Draft: April 2021

Manipulation By Mislaid Priorities

Oren Bar-Gill and Omri Ben-Shahar*

Abstract

Manipulation of consumers is widely thought to exploit people's limited rationality. This paper lays a foundation for a different theory of manipulation which relies solely on how truthful information is prioritized. It identifies a harmful manipulation tactic that is effective even when its intended audience is rational and aware of its occurrence. The paper shows that the distortions due to misprioritized information can arise not only from firms' boastful disclosures, but also, equally and surprisingly, 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.

Note to Workshop Participants: The main analysis is conducted using a formal model. But we try to make the main ideas accessible to a broad audience by offering a detailed informal summary before each segment of the technical analysis.

* 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, Holger Spamann, Kathy Spier, Rebecca Tushnet, Mark Tushnet, and workshop participants at Harvard, IDC Herzliya, Tel-Aviv 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.[2] Under these accounts, firms exploit people's mistakes and misjudgments, or inflame them, primarily by manipulating the information people receive or the format and context in which it is presented.

In this paper, we identify and study a general source of manipulation that does not fall within the traditional characterization and does not involve an exploitation of consumer irrationality. 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 disclosed information in a way that conflicts with its receipients' interests. It is an effective and harmful tactic, even when the receipients are rational and aware of its occurrence.

Consider the following example. A food product has quality dimensions that customers would regard as "good" and others that are "bad." For example, it could contain two types of flour—"whole grain" (good) and enriched white flour (bad). When designing the product's packaging, the seller could refer to both ingredients equally, for example in a list that resembles the familiar Nutrition Data disclosure. Alternatively, the seller may highlight only the good ingredient and either remain silent about the bad ingredient or present it in a manner that is more costly for buyers to review.[3] If buyers have unlimited capacity to review the 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 when it is costly for buyers to review all dimensions, and where they rationally conduct partial and selective review, the seller's priority would determine which information they get.

It is tempting to think that the only way such manipulative presentation of information could succeed is when targeting irrational customers. By making the good dimensions *salient*, the seller lulls people into focusing only on these prioritized aspects, causing them to overvalue the product. People are tricked to overlook the other, bad dimensions, oblivious to the manipulative

[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).

prioritization that is shaping their beliefs. Indeed, a large literature has focused on this problem of selective information presentation, associating it with bounded rationality (see, e.g., Bordalo, Gennaioli and Shleifer 2013; Milosavljevic, et al. 2012; Gabaix and Laibson 2006; and Jacoby 1984).

We identify a different, rational-choice understanding of "priority." In our framework, certain information is prioritized when the cost of acquiring this information is reduced; and information is de-prioritized when the cost of acquiring this information is increased. Sellers and regulators control the relative costs of acquiring information about different product dimensions by controlling the timing, location and presentation of the different bits of information. Rational buyers may succumb to the manipulated costs and acquire only partial information, but unlike irrational buyers they are fully aware of this and they make rational inferences about what they don't know.

Thus, the first contribution of this article is to show that manipulative presentation of information could also be effective and harmful when targeting *rational* consumers. Even when buyers are sophisticated and aware of how firms try to lead them on; even when they sensibly recognize that some dimensions of the products are selectively and strategically obscured; even when they make perfectly rational inferences about what is not disclosed—they can be manipulated and channeled into undesirable choices.

A key feature in the environments that exhibit 'manipulation by mislaid priorities' is the *complexity* of products. Complexity means that the product has multiple dimensions of quality, and that it is costly to inform people about all the relevant dimensions. When a product is simple—when it has a single easily disclosed dimension of quality (for example, the lifespan of a battery)—sellers of low-quality items cannot manipulate rational buyers. If they choose not to disclose the single quality dimension, buyers will infer from their silence that quality is low and will not be manipulated.

But most products are multi-dimensional, and it is often impractical to effectively communicate to buyers the quality of each dimension. Sellers may disclose the information, but rational buyers allocate their attention optimally, and this means that they will review and digest only some of the information, becoming informed only about a subset of the product's features. Anticipating buyers' review patterns, sellers could choose what true information to prioritize. While rational buyers can draw inferences about the undisclosed dimensions, recognizing that sellers may be trying to hide bad dimensions, they cannot overcome their asymmetric information. Unlike the single dimension case, information about the undisclosed features does not fully unravel. Buyers' inferences will be imperfect, causing them to make purchases they end up regretting and, due to the prudence they rationally exert, also refrain from making purchases they would value.[4]

[4] Prior literature has convincingly shown that asymmetric information is rampant in consumer markets, despite important, but only partially successful, attempts to inform consumers through reputational mechanism, private or state-sponsored certifications, ratings systems and other market and non-market mechanisms. [CITE] In this paper does not add to this literature; we take the prevalence of asymmetric information as given. Our contribution is to show how suboptimal prioritization exacerbates the harm from asymmetric information.

The second contribution of this article is to demonstrate that the distortion due to misprioritized information can arise not only from sellers' boastful 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 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 to be told about the low-quality aspects, so that they could focus their limited attention on these dimensions. Sellers, however, prefer to de-prioritize the information about these low-quality dimensions and highlight instead the high-quality ones.

In this *complements* scenario, sellers manipulate buyers by draining their capacity to rationally review the deal, and do so by prioritizing the high-quality dimensions, which have a smaller impact on buyers' utility. 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 catergory 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. Accordintly, it is more impotant 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.

The article shows 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.

Thus, we offer a theoretical framework capturing a distinct form of manipulation—mislaid disclosure priorities—and show that it applies more broadly than the standard accounts of consumer manipulation. It besets rational consumers as much as irrational ones; and it equally applies to mandated disclosures regulated by the law. This framework allows us to offer several contributions to 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]

Ours is a rational-choice model, but it can be extended to allow for imperfect rationality and behavioral biases.[6] Relaxing the rationality assumption would affect the inferences that buyers derive from their partial information and the choices they make.[7] 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 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. Digging deeper, below this abstract level, our analysis helps regulators identify the more important information, primarily through the distinction between complements versus substitutes in quality dimensions.

---

[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] In our model, buyers rationally review only a subset of product attributes due the cost of information acquisition. But there could also be other reasons, grounded in imperfect rationality, for why buyers review only a subset of product attributes. See below.

[7] 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).

This paper contributes to a growing literature on manipulation in consumer markets. Whereas much of the literature analyzes manipulation as the exploitation of an imperfectly rational consumer (see, e.g., Sunstein 2016), we study the possibility of manipulating a rational (Bayesian) consumer, through mislaid prioritization of disclosed information. 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). Buyers' choice to review information on a subset of product attributes plays a crucial role in our analysis. The reasons for, and implications of, similar choices are considered in the literature on bounded rationality and satisficing (see, e.g., Simon 1956), in the literature on information overload (see, e.g., Eppler and Mengis 2004), and in the behavioral industrial organization literature (see, e.g., Spiegler 2016; Spiegler 2006; Gabaix and Laibson 2006; Heidhues, Kőszegi and Murooka 2016).[8] Finally, 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).

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 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

[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.

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. 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.

[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.

Next, consider a *Warning* regime, where the law forces sellers to disclose information about low-quality dimensions. In this regime, if the product has two low-quality dimensions, the seller will disclose information on both dimensions (and buyers will randomly view one of them); if the product has one low-quality dimension, the seller will disclose information on that dimension, and would have to prioritize it over a voluntary disclosure of the other, high-quality dimension; only if the product has two high-quality dimensions, will the seller be free to disclose information on a high-quality dimension. This means that when no warning is viewed, buyers can identify with certainty products with 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 mirros 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: D1 and D2, with quality levels q_1 and q_2. On each dimension, quality is either 0 or 1, and we refer to 0 as low quality (or "L") and to 1 as high quality (or "H"). For example, the product could be a combination of hardware and software, and both dimensions could vary in quality. For each dimension, 50% of Sellers are H and 50% are L; the two dimensions, and their qualities, are independently distributed among Sellers.[10] We thus have four types of Sellers – HH, HL, LH, and LL (the first letter represents D1 and the second represents D2) – each comprising 25% of the market.[11] Sellers and Buyers are randomly matched.

Information structure and disclosure technology. Information is asymmetric. Sellers know both q_1 and q_2, whereas Buyers know only the distributions of quality levels across the four Seller types. Sellers can disclose information and may only do so in a truthful manner (as in Grossman and Hart 1980, Milgrom and Roberts 1986, Milgrom 2008 and Dranove and Jin 2010). We make the crucial assumption that Buyers have limited capacity to digest information: they can review information about only one dimension. If Seller discloses information about one dimension, then Buyer reviews this information; if Seller discloses information about both dimensions, then 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, and digesting them is costly to Buyers. Rational Buyers choose an optimal investment in acquiring and processing information, which in general means that they review only a subset of the disclosed dimensions. They recognize that the quality levels of the remaining dimensions are unknown to them, and form rational expectations about these dimensions. In the simplified two-dimension model, this limitation is captured by the assumption that Buyers can only become informed about one dimension.[13,14]

---

[10] This assumption is relaxed in Section 3.2, where we consider endogenous quality. [We may want to add another extension that retains exogenous quality but relaxes the assumption that quality is independently distributed across the two dimensions: When the quality dimensions are correlated, consumers will infer information on q_2 from a disclosure about q_1. For example, if Seller has a fixed R&D or production budget, then a decision to increase quality on D1 implies a decision not to increase quality on D2.]

[11] More precisely, with a sufficiently large number of firms, 25% is the expected mean share of firms in each category.

[12] The assumption is that Buyer does not know whether Seller made a disclosure on a non-reviewed dimension. If Buyer knew whether a disclosure had been made on each dimension, rational inferences would eliminate the asymmetric information problem. Asymmetric information could still be maintained if Seller always makes two disclosures, but can disclose information about the same dimension twice. Now, it is common knowledge that Seller made two disclosures, but Seller has private information about the content of those disclosures (whether they are disclosing the same dimension twice, or disclosing the two different dimensions). Another possible assumption is that Seller always discloses information on both dimensions, but Buyer can review the content of the disclosed information, i.e., whether it is High or Low, for only one dimension.

