Persuasion through Selective Disclosure: Implications for Marketing, Campaigning, and Privacy Regulation*

Florian Hoffmann† Roman Inderst‡ Marco Ottaviani§

November 2014

Abstract

By collecting personalized data, firms and political campaigners (senders) are able to better tailor their communication to the preferences and orientations of individual consumers and voters (receivers). This paper characterizes equilibrium persuasion through selective disclosure of the information that senders acquire about the preferences of receivers. We derive positive and normative implications depending on: the extent of competition among senders, whether receivers are wary of the senders' incentives to become better informed, whether firms are able to personalize prices, and whether receivers make individual or collective decisions. We find that privacy laws requiring senders to obtain consent to acquire information are beneficial when there is little or asymmetric competition among firms or candidates, when receivers are unwary, and when firms can price discriminate. Otherwise, policy intervention has unintended negative welfare consequences.

Keywords: Persuasion, selective disclosure, hypertargeting, limited attention, privacy regulation.

JEL Classification: D83 (Search; Learning; Information and Knowledge; Communication; Belief), M31 (Marketing), M38 (Government Policy and Regulation).

---

*This paper supersedes and generalizes the analysis contained in an earlier draft circulated under the title: "Hypertargeting, Limited Attention, and Privacy: Implications for Marketing and Campaigning". Hoffmann gratefully acknowledges financial support from the ERC Advanced Grant 22992, Ottaviani from the ERC Advanced Grant 295835, and Inderst from the DFG through the Research Unit "Design & Behavior" (FOR 1371).

†University of Bonn. E-mail: fhoffmann@uni-bonn.de.

‡Johann Wolfgang Goethe University Frankfurt and Imperial College London. E-mail: inderst@finance.uni-frankfurt.de.

§Bocconi University and IGIER, Milan. E-mail: marco.ottaviani@unibocconi.it.

"In the old days, everyone—Democrats, Republicans, enthusiasts, nonvoters and undecideds—saw the same television ads. Now the campaigns use 'big data' to craft highly customized and even personalized messages as people go from website to website. The campaigns test just the right ads for each voter. . . . A wealthy urban liberal sees different ads online than a working-class centrist. People who care more about jobs see different ads than people who focus on social issues." *L. Gordon Crovitz, How Campaigns Hypertarget Voters Online, Wall Street Journal, November 4, 2012.*

1 Introduction

Firms and political candidates have traditionally had two distinct ways to persuade consumers and voters. First, they could broadcast their messages through old media (leaflets, billboards, newspapers, and television), thereby achieving only a coarse segmentation of the audience, mostly along channel types and regional boundaries.[1] Alternatively, they could customize their communication strategies with direct marketing and ground-game campaigning aimed at persuading single individuals or small groups. In order to implement this second strategy, firms and candidates could hire experienced salespeople or skilled campaigners to gather critical knowledge about their audiences, making it possible to tailor their messages using face-to-face contacts or canvassing.

Nowadays, the greater availability of personally identifiable data on the internet blurs the distinction between these two traditional communication strategies. Developments in computer technology increasingly allow sellers and campaigners to systematically collect personal and detailed data about an individual's past purchasing behavior, browsing activity, and credit history, as well as the personal likes and dislikes the individual shares on social networking sites.[2] When conducting what might appear to be an impersonal

[1] For an example of coarse segmentation of the advertising audience, consider Wrangler jeans' advertising strategy: ". . . while their main product is largely the same in feel, quality and color (commercial denim is pretty much commercial denim), the experience they market that product reflecting or emulating is very different across the globe. . . . In the US, Wrangler is perceived as the cowboy brand, the go-to jeans for the workers of the Midwest. These people aren't fashionable—and they frankly don't care. . . . This is why their tagline is 'Real. Comfortable. Jeans.' And Wrangler gives them that experience by associating their denim with their lifestyle through advertising. If you look at the European and Asian sites, you'll see a very different experience. While still 'American,' the brand is giving people in those countries a taste of the limitless American freedom of an open road and exploration. . . . You'll also notice that the tagline is 'Worn Across America'—as if the jeans represent the movers and shakers and nomads on the go and exploring the last frontier of the American West." *http://philliphess.tumblr.com/post/4236832814/listen-here-sweetheart-its-all-about-marketing*

[2] Information can be either collected directly or acquired from search engines and specialized data vendors. In its privacy policy, Facebook writes: "We allow advertisers to choose the characteristics of users who will see their advertisements and we may use any of the non-personally identifiable attributes we have collected (including information you may have decided not to show to other users, such as your

transaction through the internet, a great deal of personal information may be used to finely target consumers and voters. For example, on Facebook ski resorts advertise family activities to married users with kids, but stress snowboarding and party options to younger users interested in winter sports. Behavioral targeting or hypertargeting along these lines combines features of remote broadcasting with features of personal selling/campaigning.[3]

Concerns are often raised that some consumers and voters might suffer if they remain blithely unaware of the ability of firms and candidates to collect information and communicate selectively. An active debate is underway among policymakers about reforming the regulatory framework for consumer privacy with an emphasis on the collection and use of personal data on the internet. While in this area the U.S. currently relies mostly on industry self regulation, policymakers and Congress are considering stricter regulation of consumer privacy.[4] In recent years, European legislators have intervened more directly by raising barriers to the collection and use of personally identifiable data about past purchases or recent browsing behavior, including a requirement that firms seek explicit consent to collect information.[5] The prevailing presumption—see Shapiro and Varian (1997)—is that efficiency is achieved by granting property rights over information to consumers, for example by requiring consumer consent.

We challenge this view by analyzing a model in which a set of firms or candidates (senders) attempt to persuade a set of consumers or voters (receivers). The model is designed to evaluate the private and social incentives for information collection and selective communication. In the absence of regulation, senders privately choose whether to acquire better private information about receiver preferences and then attempt to persuade by

birth year or other sensitive personal information or preferences) to select the appropriate audience for those advertisements." *https://www.facebook.com/note.php?note_id=+322194465300*

[3] "Tailor your ads and bids to specific interests: Suppose you sell cars and want to reach people on auto websites. You believe that the brand of cars you sell appeals to a wide variety of people, but some of them may react more positively than others to certain types of ads. For example, . . . you could show an image ad that associates a family-oriented lifestyle with your car brand to auto website visitors who're also interested in parenting." *Google AdWords, http://support.google.com/adwords/answer/2497941?hl=en*

[4] See American Association of Advertising Agencies (2009) for a widely adopted set of self-regulatory principles for online behavioral advertising. On the U.S. policy debate, see White House (2012), Federal Trade Commission (2012), and the extensive discussion of the Do Not Track legislation proposals on wikipedia.

[5] See the Data Protection Directive (1995/46/EC) and the Privacy and Electronic Communications Directive (2002/58/EC), also known as the E-Privacy Directive, which regulates cookies and other similar devices through its amendments, such as Directive 2009/136/EC, the so-called EU Cookie Directive, and the Privacy and Electronic Communications (EC Directive) (Amendment) Regulations 2011. The current prescription is that "cookies or similar devices must not be used unless the subscriber or user of the relevant terminal equipment: (a) is provided with clear and comprehensive information about the purposes of the storage of, or access to, that information; and (b) has given his or her consent." More recently, European authorities have been pressuring internet giants such as Facebook and Google to limit the collection of personal data without user consent.

selectively disclosing information about their horizontally differentiated offerings. Thus, our baseline model of equilibrium persuasion hinges on three key features:

(1) Senders decide whether to acquire better private information about receiver preferences.

(2) Receivers do not observe whether senders have become better informed.

(3) Senders and receivers play a disclosure game resulting in the selective revelation of the information that senders privately acquired. For illustration purposes, we cast the presentation within two classic specifications of selective disclosure:

   (a) In the first specification, each sender can increase the probability of becoming informed about a receiver's valuation of the sender's offering; then, the sender can disclose the value to the receiver. As it is well known from Dye (1985), Jung and Kwon (1998), and Shavell (1994), a sender who learns that the receiver has low values pools with the sender who does not become informed.

   (b) Alternatively, in the second specification—a continuous version of Fishman and Hagerty's (1990) model, similar to Glazer and Rubinstein (2004)—each sender decides whether to observe the receiver's valuation for two attributes of the sender's offering. However, the scope of communication is naturally restricted by factors such as airtime and screen space, or simply by the receiver's limited attention.[6] Given this limited attention, senders are able to disclose one of the two attributes, and thus select strategically the attribute to disclose so as to increase the chance of a favorable decision, as in the example about ski resorts reported in the second paragraph.

Our model of persuasion through selective disclosure is geared to address the following positive and normative questions in the economics of information privacy: Under which circumstances do senders profit from acquiring more information before selectively disclosing? Do receivers benefit as a result? What is the impact of policies aimed at protecting privacy, for example by requiring consumer consent?

[6] Senders might be unable to communicate all the attributes they know because of space or time constraints or simply because (too much) information "consumes the attention of its recipients" (Simon 1971). The limited capacity of individuals to process information is currently being investigated in a number of other areas, ranging from macroeconomics (e.g., Sims 2003) to organization economics (Dessein, Galeotti, and Santos 2012). In our model, it is the sender who must choose a particular attribute to disclose given the limitation of the communication channel, rather than the receivers having to choose how to optimally direct their limited attention and information processing capacity.

Given that additional information can be secretly acquired at no cost, we show that the sender has an incentive to become better informed. Holding fixed the receiver's expectations, a better informed sender is able to disclose information more selectively and thus induce an upward shift in the valuation of the receiver. A rational receiver should then anticipate that the sender has acquired more information, and accordingly adjust the inference for the increased selectivity in the resulting disclosure. The upshot is that rational receivers end up gaining from the additional information collected and selectively disclosed by the sender. But does the sender profit from more selective disclosure?

To answer this question, we compare the equilibrium resulting with unobserved information acquisition to the outcome that would result if the receiver were to observe the amount of information acquired by the sender. When feature (2) of our model does not hold, we show that the sender would want to commit to stay less informed when facing a receiver with a relatively favorable distribution of valuations, for example because competition is limited. In this sense, feature (2) introduces a commitment problem, akin to signal jamming à la Holmström (1999). Suppose that a policy is introduced requiring that the sender must obtain consent from the receiver to become better informed. By not asking for consent, the sender is therefore able to commit not to acquire information, to the final detriment of the receiver. Privacy regulation thus backfires through this channel.

Before turning to the applications, it is worth comparing the three key features of our baseline model to previous work on disclosure, persuasion, and information control:

- **Disclosure.** Because of feature (3), in our model lack of disclosure does not trigger complete unraveling, thus departing from the baseline models of Grossman (1981), Milgrom (1981), and Milgrom and Roberts (1986). Even when the receiver is fully rational, we do not obtain unraveling because (a) the sender is possibly uninformed in the first specification or (b) the receiver has limited attention and, thus, the sender can only disclose a subset (one "attribute") of the available information in the second specification.[7] Prior to the stage of selective disclosure, we add hidden information acquisition through features (1) and (2).

- **Persuasion.** The three key features of our model of *equilibrium* persuasion without commitment depart from *optimal* persuasion with commitment à la Rayo and Segal (2010) and Kamenica and Gentzkow (2011) where, instead, the sender (1′) chooses

[7] Our analysis abstracts away from externalities across the communication strategies of firms due to congestion effects and information overload; see, for example, Van Zandt (2004) and Anderson and de Palma (2012) for analyses in this direction using models à la Butters (1977). See also Johnson (2013) for a welfare analysis of the impact of targeted advertising in the presence of advertising avoidance by consumers.

the information structure in an unconstrained fashion, (2′) is able to fully commit to this choice, and (3′) thus completely discloses the information acquired. With these alternative assumptions, the sender always (weakly) profits from optimal persuasion, as shown by Kamenica and Gentzkow (2011). In our setting, instead, equilibrium persuasion without commitment hurts the sender when the receiver's outside option is relatively unattractive.[8] As in the case of optimal persuasion, wary receivers benefit from equilibrium persuasion in the baseline model with individual decisions at given prices; however, receivers can be hurt when prices are personalized as in our marketing application, when receivers make a collective decision as in our application to voting, or when receivers are unwary.

- **Information Control.** In our model, senders are able to control receivers' information indirectly by becoming better informed, our feature (1), and then by disclosing this information selectively, our feature (3), rather than directly and truthfully as in the literature on information control in markets à la Lewis and Sappington (1994), Johnson and Myatt (2006), and Ganuza and Penalva (2010).[9] In addition, senders in our model cannot commit to the information structure, our signal-jamming feature (2), and so can fall victim to their own incentives to secretly acquire more information; in these cases regulation helps senders achieve commitment, which in turn damages wary consumers.

Given that regulation is often motivated by the protection of unsophisticated consumers, we also analyze persuasion when receivers are not fully rational. Consider a sender facing a receiver who is unwary about the fact that the sender has acquired more information and thus discloses more selectively.[10] The sender is then able to bias the deci-

[8] Our analysis of equilibrium persuasion with multiple senders is particularly tractable given our focus on horizontal differentiation with independently distributed values. See Gentzkow and Kamenica (2012) for an analysis of optimal persuasion with multiple senders, when our three features are not satisfied. Bhattacharya and Mukerjee (2013) analyze strategic disclosure by multiple senders who share the same information; in our horizontal-differentiation model, instead, senders are endogenously informed about the values of their offerings, which are independently distributed. See also DellaVigna and Gentzkow (2010) for a survey of the literature on persuasion across economics, marketing, and political science.

[9] See also Kamenica, Mullainathan, and Thaler (2011) for a discussion of situations in which firms might know more about consumer preferences than consumers know themselves.

[10] In a disclosure setting in which the fraction of receivers who fail to update their beliefs following the lack of disclosure (analytical failure) is higher than the fraction of receivers who do not attend to the disclosure (cue neglect), Hirshleifer, Lim, and Teoh (2004) obtain an equilibrium in which the sender only discloses high realizations. Unwary consumers in our second specification, instead, attend to the disclosed attribute but fail to make the appropriate inference about the undisclosed attribute, which is chosen selectively by the sender. Thus, relative comparisons across different dimensions of information play a key role in this model. Relative comparisons across dimensions also play a role in the construction of cheap-talk equilibria by Chakraborty and Harbaugh (2007, 2010) and Che, Dessein, and Kartik (2013).

sion of the unwary receiver; at the same time, the receiver also benefits from the increased quality of information obtained. While in general there is a trade-off between the increase in bias and the increase in information, we exhibit natural examples in which the value of information trumps the bias.[11] Furthermore, we show that competition among senders in a symmetric setting has the power of completely eliminating the bias in the choices of unwary receivers, given that the bias induced by the different senders exactly cancel out under symmetry.

