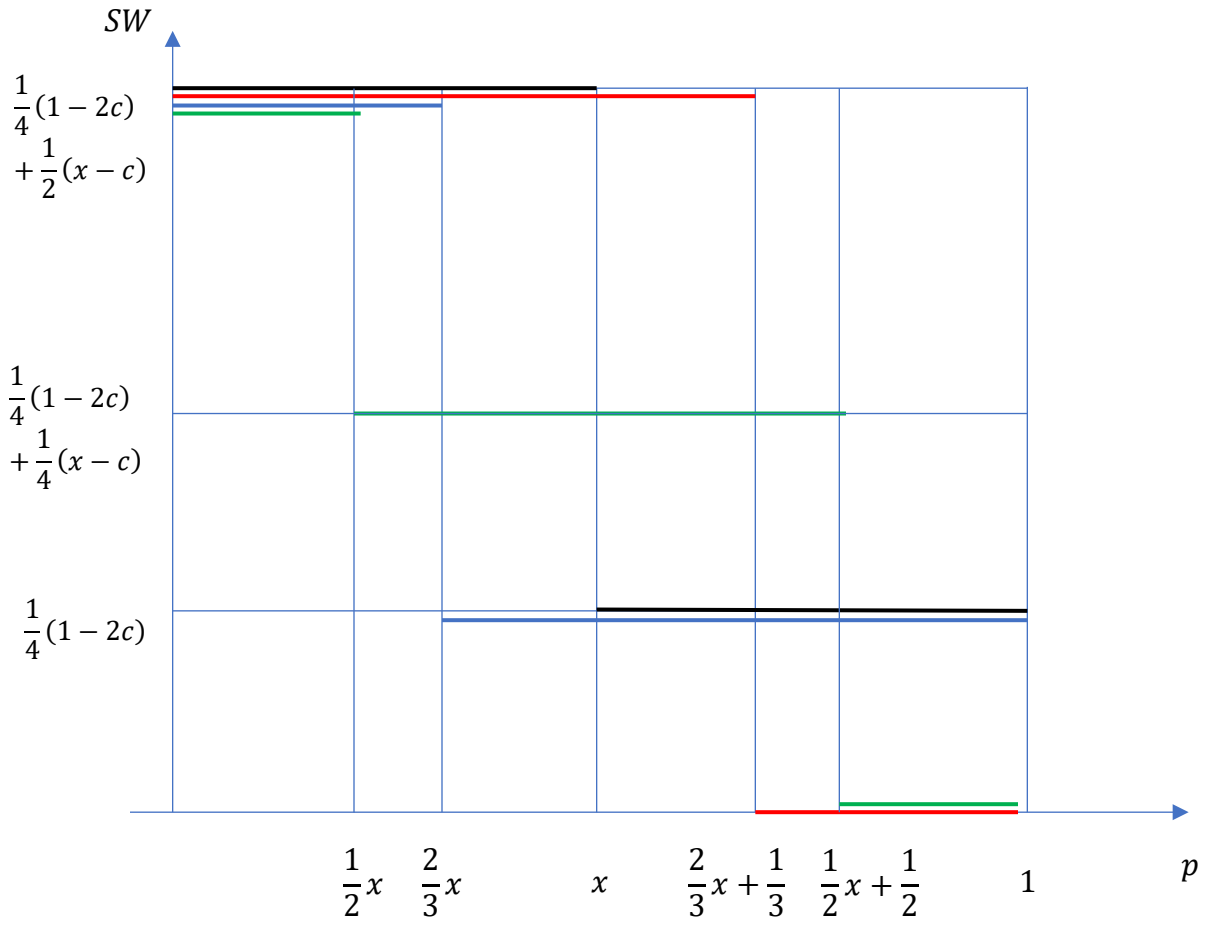
SW
$\frac{1}{4}(1-2c)$
$+\frac{1}{2}(x-c)$
$\frac{1}{4}(1-2c)$
$+\frac{1}{4}(x-c)$
$\frac{1}{4}(1-2c)$
$\frac{1}{2}x$
$\frac{2}{3}x$
x
$\frac{2}{3}x+\frac{1}{3}$
$\frac{1}{2}x+\frac{1}{2}$
1
p

Figure A1

(iv) The Section 3.3 Extension: The Static Question

Consider a 3-dimensions extension of the baseline, Section 2 model. A comprehensive analysis of the 3-dimensions model would distinguish between the utility from a single High dimension (x) and the utility from two High dimensions ($2x$) (in between zero-utility when all dimensions are L and the utility of 1 when all dimensions are H); $x \in \left[0, \frac{1}{2}\right]$. We continue to assume that buyers can review only one dimension. We assume that the quality levels are independently distributed among Sellers, such that for each dimension 50% of sellers are H and 50% are L. Accordingly, there are eight equally likely types of sellers—HHH, HHL, HLH, HLL, LHH, LHL, LLH, LLL—each comprising 12.5% of sellers.