[13] A more complete rational-choice model would derive the number of dimension that Buyer learns from an optimal information acquisition analysis. Specifically, in some cases (e.g., for some utility functions), the benefit from learning the second dimension is greater and may justify learning about that second dimension, depending on the cost of learning another dimension. Our "Buyer learns one dimension" assumption is a reduced-form representation of an optimal information acquisition model where learning another dimension is not worth the cost (perhaps because of the increasing marginal cost of learning more dimensions). Below, we comment further on the conditions that support the Buyer-learns-one-dimension assumption.

[14] We also assume that the message space, i.e., what Seller can disclose, is limited to "High" or "Low." This assumption is discussed further in Section 4.5 below.

This structure, allowing only for partial disclosure, captures a realistic information constraint. Without it, if Sellers can disclose and rational Buyers can digest all relevant information at no cost, then manipulation would be impossible in a rational choice framework. Indeed, our motivation for examining a multi-dimension model is precisely to study what happens when products are more complex and when rational Buyers can overcome some, but not all, asymmetric information.[15]

Buyers' utility and Sellers' profit. Risk-neutral Buyers each purchase at most one unit of the good. The value of the product to Buyers depends on its mix of qualities. Buyers who decide to purchase the good enjoy utility $u(q_1, q_2)$. (Buyers' "no purchase" utility is normalized to zero.) In particular, we normalize the LL utility, $u(0,0)$, to zero; and the HH utility, $u(1,1)$, to one. In the basic model, we focus on the symmetric case, where the HL and LH utilities are identical, i.e., $u(1,0) = u(0,1) = x$, where $x \in [0,1]$. 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, for the four quality combinations, are common knowledge and are summarized in Table 1 below.

	$q_2 = 1$	$q_2 = 0$
$q_1 = 1$	HH: $u(1,1) = 1$	HL: $u(1,0) = x$
$q_1 = 0$	LH: $u(0,1) = x$	LL: $u(0,0) = 0$

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—the difference between consumers' value and sellers' cost—is captured by Buyers, and thus the consumer surplus is equal to overall welfare.[16]

To motivate the assumption of quality-independent per-unit cost, think of situations where high quality is the result of a prior investment (a fixed cost) that allows the Seller to produce higher quality at the same per-unit production cost (c). An alternative motivation has to do with model design: Given our perfect competition assumption, which implies $p = c$, if cost varies with quality, then price serves as a perfect signal of quality and the informational asymmetry that motivates the analysis disappears. In Section 3, we relax the perfect competition assumption and allow the per-unit cost to vary with quality.

The legal regimes. We compare three regimes. The first regime involves no legal restriction on Sellers, and we refer to it as the Caveat Emptor (CE) regime. Sellers may decide if and what to disclose and face no sanction for failure to disclose some information. The only sanction is for lying. In particular, Sellers may, if they so wish, disclose only H-quality dimensions.

[15] 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 have scarce attention which they rationally allocate based on an 'attention budget' constraint).

[16] 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.

The second regime is a variant of mandated disclosure that requires Sellers to warn Buyers about L-quality. In this Warning (W) regime, Sellers must disclose L-quality dimensions, if such exist. Moreover, the warning must be prioritized, and it cannot be overshadowed by voluntary disclosures about the product's H-quality features. In practice, for HL and LH Sellers, this regime mandates disclosure of the L dimension and precludes the disclosure of the H dimension. (Sellers can disclose, less prominently, the H dimension, but Buyers will be able to effectively review only the highlighted L-quality warning.) For LL Sellers, the Warning regime requires the disclosure of both L dimensions, even though Buyers will be able to review only one.

The third regime is another variant of mandated disclosure, which we call Full Disclosure (FD). In this regime, Sellers must disclose information on both dimensions, regardless of their quality, even though Buyers will able to review only one. In practice, FD will require the disclosure of L-quality dimensions, but without requiring the prioritization of these L disclosures (in contrast to the Warning regime); and, since Sellers will voluntarily disclose H-quality dimensions, the result is full discosure. So understood, Full Disclosure is probably the most common regime.

To evaluate the performance of these three regimes, we begin, in the next sub-section, 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. We can thus consider the market outcome and welfare for each of the four Seller groups separately: When matched with an HH Seller, Buyer's utility is $u = \hat{u} = 1$. If $c < 1$, then Buyer purchases the good and $CS = 1 - c$; if $c \geq 1$, Buyer does not purchase the good and $CS = 0$. When matched with an HL or LH seller, Buyer's utility is $u = \hat{u} = x$. If $c < x$, then Buyer purchases the good and $CS = x - c$; if $c \geq x$, Buyer does not purchase the good and $CS = 0$. When matched with an LL seller, Buyer's utility is $u = \hat{u} = 0$, there is no purchase and $CS = 0$.

The overall consumer surplus is:

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

When the $c < x$, Buyers gain positive payoff when buying from HH Sellers (25% of cases, contributing $\frac{1}{4} \cdot (1 - c)$, to Buyers' expected surplus), as well as from HL and LH Sellers (50% of the cases, contributing $\frac{1}{2} \cdot (x - c)$ to Buyers' expected surplus). When $c \in [x, 1)$, Buyers gain

positive payoff only when buying from HH Sellers (25% of cases, contributing $\frac{1}{4} \cdot (1-c)$, to Buyers' expected surplus). When $c \geq 1$, Buyers never purchase the good.

2.4 Caveat Emptor

We now derive the equilibrium in a Caveat Emptor (CE) regime. Without a legal requirement to make any specific disclosures, and with the sole obligation not to lie, Sellers need not disclose the L-dimensions and may 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 will randomly choose one of the H-dimensions and disclose it. (An HH Seller can also disclose both H-dimensions, knowing that Buyer will randomly review one of them.) HL and LH Sellers will disclose the one H-dimension. And an LL Seller will remain silent. A rational Bayesian Buyer would draw the following inferences: First, if Seller remains silent Buyer will know with certainty that Seller is LL and that purchase utility is zero. Second, if Seller discloses an H-dimension Buyer knows that this could happen in two situations: (1) when the other dimension is L (an ex ante likelihood of $\frac{1}{4}$); and (2) when the other dimension is H, and Seller randomly chose this specific dimension for disclosure (an ex ante likelihood of $\frac{1}{4} \times \frac{1}{2}$). Thus, conditional on viewing a H-disclosure, the likelihood of the other dimension being H is $\frac{1}{3}$, and of being L is $\frac{2}{3}$.[17] The expected utility from purchasing a good from a Seller who disclosed H is: $\hat{u} = \frac{1}{3} \cdot 1 + \frac{2}{3} \cdot x$. If cost, and price, are lower than $\frac{1}{3} + \frac{2}{3} \cdot x$, Buyer will purchase the good and enjoy a consumer surplus of: $\frac{1}{3} + \frac{2}{3} \cdot x - c = \frac{1}{3} \cdot (1-c) + \frac{2}{3} \cdot (x-c)$. If cost, and price, are higher, Buyer will not purchase the good.[18]

Ex ante, there is a probability of $\frac{1}{4}$ that Buyer meets a silent LL Seller; and the expected surplus is then zero. And there is a probability of $\frac{1}{4} + \frac{1}{4} \times \frac{1}{2} = \frac{3}{8}$ that Buyer meets a Seller who discloses H on the first dimension (D1); and a probability of $\frac{1}{4} + \frac{1}{4} \times \frac{1}{2} = \frac{3}{8}$ that Buyer meets a Seller who discloses H on the second dimension (D2). Together, there is a probability of $\frac{3}{4}$ that Buyer observes

---

[17] If Buyer observes $q_1 = 1$, a rational, Bayesian Buyer will know that:

$$Pr(q_1 = q_2 = 1 | \text{B observes } q_1 = 1) = \frac{Pr(q_1 = q_2 = 1) \cdot Pr(\text{B observes } q_1 = 1 | q_1 = q_2 = 1)}{Pr(\text{B observes } q_1 = 1)} = \frac{\frac{1}{4} \cdot \frac{1}{2}}{\frac{1}{4} \cdot \frac{1}{2} + \frac{1}{4}} = \frac{1}{3}$$

and $Pr(q_1 = 1 \text{ and } q_2 = 0 | \text{B observes } q_1 = 1) = \frac{2}{3}$. Similarly, if Buyer observes $q_2 = 1$, Buyer will know that $Pr(q_1 = q_2 = 1 | \text{B observes } q_2 = 1) = \frac{1}{3}$ and $Pr(q_1 = 0 \text{ and } q_2 = 1 | \text{B observes } q_2 = 1) = \frac{2}{3}$.

[18] 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 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 $k > \frac{2}{3} \cdot (c - x)$.

a disclosure of H, and the expected surplus is then: $\frac{3}{4}\cdot\left(\frac{1}{3}\cdot(1-c)+\frac{2}{3}\cdot(x-c)\right)=\frac{1}{4}\cdot(1-c)+\frac{1}{2}\cdot(x-c)$, if $c<\frac{1}{3}+\frac{2}{3}\cdot x$. To summarize, consumer surplus under Caveat Emptor is:

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

When $c<\frac{1}{3}+\frac{2}{3}\cdot x$, Buyer purchases from HH Sellers (earning an expected surplus of $\frac{1}{4}\cdot(1-c)$) and from HL and LH Sellers (earning an expected surplus of $\frac{1}{2}\cdot(x-c)$). When $c\geq\frac{1}{3}+\frac{2}{3}\cdot x$, Buyer does not purchase from any Seller (earning an expected surplus of zero).

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, Buyers make distorted decisions. One distortion pushes for *excessive* purchases, and it occurs when the cost, and price, are in the $c\in\left[x,\frac{1}{3}+\frac{2}{3}\cdot x\right)$ range. Buyers who receive a disclosure of H-dimension are willing to buy at a price in this range, despite the uncertainty concerning the other dimension. When it turns out that the other dimension is L, Buyers end up with negative ex-post surplus. A second distortion pushes for *insufficient* purchases, and it occurs when the cost, and price, are in the $c\in\left[\frac{1}{3}+\frac{2}{3}\cdot x,1\right)$ range. Recognizing that the other dimension might be L, Buyers are unwilling to buy at this price, and lose a profitable purchase when the other dimension is (also) H.[19,20]

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 $c\in\left[x,\frac{1}{3}+\frac{2}{3}\cdot x\right)$, there is positive consumer surplus under Caveat Emptor (the red line is above 0), but less than the maximal surplus. When

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

[20] Further, at the very high end, when $c\geq 1$, Buyers should not purchase the good; and that is what happens in the Caveat Emptor equilibrium.