We then turn to an application to marketing. When firms are able to personalize prices, we show that firms always benefit from the increase in perceived differentiation that results from more selective disclosure. Wary consumers benefit only when competition is intense. Unwary receivers are exploited by price personalization if there is no competition; with competition, instead, unwary receivers underestimate differentiation and so end up paying strictly less than wary receivers. Our main conclusion is that hypertargeting—the collection and use of personally identifiable data by firms to tailor selective disclosure—benefits consumers when they are adequately protected by at least one of the following three conditions: their own wariness, symmetric competition, or the inability of firms to practice personalized pricing. A strong rationale for regulation emerges only when these three conditions are not met, that is, when few competitors exploit unwary consumers through personalized pricing. Otherwise, even seemingly light-touch regulation, such as requiring consumer consent to collect and use personal data, may backfire by giving firms a way to commit to avoid selective communication with wary consumers, who are made worse off as a result.

Finally, we extend the framework to settings in which receivers take a collective action, for example by voting for one of two candidates.[12] As in our baseline case with individual

[11] For example, the value of information exactly compensates for the bias when valuations are uniformly distributed. With exponentially distributed valuations, the value of information more than compensates the increase in bias, so even a naive receiver benefits from increased selectivity of disclosure.

[12] The setting with voting relates to the problem of persuading a group to take a collective decision considered by Caillaud and Tirole (2007); in our model voters cast their ballot simultaneously rather than sequentially. See Alonso and Câmara (2014) for a model of optimal (rather than equilibrium) persuasion of voters in a setting with vertical (rather than horizontal) differentiation.
Glaeser, Ponzetto and Shapiro (2005) also consider a model of targeted campaigning. There, in order to motivate their core supporters to show up at the polls, candidates can target their messages towards members of their own camp, which have a higher probability of learning the respective policy platform compared to the average voter. Instead, in our model, candidates communicate with each (group of) voter(s), but can tailor their messages accordingly in order to conform better with the respective political preferences.
A political economy literature has focused on the welfare economics of contribution limits for financing campaign advertising; see, for example, Coate (2004) and the survey by Prat (2007); our model abstracts away from direct costs of campaign advertising.

decision making, more selective disclosure by a single candidate still benefits each voter conditional on that voter being pivotal. However, because of the persuasion bias that is induced by the candidate who discloses more selectively, each voter is less likely to become pivotal and thus to obtain an outcome consistent with her preferences. The second effect dominates when the number of voters is sufficiently high. Thus, rational voters are made worse off when only one of the two candidates is able to disclose information more selectively. In addition, in this setting unwary voters impose a negative externality on other voters; an increase in the number of unwary voters decreases each voter's probability of becoming pivotal, and thus also reduces each voter's true expected welfare.

Overall, our analysis challenges the traditional distinction between persuasive and informative advertising.[13] In our model, advertising is at once informative and persuasive, both on the way to the equilibrium—given that improved information shifts upward consumer demand, holding fixed consumer beliefs about information quality—and in equilibrium—given that even wary receivers end up being persuaded by the selective information they are given. An important twist is that the persuasion that results in equilibrium may actually hurt senders when competition is limited; in these circumstances senders benefit from committing not to acquire and disclose information.

Section 2 sets the stage by introducing the impact of information on selective disclosure. Section 3 proceeds by analyzing our baseline model of persuasion with a single receiver who faces fixed prices. Section 4 extends the framework to marketing environments where firms can personalize pricing. Section 5 applies our model to campaigning where voters make decisions collectively. Section 6 concludes by summarizing the main insights and discussing open questions. Appendix A collects the proofs and Appendix B presents supplementary material.

2 Selective Disclosure: Impact on Ex Ante Distribution of Receiver Utility

At the heart of our analysis is a game of information acquisition and communication between senders and receivers. Senders attempt to persuade receivers to accept their offerings by tailoring their communication on the basis of what they learn about the preferences of any given receiver. Senders initially acquire information about receiver preferences. Communication then takes the form of disclosure, whereby senders truthfully

[13] See, for example, Dixit and Norman (1978) and Grossman and Shapiro (1984) on the welfare impact of persuasive and informative advertising, and Bagwell (2007) for a systematic survey on the economics of advertising.

reveal part of the information they have previously acquired. Thus, in line with the disclosure literature, the revealed information is verifiable, for example because mendacious statements result in prohibitive losses of reputation or liability.

Disclosure affects a receiver's perceived valuation of a sender's offering, relative to the alternative option given by the best among the offerings of all the other senders. A given receiver will accept sender m's offering if and only if the corresponding perceived utility U_m is sufficiently high, relative to the best alternative option. We analyze throughout settings with horizontal differentiation where the values of the offerings are independent across senders and receivers. In order to characterize the effect that a sender's information has—via selective disclosure—on the ex-ante distribution of receiver utility, this section focuses on the case of a single sender interacting with a single receiver.

To set the stage, we introduce in turn two specifications of selective disclosure. First, in the case of selective <u>non</u>-disclosure, the sender decides whether or not to disclose information at all, depending on what the sender has learned about the preferences of the receiver. But even when the sender always discloses some information, disclosure may still be selective when the sender can decide which information to disclose. Our second specification captures such selective <u>targeted</u> disclosure, where the sender learns two pieces of information about receiver preferences and decides which one to disclose.

Selective <u>Non</u>-Disclosure. Denote the receiver's *true* utility from alternative m by u_m, which is independently distributed across all M alternatives with prior logconcave distribution $F_m(u_m)$ admitting an atomless density $f_m(u_m)$. It is convenient to stipulate that the support is bounded and given by $[\underline{u}_m, \overline{u}_m]$. Denote the expectation of u_m by $E[u_m]$. Sender m learns about u_m with probability θ_m and then decides whether or not to disclose u_m to the receiver.[14] For now the probability θ_m that the sender acquires information plays the role of a parameter, as in the literature on selective disclosure à la Dye (1985);[15] θ_m is chosen by the sender in the strategic model introduced in the next section. As long as $\theta_m < 1$, the fact that the sender remains ignorant with strictly positive probability prevents full unravelling of the information about receiver preferences the sender has actually learnt. The equilibrium disclosure strategy of sender m is then characterized by a threshold value, $u_d(\theta_m)$, above which the sender discloses u_m and below

[14] When the sender remains uninformed, we may stipulate that he either cannot disclose the respective information or that it is not optimal to do so. Otherwise, there is simply no scope for *selective* non-disclosure.

[15] This baseline model of selective disclosure has been widely adopted in finance; see, for example, Guttman, Kremer, and Skrzypacz (2014) for additional references and a recent analysis in a dynamic context.

which the informed sender pools with the uninformed sender. For given θ_m, the sender only discloses when $u_m \geq u_d(\theta_m)$, which uniquely solves

$$u_d(\theta_m) = (1-\theta_m)E[u_m] + \theta_m E[u_m \mid u_m \leq u_d(\theta_m)]. \tag{1}$$

Uniqueness follows given that logconcavity of $F_m(u_m)$ implies that $dE[u_m \mid u_m \leq u']/du' \leq 1$. At $u_d(\theta_m)$ the disclosed true utility is equal to the receiver's expected utility when there is no disclosure, as given by the right-hand side of (1).

The main object of our analysis is the *ex-ante* distribution of the receiver's perceived utility from the offering of sender m, $G_m(U_m)$. For a given and known value of θ_m, this distribution is $G_m(U_m) = (1-\theta_m) + \theta_m F_m(U_m)$ for $u_d(\theta_m) \leq U_m \leq \overline{u}_m$, with a mass point at the lower bound, as explained in more detail in the proof of Proposition 1. Our analysis hinges on a comparison of $G_m(U_m)$ for two different values $0 < \theta_m^L < \theta_m^H < 1$ of the probability that sender m is informed about receiver preferences:

(i) When this probability is known to the receiver, we can apply our previous characterization for each value of θ_m. Denote by $I(U_m)$ the distribution resulting with higher probability θ_m^H that the sender is informed and by $N(U_m)$ the distribution corresponding to θ_m^L. Proposition 1(i) establishes that the transition from $N_m(U_m)$ to $I(U_m)$ represents a Single-Crossing, Mean-Preserving Spread.

(ii) Suppose now that the sender actually observes a receiver's preferences with the higher probability θ_m^H, but that the receiver wrongly believes that the probability is strictly lower, θ_m^L. This case is relevant both off-equilibrium (to determine the sender's incentives to choose the unobservable information structure when facing a receiver with rational expectations) and also when the receiver naively fails to anticipate the sender's optimal information strategy. Optimally, the sender still applies the cutoff $u_d(\theta_m^L)$, but now, when there is no disclosure, the receiver's *perceived* value U_m differs from the *true* conditional expected value by an amount equal to $(\theta_m^H - \theta_m^L)\left\{E[u_m] - E[u_m \mid u_m \leq u_d(\theta_m^L)]\right\}$. Proposition 1(ii) establishes that an increase in the actual probability above the level expected by the receiver induces a First-Order Stochastic Dominance shift in the *ex-ante* distribution of the receiver's perceived utility from $N(U_m)$ to $\widehat{I}_m(U_m)$.

As illustrated in the left-side panel of Figure 1 for a uniform example with $F_m(u_m) = u_m$ over $[0,1]$, we have the following result:

Proposition 1 *Consider the setting of selective <u>non</u>-disclosure with prior logconcave distribution $F_m(u_m)$:*

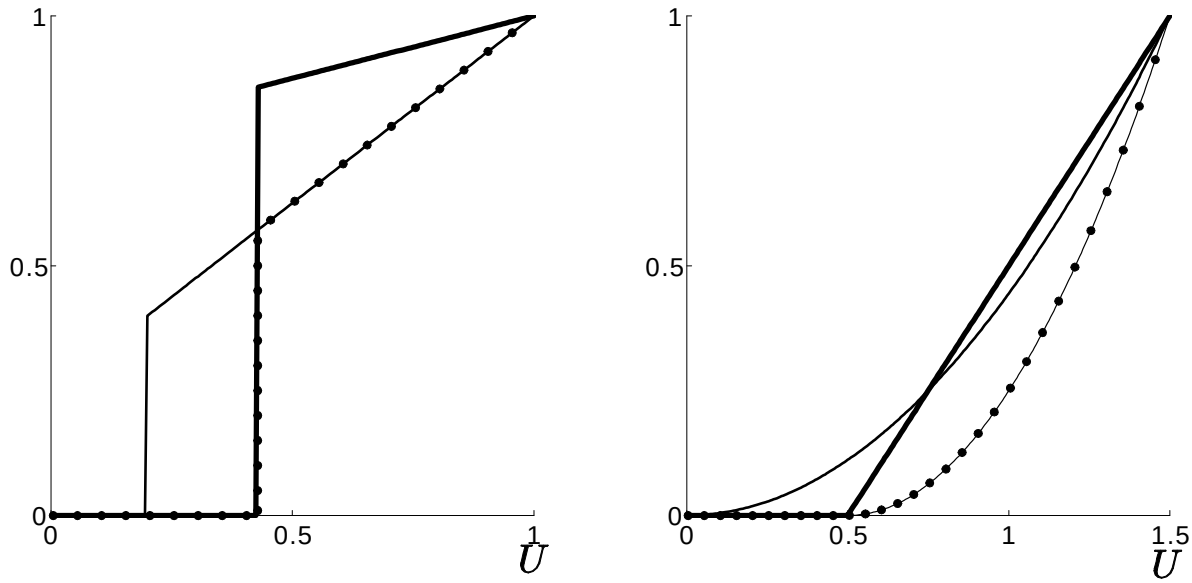

Figure 1: **Comparison of Distributions for Selective Non-Disclosure (left-hand panel) and Selective Targeted Disclosure (right-hand panel).** Both panels depict the shift from $N_m(U_m)$ (bold curve) to $I_m(U_m)$ (thin curve) and to $\widehat{I}_m(U_m)$ (dotted curve) with a uniform distribution $F_m(u_m) = u_m$ and, in case of selective non-disclosure, for $\theta_m^L = \frac{1}{4}$, $\theta_m^H = \frac{3}{4}$.

(i) If it is commonly known that the probability the sender is informed about the preferences of the receiver increases from θ_m^L to $\theta_m^H > \theta_m^L$, then the resulting shift in the ex-ante distribution of the receiver's perceived utility U_m from $G_m(U_m) = N_m(U_m)$ to $G_m(U) = I_m(U_m)$ represents a Single-Crossing, Mean-Preserving Spread (SCMPS).
(ii) If the receiver is unaware that the sender is more likely to be informed, then the new distribution of the perceived utility $G_m(U) = \widehat{I}_m(U_m)$ dominates $N_m(U_m)$ in the sense of First-Order Stochastic Dominance (FOSD).

That the shift in case (i) must be mean preserving follows immediately from Bayesian updating. The single-crossing of $I_m(U_m)$ and $N_m(U_m)$ implies that increased selectivity of disclosure by a better informed sender moves probability mass away from the center and into the tails of the distribution. Hence, more selective disclosure induces greater variability in receivers' perceived utility. According to property (ii), when the receiver is unaware of the sender's more selective information strategy, a better informed sender is able to induce in expectation a more favorable perception through selective disclosure.

The analysis to follow builds on the feature that more selective disclosure results either in a SCMPS or in a FOSD in the distribution of perceived utility depending on whether or

not the receiver is aware of this. Next, we turn to an alternative specification of selective disclosure that also satisfies these two properties.

Selective Targeted Disclosure. Under selective *non*-disclosure, full unraveling is avoided because a receiver always expects the sender to remain uninformed with positive probability. However, disclosure may still be selective and depend on the receiver's individual preferences also in situations in which the sender always discloses *some* information, provided that the sender learns the receiver's preferences. To model such selective targeted disclosure we follow Fishman and Hagerty (1990) by supposing that senders are possibly informed about two different attributes of their offering, $i = 1, 2$. If the receiver knew the values u_m^i of both attributes, the resulting true valuation would be

$$u_m = \sum_{i=1,2} u_m^i, \tag{2}$$

in the spirit of Lancaster (1966). Valuations u_m^i are independently distributed across senders m and attributes i according to the distribution function $F_m(u_m^i)$, which we assume to be uniform on $[\underline{u}_m, \overline{u}_m]$ in the main text.[16] (Our slight abuse of notation across the two applications is intentional to stress similarities.) In our application to marketing, $i = 1, 2$ may correspond to attributes of a product. The characteristics of an attribute may then represent a good match for one receiver but not for another receiver.[17]

What prevents full unraveling with *targeted* selective disclosure is the limitation that the sender can only disclose one attribute, but not both. Our key motivation for this is a receiver's limited attention; in practice such a restriction also arises naturally from limitations of time and (advertising) space, as discussed in the introduction. Given that the same distribution $F_m(u_m^i)$ applies to both u_m^1 and u_m^2, we can stipulate that sender m, when not informed about preferences, always discloses the same attribute $d_m \in \{1, 2\}$. Instead, when informed, the sender will disclose the attribute with the highest "fit", $d_m = \arg\max_{i=1,2} u_m^i$.[18] The receiver, when aware of this, will discount the value of the non-disclosed attribute accordingly when determining his perceived utility.