Under Caveat Emptor, Seller will disclose a High dimension, if such a dimension exists (and if there are multiple High dimensions, Seller will disclose one randomly); and Seller will remain silent if all dimensions are Low. A rational Buyer who observes High on will know that the product has three High dimensions with probability $\frac{1}{7}$, that the product has two High dimensions with probability $\frac{3}{7}$, and that the product has one High dimensions with probability $\frac{3}{7}$.[44] A Buyer who observes a silent Seller, will know that the product is LLL. In a Warning regime, Seller will disclosure a Low dimension, if such a dimension exists (and if there are multiple Low dimensions, Seller will disclose one randomly); and Seller will disclose High, if all dimensions are High. A rational Buyer who observes Low will know that the product has three Low dimensions with probability $\frac{1}{7}$, that the product has two Low dimensions with probability $\frac{3}{7}$, and that the product has one Low dimensions with probability $\frac{3}{7}$. A Buyer who observes High, will know that the product is HHH.

Perfect Information. Buyer will buy only from an HHH Seller, if $c < 1$, for $CS = 1 - c$ (if $c \geq 1$, Buyer will not buy and $CS = 0$). Buyer will buy from a Seller with two High dimensions, if $c < 2x$, for $CS = 2x - c$ (if $c \geq 2x$, Buyer will not buy and $CS = 0$). Buyer will buy from a Seller with one High dimension, if $c < x$, for $CS = x - c$ (if $c \geq x$, Buyer will not buy and $CS = 0$). The overall consumer surplus is: $CS = \frac{1}{8} \cdot (1 - c) + \frac{3}{8} \cdot (2x - c) + \frac{3}{8} \cdot (x - c)$, if $c < x$; $CS = \frac{1}{8} \cdot (1 - c) + \frac{3}{8} \cdot (2x - c)$, if $c \in [x, 2x)$; $CS = \frac{1}{8} \cdot (1 - c)$, if $c \in [2x, 1)$; and $CS = 0$ if $c \geq 1$.

Caveat Emptor. With probability $\frac{1}{8} \cdot \left(\frac{1}{3} + \frac{1}{2} + \frac{1}{2} + 1\right) = \frac{7}{24}$, Seller discloses High on D1; and with the same probability Seller discloses High on D2 or D3. With probability $\frac{1}{8}$, Seller remains silent. Buyer who observes High gets an expected utility of $\hat{u} = \frac{1}{7} \cdot 1 + \frac{3}{7} \cdot 2x + \frac{3}{7} \cdot x = \frac{1}{7} \cdot (1 + 9x)$, and

[44] Applying Bayes' rule, we have:

$$Pr(q_1 = q_2 = q_3 = 1 | \text{B observes } q_1 = 1) = \frac{Pr(q_1 = q_2 = q_3 = 1) \cdot Pr(\text{B observes } q_1 = 1 | q_1 = q_2 = q_3 = 1)}{Pr(\text{B observes } q_1 = 1)}$$

$$= \frac{\frac{1}{8} \cdot \frac{1}{3}}{\frac{1}{8}(\frac{1}{3} + \frac{1}{2} + \frac{1}{2} + 1)} = \frac{1}{7}$$

the consumer surplus is: $CS = \frac{1}{7} \cdot (1 + 9x) - c$ if $c < \frac{1}{7} \cdot (1 + 9x)$ and $CS = 0$ if $c \geq \frac{1}{7} \cdot (1 + 9x)$; and a Buyer who faces a silent Seller gets an expected utility of $\hat{u} = 0$, and the consumer surplus is: $CS = 0$. The overall consumer surplus is: $CS = \frac{7}{24} \cdot 3 \cdot \left(\frac{1}{7} \cdot 1 + \frac{3}{7} \cdot 2x + \frac{3}{7} \cdot x - c\right) = \frac{1}{8} \cdot (1 - c) + \frac{3}{8} \cdot (2x - c) + \frac{3}{8} \cdot (x - c)$, if $c < \frac{1}{7} \cdot (1 + 9x)$; and $CS = 0$, if $c \geq \frac{1}{7} \cdot (1 + 9x)$.