$c \in \left[\frac{1}{3} + \frac{2}{3} \cdot x, 1\right)$, there is no purchase under Caveat Emptor and surplus is zero, whereas perfect information allows for valuable purchases from HH Sellers.

This analysis suggests a key insight: As the value of *x* increases—namely, as the importance of having at least one high dimension increases—Caveat Emptor comes closer to the perfect information benchmark, delivering a more efficient outcome. Formally, when *x* is higher, two things happen: (i) the region in which Caveat Emptor is distortive, $c \in (x, 1)$, shrinks; and (ii) the $\frac{1}{3} + \frac{2}{3}x$ threshold for making a purchase goes up, such that Buyers are willing to purchase at a higher price. The combined effects reduce both distortions, but the insufficient-purchases region, $1 - \left(\frac{1}{3} + \frac{2}{3} \cdot x\right) = \frac{2}{3} \cdot (1 - x)$, shrinks faster than the excessive-purchases region, $\frac{1}{3} + \frac{2}{3} \cdot x - x = \frac{1}{3} \cdot (1 - x)$. Graphically, when *x* is higher the region in which the red line merges with the social optimum black line is larger. Informally, a higher *x* makes the difference between one versus two H-dimensions smaller, and it becomes all the more important for Buyers to know that there is *at least* one H-dimension. Caveat Emptor delivers this information to Buyers.

2.5 Warning

Caveat Emptor allows Sellers to prioritize information about H-dimensions, to the detriment of Buyers whose priority is to know whether the undisclosed dimension is Low. This mislaid priority in disclosure is manipulative. Since the problem is due to the Seller's non-disclosure of the L-dimension, a straightforward legal intervention would mandate different disclosure priorities. Two such responses are typical. The first, which we call a "Warning" regime and examine in this subsection, requires Sellers to disclose the L-dimensions. The second, which we call a "Full Disclosure" regime and examine in the next subsection, requires Sellers to disclose all dimensions fully and equally. We show that these mandated disclosure regimes do indeed prevent the manipulation created by Caveat Emptor, but introduce a different, unintended, distortion.

The Warning regime requires that Sellers disclose any L-dimension, and not overshadow the warning by disclosing an H dimension.[21] LL Sellers will have to disclose both L-dimensions, and Buyers will randomly review one. HL and LH Sellers will have to disclose only the one L-dimension. And an HH Seller, free from a warning obligation, will voluntarily disclose one or both H-dimensions (and if both are disclosed, Buyer will review one of them randomly.)

Rational, Bayesian Buyers draws the following inferences from these disclosures: First, if they observe an H-dimension, Buyers will infer that she is facing an HH Seller and that the purchase utility is 1. Buyers will purchase the good and enjoy a consumer surplus of: $CS = 1 - c$, as long as the cost, and price, are below 1. Second, if Buyers observe a warning about an L-dimension, they know that this could happen in two situations: (1) when the other dimension is H (an ex ante

[21] Generally, a Warning regime could allow sellers to also disclose, if they so choose, any High dimension, but it requires that such voluntary High disclosures not compete for attention with the Low-quality warnings. Thus, HL and LH Sellers could disclose in a secondary manner the High dimension, but since Buyers have limited capacity to review and will only observe the Low dimension, such additional High disclosure would be futile for Sellers. For simplicity, then, we assume that under a Warning regime only HH sellers will disclose a High dimension.

likelihood of $\frac{1}{4}$); and (2) when 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 disclosure of L-dimension, 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).[22] The expected utility from purchasing a good in this case is: $\hat{u} = \frac{2}{3} \cdot x$. If cost, and price, are lower than $\frac{2}{3} \cdot x$, Buyer who observes a Low warning will purchase the good and enjoy a consumer surplus of: $CS = \frac{2}{3} \cdot x - c = \frac{2}{3} \cdot (x - c) - \frac{1}{3} \cdot c$. 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 an H-dimension; the expected surplus is: $\frac{1}{4} \cdot (1 - c)$. And there is a probability of $\frac{1}{4} + \frac{1}{4} \times \frac{1}{2} = \frac{3}{8}$ that Buyer meets a Seller who discloses L-quality on the first dimension (D1); and a probability of $\frac{1}{4} + \frac{1}{4} \times \frac{1}{2} = \frac{3}{8}$ that Buyer meets a Seller who discloses L-quality on the second dimension (D2). Together, there is a probability of $\frac{3}{4}$ that Buyer observes a disclosure about an L-dimension; the expected surplus is: $\frac{3}{4} \cdot \left(\frac{2}{3} \cdot (x - c) - \frac{1}{3} \cdot c\right) = \frac{1}{2} \cdot (x - c) - \frac{1}{4} \cdot c$ Thus, consumer surplus in the Warning regime is:

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

When $c < \frac{2}{3} \cdot x$, Buyer purchases from HH Sellers (earning an expected surplus of $\frac{1}{4} \cdot (1 - c)$), from HL and LH Sellers (earning an expected surplus of $\frac{1}{2} \cdot (x - c)$) and from LL Sellers (earning an expected surplus of $-\frac{1}{4} \cdot c$). When $c \in \left[\frac{2}{3} \cdot x, 1\right)$, Buyer purchases only from HH Sellers (earning an expected surplus of $\frac{1}{4} \cdot (1 - c)$). When $c \geq 1$, Buyer does not purchase from any Seller (earning an expected surplus of zero).

Comparison to the Perfect Information Benchmark. At the higher range of costs (and 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-dimension

[22] If Buyer observes $q_1 = 0$, a rational, Bayesian Buyer will know that:

$$Pr(q_1 = q_2 = 0 | \text{B observes } q_1 = 0) = \frac{Pr(q_1 = q_2 = 0) \cdot Pr(\text{B observes } q_1 = 0 | q_1 = q_2 = 0)}{Pr(\text{B observes } q_1 = 0)} = \frac{\frac{1}{4} \cdot \frac{1}{2}}{\frac{1}{4} \cdot \frac{1}{2} + \frac{1}{4}} = \frac{1}{3}$$

and $Pr(q_1 = 0 \text{ and } q_2 = 1 | \text{B observes } q_1 = 0) = \frac{2}{3}$. Similarly, if Buyer observes $q_2 = 0$, Buyer will know that $Pr(q_1 = q_2 = 0 | \text{B observes } q_2 = 0) = \frac{1}{3}$ and $Pr(q_1 = 1 \text{ and } q_2 = 0 | \text{B observes } q_2 = 0) = \frac{2}{3}$.

warning. Buyers make the optimal decision in this range of high costs and high prices: they purchase only if a warning is not observed. 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 $c < x$, things don't look as good for the Warning regime. In the perfect information benchmark, Buyers distinguish between LL Sellers and HL/LH Sellers and, in this range of prices, make purchases only from the HL/LH Sellers. In the Warning regime, Buyers cannot distinguish between LL from HL/LH Sellers and cannot make the selective purchase decision. As a result, Buyers that observe an L-disclosure are subject to two distortions: One distortion pushes for *excessive* purchases, and it occurs when the cost, and price, are in the $c < \frac{2}{3} \cdot x$ range. At such a low price, even the Buyer who is warned about an L-dimension purchases the product. If it turns out that the other dimension is also L, Buyer ends up with negative ex-post surplus. A second distortion pushes for *insufficient* purchases, and it occurs when the cost, and price, are in the $c \in \left[\frac{2}{3}x, x\right)$ range. Recognizing the risk that the other dimension might (also) be L, Buyer does not buy at this price, forgoing the value of a profitable purchase when the other dimension turns out to be H.[23,24]

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): 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. Formally, when x is lower, two things happen: (i) the region in which Warning is distortive, $c \in (0, x)$, shrinks; and (ii) the $\frac{2}{3}x$ threshold for making a purchase when a warning of an L-dimension is observed decreases. The combined effects reduce both distortions, but the excessive-purchases region, $\frac{2}{3}x - 0 = \frac{2}{3}x$, shrinks faster than the insufficient-purchases region, $x - \frac{2}{3}x = \frac{1}{3}x$. Graphically, when x is lower, the region in which the blue line merges with the social optimum black line is larger. 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 second mandated disclosure regime: Full Disclosure (FD). It requires Sellers to disclose the quality of both dimensions. Under this regime, Sellers have no discretion about the

[23] If $c < \frac{2}{3} \cdot x$, 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 $\frac{1}{2} \cdot (x - c)$, as compared to the perfect information benchmark.

[24] If $c \geq 1$, then Buyer should never purchase the good; and that is what happens in the Warning equilibrium.

content of their disclosures and must equally disclose both H and L dimensions. Buyers, limited in their ability to review, randomly view one of the two disclosed dimensions.

Rational, Bayesian Buyers draw the following inferences from Sellers' disclosures. 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 (each with a likelihood of 50%), which promises an expected utility of $\hat{u} = \frac{1}{2} \cdot (1 + x)$. Buyers make a purchase when the cost, and price, is lower than $\frac{1}{2} \cdot (1 + x)$, and enjoy a payoff of $CS = \frac{1}{2} \cdot (1 + x) - c = \frac{1}{2} \cdot (1 - c) + \frac{1}{2} \cdot (x - c)$. The likelihood of viewing H is $\frac{1}{2}$ (a $\frac{1}{4}$ probability of facing an HH Seller plus a $\frac{1}{2}$ probability of facing an HL or LH Seller and then viewing the one H dimension with probability $\frac{1}{2}$); and thus the contribution of the "view H" payoff to Buyers' ex-ante surplus is $\frac{1}{2} \cdot \left(\frac{1}{2} \cdot (1 - c) + \frac{1}{2} \cdot (x - c)\right) = \frac{1}{4} \cdot (1 - c) + \frac{1}{4} \cdot (x - c)$. The other scenario occurs if Buyers view an L-dimension. They expect the product to be either LL of LH (each with a likelihood of 50%), which promises an expected utility of $\hat{u} = \frac{1}{2} \cdot x$. Buyers make a purchase when the cost, and price, is lower than $\frac{1}{2} \cdot x$, and enjoy a payoff of: $CS = \frac{1}{2} \cdot x - c = \frac{1}{2} \cdot (x - c) - \frac{1}{2} \cdot c$. The likelihood of viewing L is $\frac{1}{2}$; and thus the contribution of the "view L" payoff to Buyers' ex-ante surplus is $\frac{1}{2} \cdot \left(\frac{1}{2} \cdot (x - c) - \frac{1}{2} \cdot c\right) = \frac{1}{4} \cdot (x - c) - \frac{1}{4} \cdot c$. In sum, consumer surplus under Full Disclosure is:

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

When $c < \frac{1}{2} \cdot x$, Buyers purchase regardless of the observed disclosure, earning an expected surplus of $\frac{1}{4} \cdot (1 - c)$ from HH Sellers, $\frac{1}{2} \cdot (x - c)$ from HL/LH Sellers, and $-\frac{1}{4} \cdot c$ from LL Sellers. When $c \in \left[\frac{1}{2} \cdot x, \frac{1}{2} \cdot (1 + x)\right)$, Buyers purchase only when observing a disclosure of an H-dimension, earning an expected surplus of $\frac{1}{4} \cdot (1 - c)$ from HH Sellers and $2 \cdot \frac{1}{4} \cdot \frac{1}{2} \cdot (x - c) = \frac{1}{4} (x - c)$ from HL/LH Sellers. When $c \geq \frac{1}{2} \cdot (1 + x)$, Buyers do not make a purchase, irrespective of the observed disclosure, earning zero.