[16] See Appendix B for a more general analysis where $F_m(u_m^i)$ can be any logconcave distribution with an atomless density $f_m(u_m^i)$ over $[\underline{u}_m, \overline{u}_m]$. There we also consider cases with $\overline{u}_m \to \infty$.

[17] An alternative interpretation is that disclosure by a firm allows a consumer to learn the distances between the product's true characteristics and her own preferred characteristics. In an earlier version of this paper, using a Salop circle for each attribute, we have explicitly disentangled the actual location of product attributes from the preferred location according to the preferences of a particular consumer. The subsequently derived properties then still hold despite the fact that the distribution of utility does not necessarily inherit the properties of the resulting distribution of distances. In Appendix B we also consider the case where the receiver places different weights on the utilities u_m^i in (2).

[18] A strategy where the sender does not disclose any attribute, but where the receiver knows that the sender is informed, would also not arise in equilibrium due to a standard unraveling argument.

Consider the shift in the distribution $G_m(U_m)$ from $N_m(U_m)$ (when the sender is not informed about the receiver's preferences) to $I_m(U_m)$ (when the sender is informed and the receiver knows that the sender is informed). Targeted selective disclosure, when the sender known to be informed, increases probability mass at the lower tail given that the support of U_m shifts downward: the lower bound of the support of $N_m(U_m)$ is $\underline{u}_m + E[u_m^i]$, while that of $I_m(U_m)$ is $2\underline{u}_m$, given that the lowest perceived value $U_m = 2\underline{u}_m$ is realized when $u_{d_m} = \underline{u}_m$ is disclosed. At the upper tail of the distribution two conflicting forces are now at work under selective targeted disclosure: on the one hand, it is more likely that higher values of u_m^i are disclosed when disclosure is selective; on the other hand, the perceived value of the non-disclosed attribute is then discounted by the receiver. Appendix B shows that with logconcavity of $F_m(u_m^i)$ the first effect dominates at the upper tail, because log-concavity sufficiently bounds the change in the conditional expected value of the non-disclosed attribute. Instead, when the receiver is unaware that the sender is informed, the resulting distribution $\widehat{I}_m(U_m)$ FOSD dominates $N_m(U_m)$, given that the receiver then holds inflated expectations about the undisclosed attribute.

The right-side panel of Figure 1 illustrates that selective targeted disclosure in the uniform example (with $F_m(u_m^i) = u_m^i$ for $u_m^i \in [0,1]$) impacts the distributions of perceived utility through SCMPS and FOSD shifts, in analogy with what Proposition 1 established for selective non-disclosure. These specifications motivate the following general definition of the impact of a sender's information strategy, $s_m \in \{L, H\}$:

Definition 1 (More Selective Disclosure). *We say that sender m's information strategy $s_m = H$ leads to "more selective disclosure" than $s_m = L$ when it impacts the ex-ante distribution of a receiver's perceived utility $G_m(U_m)$ in the following two ways:*
(i) When the receiver believes $s_m = H$, the resulting shift from $G_m(U_m) = N_m(U_m)$ (for $s_m = L$) to $G_m(U) = I_m(U_m)$ (for $s_m = H$) represents a <u>Single-Crossing, Mean-Preserving Spread</u> (SCMPS): There exists $\widetilde{U}_m$ so that for all U_m in the interior of at least one support it holds that $I_m(U_m) > N_m(U_m)$ when $U_m < \widetilde{U}_m$ and $I_m(U_m) < N_m(U_m)$ when $U_m > \widetilde{U}_m$.
(ii) When the receiver believes that $s_m = L$, the resulting shift from $G_m(U_m) = N_m(U_m)$ (for $s_m = L$) to $G_m(U) = \widehat{I}_m(U_m)$ (for $s_m = H$) is such that $\widehat{I}_m(U_m)$ dominates $N_m(U_m)$ in the strict <u>First-Order Stochastic Dominance</u> (FOSD) order.

In the two specifications developed in this section, the choice of $s_m = L$ or $s_m = H$ corresponds to the choice of probability θ_m^L or θ_m^H with which the sender is informed about receiver preferences. The analysis that follows from now on directly assumes the general properties stated in Definition 1, and thus applies beyond these specifications. In light of our application to selective non-disclosure, we allow the distributions $G_m(U_m)$ to

have an arbitrary (countable) number of mass points, though it is convenient to suppose that the respective support is convex and bounded. Further, our analysis extends to situations in which the receiver's utility depends on additional, possibly only privately known determinants.[19] To see this, suppose that the receiver's *total* perceived utility under option m is equal to U_m *plus* some additional, independently distributed component V_m, i.e., equal to $\widehat{U}_m = U_m + V_m$, now with respective distribution $\widehat{G}_m(\widehat{U}_m)$. In particular, as long as V_m has a logconcave density, when G_m undergoes a SCMPS, the same holds for $\widehat{G}_m$ so that Definition 1 now applies to $\widehat{U}_m$.[20]

3 Equilibrium Selective Disclosure

3.1 Baseline Model

Let $M \geq 2$ be the set, as well as the number, of the senders and their corresponding offerings. We assume throughout that receivers' values for the senders' offerings are independently distributed across senders and across receivers. For this baseline model we consider receivers who each individually choose one of the senders' offerings, so that we can focus on the case with a single receiver; Section 5 extends the analysis to the case with multiple receivers who make a collective choice through voting. Our setup is general enough to allow for "non-strategic" senders who are known to always choose the same information strategy, e.g., $s_m = L$. One could then think of such a non-strategic sender as an outside option $U_m = R$ with fixed (possibly degenerate) distribution $N_m(R)$.[21] Thus, our setting also encompasses the monopoly case of a single strategic sender who wants to persuade a receiver to take his offering over an outside option of fixed value.

At $t = 1$ senders choose their information strategy $s_m \in \{L, H\}$. The baseline assumption is that the choice of s_m is an unobservable hidden action; this case without commitment is natural in the absence of regulation. Section 3.4 considers the commitment case where the choice of s_m is observable.

At $t = 2$, senders can disclose information to a receiver, and the ex-ante distribution of the perceived value U_m is given by $N_m(U_m)$ when the choice of $s_m = L$ is commonly known, by $I_m(U_m)$ when the choice of $s_m = H$ is commonly known, and by $\widehat{I}_m(U_m)$ when $s_m = H$ but the receiver wrongly believes that $s_m = L$.[22] Recall that we refer to $s_m = H$,

[19] All results generalize apart from the results with personalized pricing derived in Section 4, because the addition of receiver heterogeneity could interfere with optimal pricing.

[20] Cf. e.g. Theorem 4 in Jewitt (1987).

[21] In what follows we assume that all senders are strategic unless we explicitly state otherwise.

[22] Thus, the case where $s_m = L$ but a receiver wrongly believes that $s_m = H$ is not relevant.

compared to $s_m = L$, as more selective disclosure according to Definition 1.

At $t = 3$, the receiver chooses sender $m \in M$ that results in the highest perceived utility, U_m, and randomizes with equal probability in case of a tie. With respect to the payoffs that are then realized, for the baseline specification of the model we only need to specify that each sender m is strictly better off when the receiver chooses their offering m.[23] The receiver in the baseline is fully rational; Section 3.5 investigates the policy-relevant case with unwary receivers. Section 4 extends the model to allow for personalized pricing.

3.2 Receiver Preferences

Before proceeding to analyze the equilibrium, it is helpful to consider the preferences of a receiver. The receiver realizes $U^{(1)} = \max_{m \in M} U_m$. Suppose the receiver is aware of the strategies (s_m) of all senders. Then, pick some $m \in M$ and denote by $U^{(1:M \setminus m)}$ the maximum over all remaining senders, with distribution $G^{(1:M \setminus m)}(U^{(1:M \setminus m)})$. For given U_m, the receiver accordingly realizes $E\left[\max\left\{U^{(1:M \setminus m)}, U_m\right\}\right]$. Given that the expression in brackets is a convex function of U_m, the receiver obtains a higher expected utility after a mean-preserving spread in $G_m(U_m)$:

Proposition 2 *More selective disclosure by any sender m, resulting in a switch of $G_m(U)$ from $N_m(U_m)$ to $I_m(U_m)$, benefits a receiver who is aware of this.*

The intuition for Proposition 2 is straightforward. For each alternative m the receiver has a binary decision to make, namely whether to take this alternative or not. Once the receiver decides against this alternative, the actual value of the resulting utility is inconsequential because the receiver obtains the value of her next best alternative. Each alternative m thus represents an option for the receiver, whose expected value increases with a mean-preserving spread under $I_m(U_m)$. Note that in three extensions of this baseline setting receivers will no longer unambiguously benefit from more selective disclosure: (i) when a receiver is unwary of the increased selectivity of disclosure (Section 3.5); (ii) when selective disclosure is combined with personalized pricing (Section 4); and (iii) when decision-making is collective (Section 5).

3.3 Equilibrium without Commitment

Recall our assumption that the choice of s_m is an unobservable hidden action. For our baseline analysis we now further assume that there is also no other way for the senders to

[23] Given that the distributions of preferences across alternatives are independent, we also need not specify whether sender m receives a different (or the same) payoff when alternatives $m' \neq m$ or $m'' \neq m$ are chosen.

commit to a certain information strategy. We can then formalize the senders' incentives to become better informed in order to disclose more selectively as follows. For a given realization of U_m, denote by $w_m(U_m)$ the "winning" likelihood with which the receiver chooses alternative m. Hence, from an ex-ante perspective, alternative m is chosen with probability

$$q_m = \int w_m(U_m) dG_m(U_m).$$

The following is an immediate implication of the fact that $\widehat{I}_m(U_m)$ dominates $N_m(U_m)$ in the strict First-Order Stochastic Dominance (FOSD) order and that $w_m(U_m)$ is non-decreasing.[24]

Proposition 3 *Suppose senders are unable to commit to the information strategy s_m. Then, in the only equilibrium all senders choose to disclose more selectively ($s_m = H$).*

The intuition is straightforward. Suppose that the receiver believes that $s_m = L$. Then, by the FOSD property, the sender has a strict incentive to disclose more selectively ($s_m = H$), because this induces a favorable upward shift of the distribution of the receiver's perceived utility from $N_m(U_m)$ to $\widehat{I}_m(U_m)$. When all senders have the option to disclose more selectively, in equilibrium the receiver should thus anticipate $s_m = H$ for all $m \in M$.

3.4 Equilibrium with Commitment

We now turn to circumstances in which senders have the ability to credibly commit to a certain information strategy. This is so, in particular, when information acquisition requires either the direct cooperation or at least the consent of receivers, in which case receivers in fact directly observe a sender's attempt to become better informed about receiver preferences in order to disclose more selectively.

Sender's Incentives. In contrast to the preceding analysis without commitment, now a given sender m's preferred choice depends crucially on the distribution of a receiver's next best alternative, which we denoted by $G^{(1:M\setminus m)}(U^{(1:M\setminus m)})$. To illustrate, we first consider a simple, though somewhat extreme example and then generalize the insight by mapping changes in $G^{(1:M\setminus m)}(U^{(1:M\setminus m)})$ into the model's primitives, notably the number of competing senders and offers.

[24] The latter property follows from an argument by contradiction. When all distributions $G_{m'}(U_{m'})$ are atomless, it is straightforward to derive $w_m(U_m) = \prod_{m' \in M\setminus m} G_{m'}(U_m)$, because then the likelihood of a draw is zero. The expression for the general case, allowing for mass points of all $G_m(U_m)$ and for asymmetries across senders, has a more unwieldy expression but can also be easily derived.

Consider first the degenerate case where the receiver's best alternative has a deterministic value: $U^{(1:M\backslash m)} = R$. The receiver would then choose sender m's preferred alternative with probability $1 - I_m(R)$ when the sender discloses more selectively, rather than with probability $1 - N_m(R)$. By the SCMPS property, the probability that the receiver accepts the sender's offering is strictly higher under more selective disclosure if and only if R lies to the right of the intersection of $I_m(U_m)$ and $N_m(U_m)$, while otherwise it is strictly lower. That is, in this example the sender prefers more selective disclosure if the receiver's preferred alternative is sufficiently attractive (high R), while he otherwise prefers less selective disclosure. We now generalize this insight to the case where $U^{(1:M\backslash m)}$ is stochastic. To this end, first restrict consideration to distributions $G^{(1:M\backslash m)}(U^{(1:M\backslash m)})$ that are ordered in terms of their Likelihood Ratio (LR).[25] Below, we relate such shifts to changes in the model's primitives and also extend the result. To obtain a non-trivial comparison, we only consider the case where the supports of $N_m(\cdot)$, $I_m(\cdot)$, and $G^{(1:M\backslash m)}(\cdot)$, before and after the LR shift, overlap sufficiently to guarantee that the probability of the receiver choosing option m is strictly between zero and one.

Lemma 1 *Consider a single sender m and suppose the distribution of the receiver's next best alternative, $U^{(1:M\backslash m)}$, undergoes a shift resulting in a distribution that dominates in the LR order. Then, if the sender weakly preferred $s_m = H$ to $s_m = L$ before the shift, he strictly does so after the shift. Likewise, if the sender weakly prefers $s_m = L$ to $s_m = H$ after the shift, he strictly did so before the shift.*

Below we will also consider weaker conditions for when sender m prefers $s_m = H$, though then the converse does not hold as is the case for the LR shift in Lemma 1.

To close the analysis of the sender's preferences, we need to map the comparative analysis into the model's primitives. An immediate way to do so is to consider directly a switch in the respective distribution $G_{m'}$ of any other sender.[26] This could then be further linked to a change in the distribution of a consumer's utility. To illustrate, consider selective targeted disclosure with $M = 2$ and u_1^i distributed uniformly over $[0, 1]$, as in Figure 1. For sender $m = 2$ take $G_2(U_2) = N_2(U_2)$ and suppose that u_2^i is also distributed uniformly but with support $[0, \overline{u}_2]$, for $0 < \overline{u}_2 < 3$. Then $m = 1$ prefers to become informed and target his disclosure, resulting in $G_1(U_1) = I_1(U_1)$, if $\overline{u}_2 > \frac{3}{4}$ (i.e., if the

[25] That is, we assume here that $G^{(1:M\backslash m)}(\cdot)$ admits a density $g^{(1:M\backslash m)}(\cdot)$. Then $\widehat{G}^{(1:\widehat{M}\backslash m)}(\cdot)$ dominates $G^{(1:M\backslash m)}(\cdot)$ in the likelihood ratio order if and only if the ratio of the respective densities $\widehat{g}^{(1:\widehat{M}\backslash m)}(\cdot)/g^{(1:M\backslash m)}(\cdot)$ is an increasing function. Below we also treat the case with jumps when we relate a change of $G^{(1:M\backslash m)}(\cdot)$ directly to the model's primitives.