We now compare the Caveat Emptor outcome to the perfect information benchmark. There are two cases: The first case is when $x < \frac{1}{5}$, which implies $x < 2x < \frac{1}{7} \cdot (1 + 9x)$. In this case, if $c < x$, then Caveat Emptor achieves the perfect information outcome. If $c \in [x, 2x)$, then Caveat Emptor induces inefficient purchases from Sellers with only one High dimension. If $c \in [2x, \frac{1}{7} \cdot (1 + 9x))$, then Caveat Emptor induces inefficient purchases from Sellers with either one or two High dimensions. If $c \in \left[\frac{1}{7} \cdot (1 + 9x), 1\right)$, then under Caveat Emptor Buyers inefficiently fail to purchase from HHH Sellers. If $c \geq 1$, then there are no purchases, as in the perfect information benchmark. The second case is when $x > \frac{1}{5}$, which implies $x < \frac{1}{7} \cdot (1 + 9x) < 2x$. In this case, if $c < x$, then Caveat Emptor achieves the perfect information outcome. If $c \in [x, \frac{1}{7} \cdot (1 + 9x))$, then Caveat Emptor induces inefficient purchases from Sellers with only one High dimension. If $c \in \left[\frac{1}{7} \cdot (1 + 9x), 2x\right)$, then under Caveat Emptor Buyers inefficiently fail to purchase from Sellers with two High dimensions and from HHH Sellers. If $c \in [2x, 1)$, then under Caveat Emptor Buyers inefficiently fail to purchase from HHH Sellers. If $c \geq 1$, then there are no purchases, as in the perfect information benchmark.

Warning. With probability $\frac{1}{8} \cdot \left(\frac{1}{3} + \frac{1}{2} + \frac{1}{2} + 1\right) = \frac{7}{24}$, Seller discloses Low on D1; and with the same probability Seller discloses Low on D2 or D3. With probability $\frac{1}{8}$, Seller discloses High. Buyer who observes Low gets an expected utility of $\hat{u} = \frac{1}{7} \cdot 0 + \frac{3}{7} \cdot x + \frac{3}{7} \cdot 2x = \frac{1}{7} \cdot 9x$, and the consumer surplus is: $CS = \frac{1}{7} \cdot 9x - c$ if $c < \frac{1}{7} \cdot 9x$ and $CS = 0$ if $c \geq \frac{1}{7} \cdot 9x$; and a Buyer who observes High gets an expected utility of $\hat{u} = 1$, and the consumer surplus is: $CS = 1 - c$ if $c < 1$ and $CS = 0$ if $c \geq 1$. The overall consumer surplus is: $CS = \frac{7}{24} \cdot 3 \cdot \left(\frac{1}{7} \cdot 0 + \frac{3}{7} \cdot x + \frac{3}{7} \cdot 2x - c\right) = \frac{1}{8} \cdot (1 - c) + \frac{3}{8} \cdot (2x - c) + \frac{3}{8} \cdot (x - c) + \frac{1}{8} \cdot (-c)$, if $c < \frac{1}{7} \cdot 9x$; $CS = \frac{1}{8} \cdot (1 - c)$, if $c \in \left[\frac{1}{7} \cdot 9x, 1\right)$; and $CS = 0$, if $c \geq 1$.

We now compare the Warning outcome to the perfect information benchmark. If $c < x$, then Warning induces inefficient purchases from LLL Sellers. If $c \in [x, \frac{1}{7} \cdot 9x)$, then Warning induces inefficient purchases from LLL Sellers and from Sellers with only one High dimension. If $c \in \left[\frac{1}{7} \cdot 9x, 2x\right)$, then under Warning Buyers inefficiently fail to purchase from Sellers with two High dimensions. If $c \geq 2x$, then Warning achieves the perfect information outcome.

Comparison. As in the baseline model, Warning is better in the complements case, i.e., when the cost, and price, are high; and Caveat Emptor is better in the substitutes case, i.e., when the cost,

and price, are low. But here, with three dimensions, the harm inflicted upon Buyers is greater. With two dimensions, Caveat Emptor achieves the perfect information outcome when cost, and price, are below the value of a product with one High dimension *out of two*. With three dimensions, Caveat Emptor achieves the perfect information outcome when cost, and price, are below the value of a product with one High dimension *out of three*, which will usually be a narrower range. Symmetrically, with two dimensions, Warning achieves the perfect information outcome when cost, and price, are above the value of a product with *one* High dimension *out of two*. With three dimensions, Warning achieves the perfect information outcome when cost, and price, are above the value of a product with *two* High dimensions *out of three*, which will usually be a narrower range.