Comparison to the Perfect Information Benchmark. A Buyer who observes an H-dimension cannot distinguish between an HH Seller and an HL or LH Seller. This uncertainty leads to two distortions: One distortion is *insufficient* purchases, and it occurs when the cost, and price, are in the $c \in \left[\frac{1}{2} \cdot (1 + x), 1\right)$ range. Because they recognize the risk that the other dimension might be Low, Buyers are unwilling to buy at this price, and lose the value of a profitable purchase when

the other dimension turns out to be (also) H. Another distortion is *excessive* purchases, and it occurs when the cost, and price, are in the $\left[x, \frac{1}{2} \cdot (1+x)\right)$ range. Buyers who observe High are willing to buy at this price, despite the uncertainty concerning the other dimension. If it turns out that the other dimension is L, they end up with negative surplus.

A Buyer who observes an L-dimension cannot distinguish between an LL Seller and an HL or LH Seller. This uncertainty, too, leads to two distortions: One distortion is *insufficient* purchases, and it occurs when the cost, and price, are in the $c \in \left[\frac{1}{2} \cdot x, x\right)$ range. Because they recognize the risk that the other dimension might be (also) L, buyers are unwilling to buy at this price, and lose the value of a profitable purchase when the other dimension turns out to be H. Another distortion is *excessive* purchases, and it occurs when the cost, and price, are in the $c < \frac{1}{2} \cdot x$ range. Buyers are willing to buy at this price, despite the uncertainty concerning the other dimension. If it turns out that the other dimension is (also) L, they end up with negative surplus.

Figure 1 below depicts these distortions. The black line represents the maximal, perfect-information surplus and the green line represents the surplus in the Full Disclosure regime. The FD line is always below the perfect-information line (when $c < 1$).

2.7 Comparing the Three Regimes

The last three subsections examined the outcome under three familiar legal regimes: Caveat Emptor, Warning, and Full Disclosure. We saw that each regime is associated with inefficient purchase decisions by Buyers, and that the type and magnitude of the distortions depend on the parameter *x*, which measures the importance of having at least one High dimension. We can now compare the three regimes, in terms of consumer surplus.[25]

Under the Caveat Emptor regime, distortions occur when $c > x$. The surplus loss is due to Buyer's inability to distinguish between Sellers with one H-dimension versus Sellers with two. (Buyer can identify an LL Seller with certainty and never purchases from this Seller.) The uncertainty drives the rational Buyer to purchase only if the cost, and price, are below the threshold $\frac{2}{3}x + \frac{1}{3}$. Unable to separate the HH from the HL/LH Sellers, Buyer inefficiently purchases the product from HL and LH Sellers when $c \in \left[x, \frac{2}{3}x + \frac{1}{3}\right)$, for a negative expected payoff of $\frac{1}{2}(x - c)$; and Buyer inefficiently fails to purchase the product from HH Sellers when $c \in \left[\frac{2}{3}x + \frac{1}{3}, 1\right)$, forgoing a positive expected payoff of $\frac{1}{4}(1 - c)$.

Under the Warning regime, distortions occur when $c < x$. The surplus loss is due to Buyer's inability to distinguish between Sellers with one L-dimension versus Sellers with two. (Buyer can identify the HH Seller with certainty and purchases from this Seller whenever the purchase is efficient, i.e., whenever $c < 1$.) The uncertainty drives Buyer to purchase from a Seller who

[25] 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.

discloses Low if and only if the cost, and price, are below the threshold $\frac{2}{3}x$. Unable to separate the LL from the HL/LH Sellers, Buyer inefficiently purchases the product from LL Sellers when $c < \frac{2}{3}x$, for a negative expected payoff of $-\frac{1}{4}c$; and Buyer inefficiently fails to purchase the product from HL and LH Sellers when $c \in \left[\frac{2}{3}x, x\right)$, forgoing a positive expected payoff of $\frac{1}{2}(x - c)$.

As we see, Caveat Emptor enjoys an advantage when the relative cost, and price, of the product is low, whereas Warning enjoys an advantage when the relative cost, and price, of the product is high. Because the distortions under Caveat Emptor are due to Buyer's inability to distinguish a product with only one H-dimension from a product with two, the uncertainty is irrelevant when the price is low enough to make either purchase worthwhile. Conversely, because the distortions under Warning are due to Buyer's inability to distinguish a product with only one H-dimension from a product with none, 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 on both dimensions and even one Low significantly lowers Buyers' payoffs. In these scenarios, the parameter *x* (the value of the product with only one H-dimension) is low. When the two dimensions are complements it is important for Buyers to distinguish HH Sellers from all other Sellers, 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. The complements case captures products where all dimensions are critical and the presence of any low dimension can significantly detract from the value of the product. 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 Buyer to know that there is at least one H-dimension, namely, to separate the LL Sellers. Having a second H-dimension is of relatively less value. In these scenarios, the parameter *x* takes on a high value. The fact that Caveat Emptor induces pooling among HH, HL and LH Sellers matters less to Buyers who care primarily about securing at least one H-dimension. The substitutes case captures products with dimensions that overlap, in the sense that they have more features than Buyers have the capacity to enjoy, and could thus be highly valuable even if not all are high-quality. Caveat Emptor's attention to these H-dimensions accounts for its superiority.

Unlike the Caveat Emptor and Warning regimes, the Full Disclosure regime does not allow Buyer to identify any Seller with certainty. When the disclosure observed is an L-dimension, Buyer cannot distinguish between LL, HL and LH (but can rule out HH). The uncertainty drives Buyer to purchase if and only if the cost, and price, are below $\frac{1}{2} \cdot x$. When $c < x$, this "pooling" leads to both inefficient purchases (from LL Sellers when $c < \frac{1}{2}x$, for a loss of $-\frac{1}{4}c$) and an inefficient failure to purchase (from HL/LH Sellers when $c \in \left[\frac{1}{2}x, x\right)$, forgoing a gain of $\frac{1}{4}(x - c)$.) Symmetrically, when the disclosure observed is an H-dimension, Buyer rules out LL but cannot distinguish between HH, HL and LH, and thus purchases if and only if the cost, and price, are

below $\frac{1}{2}\cdot(1+x)$. Here, when $c > x$, there are both inefficient purchase (from HL/LH Sellers when $c \in [x, \frac{1}{2}(1+x))$, for a loss of $\frac{1}{4}(x-c)$) and an inefficient failure to purchase (from HH Sellers when $c \in [\frac{1}{2}(1+x), 1)$, forgoing a gain of $\frac{1}{4}(1-c)$.)

How does Full Disclosure rank relative to Caveat Emptor and Warning? 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 could 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 following a High disclosure; 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 High only when randomly viewing the H-dimension. As a result, the likelihood of observing High, and thus making an inefficient purchase, is lower under FD. (Full disclosure also dominates Caveat Emptor in the $c \in [\frac{1}{3}+\frac{2}{3}x, \frac{1}{2}(1+x))$ 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.)[26]

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* is at the low range, Caveat Emptor is the most efficient rule. Indeed, it achieves the first-best, perfect-information outcome (due to the fact that, when the product is cheap, Buyers are not chilled by the concern over an undisclosed L-dimension). And, when *c* is at the high range, Warning is the most efficient rule, inducing the first-best, perfect-information outcome (due to the fact that, when the product is expensive, Buyer's uncertainty over a potential undisclosed H-dimension does not induce risky purchases). There are no *c* values for which Full Disclosure is optimal.[27] 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.[28]

---

[26] Full Disclosure dominates Warning in the $c \in [\frac{2}{3}x, x)$ 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 $c \in [\frac{1}{2}x, \frac{2}{3}x)$ 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.)

[27] When $c < x$, FD (weakly) dominates Warning, but is always dominated by CE; and when $c > x$, FD (weakly) dominates CE, but is always dominated by Warning.