[26] Obviously, a switch from $s_{m'} = L$ to $s_{m'} = H$ alone does not satisfy the condition of Lemma 1. See, however Proposition 4 below.

receiver's alternative is relatively attractive), but prefers not to become informed if $\overline{u}_2 < \frac{3}{4}$ (i.e., if the receiver's alternative is relatively unattractive).

In what follows we focus on a comparative analysis in terms of competition, as expressed by the number of senders M. Importantly, our results then hold irrespective of the specific application and the underlying distribution of receiver preferences. For the following result we first impose symmetry across senders and their alternatives, as well as a restriction on admissible mass points of the respective distributions, so as to directly appeal to Lemma 1 (and the LR shift).

Proposition 4 *Suppose that the distributions $N(U_m)$ and $I(U_m)$ are identical across senders and admit a density with at most a mass point at the lower support. Suppose also that senders can commit to an information strategy s_m. Then there is a finite threshold M' such that for all $M \geq M'$ there exists an equilibrium where all senders disclose more selectively, $s_m = H$, while for $M < M'$ this outcome is not an equilibrium.*

Proposition 4 thus provides an equilibrium characterization for large M. Intuitively, as the number of senders increases, it becomes more and more likely that the receiver obtains a highly valuable offer elsewhere, so that from the perspective of sender m, more and more probability mass shifts to the upper tail in the distribution of the receiver's best "outside option", $U^{(1:M\backslash m)}$. As a consequence, each sender m prefers to disclose more selectively in order to increase the probability of particularly high realizations of U_m. Below we show that this intuition also extends to the case where no symmetry across senders is imposed. Note that in the two specifications from Section 2 the conditions of Proposition 4 hold whenever the distribution of consumer preferences ($F_m(u_m)$ and $F_m(u_m^i)$, respectively) is continuous.

Proposition 4 also covers the converse case: When M falls below the threshold M', there no longer is an equilibrium where all senders choose to disclose more selectively. There, the same intuition applies as in case of a deterministic outside option $U^{(1:M\backslash m)} = R$ with a low value of R: A sender does not want to disclosure more selectively when, given a low M, the receiver's alternative is likely to be relatively unattractive.[27]

Before proceeding to a further interpretation, we extend the result for high M. In fact, the above intuition suggests that the result should hold more generally, notably without invoking symmetry as in Proposition 4. This is indeed the case, and we only need some milder technical restrictions that essentially ensure that the addition of more senders exerts

[27] We do not provide a complete characterization of equilibrium for the case with low M, when existence can only be guaranteed by allowing for mixed strategies.

competitive pressure on a given sender M.[28]

Proposition 5 *Suppose senders can commit to an information strategy s_m. If competition is sufficiently intense (high M), any sender m benefits when he is known to disclose more selectively, so that the unique equilibrium has $s_m = H$ for all $m \in M$, as long as the following conditions are met:*
(i) Distributions $N_m(U_m)$ and $I_m(U_m)$ have the same support across senders and
(ii) Sender types as characterized by distributions $G_m(U_m)$ are drawn from a finite set Ω.

Discussion. Propositions 4 and 5 assume that receivers are aware of a sender's attempt to become better informed about their preferences in order to disclose more selectively. Absent regulation, consumers may not be able to control the extent to which firms collect personal data, such as past purchases that reflect also on consumer preferences for the current offering. By Proposition 3, all firms would then indeed want to disclose more selectively. By Proposition 2, consumers would strictly benefit from increased selective disclosure and, consequently, would be harmed by regulation that prohibits firms from collecting and using personal information.

Interestingly, even a less restrictive regulation requiring firms to seek consumer *consent* can backfire and lead to a reduction in consumer surplus. This is the case when firms would like to commit not to become better informed about individual consumers' preferences before disclosing selectively, but cannot do so because consumers do not observe whether firms collect and use personalized data. Regulation that prescribes consumer consent then provides such commitment, which is in the interest of firms but not in the interest of consumers. From Proposition 4 we know that a firm is more likely to benefit from such a commitment, as provided by regulation, when competition is limited (small M). Then, even regulation that, rather than prohibiting the collection and use of personal data, just requires consumer permission can backfire and decrease consumer welfare and efficiency. We return to a more complete discussion of policy after analyzing the case with unwary receivers.

3.5 Unwary Receivers

Given the novelty of hypertargeting technology, receivers may well be unaware of senders' capability to collect and use data for selective disclosure. To analyze this case, suppose that

[28] For instance, this would clearly not be the case when for all newly added senders m' the mass of $G_{m'}(U_{m'})$ was confined to values $U_{m'}$ that are below the lower bound of the support of $G_m(U_m)$ (for a given choice s_m). We thus require that across all senders m, the distributions $N_m(U_m)$ have the same support, and that also the distributions $I_m(U_m)$ have the same support.

a sender chooses $s_m = H$, but a receiver wrongly believes that $s_m = L$.[29] Recall that the distribution of perceived valuation of a receiver who wrongly believes that $s_m = L$ shifts in the FOSD order from $N_m(U_m)$ to $\widehat{I}_m(U_m)$. Directly working with $\widehat{I}_m(U_m)$ was sufficient to derive the senders' expected payoff and thus preferences over communication strategies in Section 3.3; however, to derive the impact on a receiver's *true* expected payoff we must return to the primitives that give rise to $\widehat{I}_m(U_m)$. Thus, we consider our two specifications of selective non-disclosure and selective targeted disclosure.

When a receiver remains unwary, the perceived value U_m risks being inflated, which may distort decisions and decrease consumer welfare.[30] However, there are two, less immediate effects that work in the opposite direction and tend to benefit unwary receivers through the increased quality of information collected, even though this information is selectively disclosed. One effect, which we illustrate first, works through (symmetric) competition when the perception of other senders' options is equally inflated. More concretely, when the receiver has to decide between some option m and the next best alternative m' whose perceived value is *also* inflated under competition, then these errors can disappear altogether, so that ultimately also an unwary receiver becomes unambiguously better off when (now all) senders disclose more selectively. Intuitively, under symmetry, when all senders choose the same disclosure strategy and when also the respective utilities are chosen from the same distributions, the distortions in perceptions fully cancel out. Then, an unwary receiver ends up making the same decision that a wary receiver would have made, for given choices of disclosure strategies by senders.

Proposition 6 *In our two specifications of selective non-disclosure and selective targeted disclosure, also an unwary receiver benefits when all senders choose to disclose more selectively ($s_m = H$) if preferences over all options M are symmetric.*

Combining Proposition 6 for unwary receivers with the results for wary receivers and firm preferences obtained above, we can conclude more generally that regulation that restricts the collection of personal data, from which senders can learn about receivers' preferences, always harms receivers in our model with symmetric competition. This holds despite the fact that unwary receivers' perceived valuations are inflated, precisely because they are protected by competition.

[29] In the selective targeted disclosure specification, unwary receivers do not adjust for the the fact that the informed sender discloses the most favorable attribute. Thus, they are effectively "cursed", as in Eyster and Rabin (2005).

[30] With selective targeted disclosure (respectively selective non-disclosure), the receiver's perceived value U_m exceeds the *true* conditional expected value by $E[u_m^i] - E[u_m^i \mid u_m^i \leq u_{d_m}]$ (respectively, $E[u_m] - E[u_m \mid u_m \leq u_{d_m}]$) multiplied with the difference in the respective probabilities with which the sender becomes informed, 1 (respectively, $\theta_m^H - \theta_m^L$).

The second effect through which unwary receivers may benefit from more selective disclosure is independent of competition between senders and the symmetry imposed in Proposition 6. It is based on the observation that, with incomplete information, an inflated perception of U_m not only increases the error of inefficiently opting for option m, but also reduces the error of inefficiently rejecting option m. In the rest of this section, we illustrate this separately for our two specifications.

Selective Targeted Disclosure. To abstract from the effects of competition, it is particularly illustrative to first take the auxiliary case with a single monopolistic sender and an outside option of known value R. By biasing upward an unwary receiver's perceived utility relative to the known value R, the mistake of erroneously choosing the sender's preferred option, even though $u^1 + u^2 < R$, evidently becomes larger when the sender discloses selectively. However, at the same time it becomes less likely that the receiver erroneously decides against the sender's preferred alternative, namely when actually $u^1 + u^2 > R$ holds. How these two errors trade off depends on the distribution $F(u^i)$. Interestingly, with a uniform distribution, as in Figure 1, it is straightforward to establish that these two errors exactly cancel out. With an exponential distribution, as also considered explicitly in Appendix B, unwary consumers are in fact strictly better off when $G(U)$ switches from $N(U)$ to $\widehat{I}(U)$, even though in the case of a single monopolistic sender their decision is biased. As these results for both the uniform and the exponential distribution hold for any choice of R, the observations extend to the case where R is itself random and, thereby, also to when R again represents the best alternative choice from $M - 1 > 0$ other senders (provided that the respective perception is not biased).

Selective Non-Disclosure. To illustrate, consider again the case with a single monopolistic sender and an outside option of known value. In this case also the unwary receiver is unambiguously better off under selective disclosure independent of the distribution $F(u)$. To see this, note that as the unwary receiver does not update his beliefs when there is no disclosure, the sender's strategy only affects his behavior when there is disclosure, which happens when $u \geq R$. In this case the receiver realizes indeed a strictly higher utility by choosing the sender's preferred option. The result then follows immediately because the case where $u \geq R$ and there is disclosure is strictly more likely when the sender is better informed.

4 Personalized Pricing

Scope for Personalized Pricing. When applying our model to marketing, senders represent firms and receivers represent individual consumers deciding which product to purchase. A distinctive feature of our baseline analysis is that firms do not adjust prices individually, based on their knowledge of consumer preferences. Given that each firm offers all consumers the same product, even when it selectively gives them different information, such personalized pricing may be difficult with physical goods that can be easily resold. Price discrimination would then create scope for arbitrage, either through a grey (or parallel) market between consumers or through the activity of intermediaries.[31] Beyond these circumstances for which our baseline analysis without price discrimination is suitable, in other markets transaction costs could impede arbitrage across consumers. This section turns to situations in which firms not only are able to learn about the preferences of consumers and target their communication accordingly, but also are able to charge personalized prices to customers.[32]

Firm Preferences with Personalized Pricing. With competition, we stipulate that firms learn the utility that the consumer perceives for each product, for example, on the basis of some commonly collected information. When no firm chooses weakly dominated prices, this ensures that, first, the consumer still purchases the product with the highest perceived utility $U^{(1)}$, and that, second, the price that the consumer pays is equal to the incremental utility relative to the second-highest such value. Consequently, a consumer realizes the second order-statistic, denoted by $U^{(2)}$. We first establish that with personalized pricing all firms prefer to disclose more selectively, now regardless of whether this is observed by consumers or not and irrespective of the intensity of competition (and thus in contrast to our previous findings without personalized pricing; see Propositions 4 and 5).

Recall our notation $U^{(1:M\backslash m)}$ for the highest expected utility over all other $M\backslash m$ firms. Then, the expected profit of firm m is given by

$$\int \left[\int \max\left\{U_m - U^{(1:M\backslash m)}, 0\right\} dG^{(1:M\backslash m)}(U^{(1:M\backslash m)}) \right] dG_m(U_m).$$

[31] Also, price discrimination may be limited when consumers are concerned about fairness. Price (or rate) parity has become a major objective for firms, e.g., hotels, given the increasing transparency via online channels. Furthermore, when the considered channel may only represent one among several (online or offline) distribution channels, the firm's pricing flexibility for this channel may be seriously compromised, so that we may indeed abstract away from pricing differences depending on the firm's disclosure policy.

[32] The industrial organization literature on behavior-based price discrimination has focused on personalized pricing where, in particular, the past purchasing history of consumers is used; see, for example, Villas-Boas (1999) and Acquisti and Varian (2005). As we abstract from this dynamic feature, our analysis will be quite different.

Because the term in rectangular brackets is a convex function of U_m,[33] it is higher after a mean-preserving spread in $G_m(U_m)$. With personalized pricing, a firm that offers a consumer's preferred choice—and can thus make a profit—wants to maximize the distance between the consumer's expected utility for the firm's own product and the utility for the product of its closest rival, because the firm extracts exactly this difference. From an ex-ante perspective, the firm thus prefers a greater dispersion of U_m. Given that the firm also benefits from a FOSD shift in $G_m(U_m)$, we have the following result:

Proposition 7 *With personalized pricing, any given firm m that can do so prefers to disclose more selectively, irrespective of whether this is observed by the consumer or not. Thus, with personalized pricing in the unique equilibrium outcome all firms that can do so disclose more selectively ($s_m = H$).*

Consumer Preferences with Personalized Pricing. From the perspective of consumers, the effect of more selective disclosure depends now crucially on the degree of competition, in stark contrast to the baseline case without personalized pricing. As a starting point, consider a duopoly with $M = 2$, where the differences between the two cases are particularly pronounced. While without personalized pricing a consumer realized the maximum of the two expected utilities $U^{(1)} = \max\{U_1, U_2\}$, with personalized pricing the consumer realizes the second-highest value, which for $M = 2$ is the minimum $U^{(2)} = \min\{U_1, U_2\}$. The consumer is now strictly worse off when any of the two firms discloses more selectively. Formally, this can be seen in complete analogy to the argument to why without personalized pricing the consumer was strictly better off. With personalized prices the consumer's expected utility is

$$E[U^{(2)}] = \int \left[\int \min\{U_m, U_{m'}\} \, dG_{m'}(U_{m'}) \right] dG_m(U_m).$$

The expression in rectangular brackets is now a strictly concave function of U_m,[34] while it was a strictly convex function when without personalized pricing we applied the maximum. Thus a mean-preserving spread in $G_m(U_m)$ reduces the consumer's expected utility. The intuition is that as firm m discloses more selectively, a consumer's updating makes firms more differentiated from an ex-ante perspective. This ensures that in expectation the (winning) firm with the highest perceived value can extract a higher price. As we

[33] Its derivative with respect to U_m is $G^{(1:M\setminus m)}(U_m)$ and thus increases in U_m.

[34] It can be written as $U_m - \int_{\underline{U}}^{U_m} [U_m - U_{m'}] dG_{m'}(U_{m'})$, which is a.e. differentiable with first derivative $1 - G_{m'}(U_m)$, which is decreasing.

show next, this negative effect of increased differentiation for consumers is, however, subdued when M increases, so that it becomes increasingly likely that each firm has a close competitor, in which case the efficiency benefits, as in Proposition 2, again accrue to the consumer.