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

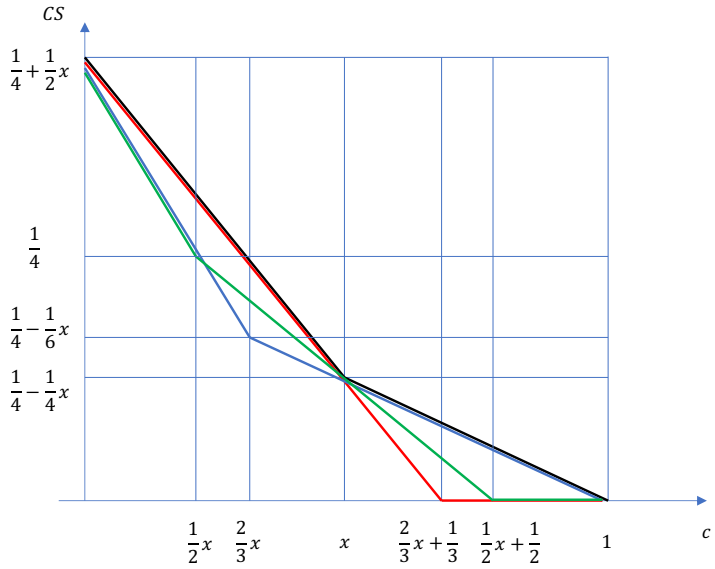

Figure 1

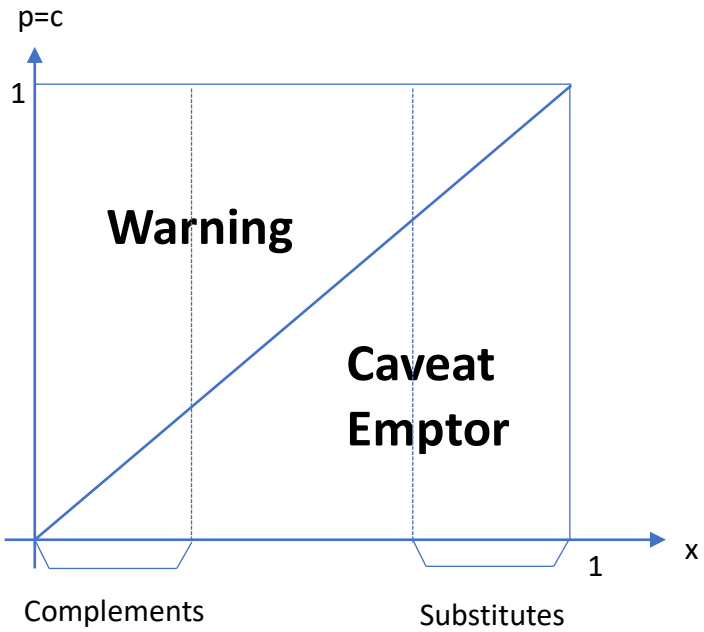

Figure 2

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). This principle supports a conditional

disclosure regime: Sellers must warn about L-dimensions when such information is critical to Buyers; otherwise, when it is more important for Buyers to know about an H-dimension, Sellers are under no obligation to disclose. Specifically, the ideal PIP regime would require a warning any time $x < c$. Otherwise, no warning is necessary and sellers should be free to boast High quality dimensions.

PIP rejects a one-size-fits-all priority of information. The Warning regime always prioritizes L-dimensions for disclosure, whereas Caveat Emptor always prioritizes H-dimensions. Buyer's priorities, however, vary with the product. They need warnings, but only when their priority is to be alerted to aspects of low quality.

The PIP conditional mandate "issue a warning if and only if $x < c$" requires courts or lawmakers to verify these 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 were quality dimensions are *complements*—where Buyers need to know that both are H, namely, that none are L. Forcing Sellers to warn about an L-dimension is the surest way to give Buyers this needed information. Conversely, when x is high, the PIP warning condition is less likely to hold. This is the case where quality dimensions are *substitutes*—where Buyers need to know that at least one dimension is H. Sellers can be counted on to give this information when they are free to prioritize high quality aspects.

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 such products, it is less likely that a single H-dimension will give Buyer enough utility. A warning regarding the presence of an L-dimension is thus critical. Conversely, when c is low, a product with even one H-dimension is more likely to generate enough utility to justify its cost. A Caveat Emptor regime that enables Buyers to know if there is at least one H dimension provides the needed information.

These practical criteria are, of course, crude approximations for the underlying condition "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, allowing us to investigate the generality of the legal implications from the model and to identify additional refinements for anti-manipulation law. We focus on three aspects. First, we introduce asymmetry in the quality dimensions, where some product attributes 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 relax the assumption that products' costs are independent of their quality. Once it becomes costlier to make high-quality products, we can examine the incentives of sellers to invest in improved quality. We show how the prospect of ex post manipulation distorts sellers' ex ante incentives to invest in quality. Third, we relax the simplifying assumption that products have only two dimensions and revisit the comparison between the three legal regimes when there are more dimensions. This extension also allows us to analyze sellers' incentives to increase product complexity, and to identify manipulative strategies in this domain. 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

The analysis in Section 2 captured the complexity of products by assuming that there are more dimensions than people can effectively review. Products had two dimensions, while people were able to review only one. Manipulation occurred by mislaid priorities in the presentation of information, and we saw that it could occur even when the two dimensions were symmetric in their contribution to the product's value. Sometimes people need more than anything to know if a product has any low-quality dimension, and yet sellers prefer to tell them about some high-quality dimension. Other times people's priority is to know about the presence of even a single high-quality dimension, and yet the law requires that sellers draw buyers' attention to a low-quality dimension.

We now expand the lens and consider a product that has two dimensions that are not symmetrically valued. In the basic model, the symmetry assumption meant that if a product had only one high-quality dimension, it didn't matter to buyers which of the two dimensions was high. But products have some attributes that are more important than others. We capture this possibility by assuming that among the two dimensions, one is "primary" and the other "secondary." While buyers still prefer products to have high quality on both dimensions, if the product is to have only one high-quality dimension, buyers prefer that it would be in the primary dimension.

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

Consider *Caveat Emptor*. Buyers know that sellers with a perfect product (high quality in both dimensions), who recognize that their buyers can only review one dimension, will prioritize disclosure of the primary dimension. And, of course, if the product is high-quality only in the primary dimension, sellers again will disclose only that dimension. Accordingly, if buyers observe a seller disclosing high quality in the secondary dimension, they know for sure that the primary dimension is 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 have no 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. Because of this remaining uncertainty, some distortion remains, but it is negligible when the cost of the product is low, in which case buyers purchase the product despite the uncertainty.

The presence of a primary dimension also helps design a superior warning regime. In the symmetric case, sellers were required to disclose all low-quality dimensions, and when buyers viewed a warning with regard to a certain dimension, they 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 than the secondary dimension warning.) Under this *Priority Warning* regime, if buyers observe a warning on the secondary dimension, they know for sure that the primary dimension is high quality—otherwise it would have been prioritized for a warning. As a result, buyers are able to infer more information than under the (regular) Warning regime that applied in the symmetric case, and they end up with perfect information about the primary dimension. If it is low-quality, sellers would warn about it. Any other disclosure by sellers means that the primary dimension is high quality. The only remaining uncertainty under Priority Warning surrounds the secondary dimension, and it occurs when the primary dimension is (warned to be) low-quality. Thus, some distortion remains, but it is negligible when the cost of the product is high, in which case buyers do not purchase the product regardless of the secondary dimension's quality.

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. This is true in a Caveat Emptor regime—sellers have less room to manipulate; and in a Priority Warning regime—which limits the concern about manipulation by disclosure regulation. Second, the problem of manipulation does not go away, requiring a thoughtful regulatory intervention. The spirit of the optimal intervention we saw 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 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. 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.[29] In

[29] Under Caveat Emptor, when buyers observe High on the primary dimension, they don't know if the product is High on both dimensions or only on the primary dimension; but as long as the price is below the value contribution of the primary dimension they make the optimal decision to purchase. (When buyers observe High on the secondary dimension, they know for sure that the product is Low on the primary dimension and make an optimal decision to purchase if and only if the price is below the value contribution of the secondary dimension.) Under Priority Warning,

the asymmetric dimensions case, the *principle of information priority* (that we characterized in Section 2) 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. With perfect information, Buyers know with certainty which type of Seller they face. When matched with an HH Seller, Buyers purchase the good whenever $c < 1$. When matched with an HL Seller, Buyers purchase the good whenever $c < x_1$. When matched with an LH seller, Buyers purchase the good whenever $c < x_2$. And when matched with an LL seller, Buyers never purchase the good.

3.1.3 Caveat Emptor

With asymmetric quality dimensions, there is one key difference in Sellers' disclosure strategy: Whereas in the baseline, symmetric-dimensions model, an HH Seller randomly chose to disclose one of the two quality dimensions, here an HH Seller discloses the more important dimension, D1. Thus, Caveat Emptor entails inherent prioritization. This means that Buyers 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 inherent prioritization allows Buyers to infer more information than they could in the symmetric-dimensions scenario. The only remaining uncertainty is due to Buyers' inability to distinguish between HH and HL Sellers, who both disclose H on D1. This uncertainty accounts for some loss of surplus: it drives the rational Buyer who observes H on D1 to purchase only if the cost, and price, are below the threshold $\frac{1}{2} \cdot (1 + x_1)$. Unable to separate the HH from the HL Sellers, 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 Buyer 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 $c > x_1$.

---

when buyers observe Low on the primary dimension, they don't know if the product is Low on both dimensions or only on the primary dimension; but as long as the price is above the value contribution of the secondary dimension they make the optimal decision not to purchase. (When buyers observe Low on the secondary dimension, they know for sure that the product is High on the primary dimension and make an optimal decision to purchase if and only if the price is below the value contribution of the primary dimension; and when they observe a High disclosure they know that the product is High on both dimensions and make an optimal decision to purchase if and only if the price is below the value of a perfect product.) We thus see that under both regimes buyers make optimal decisions when the price is between the value contribution of the primary dimension and the value contribution of the secondary dimension.

3.1.4 Warning

In the asymmetric-dimensions extension, we need to distinguish between two versions of the Warning regime: The (regular) Warning regime and the Priority Warning regime. The (regular) Warning regime is the one described in the baseline model; it requires that Sellers disclose all L-dimensions, and—if an L dimension exists—Sellers cannot disclose any H dimension. The Priority Warning regime qualifies the (regular) Warning regime, such that LL Sellers will have to disclose only the important, D1 L-dimension, and will be prohibited from disclosing the other L-dimension. We focus on the Priority Warning regime and show that it is superior to other disclosure regimes; we then compare Priority Warning to Caveat Emptor.

Whereas Caveat Emptor entails inherent prioritization, the Priority Warning regime imposes legally-mandated prioritization. Under Priority Warning, Buyers know with certainty when they are facing an HH Seller, and when they are facing HL Seller (if Seller discloses L on D2). (Recall that, in the baseline, symmetric-dimensions case, under Warning, Buyers know with certainty only when they are facing an HH Seller.) The legally-mandated prioritization gives Priority Warning a distinct advantage. This regime does lead to surplus loss, due to Buyer's inability to distinguish between LL and LH Sellers who both disclose L on D1. The uncertainty drives Buyer to purchase from a Seller who discloses L on D1 if and only if the cost, and price, are below the threshold $\frac{1}{2} \cdot x_2$. Unable to separate the LL from the LH Sellers, Buyer inefficiently purchases the product from LL Sellers when $c < \frac{1}{2} \cdot x_2$, for a negative expected payoff of $-\frac{1}{4}c$; and Buyer inefficiently fails to purchase the product from LH Sellers when $c \in \left[\frac{1}{2} \cdot x_2, x_2\right)$, forgoing a positive expected payoff of $\frac{1}{4}(x_2 - c)$. Both distortions occur when $c < x_2$.