More precisely, Proposition 8 establishes that regardless of what all other firms do, when there are sufficiently many firms a consumer strictly benefits when a particular firm discloses more selectively, so that $I_m(U_m)$ instead of $N_m(U_m)$ applies. In case all distributions are atomless, as in the case of selective targeted disclosure, we have a particularly clear-cut result as to when a consumer benefits from a switch to $I_m(U_m)$. As Proposition 8 shows, this is the case if and only if

$$\int \left[N_m(U_m) - I_m(U_m)\right] \left[G^{(2:M\setminus m)}(U_m) - G^{(1:M\setminus m)}(U_m)\right] dU_m > 0. \tag{3}$$

Given single-crossing of $N_m(U_m)$ and $I_m(U_m)$, in Proposition 8 we can sign expression (3) unambiguously to be positive whenever there are sufficiently many firms (high M). While we obtain expression (3) only after some transformations, it intuitively captures the fact that, from a consumer's perspective, with personalized pricing the *precise* realization of U_m only matters when it falls between the first and second highest realizations of all other $M-1$ utilities, which happens with probability $G^{(2:M\setminus m)}(U_m) - G^{(1:M\setminus m)}(U_m)$. Expression (3) also makes it intuitive why in the following result we need to exclude one particular case in which receivers do not benefit from a shift of probability mass into the tails of $G_m(U_m)$.

Proposition 8 *When competition is sufficiently intense (high M), then even with personalized pricing a consumer strictly benefits when a given sender m discloses more selectively ($s_m = H$), as long as conditions (i) and (ii) from Proposition 5 together with at least one of the following hold: Either the upper supports of $I_m(U_m)$ and $N_m(U_m)$ coincide or at least one other sender $m' \neq m$ chooses $s_{m'} = H$.*

Proposition 8 relates our paper to results by Board (2009) and Ganuza and Penalva (2010) on the effect of providing bidders with private information in a private-values second-price auction. With personalized pricing, a comparison of consumer surplus in our model effectively amounts to comparing the expectation of the second-order statistic $E[U^{(2)}]$, as in a second-price auction. In these papers, however, the question that is asked is whether providing more information to *all* bidders increases the auctioneer's expected payoff, while for Proposition 8 we ask whether more information held by a *single* firm benefits the consumer.

Implications for Regulation. Next, we derive implications for the regulation of consumer privacy. Consumers can be worse off when more selective disclosure is combined with personalized pricing, contrary to our finding in the absence of personalized pricing. A second difference arises from firms' preferences. Because with personalized pricing firms will always want to collect personal data and disclose selectively, regulation that requires consumer consent can no longer backfire by granting firms commitment power, as in the baseline case. Regulation that strictly prohibits the collection and use of personal data will, however, reduce efficiency and consumer welfare when competition is sufficiently intense.

Suppose now again that consumers remain unwary of firms' capability to collect and use personally identifiable data for selective disclosure. Then, as noted previously, their perceived value U_m risks being inflated. While for the baseline analysis we showed that consumers may still be better off even with a monopolistic firm, with personalized pricing this is clearly no longer the case. Personalized pricing allows the firm to extract a consumer's perceived incremental utility relative to her next best choice. If the perceived utility is inflated, this generates the potential for consumer exploitation. Facing a monopolistic firm, an unwary consumer would then indeed be better off staying out of the market.

Competition between senders, however, protects also unwary consumers, both with and without personalized pricing. As in the baseline case, this is most immediate when, under symmetry, the perception of the different offers is equally inflated, so that the decision of an unwary consumer fully matches that of a wary consumer. Interestingly, in the case with personalized pricing an unwary consumer pays a lower price and thus ends up being even strictly better off than a wary receiver. This follows because unwary consumers do not adjust expectations for the quality of information underlying firms' selective disclosure, which works towards reducing the perceived difference between the first-best and second-best alternative. With selective targeted disclosure, if $u^{(1)}$ is the highest disclosed fit (attribute) and $u^{(2)}$ is the second highest, an unwary consumer pays the price $\widehat{p} = u^{(1)} - u^{(2)}$, given that the expectations about the non-disclosed attribute of either firm wrongly remain unchanged at $E[u^i]$. A wary consumer pays, instead, the strictly higher price

$$p = u^{(1)} - u^{(2)} + \left\{ E[u^i \mid u^i \leq u^{(1)}] - E[u^i \mid u^i \leq u^{(2)}] \right\},$$

given that the term in braces, equal to the difference in the updated conditional expectations for the attributes not disclosed by firms 1 and 2, is strictly positive. With selective non-disclosure, unwary consumers equally pay a strictly lower price than wary consumers whenever this is determined by the alternative of a firm that does not disclose. As in this

case an increase in θ_m also increases the likelihood of disclosure, which (weakly) reduces the expected price and the likelihood of making the wrong choice, together with our previous results we have:

Proposition 9 *When firms are able not only to selectively disclose information but also to price discriminate based on collected personal data, then regulation that requires firms to obtain consumer consent is always (weakly) beneficial. Instead, regulation that prohibits the collection of such data will harm consumers and also reduce efficiency when competition is sufficiently intense. At least when firms are symmetric, the latter conclusion holds also when consumers remain unwary of firms' capability to collect and use personal data.*

5 Political Campaigning and Collective Decision-Making

We now extend the model to political campaigning by supposing that $M = 2$ candidates compete for voters.[35] Equivalently, the model also captures voting for or against a motion brought forward in a debate or referendum. For convenience only, suppose there is an odd number $V \geq 3$ of voters; the case with $V = 1$ coincides with our previous analysis. Denote by $U_m(v)$ voter v's expected utility when candidate (or motion) m wins, which now (for a given strategy s_m) is independently and symmetrically distributed across candidates. In terms of motivation, a candidate could generally be more or less forthcoming with respect to his views, i.e., more or less willing to disclose them (rather than not disclosing). With respect to selective targeted disclosure, candidates' platforms could comprise issues on which a candidate's stance can more or less coincide with the preference and political orientation of a particular voter.

Each voter casts the ballot for the candidate whose future decisions promise to deliver the highest utility value, and the candidate m who obtains the highest number of votes wins the election. When V^1 is the number of votes cast for candidate $m = 1$ out of a total of V votes, the ex-ante probability with which candidate $m = 1$ is elected is $Q = 1 - \Pr[V^1 \leq (V-1)/2]$. Denote the likelihood of a vote for candidate $m = 1$ by $q_1 = q$. (Of course, this will differ depending on candidates' choice of strategy.) Then, given that V is odd, the expression for Q conveniently becomes

$$Q = 1 - \sum_{k=0}^{\frac{V-1}{2}} \binom{V}{k} q^k (1-q)^{V-k}.$$

[35] In recent field experiments, Dewan, Humphreys, and Rubenson (2014) and Kendall, Nannicini, and Trebbi (2013) find causal evidence of the effect of campaign information on intention to vote and electoral outcomes, respectively. For a view from the advertising industry, see Abse's (2013) account of hypertargeting strategies pursued in the 2012 U.S. presidential election.

Further, let y be the probability with which any given voter will become pivotal. Using again that V is odd, we have

$$y = \binom{V-1}{\frac{V-1}{2}} q^{\frac{V-1}{2}} (1-q)^{\frac{V-1}{2}} . \tag{4}$$

Under collective decision-making, each voter's ex-ante utility now comprises two terms. If the voter ends up not being pivotal, the vote does not influence the decision. As we now focus on the new effect arising from collective decision-making, we impose symmetry across candidates, implying in particular that when not being pivotal, a voter's unconditional expected utility $E[U]$ is independent of which candidate is chosen. However, when this voter is pivotal, the conditional expected utility equals $E[U^{(1)}]$. Multiplied with the respective probabilities, y for becoming pivotal and $1-y$ otherwise, a voter's ex-ante expected utility becomes

$$E[U] + y \left\{ E[U^{(1)}] - E[U] \right\} . \tag{5}$$

From Proposition 2, $E[U^{(1)}]$, and thus the term in braces, is strictly higher when one of the candidates discloses more selectively, and even more so when both candidates disclose more selectively. It is in this sense that our previous analysis can be applied. However, in this application to elections, a decision is no longer determined by the preferences of an individual receiver alone, as in our baseline case, but instead by the aggregate preferences of all receivers (voters). Therefore, we have to take into account a second effect. The strategies of candidates now also affect a voter's likelihood y of becoming pivotal and, thereby, affect voter utility according to expression (5). This is what we explore in the rest of this section.

The likelihood of becoming pivotal decreases when, instead of $q = 1/2$, asymmetry in candidates' strategies leads to $q \neq 1/2$, notably when one candidate alone, say $m = 1$, has a sufficiently sophisticated campaign team to collect the necessary information to "hyper-target" voters ($s_1 = H$). He will then indeed do so if this increases the likelihood of a vote in his favor, to $q > 1/2$. In this case, a trade-off results given that the improved informativeness increases a voter's utility conditional on being pivotal, while the probability that any given voter becomes pivotal decreases. We show next how this trade-off is always resolved unambiguously when the number of voters becomes sufficiently high.

Denote by q_I and y_I a voter's probability of voting for $m = 1$ and for being pivotal when only one candidate, here candidate 1, can disclose more selectively ($s_1 = H$, $s_2 = L$). Denote by $q_N = 1/2$ and y_N the corresponding probabilities when no candidate is better placed in this way ($s_1 = s_2 = L$). With a slight abuse of notation, $E_I[U^{(1)}]$ and $E_N[U^{(1)}]$ denote the respective conditional utilities of a pivotal voter. Then, a voter's ex-ante

expected utility, as given by (5), is higher when candidate 1 discloses more selectively if

$$\frac{y_I}{y_N} = [4q_I(1-q_I)]^{\frac{V-1}{2}} > \frac{E_N[U^{(1)}] - E[U]}{E_I[U^{(1)}] - E[U]}. \tag{6}$$

While the right-hand side of the inequality (6) does not depend on V (and is strictly smaller than one), the left-hand side is strictly decreasing in V (and goes to zero) as long as $q_I \neq 1/2$. Hence, either condition (6) does not hold for all V, including $V = 3$, or there exists a cutoff value $\widetilde{V}$ such that it is only for $V < \widetilde{V}$ that a voter is better off when (only) candidate 1 discloses more selectively. While for each voter the benefits conditional on being pivotal remain the same when the number of voters increases, this makes it relatively less likely that a voter becomes pivotal when candidates choose different disclosure strategies (even though both y_I and y_N approach zero as $V \to \infty$).[36] As a result, voters are no longer unambiguously better off when only one candidate is in a position to disclose information more selectively, in contrast to the baseline analysis with individual decision making:

Proposition 10 *Suppose V voters decide by majority rule over $M = 2$ ex-ante symmetric candidates. Voters are strictly better off when <u>both</u> candidates disclose more selectively ($G_m(U_m) = I(U_m)$), as in our baseline analysis. When only <u>one</u> candidate is in a position to disclose more selectively, asymmetric selective disclosure still benefits each voter conditional on being pivotal, but it decreases the likelihood with which individual voters become pivotal and can thus decide the outcome based on their own preferences. There exists a cutoff $\widetilde{V}$ on the number of voters such that voters are strictly worse off with asymmetric selectivity of disclosure when $V > \widetilde{V}$, and better off when $V \leq \widetilde{V}$.*

Groups of Voters. Our application to voting can apply to various settings, from small committees to larger elections, and the vote could be cast for individual candidates as well as for or against a particular motion. In the application thus far, V has represented the number of voters. For each voter, the respective fit with a candidate's orientation (or the content of a motion) was chosen independently. Such preferences may also be shared across different voters in an electorate. One way to extend our results is to now suppose that there are V voter groups, each composed of z_v voters with the same preferences or political orientation. A campaign in this case would target groups of voters rather than individual voters.[37] In an increasingly fragmented media landscape this could be achieved

[36] With a uniform distribution over preferences, the threshold can be obtained explicitly. For instance, irrespective of the support, for the case with selective targeted disclosure we have $\widetilde{V} = 153$.

[37] Indeed, candidates traditionally make highly targeted speeches at private events, such as the behind-closed-doors fundraiser where shortly before the 2012 election Mitt Romney was unwittingly recorded

by tailoring the campaign message to different channels that are frequented by voters with a particular orientation. Our previous analysis immediately extends to the case in which each of these V groups has identical size $z_v = z$.[38]

Unwary Voters. With collective decision-making, as in the case of voting, the consideration of unwary receivers has a particular, novel twist. Notably, we show now how unwary voters can have a negative *externality* on other voters, even though in our setting voters' preferences are independently drawn. As voters, in our setting, have also ex-ante symmetric preferences across candidates, an asymmetry is generated when only one candidate ($m = 1$) is sufficiently sophisticated to gather more information about voters and disclose (more) selectively. When a voter is unwary of this, from the resulting FOSD shift in the ex-ante distribution of his preferences it is immediate that this increases the likelihood with which she will vote for the respective candidate, i.e., $m = 1$, which we denote by $\widehat{q} > 1/2$.[39] As this further reduces the likelihood with which any voter becomes pivotal, compared to the case where the considered voter is wary, we have:

Proposition 11 *Consider the application to voting, where voters have a priori symmetric preferences over two candidates, and now suppose that V_u out of V voters remain unwary when one candidate discloses more selectively, while the other candidate cannot do so. Then, an increase in the number of unwary voters V_u decreases each voter's probability of becoming pivotal, and thus also reduces each voter's true expected welfare.*

suggesting that some 47% of Americans are government-dependent "victims" who do not pay taxes or take responsibility for their lives, and about whom "it's not my job to worry".

[38] As suggested by Mitt Romney's gaffe, an important drawback of targeted campaigns is the risk that voters exchange information about the different messages they receive. We leave this extension to future research; see, for example, Galeotti and Mattozzi (2011) for a model in which information sharing among voters reduces the incentives for information disclosure by candidates.

[39] Precisely, for selective targeted disclosure we have

$$\begin{aligned}\widehat{q} &= \int_{\underline{U}}^{\overline{U}} \left(1 - \widehat{I}(U)\right) dN(U) = \int_{\underline{u}+E[u^i]}^{\overline{u}+E[u^i]} \widehat{I}'(U)\, N(U) dU \\ &= \int_{\underline{u}+E[u^i]}^{\overline{u}+E[u^i]} 2f(U - E[u^i])F^2(U - E[u^i])dU = \left[\frac{2}{3}F^3(U - E[u^i])\right]_{\underline{u}+E[u^i]}^{\overline{u}+E[u^i]} = \frac{2}{3},\end{aligned}$$

so that, given ex-ante symmetry, when only one candidate can hypertarget in this way, an unwary voter's probability to vote for this candidate would thus always increase by $\frac{2}{3} - \frac{1}{2} = \frac{1}{6}$. For selective non-disclosure we have

$$q = \frac{1}{2}\frac{\theta^H}{\theta^L}\left[1 - N^2(u_d(\underline{\theta}))\right]$$

and

$$\widehat{q} = q + \frac{1}{2}N(u_d(\theta^L))\widehat{I}(u_d(\theta^L)).$$

6 Conclusion

The greater availability of personally identifiable data opens up new opportunities for tailoring advertising messages to the perceived preferences of particular consumers or voters. Online marketing is thus similar to more traditional personalized channels, where face-to-face interaction with consumers allows salespeople to learn about consumer preferences and to tailor their communication accordingly. Our simple model of equilibrium persuasion through selective disclosure can then be applied to a broad range of settings:

- First, the model applies to traditional communication strategies. As we mention in the opening paragraph of the introduction, old media only allow for a segmentation of receivers into coarse groups; nevertheless, different messages can be (and often are) sent to groups with different preferences.

- Second, consider a face-to-face interaction between a salesperson or campaigner and an individual consumer or voter. Even when meeting a consumer for the first time, an experienced salesperson should be able to draw inferences about the consumer's needs and preferences and use the limited time available (or the consumer's limited attention) to communicate only those product attributes that dovetail nicely with those preferences.

- A third relevant setting is distance selling and campaigning through communication channels that were previously anonymous but now allow for increased personalization, given the ability of senders to collect personally identifiable data on the internet. Based on an individual receiver's profile, a firm may choose how to best use the limited amount of time or space to selectively convey the attributes of an offering.

We now summarize our main findings. In our baseline case with individual decisions at fixed prices, wary receivers benefit from an increase in information quality, given that they rationally adjust for the resulting increased selectivity of disclosure (Proposition 2). However, senders' incentives to become better informed—as a basis of (more) selective disclosure—are more subtle. In the absence of policy intervention, we naturally assume that receivers *do not observe* the senders' choice of information acquisition. Thus, for given expectations by the receivers, senders have an incentive to become better informed because, off-equilibrium, they would increase the chances that their offering is chosen (Proposition 3). Even though senders' own incentives force them to become better informed, (more) selective disclosure ends up either benefitting or hurting them in equilibrium. Notably, senders benefit from increased information when competition is intense, but they are hurt

when competition is limited and their offerings are initially attractive (Propositions 4 and 5). In the latter case, policy intervention that makes information acquisition *observable*, for example by requiring consumer consent, hurts consumers because it allows senders to commit not to acquire information by not requiring consent.

The introduction of personalized pricing changes the outcome of our baseline analysis in important ways. The extent to which the efficiency gains associated with more informative communication are shared between firms and consumers depends on whether firms can price discriminate according to the perceived expected valuation of a particular consumer.[40] In the presence of competition, perceived product differentiation matters. Selective disclosure based on better information dampens competition by increasing perceived differentiation, from an ex-ante perspective. Firms always benefit as a result (Proposition 7), but the impact on consumer welfare depends on the intensity of competition (Proposition 8).

Next, consider political campaigning by two candidates vying for a majority of voters. While personal purchasing decisions depend on the preferences of each individual consumer in isolation, collective voting decisions depend on the aggregation of individual preferences. This key difference allows us to obtain a number of new insights. From the perspective of an individual voter, for instance, what matters now is the product of the probability with which the voter becomes pivotal and the utility conditional on this event. While the conditional utility always increases when more candidates are better informed about voters' preferences before communicating selectively, at least when voters are wary, asymmetric abilities in information acquisition across candidates can tilt the vote shares of different candidates, thus reducing the probability that any given voter becomes pivotal. As a result, Proposition 10 finds that voters are made better off when any given candidate is better informed also in the asymmetric case where his competitor is not equally sophisticated, only as long as the number of voters is not too large. The model also applies when there are differences of preferences or political orientation across groups of voters, so that selective communication targets different groups instead of each voter individually.

For policy purposes we also consider the case of receivers who remain unwary of the senders' ability and incentives to become better informed. When unwary receivers wrongly believe that senders are less informed, they do not properly discount the valuation for the adverse selection that is implicit in the fact that the disclosed attribute is the most favorable (in the case of selective targeted disclosure) or that lack of disclosure is more likely to result with low realizations (in the case of selective non-disclosure). We find that:

[40] Such price discrimination may only be feasible for services or low-value products, when customers or intermediaries have little scope for arbitrage.

- Unwary receivers can partly benefit from the increased information that sellers now *always* want to acquire, but may also lose because their choices become biased. However, in a fully symmetric setting, competition benefits unwary receivers particularly through the following mechanism (Proposition 6). Given that unwary receivers have inflated perceptions about the offerings by all senders, the distortion in the processing of information obtained from one sender is compensated by a similar distortion for competing firms, and exactly so under symmetry. Through this channel, competition eliminates bias and thus protects unwary consumers.

- In the case of personalized pricing, consumer ignorance about firms' capabilities to acquire detailed information about customer preferences reduces differentiation and can thereby spur rivalry among firms for a particular consumer. As a consequence, with competition and price discrimination, consumer unwariness can even become a blessing (Proposition 9). Overall, competition can protect unwary consumers from exploitation.

- For the model of campaigning, the presence of unwary voters has a negative externality on other voters (Proposition 11). When candidates differ in their capabilities to acquire information about individual voters' preferences, an increase in the number of unwary voters decreases each voter's probability of becoming pivotal, and thus also reduces each voter's true expected welfare.

Our results are obtained in a stylized model where senders can freely acquire and divulge information only about their own offerings in a private-value environment. First, we have abstracted away from learning about common values, which has been the focus of the literature on strategic voting (e.g., Feddersen and Pesendorfer 1996 and McMurray 2013). Second, an extension could allow senders to disclose information about their competitors. Third, information could be costly as in the law and economics literature on transparency,[41] or it could be sold by an information broker as in Taylor (2004). Overall, our analysis reveals a twist to Shapiro and Varian's (1997) policy of granting property rights over information. We show that requiring consumer consent may allow firms to commit to abstain from selective communication, which would have benefited consumers.[42]

[41] When the prime purpose of information is to affect the distribution of surplus, the incentives to collect information may be too high (Hirshleifer 1971). The literature on law and economics has also discussed more broadly the benefits of greater transparency for expanding efficiency-enhancing trade (Stigler 1980, Posner 1981). Hermalin and Katz (2006) show, however, that trade efficiency may not monotonically increase with information.

[42] A different twist on the costs of transparency has been recently offered in the marketing literature on targeted advertising, which allows firms to better restrict the scope of their marketing to those consumers

References

Abse, Nathan. 2013. *Big Data and Microtargeted Political Advertising in Election 2012: The Challenge Ahead.* Interactive Advertising Bureau, http://www.iab.net/political advertising

Acquisti, Alessandro, and Hal R. Varian. 2005. Conditioning Prices on Purchase History. *Marketing Science*, 24(3), 367–381.

Alonso, Ricardo, and Odilon Câmara. 2014. Persuading Voters. Unpublished.

American Association of Advertising Agencies. 2009. *Self-Regulatory Principles for Online Behavioral Advertising.* Report.

Anderson, Simon, and André de Palma. 2012. Competition for Attention in the Information (Overload) Age. *RAND Journal of Economics*, 43(1), 1–25.

Athey, Susan, and Joshua Gans. 2010. The Impact of Targeting Technology on Advertising and Media Competition. *American Economic Review*, 100(2), 608–613.

Bagnoli, Mark, and Ted Bergstrom. 2005. Log-concave Probability and its Applications. *Economic Theory*, 26(2), 445–469.

Bagwell, Kyle. 2007. The Economic Analysis of Advertising. In: Mark Armstrong and Robert Porter (eds.), *Handbook of Industrial Organization, Volume 3*, Chapter 28, 1701–1844.

Bhattacharya, Sourav, and Arijit Mukerjee. Strategic Information Revelation When Experts Compete to Influence. *RAND Journal of Economics*, 44(3), 522–544.

Board, Simon. 2009. Revealing Information in Auctions: The Allocation Effect. *Economic Theory*, 38(1), 125–135.

Butters, Gerard R. 1977. Equilibrium Distributions of Sales and Advertising Prices. *Review of Economic Studies*, 44(3), 465–491.

who are likely to purchase in the first place (see Athey and Gans 2010 for its impact on media competition). Several recent papers in marketing (e.g., Goldfarb and Tucker 2011; Campbell, Goldfarb, and Tucker 2011) analyze, both theoretically and empirically, how more restrictive privacy rights affect competition and welfare by potentially making advertising campaigns less cost-effective. Combined with the insights from our analysis, the protection of privacy rights should thus always be considered while taking into account competition and its benefits to consumers.

Caillaud, Bernard, and Jean Tirole. 2007. Consensus Building: How to Persuade a Group. *American Economic Review*, 97(5), 1877–1900.

Campbell, James, Avi Goldfarb, and Catherine Tucker. 2011. Privacy Regulation and Market Structure. Unpublished.

Chakraborty, Archishman, and Rick Harbaugh. 2007. Comparative Cheap Talk. *Journal of Economic Theory*, 132(1), 70–94.

Chakraborty, Archishman, and Rick Harbaugh. 2010. Persuasion by Cheap Talk. *American Economic Review*, 100(5), 2361–2382.

Che, Yeon-Koo, Wouter Dessein, and Navin Kartik. 2013. Pandering to Persuade. *American Economic Review*, 103(1), 47–79.

Coate, Stephen. 2004. Political Competition with Campaign Contributions and Informative Advertising. *Journal of the European Economic Association*, 2(5), 772–804.

DellaVigna, Stefano, and Matthew Gentzkow. 2010. Persuasion: Empirical Evidence. *Annual Review of Economics*, 2, 643–669.

Dessein, Wouter, Andrea Galeotti, and Tano Santos. 2012. Rational Inattention and Organizational Focus. Mimeo.

Dewan, Torun, Macartan Humphreys, and Daniel Rubenson. 2014. The Elements of Political Persuasion: Content, Charisma and Cue. *Economic Journal*, 124(574), F257–F292.

Dixit, Avinash and Victor Norman. 1978. Advertising and Welfare. *Bell Journal of Economics*, 9(1), 1–17.

Dye, Ronald A. 1985. Disclosure of Nonproprietary Information. *Journal of Accounting Research,* 23(1), 123–145.

Eyster, Erik, and Matthew Rabin. 2005. Cursed Equilibrium. *Econometrica*, 73(5), 1623–1672.

Feddersen, Timothy J., and Wolfgang Pesendorfer. 1996. The Swing's Voter Curse. *American Economic Review*, 86(3), 408–424.

Federal Trade Commission. 2012. *Protecting Consumer Privacy in an Era of Rapid Change. Recommendations for Businesses and Policy Makers.* http://www.ftc.gov/os/2012/03/120326privacyreport.pdf

Fishman, Michael J., and Kathleen M. Hagerty. 1990. The Optimal Amount of Discretion to Allow in Disclosure. *Quarterly Journal of Economics*, 105(2), 427–444.

Galeotti, Andrea, and Andrea Mattozzi. 2011. 'Personal Influence': Social Context and Political Competition. *American Economic Journal: Microeconomics*, 3(1), 307–327.

Ganuza, Juan-José, and José S. Penalva. 2010. Signal Orderings based on Dispersion and the Supply of Private Information in Auctions. *Econometrica*, 78(3), 1007-1030.

Gentzkow, Matthew, and Emir Kamenica. 2012. Competition in Persuasion. Unpublished. University of Chicago Booth School of Business.

Glazer, Jacob, and Ariel Rubinstein. 2004. On Optimal Rules of Persuasion. *Econometrica*, 72(6), 1715–1736.

Glaeser, Edward J., Giacomo A.M. Ponzetto and Jesse M. Shapiro. 2005. Strategic Extremism: Why Republicans and Democrats Divide on Religious Values. *Quarterly Journal of Economics*, 120(4), 1283–1330.

Goldfarb, Avi, and Catherine Tucker. 2011. Privacy Regulation and Online Advertising. *Management Science*, 57(1), 57–71.

Grossman, Gene M. and Carl Shapiro. 1984. Informative Advertising with Differentiated Products. *Review of Economic Studies*, 51(1), 63–81.

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

Guttman, Ilan, Ilan Kremer and Andrzej Skrzypacz. 2014. Not Only What but also When: A Theory of Dynamic Voluntary Disclosure. *American Economic Review*, 104(8), 2400–2420.

Hermalin, Benjamin, and Michael Katz. 2006. Privacy, Property Rights, and Efficiency: The Economics of Privacy as Secrecy. *Quantitative Marketing and Economics,* 4(3), 209–239.

Hirshleifer, David A., Sonya S. Lim, and Siew Hong Teoh. 2004. Disclosure to an Audience with Limited Attention. SSRN Working Paper.

Hirshleifer, Jack. 1971. The Private and Social Value of Information and the Reward to Inventive Activity. *American Economic Review*, 61(4), 561–574.

Holmström, Bengt. 1999. Managerial Incentive Problems: A Dynamic Perspective. *Review of Economic Studies*, 66(1), 169–182.

Jewitt, Ian. 1987. Risk Aversion and the Choice Between Risky Prospects: The Preservation of Comparative Statics Results. *Review of Economic Studies*, 54(1), 73–85.

Johnson, Justin P. 2013. Targeted Advertising and Advertising Avoidance. *RAND Journal of Economics*, 44(1), 128–144.

Johnson, Justin P., and David P. Myatt. 2006. On the Simple Economics of Advertising, Marketing, and Product Design. *American Economic Review*, 96(3), 756–784.

Jung, Woon-Ho and Young K. Kwon. 1998. Disclosure When the Market Is Unsure of Information Endowment of Managers. *Journal of Accounting Research*, 26(1), 146–153.

Kamenica, Emir, and Matthew Gentzkow. 2011. Bayesian Persuasion. *American Economic Review*, 101(6), 2590–2615.

Kamenica, Emir, Sendhil Mullainathan, and Richard Thaler. 2011. Helping Consumers Know Themselves. *American Economic Review: Papers and Proceedings*, 101(3), 417–422.

Kendall, Chad, Tommaso Nannicini, and Francesco Trebbi. 2013. How Do Voters Respond to Information? Evidence from a Randomized Campaign. NBER Working Paper No. 18986.

Lancaster, Kelvin. 1966. A New Approach to Consumer Theory. *Journal of Political Economy*, 74(2), 132–157.

Lewis, Tracy R., and David E. M. Sappington. 1994. Supplying Information to Facilitate Price Discrimination. *International Economic Review*, 35(2), 309–327.

McMurray, Joseph C. 2013. Aggregating Information by Voting: The Wisdom of the Experts versus the Wisdom of the Masses. *Review of Economic Studies*, 80(1), 277–312.

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

Milgrom, Paul, and John Roberts. 1986. Relying on Information of Interested Parties. *RAND Journal of Economics*, 17(1), 18–32.

Posner, Richard A. 1981. The Economics of Privacy. *American Economic Review*, 71(2), 405–409.