Priority Warning is superior to the two disclosure rules examined earlier, the (regular) Warning and the Full Disclosure regimes. Under the (regular) Warning regime, Buyers can only know with certainty when they are facing HH Sellers, whereas the Priority Warning allows them to also identify HL Sellers. Thus, the distortion under Priority Warning is smaller and occurs in a narrower band of *c* values. Priority Warning is also superior to Full Disclosure, which does not allow Buyers to identify any Seller with certainty. As in the symmetric-dimensions case, Full Disclosure is superior to (regular) Warning for some *c* values; but it is always inferior to Priority Warning.

3.1.5 Caveat Emptor v. Priority Warning

In the asymmetric-dimensions case, the Caveat Emptor regime with its inherent prioritization of D1 is symmetric to the Priority Warning regime with its legally-mandated priority of D1. Caveat Emptor enjoys an advantage when the relative cost, and price, of the product are low, and Priority Warning enjoys an advantage when the relative cost, and price, of the product are high. Intuitively, the distortions under Caveat Emptor are due to Buyer's inability to distinguish HH from HL, and this uncertainty is irrelevant when the price is low and Buyer makes the efficient decision to purchase from both of these Seller types. The distortions under Priority Warning are due to Buyer's inability to distinguish LL from LH, and this uncertainty is irrelevant when the price is high and Buyer makes the efficient decision not purchase from both of these Seller types.

When quality dimensions are asymmetric, featuring a primary and a secondary dimension, the welfare loss due to Buyers' asymmetric information decreases. Recognizing the inherent or the legally-mandated priority in disclosure—namely, that Sellers are either eager or required to disclose information about the primary dimension—Buyers are able to infer more information about the undisclosed dimension, and are thus less vulnerable to manipulation. Some distortion remains: it occurs under Caveat Emptor when the cost, and price, are high (above x_1); and it occurs under Priority Warning when the cost, and price, are low (below x_2). Caveat Emptor is superior when $c < x_2$; and Priority Warning is superior when $c > x_1$. Thus, 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.

The asymmetric case also exhibits an intermediate range, in which $c \in [x_2, x_1]$. In this range, 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.[30] See Figure 3 below.

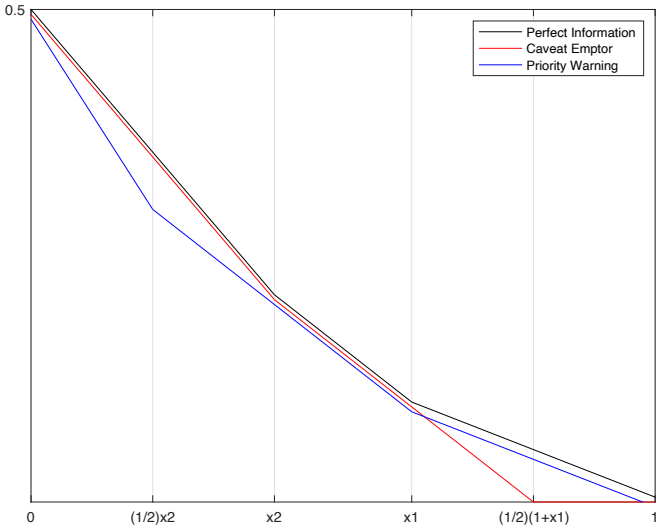

Figure 3

The Principle of Information Priority (PIP), which we introduced in Section 2 as a criterion for the optimal legal intervention, has a new valance in the asymmetric-dimensions case. In the asymmetric case, PIP has "more room" to operate. Here, PIP tells us not only when information about a high dimension is more important than information about a low dimension, or vice versa, but also which high dimension, or which low dimension, should be prioritized. In the hands-off, Caveat Emptor regime, Sellers voluntarily implement this inter-dimensional prioritization (albeit only for HH products). The Priority Warning regime implements inter-dimensional prioritization

[30] 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.

by legal design (albeit only for LL products). PIP also guides the optimal regime choice. In the baseline, symmetric-dimensions case, we said that a conditional legal rule "warning if and only if $x < c$" eliminates the distortions and achieves the perfect information benchmark. In the asymmetric case, this criterion becomes "warning only if $x_1 < c$." (PIP would work equally well if the criterion is expanded, "warning only if $x_2 < c$," but expanding the warning requirement is not necessary to achieve the perfect information benchmark.) Namely, the law should mandate a (prioritized) warning only when the cost, and price, exceed the value of the more important dimension. The inherent prioritization under Caveat Emptor reduces the need for a warning mandate.

Finally, if a conditional rule based on PIP is infeasible and either Caveat Emptor or Priority Warning have to be selected, the values of x_1 and x_2 would determine the optimal one-size rule. Caveat Emptor creates a distortion in the upper interval, $c \in (x_1, 1]$, and as x_1 draws closer to 1 the distortion shrinks. Priority Warning creates a distortion in the lower interval $c \in [0, x_2)$, and as x_2 draws closer to 0 the distortion shrinks. A very large x_1 makes Caveate Emptor the better one-size rule; a very low x_2 makes Priority Warning the superior rule. But more can be said.

In the baseline, symmetric-dimensions model, PIP tracked the complements v. substitutes distinction. A low *x* represented the complements case, where the Warning regime was optimal; and a high *x* represented the substitutes case, where Caveat Emptor was optimal. In the asymmetric case, PIP requires some fine-tuning of 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.

3.2 Investment in Quality

3.2.1 Informal Summary

The analysis in Section 2 made the simplifying assumption that the product quality is exogenously given, and that the share of products with different quality combinations is fixed and predetermined. This assumption allowed us to focus on ex post allocative efficiency and to show how rational buyers can be manipulated. But the assumption of exogenous quality is unrealistic and limiting. Sellers invest in quality, when higher-quality products generate higher profits. When manipulation is possible, sellers will prefer to invest in certain quality dimensions and not in others; and the profit-maximizing quality investment might not align with the welfare-maximizing choice.

We now expand the Section 2 framework, in two steps. First, we relax the assumption that cost is independent of product quality and adopt the more reasonable assumption that the cost of producing a good with two high-quality dimensions is higher than the cost of producing a good with one high-quality dimension, which in turn is higher than the cost of producing a good with

two low-quality dimensions.[31] In this first step, however, we retain the assumption of exogenous quality. The move to quality-dependent costs does not change our analysis or results, as long as quality remains exogenous.[32]

The second, and more critical, step is to introduce endogenous quality. We allow sellers to invest in quality, and let these investment decisions determine the distribution of quality levels in the market. We return to the symmetric-dimensions scenario. To provide incentives to invest in higher quality, we must allow sellers to capture at least some of the benefit from their investment. Therefore, we depart from Section 2's perfect competition assumption, and instead assume that sellers have enough bargaining power to set prices that expropriate buyers' entire expected value from the product. Adding this second step – the endogenous quality step – significantly changes the analysis.

Under *Caveat Emptor*, the key new insight is that sellers would never produce products with two high-quality dimensions. As we have seen, under Caveat Emptor buyers cannot distinguish between products with one versus two high-quality dimensions. Since buyers are able to review information on only one quality dimension, 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. Rational buyers will realize that there will be no products with two high-quality dimensions, and thus if they observe a disclosure of a high-quality dimension they will correctly infer that the other dimension is low quality. Therefore, there is no manipulation in equilibrium. But while there is no manipulation, there will still be a welfare loss, whenever it is socially desirable to produce products with two high-quality dimensions. We see that buyers' fear of manipulation, and the rational prudence they exercise, destroys the incentives of sellers to invest in high quality on both dimensions, and, ultimately, guarantees that buyers will not be manipulated. Buyers are now hurt by being sold lower quality products.

A symmetric distortion occurs in the *Warning* regime. Here, buyers are unable to distinguish between products with one versus two low-quality dimensions. Both of these products would be accompanied by a warning, buyers will recognize that one of the dimensions is low-quality and remain unsure about the other. Recognizing this, sellers will never offer a product with one high-quality dimension. 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-

[31] The Section 2 analysis assumed a uniform price, which was equal to the uniform cost. Here, to retain the assumption that the price does not vary with product quality, we assume a (uniform) price that is set above the cost of producing a good with two high-quality dimensions, in accordance with sellers' relative market power. 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. An alternative approach, which we explore below, gives sellers all of the bargaining power, so that they can set a price that extracts 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.

[32] The move to quality-dependent costs qualifies the Section 2 analysis in one regard: We have seen that buyers might decide on a purchase that leaves them with a net ex post loss. In the Section 2 analysis, a loss for buyers was necessarily also a social loss. Here, we assume that prices are above cost, and thus buyers' ex-post loss may not be a social loss. With perfect information, above-cost prices result in insufficient purchasing. The additional purchases caused by the informational asymmetry mitigate the insufficient purchasing problem; buyers' loss is offset by sellers' gain. This qualification applies to the analysis of all legal regimes. A more detailed analysis of this extension is provided in the Appendix.

high, or all-low. Buyers will not be manipulated, 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.

While this analysis reveals a new source of productive inefficiency, 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 in 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 guarantees the existence of such products.

3.2.2 Revised Formal Framework

We employ the setup from Section 2, subject to the following modifications: We allow sellers to choose the quality of their product (on both dimensions), thus replacing the exogenous division of Sellers into the four, equal-sized groups – LL, LH, HL and HH – with an endogenous division. Specifically, we let ρ_{LL}, ρ_{LH}, ρ_{HL} and ρ_{HH} represent the shares of LL, LH, HL and HH Sellers, respectively. The cost, to Seller, of producing LL, LH, HL and HH products product is c_{LL}, c_{LH}, c_{HL} and c_{HH}, respectively. 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 highlights incentives to invest, we adopt the opposite assumption: We give Sellers all the bargaining power, such that $p = \hat{u}$.[33] 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 (HH, HL, LH and LL). Under Caveat Emptor, if Seller remains silent, Buyer will know that she is facing an LL Seller. And if Seller discloses High, Buyer will know 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 High, Buyer will know that she is facing an HH Seller. An if Seller discloses Low, Buyer will know 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 $\frac{\rho_{1H}}{\frac{1}{2}\rho_{LL}+\rho_{1H}}$.