Prat, Andrea. 2006. Rational Voters and Political Advertising. In: Barry R. Weingast and Donald A. Wittman (eds.). *The Oxford Handbook of Political Economy*, Chapter 3, 50–63.

Rayo, Luis, and Ilya Segal. 2010. Optimal Information Disclosure. *Journal of Political Economy*, 118(5), 949–987.

Shavell, Steven. 1994. Acquisition and Disclosure of Information Prior to Sale. *RAND Journal of Economics*, 25(1), 20–36.

Shapiro, Carl, and Hal R. Varian. 1997. US Government Information Policy. Unpublished. University of California, Berkeley.

Simon, Herbert A. 1971. Designing Organizations for an Information-Rich World. In: Martin Greenberger (ed.), *Computers, Communication, and the Public Interest*, Baltimore, MD: The Johns Hopkins Press.

Sims, Christopher A. 2003. Implications of Rational Inattention. *Journal of Monetary Economics*, 50(3), 665–690.

Stigler, George J. 1980. An Introduction to Privacy in Economics and Politics. *Journal of Legal Studies*, 9(4), 623–644.

Taylor, Curtis. 2004. Consumer Privacy and the Market for Customer Information. *RAND Journal of Economics*, 35(4), 631–651.

Van Zandt, Timothy. 2004. Information Overload in a Network of Targeted Communication. *RAND Journal of Economics*, 35(3), 542–560.

Villas-Boas, Miguel J. 1999. Dynamic Competition with Customer Recognition. *RAND Journal of Economics*, 30(4), 604–631.

White House. 2012. *Consumer Data Privacy in a Networked World: A Framework for Protecting Privacy and Promoting Innovation in the Global Digital Economy.* Washington, DC, February. http://www.whitehouse.gov/sites/default/files/privacy-final.pdf

Appendix A: Proofs and Derivations

Proof of Proposition 1. We now suppress the subscript m denoting the respective sender. Note, first, that uniqueness of $u_d(\theta)$, as given in (1), follows from

$$\frac{dE[u \mid u \leq u']}{du'} \leq 1, \tag{7}$$

which holds if and only if $H(u') = \int_{\underline{u}}^{u'} F(u)du$ is logconcave, which in turn is implied by logconcavity of $F(u)$; see Bagnoli and Bergstrom (2005), in particular their Lemma 1 and Theorem 1. From (7) we also have that

$$\frac{du_d(\theta)}{d\theta} = -\frac{E[u] - E[u \mid u \leq u_d(\theta)]}{1 - \theta \frac{dE[u|u \leq u_d(\theta)]}{du_d(\theta)}} < 0. \tag{8}$$

To characterize $N(U)$ and $I(U)$, where θ is known to the receiver, note that in either case, with $\theta = \theta^L$ or $\theta = \theta^H$, we have for $u_d(\theta) \leq U \leq \overline{u}$ that $G(U) = (1-\theta) + \theta F(U)$, while at the lower bound there is a jump $j(\theta) = G(u_d(\theta))$, satisfying

$$\frac{dj(\theta)}{d\theta} = \theta \frac{du_d(\theta)}{d\theta} f(u_d(\theta)) - (1 - F(u_d(\theta)) < 0, \tag{9}$$

where we made use of (8). As we move from $\theta = \theta^L$ to $\theta = \theta^H$, single crossing thus follows from (8), (9), and from the fact that $G(U)$ is strictly decreasing in θ (with slope $F(U) - 1$) for all $U > u_d(\theta)$. Note also that the single crossing point is at $u_d(\theta^L)$, the lower support of $N(U)$.

Finally, when the receiver is unwary of the switch to $\theta = \theta^H$, we have that the distribution $\widehat{I}(U)$ is zero up to $u_d(\theta^L)$, then jumps to $1 - \theta^H \left(1 - F(u_d(\theta^L))\right)$, and is equal to $\left(1 - \theta^H\right) + \theta^H F(U)$ for $[u_d(\theta^L), \overline{u}]$. Thus FOSD holds, given that $\widehat{I}(U) < N(U)$ holds strictly for all $u \in [u_d(\theta^L), \overline{u})$, i.e., at all points of their joint support apart from the upper bound.[43] **Q.E.D.**

Proof of Proposition 2. The receiver's expected utility can be written as the following Lebesgue integral

$$E[U^{(1)}] = \int \left[\int \max\left\{U^{(1:M\backslash m)}, U_m\right\} dG^{(1:M\backslash m)}(U^{(1:M\backslash m)}) \right] dG_m(U_m). \tag{10}$$

Given that the expression in brackets is a convex function of U_m,[44] it is higher after a mean-preserving spread in $G_m(U_m)$. **Q.E.D.**

[43] All the results extend also to the case with an unbounded upper support $\overline{u} \to \infty$.

[44] It can be written as $U_m + \int_{U_m}^{\infty} [U^{(1:M\backslash m)} - U_m] dG^{(1:M\backslash m)}(U^{(1:M\backslash m)})$, which is a.e. differentiable with first derivative $G^{(1:M\backslash m)}(U_m)$, which is increasing.

Proof of Lemma 1. For any given distribution of the receiver's next best alternative $G^{(1:M\backslash m)}(\cdot)$, the difference in the likelihood q_m that option m is chosen when sender m switches from $N_m(\cdot)$ to $I_m(\cdot)$ is given by[45]

$$\begin{aligned}\Delta q_m &= \int_{\underline{U}}^{\overline{U}} G^{(1:M\backslash m)}(U_m) d\left[I_m(U_m) - N_m(U_m)\right] \\ &= \int_{\underline{U}}^{\overline{U}} \left[N_m(U_m) - I_m(U_m)\right] dG^{(1:M\backslash m)}(U_m),\end{aligned}$$

where the second line follows from integration by parts, which we can apply as $G^{(1:M\backslash m)}(\cdot)$ is presently assumed to be absolutely continuous. Now consider two choices for the distribution $G^{(1:M\backslash m)}(\cdot)$: $H'(\cdot)$ and $H''(\cdot)$, where the latter dominates in the likelihood ratio order. Further, denote the support of H' by $\left[\underline{U}', \overline{U}'\right]$ and by $\left[\underline{U}'', \overline{U}''\right]$ the respective support of H'', where we must have, from the LR shift, that $\underline{U}' \leq \underline{U}''$ as well as $\overline{U}' \leq \overline{U}''$. In what follows, we now only consider the non-trivial case where both $N_m(\cdot)$ as well as $I_m(\cdot)$ have strictly positive mass for $U_m \in \left[\underline{U}'', \overline{U}'\right]$, such that the probability of winning satisfies, independently of the information strategy, $q_m \in (0,1)$ both before and after the shift in the outside option.

Then, using $Z_m(U_m) = N_m(U_m) - I_m(U_m)$, which by the SCMPS is, in the interior of the respective supports, strictly negative for $U_m < \widetilde{U}_m$ and strictly positive for $U_m > \widetilde{U}_m$, we have, for the non-trivial case where $\widetilde{U}_m$ satisfies $\underline{U}'' \leq \widetilde{U}_m \leq \overline{U}'$,[46] that

$$\begin{aligned}\Delta q_m'' &= \int_{\underline{U}'}^{\overline{U}''} Z_m(U_m) dH''(U_m) \\ &= \int_{\underline{U}'}^{\widetilde{U}_m} Z_m(U_m) \frac{h''(U_m)}{h'(U_m)} dH'(U_m) \\ &\quad + \int_{\widetilde{U}_m}^{\overline{U}'} Z_m(U_m) \frac{h''(U_m)}{h'(U_m)} dH'(U_m) + \int_{\overline{U}'}^{\overline{U}''} Z_m(U_m) dH''(U_m) \\ &> \frac{h''(\widetilde{U}_m)}{h'(\widetilde{U}_m)} \left[\int_{\underline{U}'}^{\overline{U}'} Z_m(U_m) dG^{(1:M\backslash m)}(U_m)\right] = \frac{h''(\widetilde{U}_m)}{h'(\widetilde{U}_m)} \Delta q_m'\end{aligned}$$

where we have used that the likelihood ratio $\frac{h''(U_m)}{h'(U_m)}$ is increasing in U_m. This implies that whenever $\Delta q_m' \geq 0$ it must hold that $\Delta q_m'' > 0$. The converse result that $\Delta q_m'' \leq 0$ implies $\Delta q_m' < 0$ follows similarly. **Q.E.D.**

[45] The upper and lower bounds of the integral are chosen to contain the supports of both $N_m(\cdot)$ and $I_m(\cdot)$.

[46] If $\underline{U}'' > \widetilde{U}_m$, then $I_m(\cdot)$ first-order stochastically dominates $N_m(\cdot)$ on the relevant support, such that we must have $\Delta q_m'' > 0$. Similarly, if $\overline{U}' < \widetilde{U}_m$, we must have from the SCMPS that $\Delta q_m' < 0$.

Proof of Proposition 4. Assume that $N_m(U_m) = N(U_m)$ and $I_m(U_m) = I(U_m)$ for all m and consider, first, the case without mass points. For ease of exposition only we consider without loss of generality the choice of sender $m = 1$. Then, when all competing $M\backslash 1$ senders disclose more selectively the receiver's best alternative is distributed according to $G^{(1:M\backslash 1)}(U_1) = I^{M-1}(U_1)$. Thus, as the number of senders increases from M to $M+1$ we obtain for the likelihood ratio of receiver's "outside option" that

$$\frac{g^{(1:(M+1)\backslash 1)}(U_1)}{g^{(1:M\backslash 1)}(U_1)} = \frac{M}{(M-1)} I(U_1),$$

which is increasing in U_1. The result then follows from Lemma 1.

We turn now to the case where $G_m(U_m)$ can have a mass point, albeit at most at the lower boundary of the support. Choose again $G_m(U) = I(U)$ for all $m \in M\backslash 1$ and recall that from the properties of the SCMPS the lower support of $I(U)$, which we denote by $\underline{U}^I$, is (weakly) smaller than the respective lower support for $N(U)$, denoted by $\underline{U}^N \geq \underline{U}^I$. Defining $j = I(\underline{U}^I)$ and $k = N(\underline{U}^N)$, for the case where $\underline{U}^N = \underline{U}^I$, such that from the properties of SCMPS we must have $j \geq k$. Then, denoting the difference in the likelihood q_1 that option $m = 1$ is chosen when sender $m = 1$ chooses $I_1(U_1)$ instead of $N_1(U_1)$ as a function of the number of firms by $\Delta q_1(M)$, we can write

$$\begin{aligned}
\Delta q_1(M+1) &= \int_{\underline{U}}^{\overline{U}} Z(U) \frac{M}{(M-1)} I(U) dI^{M-1}(U) - (j-k) \frac{M}{M+1} j^M \\
&= \int_{\underline{U}}^{\overline{U}} Z(U) \frac{M}{(M-1)} I(U) dI^{M-1}(U) - \left(\frac{M}{M-1} \frac{M}{M+1} j \right) \frac{M-1}{M} (j-k) j^{M-1} \\
&> \frac{M}{M-1} I(\widetilde{U}) \left(\int_{\underline{U}}^{\widetilde{U}} Z(U) dI^{M-1}(U) + \int_{\widetilde{U}}^{\overline{U}} Z(U) dI^{M-1}(U) - \frac{M-1}{M} (j-k) j^{M-1} \right) \\
&= \frac{M}{M-1} I(\widetilde{U}) \Delta q_1(M),
\end{aligned}$$

where we have used the properties of the SCMPS and, from $j = I(\underline{U}^I)$, that $\frac{M}{M-1}\frac{M}{M+1} j < \frac{M}{M-1}\frac{M}{M+1} I(\widetilde{U}) < \frac{M}{M-1} I(\widetilde{U})$. For the remaining case where $\underline{U}^N > \underline{U}^I$, we similarly have

$$\begin{aligned}
\Delta q_1(M+1) &= \int_{\underline{U}}^{\overline{U}} Z(U) \frac{M}{(M-1)} I(U) dI^{M-1}(U) - \left(\frac{M}{M-1} \frac{M}{M+1} j \right) \frac{M-1}{M} j^M \\
&> \frac{M}{M-1} I(\widetilde{U}) \left(\int_{\underline{U}}^{\widetilde{U}} Z(U) dI^{M-1}(U) + \int_{\widetilde{U}}^{\overline{U}} Z(U) dI^{M-1}(U) - \frac{M-1}{M} j^M \right) \\
&= \frac{M}{M-1} I(\widetilde{U}) \Delta q_1(M).
\end{aligned}$$

For both cases the result then follows as $\Delta q_1(M) \geq 0$ implies $\Delta q_1(M+1) > 0$. That the threshold M' is finite follows from the more general argument in the proof of Proposition 5, where symmetry is no longer imposed. **Q.E.D.**

Proof of Proposition 5. For ease of exposition only we consider the choice of sender $m = 1$. Note, first, that when sender $m = 1$ switches from $N_1(U_1)$ to $I_1(U_1)$, the resulting difference in the likelihood q_1 is given by the following Lebesgue integral

$$\int_{\underline{U}}^{\overline{U}} w_1(U_1) d\left[I_1(U_1) - N_1(U_1)\right], \tag{11}$$

where we choose the two values $\underline{U}$ and $\overline{U}$ large enough so that the (finite) supports of both $N(\cdot)$ and $I(\cdot)$ are contained in $[\underline{U}, \overline{U}]$. Recall that $\widetilde{U}_1$ denotes the point of intersection of $I_1(U_1)$ and $N_1(U_1)$ and take now some value $\widetilde{U}_1 < \widehat{U} \leq \overline{U}$. Then, by extending (11) with $w_1(\widehat{U})$, it can be written as[47]

$$w_1(\widehat{U}) \left[\int_{[\underline{U}, \widehat{U})} \frac{w_1(U_1)}{w_1(\widehat{U})} d\left[I_1(U_1) - N_1(U_1)\right] + \int_{[\widehat{U}, \overline{U}]} \frac{w_1(U_1)}{w_1(\widehat{U})} d\left[I_1(U_1) - N_1(U_1)\right] \right]. \tag{12}$$

We will next show that (12) is strictly positive for all sufficiently large M. In doing so, we will distinguish between the following three cases: (1) The two upper supports of N and I coincide, i.e., $\overline{U}^N = \overline{U}^I = \overline{U}$, and I_1 has no mass point at $\overline{U}^I$; (2) $\overline{U}^N = \overline{U}^I = \overline{U}$ but I_1 has a mass point at $\overline{U}^I$; (3) $\overline{U}^N < \overline{U}^I$.