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

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; namely, all sellers are HH Sellers. 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; namely, all sellers are HL or LH Sellers. And when $max(1 - c_{HH}, x - c_{1H}, 0) = 0$, Sellers do not invest in quality; namely, all sellers are LL Sellers. In all three scenarios, the first best welfare level is achieved.

Caveat Emptor. Under Caveat Emptor, the consumer cannot distinguish between High on two quality dimensions (HH Sellers) and High on one quality dimension (HL or LH Sellers). Therefore, Sellers will have no incentive to ensure High on both dimensions, and we have: $\rho_{HH} = 0$.[34] In a rational-expectations equilibrium, Buyers will realize that, when they observe a High disclosure, it must be from an HL or 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 a 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 $max(1 - c_{HH}, x - c_{1H}, 0) = 0$.

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. Here, the rationality of Buyers guarantees that they will not be misled, and instead reduces quality.[35] 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 Low on two quality dimensions (LL Sellers) and Low on one quality dimension (HL or 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: $\rho_{1H} = 0$.[36] In a rational-expectations equilibrium, Buyers will realize

---

[34] 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: $\rho_{HH} = 0$.

[35] 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"). Of course, in the real-world, complete unravelling is unlikely to occur, and there will be manipulation in equilibrium.

[36] 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

that, when they observe a Low disclosure, 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, because Buyers will not pay more for these products than for HL or LH products. If HH products create the largest net benefit, then Caveat Emptor is inefficient. Sellers will choose between HL or LH products and LL products. If HL, LH or LL products create the largest net benefit, then Caveat Emptor is efficient. In the Warning regime, Sellers will not invest in high quality on only one dimension, because Buyers will not pay more for HL or LH products as compared to LL products. If HL or LH products create the largest net benefit, then Warning is inefficient. Sellers will choose between HH products and LL products. If HH or LL products create the largest net benefit, then Warning is efficient.

We see that both Caveat Emptor and Warning achieve the optimal outcome, when LL products create the largest net benefit. 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, then Warning is the better rule. If it is cost-effective to have high quality on only one dimensions, then Caveat Emptor is the better rule.

3.3 Multiple 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 more than two dimensions. It allows us to examine how the extent of buyers' uncertainty affects the problem of manipulation. 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: We allow sellers to determine the number of quality dimensions, and make the product more or less complex for buyers to review. We show that some sellers will have incentives to create spurious complexity in order to profit from manipulation, and study the effects of the different legal regimes on these incentives.

The static question. What happens when the number of dimensions increases, while buyers' ability to review information remains unchanged? Unsurprisingly, things get worse for buyers. The underlying problem is buyers' incomplete information, and this problem is aggravated when the complexity of products increases. With more dimensions, the information content of the

---

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: $\rho_{1H} = 0$.

disclosures goes down. As a result, 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 show 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 suggests that low-quality sellers may want to increase the number of dimensions, when they have a good chance of producing high quality on the new dimension. 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 both the old and new 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 both the old and new dimensions.

3.3.2 The Static Question: More Dimensions, More Harm

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 the zero utility when all dimensions are Low and the utility of one when all dimensions are High); $x \in \left[0, \frac{1}{2}\right]$.

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. We thus have eight groups of sellers—HHH, HHL, HLH, HLL, LHH, LHL, LLH, LLL—each comprising 12.5% of sellers.

Under Caveat Emptor, Seller will disclosure 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 $\frac{3}{7}$.[37] A Buyer who

[37] Applying Bayes' rule, we have:

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 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 $c \in \left[2x, \frac{1}{7} \cdot (1 + 9x)\right)$, 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 $c \in$

$$Pr(q_1 = q_2 = q_3 = 1 | \text{B observes } q_1 = 1) = \frac{Pr(q_1 = q_2 = q_3 = 1) \cdot Pr(\text{B observes } q_1 = 1 | q_1 = q_2 = q_3 = 1)}{Pr(\text{B observes } q_1 = 1)}$$
$$= \frac{\frac{1}{8} \cdot \frac{1}{3}}{\frac{1}{8}(\frac{1}{3} + \frac{1}{2} + \frac{1}{2} + 1)} = \frac{1}{7}$$

$[x, \frac{1}{7} \cdot (1+9x))$, then Caveat Emptor induces inefficient purchases from Sellers with only one High dimension. If $c \in [\frac{1}{7} \cdot (1+9x), 2x)$, 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 $c \in [\frac{1}{7} \cdot 9x, 1)$; 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 $c \in [\frac{1}{7} \cdot 9x, 2x)$, 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.

3.3.3 The Dynamic Question: Sellers' Incentives to Increase the Number of Dimensions

We have seen that a larger number of quality dimensions can increase the social cost of manipulation. 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 x denote the value of High on the single dimension, we see that H Sellers get $\pi = x - c$ if $c < x$, and L sellers get $\pi = 0$.

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 of H Sellers who did not add a new dimension, a share ρ_L of L Sellers 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.[38] Therefore, Buyer can identify H Sellers and L Sellers with certainty.

Under Caveat Emptor, if Seller remains silent, Buyer will know that she is 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 $\frac{\rho_{LH}}{\rho_{HH}+\rho_{HL}+\rho_{LH}}$. The expected profit of HH, HL and LH Sellers is: $\pi = \frac{\rho_{HH}}{\rho_{HH}+\rho_{HL}+\rho_{LH}} \cdot 1 + \frac{\rho_{HL}+\rho_{LH}}{\rho_{HH}+\rho_{HL}+\rho_{LH}} \cdot x - c$ if $c < \frac{\rho_{HH}}{\rho_{HH}+\rho_{HL}+\rho_{LH}} \cdot 1 + \frac{\rho_{HL}+\rho_{LH}}{\rho_{HH}+\rho_{HL}+\rho_{LH}} \cdot x$.

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: $\pi = x - c$; and an HH Seller's expected profit will be: $\pi = 1 - c$. 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 $\frac{\rho_{LH}}{\rho_{LL}+\rho_{HL}+\rho_{LH}}$. The expected profit of LL, HL and LH Sellers is: $\pi = \frac{\rho_{HL}+\rho_{LH}}{\rho_{LL}+\rho_{HL}+\rho_{LH}} \cdot x - c$ if $c < \frac{\rho_{HL}+\rho_{LH}}{\rho_{LL}+\rho_{HL}+\rho_{LH}} \cdot x$.

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 $\frac{1}{2}$).[39]

[38] 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.

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

Caveat Emptor. 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 $\pi = \frac{\rho_{HH}}{\rho_{HH}+\rho_{HL}+\rho_{LH}} \cdot 1 + \frac{\rho_{HL}+\rho_{LH}}{\rho_{HH}+\rho_{HL}+\rho_{LH}} \cdot x - c$, if $c < \frac{\rho_{HH}}{\rho_{HH}+\rho_{HL}+\rho_{LH}} \cdot 1 + \frac{\rho_{HL}+\rho_{LH}}{\rho_{HH}+\rho_{HL}+\rho_{LH}} \cdot x$. Without adding a new dimension, L Seller gets zero. From a social perspective, looking at the Buyers who are matched with this Seller: If $c < x$, the resulting transactions generate welfare of $\frac{1}{2} \cdot (x - c)$ if a new dimension is added, rather than zero if Seller sticks with one dimension. Therefore, L Sellers have a socially excessive incentive to add a new dimension. If $c > x$, the resulting transactions generate welfare of zero, as when Seller sticks with one dimension. When $c \in [x, x\frac{\rho_{HH}}{\rho_{HH}+\rho_{HL}+\rho_{LH}} \cdot 1 + \frac{\rho_{HL}+\rho_{LH}}{\rho_{HH}+\rho_{HL}+\rho_{LH}} \cdot x\Big)$, L Sellers have a socially excessive incentive to add a new dimension.

If an H Seller adds a dimension: With probability $\frac{1}{2}$, he will become HH; and with probability $\frac{1}{2}$ he will become HL. Either way he will disclose High and get $\pi = \frac{\rho_{HH}}{\rho_{HH}+\rho_{HL}+\rho_{LH}} \cdot 1 + \frac{\rho_{HL}+\rho_{LH}}{\rho_{HH}+\rho_{HL}+\rho_{LH}} \cdot x - c$, if $c < \frac{\rho_{HH}}{\rho_{HH}+\rho_{HL}+\rho_{LH}} \cdot 1 + \frac{\rho_{HL}+\rho_{LH}}{\rho_{HH}+\rho_{HL}+\rho_{LH}} \cdot x$. Without adding a new dimension, H Seller gets $\pi = x - c$. So, the private benefit from adding a dimension is: $\frac{\rho_{HH}}{\rho_{HH}+\rho_{HL}+\rho_{LH}} \cdot (1 - x)$. From a social welfare perspective, looking at the Buyers who are matched with this Seller: If $c < x$, the resulting transactions generate welfare of $\frac{1}{2} \cdot (1 - c) + \frac{1}{2} \cdot (x - c)$, rather than $x - c$ without the new dimension. The social benefit is $\frac{1}{2} \cdot (1 - x)$. Therefore, an H Seller has insufficient incentives to add a dimension, if $\frac{\rho_{HH}}{\rho_{HH}+\rho_{HL}+\rho_{LH}} < \frac{1}{2}$; and an excessive incentive, if $\frac{\rho_{HH}}{\rho_{HH}+\rho_{HL}+\rho_{LH}} > \frac{1}{2}$. Since L Sellers have the stronger incentive to add dimensions, it is likely that ρ_{LH} will be larger, and so H Sellers will have an insufficient incentive to add dimensions. If $c > x$, the resulting transactions generate welfare of $\frac{1}{2} \cdot (1 - c)$, rather than zero. The social benefit is $\frac{1}{2} \cdot (1 - c)$, which means, once again, that an H Seller has insufficient incentives to add a dimension.

Warning. If an L Seller adds a new dimension, then: With probability $\frac{1}{2}$, he will become LL; and with probability $\frac{1}{2}$, he will become LH. Either way he will disclose Low and get $\pi = \frac{\rho_{HL}+\rho_{LH}}{\rho_{LL}+\rho_{HL}+\rho_{LH}} \cdot x - c$, if $c < \frac{\rho_{HL}+\rho_{LH}}{\rho_{LL}+\rho_{HL}+\rho_{LH}} \cdot x$. Without adding a new dimension, L Seller will get zero. From a social welfare perspective, looking at the Buyers who are matched with this Seller: If $c < x$, the resulting transactions generate welfare of $\frac{1}{2} \cdot (x - c)$, rather than zero without the new dimension. The social benefit is $\frac{1}{2} \cdot (x - c)$. An L Seller will have excessive incentives to add a dimension, when $c < \frac{\rho_{HL}+\rho_{LH}-\rho_{LL}}{\rho_{LL}+\rho_{HL}+\rho_{LH}} \cdot x$; and an insufficient incentive, when $c \in \left[\frac{\rho_{HL}+\rho_{LH}-\rho_{LL}}{\rho_{LL}+\rho_{HL}+\rho_{LH}} \cdot x, x\right)$. If $c > x$, the resulting transactions generate welfare of zero, as when Seller sticks with one dimension. An L Seller will have socially optimal incentives to add a dimension.

If an H Seller adds a dimension: With probability $\frac{1}{2}$, he will become HH, disclose High, and get $\pi = 1 - c$; and with probability $\frac{1}{2}$ he will become HL, disclose Low, and get $\pi = \frac{\rho_{HL}+\rho_{LH}}{\rho_{LL}+\rho_{HL}+\rho_{LH}} \cdot$

$x - c$ if $c < \frac{\rho_{HL}+\rho_{LH}}{\rho_{LL}+\rho_{HL}+\rho_{LH}} \cdot x$. Without adding a new dimension, H Seller will get $x - c$. The private benefit from adding a dimension is: $\frac{1}{2} \cdot (1 - x) + \frac{1}{2} \cdot \left(-\frac{\rho_{LL}}{\rho_{LL}+\rho_{HL}+\rho_{LH}} \cdot x\right)$. From a social welfare perspective, looking at the Buyers who are matched with this Seller: If $c < x$, the resulting transactions generate welfare of $\frac{1}{2} \cdot (1 - c) + \frac{1}{2} \cdot (x - c)$, rather than $x - c$ without the new dimension. The social benefit is $\frac{1}{2} \cdot (1 - x)$, greater than the private benefit. H Sellers have socially insufficient incentives to add a dimension. If $c > x$, the resulting transactions generate $\frac{1}{2} \cdot (1 - c)$, identical to the social benefit. Here, H Sellers have optimal incentives to add a dimension.

Comparison. If $c < x$, L Sellers have socially excessive incentives to add dimensions under Caveat Emptor; under Warning their incentives to add a dimension can be either excessive (for lower cost products) or insufficient (for higher cost products). H Sellers have socially insufficient incentives to add dimensions under Caveat Emptor (probably) and under Warning (definitely). If $c > x$, L Sellers have socially excessive incentives to add dimensions under Caveat Emptor; under Warning their incentives to add a dimension are socially optimal. H Sellers have socially insufficient incentives to add dimensions under Caveat Emptor; under Warning their incentives to add a dimension are socially optimal.

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 Revisiting the Rationality Assumption

The analysis in this paper departed from much of the academic literature on consumer manipulation by assuming rationality. It is, of course, well documented that people display various forms of irrational conduct and that firms exploit these cognitive lapses. The assumption of rationality enabled us to show that problems akin to behavioral manipulations arise even if people are sophisticated, and even if they know that they are being manipulated. It is not necessary for people to make mistakes or draw naïve, irrational inferences, nor is it necessary to base legal intervention on the showing of a particular cognitive bias. Given the complexity of the decisions and choices people make, they have to allocate their attention wisely and sparingly, thus remain consciously uncertain about some product features. It is this uncertainty that sellers can exploit and that the law can try to redress.

Our rational choice approach not only expands the domain of manipulation; it also generates novel insights: We have seen that the same legal intervention can sometimes improve people's decisions, but other times disrupt them. In the behavioral literature, the typical solution to manipulation is a prominent disclosure, a warning. We showed that this solution works only under certain conditions. It could fail, for two reasons: the deficiencies warned about are not a priority; and the warnings accumulate. By identifying the complements-versus-substitutes dichotomy, our analysis provided guidance how to use warnings only when they are valuable.

Our analytical framework can be extended to allow for imperfect rationality and behavioral biases. 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. The uncertainty about the undisclosed dimensions leads to both excessive and insufficient purchases, but the rational, Bayesian value estimate optimally balances the two distortions to minimize consumer harm. In contrast, 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 preceding comparison between rational and 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.[40] 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.

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

To illustrate, consider the privacy disclosures typically appearing as prominent upfront warnings upon entry into websites.[41] 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 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.[42]

Disclaimers. Another lesson from the analysis concerns the impact of disclaimers—mandated disclosures that accompany and qualify sellers' claims of high quality.[43] 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 low-quality 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.[44]

---

[41] 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).

[42] These data-privacy entry warnings create an additional problem highlighted by our analysis of asymmetric dimensions (Sec. 3.1). Prioritizing data-privacy concerns harms users, when the service includes other low-quality attributes with a greater negative impact. Having attended to the data-privacy warnings and having decided to go forward, users might not spend time they otherwise would on examining other low-quality aspects. When the quality dimensions are complements, the law should prioritize disclosure of the dimension that has the greatest negative impact on value.

[43] 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).

[44] 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.

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. (Of course, a price dimension would be regarded as high-quality when the price is low; and a price dimension would be regarded as low-quality when the price is high.) 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 prices 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, both quality and price dimensions have to be high valued 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 satisfaction 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.[45]

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 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.

[45] 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 food safety and 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.[46] 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 and only 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. The impact that other inexpensive dimensions have on the score should depend on whether they are substitutes or complements to the expensive component.

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

The underlying problem is information asymmetry in consumer markets. Scores are one possible method for reducing information asymmetry. Other market solutions include reputational mechanisms, ratings systems and quality certification. We acknowledge the value of these market solutions. And, yet, as with scores, other market solutions do not eliminate the information

[46] 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).

asymmetry problem. Indeed, a large literature has convincingly demonstrated the prevalence of asymmetric information in consumer markets (see, e.g., [CITE]), despite the mitigating effect of various market solutions. It bears emphasis that we do not purport to add to this literature and 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, we limit the message space, i.e., what sellers can disclose (and what regulators can force sellers to disclose): we allow sellers to disclose either High or Low. Perhaps a broader message space could solve the information asymmetry problem. For example, under Caveat Emptor, the concern was that buyers cannot distinguish between HH Sellers and HL/LH Sellers, since any Seller with at least one High dimension would disclose High. What if an HH Seller can separate itself by communicating "nothing to warn about" (rather than disclosing High)? Such a disclosure would eliminate the asymmetric information problem. But this solution becomes less feasible as the number of quality dimensions increases. If there are, say, 15 dimensions and a seller with 14 High dimensions wants to distinguish itself from a seller with 13 High dimensions, the "nothing to warn about" statement no longer works (recall that sellers cannot lie). To affect the desired separation, the seller with 14 High dimensions would need a more complex communication, which would be costly for buyers to review.

A second, market-based solution would allow buyers to be more strategic in selecting the disclosures that they review. If the quality dimensions are substitutes, such that information about high-quality attributes is more important, buyers could ignore mandated warnings and focus on sellers' voluntary disclosures about high-quality attributes. For example, some internet users may be able to ignore GDPR-madated warnings. Similarly, if quality dimensions are complements, such that information about low-quality attributes is more important, buyers could ignore sellers' (voluntary) disclosures about the products high-quality attributes and focus on the mandated warnings. For example, some consumers may be able to look past sellers' claims on packaged foods and focus on the nutrition table. And yet it is not always poaaible, or easy, to ignore some information and focus on other information. Indeed, it is often unclear (before reading) if a disclosure is a volunarilty provided High disclosure or a legally mandated Low disclosure, 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 content of the disclosure.) Also, if buyers were able to ignore sellers' voluntary disclosures just by their "look" without having to read them, then sellers would change the format to make their disclosures look like mandated warnings. For example, consumers often get unsolicited mail that looks official – formatted to undermine consumers' attempts to filter seller-provided information without reading it. Or, to take another example, looking at the labels attached to dietary supplements, it is often hard to distinguish between the mandated disclosures and the voluntarily provided disclosures. Finally, considering video ads, it is impractical to mute the ad and listen only to the disclaimer.

Another market solution invokes product warranties. Warranties are a well-known solution to the asymmetric information problem (see Grossman 1981). We acknowledge the important role of warranties; and only note their more limited impact when products are multi-dimensional. Indeed, we suggest that a warranty should be considered as a quality dimension itself. Moreover, a warranty can itself be multi-dimensional, when we consider the scope of the warranty – what parts or defects are covered and for how long. Finally, our model allows only for disclosue by a seller about the quality of that seller's product. In some cases, 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—an outcome that 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 information v. bad information—that should be prioritized for disclosure. Throughout the paper we gave examples for products that fall into each of the two categories. In our baseline model, the two quality dimensions could be either complements or substitutes. With more than two dimensions, we can have some subsets of dimensions that are substitutes and others that are complements. A restaurant may offer a complex menu with items that are substitutes, but it has other aspects—service quality, ambience, food safety—that are complements (among themselves and vis-à-vis the menu items).

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 now define 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 on 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), some buyers view D1 to be the more important dimension, whereas for other buyers D2 may be the more important dimension. 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 choice—whether to demand a warning about the low-quality dimension or to permit prioritized disclosure of the high-quality dimension—could depend on which group of buyers has lower information costs. We should prioritize information that is most critical to buyers who typically acquire less information. Other buyers can make up for the less-than-ideal disclosure by exerting more effort to review other dimensions.[47] If both groups of buyers have the same information-acquisition costs, then policymakers should follow a majoritarian approach and choose the legal regime that protects the larger group. Similarly, if the larger group of buyers cares more about D1, a Priority Warning rule should mandate disclosure of Low on D1 (if D1 is low-quality). Buyers who care more about D2 would recognize that D1 is prioritized for disclosure and would draw the same inferences, but receive less protection, since there will remain uncertainty about the dimension that matters to them more.

[47] 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.

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.

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 the one hand, courts are the right place to issue product-specific commands. 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 could thwart the goal of deprioritizing less critical information. 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.

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).

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

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.

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.

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

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." *The 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." *The 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.

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

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.