Case 1: Define $G_m^-(U_1) = \lim_{U \to U_1^-} G_m(U)$ and note that monotonicity of $G_m(U)$ implies that

$$\prod_{m' \in M \setminus 1} G_{m'}^-(U_1) \leq w_1(U_1) \leq \prod_{m' \in M \setminus 1} G_{m'}(U_1).$$

Hence, for any two $U_1'' > U_1'$ it holds that

$$\frac{w_1(U_1'')}{w_1(U_1')} \geq \prod_{m' \in M \setminus 1} \frac{G_{m'}^-(U_1'')}{G_{m'}(U_1')}. \tag{13}$$

From the SCMPS and the properties of Case 1, there exists $\widehat{U} < \overline{U}$ so that the following holds. First, we have for $U_1 \geq \widehat{U}$ that $dI(U_1) > dN(U_1)$. Second, defining for any $U_1 < \widehat{U}$

$$\varphi_\omega(U_1) = \frac{G_\omega(U_1)}{G_\omega^-(\widehat{U})},$$

we have $\varphi(U_1) = \max_{\omega \in \Omega} \varphi_\omega(U_1) < 1$. So, applying (13), it surely must hold for any $U_1 < \widehat{U}$ that

$$\frac{w_1(U_1)}{w_1(\widehat{U})} \leq \left[\varphi(U_1)\right]^{M-1}.$$

[47] Note that both $I_1(\cdot)$ and $N_1(\cdot)$ are right-continuous functions.

That $w_1(U_1)/w_1(\widehat{U}) \to 0$ uniformly as $M \to \infty$ follows then from the observations that $\varphi(U_1) \in [0,1)$ and that U_1 is taken from a bounded set. As a consequence, if we now consider a sequence of markets, indexed by M, we have

$$\lim_{M\to\infty} \int_{[\underline{U},\widehat{U})} \frac{w_{1,M}(U_1)}{w_{1,M}(\widehat{U})} d\left[I_1(U_1) - N_1(U_1)\right] \to 0. \tag{14}$$

Making now use of $\widehat{U} < \overline{U}$, $dI(U_1) > dN(U_1)$ for all $\widehat{U} \leq U_1 < \overline{U}$, as well as obviously $\frac{w_1(U_1)}{w_1(\widehat{U})} \geq 1$ over this range, there surely exists a value $\varepsilon > 0$ such that along any considered sequence we have

$$\int_{[\widehat{U},\overline{U}]}^{\overline{U}} \frac{w_{1,M}(U_1)}{w_{1,M}(\widehat{U})} d\left[I_1(U_1) - N_1(U_1)\right] > \varepsilon. \tag{15}$$

Substituting (14) and (15) into (12) proves the assertion for Case 1.

Case 2: With $\overline{U}^N = \overline{U}^I = \overline{U}$ but $I(U_1)$ having a mass point at the upper bound, we choose $\widehat{U} = \overline{U}$. From the SCMPS, (15) is then again immediate. (In particular, in case also $N(U_1)$ has a mass point at $\overline{U}$, then this is strictly smaller.) Next, as $\overline{U}$ is the upper bound of both supports, also the preceding argument for all $U_1 < \widehat{U}$ applies, so that also (14) holds.

Case 3: With $\overline{U}^N < \overline{U}^I$, we set $\widehat{U} = \overline{U}^N$. The first part of the argument, as in Case 1, now applies as $w_1(U_1)/w_1(\widehat{U}) \to 0$ holds, using convexity of the supports, for all $U_1 < \widehat{U}$. Hence, (14) again applies. We come next to the second part of the argument. Using monotonicity of $w_1(U_1)$, we can write

$$\int_{[\widehat{U},\overline{U}]} \frac{w_1(U_1)}{w_1(\overline{U}^N)} d\left[I_1(U_1) - N_1(U_1)\right] \geq \left[1 - I_1^-(\overline{U}^N)\right] - \left[1 - N_1^-(\overline{U}^N)\right],$$

which is strictly positive from the SCMPS. Thus, also (15) holds along any considered sequence with $M \to \infty$. This completes the proof. **Q.E.D.**

Proof of Proposition 6. We will prove the result separately for the two cases with selective targeted disclosure and selective non-disclosure. Consider first selective targeted disclosure. In this case senders always disclose $u_{d_m} = \max\{u_m^1, u_m^2\}$ irrespective of whether they face a wary or an unwary receiver. As both the true expected valuation, $u_{d_m} + E[u_m^i | u_m^i \leq u_{d_m}]$, as well as the one perceived by the unwary receiver, $u_{d_m} + E[u_m^i]$, are strictly increasing in u_{d_m}, and, by symmetry, for a given u_{d_m} constant over m, an unwary receiver's decision rule is the same as that of a wary receiver, namely to choose the option where the respective disclosed value u_{d_m} is maximal. Thus, wary and unwary receivers realize the same expected utility when all senders disclose an attribute selectively. Hence,

we conclude from Proposition 2 that unwary consumers, like wary consumers, are better off with selective disclosure.

Next, consider the case of selective non-disclosure. Under $G_m(U_m) = \widehat{I}(U_m)$, with $s_m = H$ senders still disclose when $u_m \geq u_d(\theta^L)$. Also, given symmetry, when no firm chooses to disclose, the unwary receiver chooses one of the M options at random, while when at least one sender chooses to disclose, she correctly chooses the alternative that gives her the highest true expected utility. The claim then follows immediately as the probability of disclosure increases for each sender if θ switches from θ^L to θ^H. **Q.E.D.**

Proof of Proposition 8. To simplify the exposition, without loss of generality we consider again the choice of firm $m = 1$ and thus a switch from $G_1(U_1) = N_1(U_1)$ to $G_1(U_1) = I_1(U_1)$. Note also that the case with $M = 2$ was already fully solved in the main text and that, presently, we are interested in the case for high M, which is why without loss of generality we can assume that $M \geq 3$. Denote by $v_1(U_1)$ a consumer's expected utility for given U_1, so that the respective difference in ex-ante utility is given by

$$\int_{\underline{U}}^{\overline{U}} v_1(U_1) d\left[I_1(U_1) - N_1(U_1)\right]. \tag{16}$$

Next, note that $v_1(U_1)$ is continuous and almost everywhere differentiable with left-hand side derivative $\lim_{U \to U_1^-} v_1'(U) = \lim_{U \to U_1^-} G^{(2:M\setminus 1)}(U_1) - G^{(1:M\setminus 1)}(U_1) = \eta(U_1)$, where

$$\eta(U_1) = \sum_{m \in M\setminus 1} \left(\left[1 - G_m^-(U_1)\right] \prod_{m' \notin \{1,m\}} G_{m'}^-(U_1) \right).$$

Hence, we can transform (16), using integration by parts, to obtain

$$\int_{\underline{U}}^{\overline{U}} \eta(U_1)[N_1(U_1) - I_1(U_1)] dU_1. \tag{17}$$

This generalizes expression (3) in the main text, where for ease of exposition it was assumed that all distributions are continuous. In order to show (17) to be strictly positive for high M, we distinguish two cases: 1) $\widetilde{U}_1 < \overline{U}^N$ and 2) $\widetilde{U}_1 = \overline{U}^N$.

Case 1: First, assume that $\widetilde{U}_1 < \overline{U}^N$ and define $\phi(U_1) = \eta(U_1)/\eta(\widetilde{U}_1)$. Then, extending the expression in (17), for each U_1, by multiplying and dividing with the term $\eta(\widetilde{U}_1)$, and noting that $\eta(\widetilde{U}_1) > 0$, a sufficient condition for (17) to be greater than zero is that

$$\int_{[\underline{U},\widetilde{U}_1)} \left[N_1(U_1) - I_1(U_1)\right] \phi(U_1) dU_1 + \int_{[\widetilde{U}_1,\overline{U}]} \left[N_1(U_1) - I_1(U_1)\right] \phi(U_1) dU_1 > 0. \tag{18}$$

Next, define, using finiteness of Ω, for any U_1, $\overline{\varsigma} = \max_{\omega\in\Omega}[1-G_\omega(U_1)]$ and $\underline{\varsigma} = \min_{\omega\in\Omega}[1-G^-_\omega(\widetilde{U}_1)] > 0$ and note that for $U_1 < \widetilde{U}_1$ we have

$$\phi(U_1) \leq \left(\frac{\overline{\varsigma}}{\underline{\varsigma}}\right) \frac{\sum_{m\in M\backslash 1} \prod\limits_{m'\notin\{1,m\}} G^-_{m'}(U_1)}{\sum_{m\in M\backslash 1} \prod\limits_{m'\notin\{1,m\}} G^-_{m'}(\widetilde{U}_1)}.$$

Then, using $\varkappa(U_1) = \max_{\omega\in\Omega} \frac{G_\omega(U_1)}{G^-_\omega(\widetilde{U}_1)} < 1$ for $U_1 < \widetilde{U}_1$, it holds that $\phi(U_1) \leq \left(\frac{\overline{\varsigma}}{\underline{\varsigma}}\right) \varkappa^{M-2}(U_1)$. From this, together with the observations that U_1 is chosen from a bounded set and $\varkappa(U_1) \in [0,1)$, it follows that $\phi(U_1) \to 0$ uniformly as $M \to \infty$. It then follows by the same arguments as in the proof of Proposition 5 that the first integral in (18) goes to zero as $M \to \infty$. It remains to show that the second integral remains bounded away from zero (and positive). Clearly, for $U_1 > \widetilde{U}_1$ we have from the SCMPS that $N_1(U_1) > I_1(U_1)$ such that it is sufficient to show that the weights $\phi(U_1)$ do not go to zero on a set of positive measure. But this follows immediately from the fact that, for any $\widetilde{U}_1 < U_1 < \overline{U}^N$ in the interior of all supports (Case 1), we have

$$\phi(U_1) \geq \left(\frac{\underline{\varsigma}'(U_1)}{\overline{\varsigma}'(U_1)}\right) \frac{\sum_{m\in M\backslash 1} \prod\limits_{m'\notin\{1,m\}} G^-_{m'}(U_1)}{\sum_{m\in M\backslash 1} \prod\limits_{m'\notin\{1,m\}} G^-_{m'}(\widetilde{U}_1)},$$

where $\underline{\varsigma}'(U_1) = \min_{\omega\in\Omega}[1-G^-_\omega(U_1)] > 0$ and $\overline{\varsigma}'(U_1) = \max_{\omega\in\Omega}[1-G^-_\omega(\widetilde{U}_1)]$. Then, using $\kappa(U_1) = \min_{\omega\in\Omega} \frac{G^-_\omega(U_1)}{G^-_\omega(\widetilde{U}_1)} > 1$, it holds that $\phi(U_1) \geq \frac{\underline{\varsigma}'(U_1)}{\overline{\varsigma}'(U_1)}\kappa^{M-2}(U_1) > \frac{\underline{\varsigma}'}{\overline{\varsigma}'} > 0$.

Case 2: Here it holds that $\widetilde{U}_1 = \overline{U}^N$. From the SCMPS, this case can only arise when $\overline{U}^N < \overline{U}^I$ and N_1 has a jump at $\widetilde{U}_1 = \overline{U}^N$. Then, assume, first, that $G_m = N_m$ for all $M\backslash m$, implying that $\eta(U_1) = 0$ for all $U_1 > \overline{U}^N$. Hence, we can write (17) as

$$\int_{[\underline{U},\widetilde{U}_1)} \eta(U_1)[N_1(U_1) - I_1(U_1)]dU_1 \leq 0.$$

A necessary condition for (17) to be strictly positive for high M is that the set $M^* \subseteq M\backslash 1$ where $G_{m^*} = I_{m^*}$ for all $m^* \in M^*$ is non-empty. We will show next that this is also sufficient. To see this take any $\widetilde{U}_1 = \overline{U}^N < \widehat{U} < \overline{U}^I$, and extend the expression in (17), for each U_1, by multiplying and dividing with the term $\eta(\widehat{U}) > 0$. Then, in analogy to (18), a sufficient condition for (17) to be greater than zero is that

$$\int_{[\underline{U},\widehat{U})} [N_1(U_1) - I_1(U_1)]\, \phi(U_1)dU_1 + \int_{[\widehat{U}_1,\overline{U}]} [N_1(U_1) - I_1(U_1)]\, \phi(U_1)dU_1 > 0. \tag{19}$$

Clearly, the first integral in (19) still goes to zero as $M \to \infty$ by the same argument as in the proof of case 1. To show that the second integral is bounded away from zero (and positive), from $N_1(U_1) - I_1(U_1) > 0$ for all $\widehat{U} < U_1 < \overline{U}^I$, it suffices to show that the weights $\phi(U_1)$ do not go to zero on a set of positive measure. But, using $\underline{\varsigma}''(U_1) = \min_{m \in M^*}[1 - I_m^-(U_1)] > 0$ and $\overline{\varsigma}''(U_1) = \max_{m \in M^*}[1 - I_m^-(\widehat{U})]$, this follows immediately from

$$\phi(U_1) \geq \left(\frac{\underline{\varsigma}''(U_1)}{\overline{\varsigma}''(U_1)}\right) \frac{\sum_{m \in M^*} \prod\limits_{m' \in M \setminus \{1,m\}} G_{m'}^-(U_1)}{\sum_{m \in M^*} \prod\limits_{m' \in M \setminus \{1,m\}} G_{m'}^-(\widehat{U})} \geq \left(\frac{\underline{\varsigma}''}{\overline{\varsigma}''}\right) \nu^{M-2}(U_1),$$

where $\nu(U_1) = \min_{\omega \in \Omega} \frac{G_\omega^-(U_1)}{G_\omega^-(\widehat{U})} \geq 1$. **Q.E.D.**

Proof of Proposition 11. Note first that the expected utility of a pivotal voter, no matter whether she is wary or unwary, is independent of the number of unwary voters V_u. Hence, V_u can only affect expected voter welfare through its effect on the probability of becoming pivotal. Recall that we denote the probability that a wary voter elects candidate $m = 1$ by q and the respective probability for an unwary voter by $\widehat{q}$. As $\widehat{I}(U)$ dominates $N(U)$ in the sense of strict FOSD, we must have that $\widehat{q} > \frac{1}{2}$.

We show finally that $\widehat{q} > q \geq 1/2$ implies that the likelihood of becoming pivotal decreases with V_u. To see this, pick an arbitrary voter v' and determine how the probability of the remaining $V - 1$ voters generating a draw changes when a single voter $v \neq v'$ is wary versus when v is unwary. To do so, decompose the probability of a draw among these $V - 1$ voters according to the voting decision of voter v; denote the number of votes out of $V \backslash \{v, v'\}$ cast for candidate $m = 1$ by V_1 and the number of votes cast for $m = 2$ by V_2. Then, if v votes for $m = 1$, the probability of a draw is given by $X = \Pr\left(V_1 = \frac{V-1}{2} - 1, V_2 = \frac{V-1}{2}\right)$. If v votes for $m = 2$, the respective probability is given by $Y = \Pr\left(V_1 = \frac{V-1}{2}, V_2 = \frac{V-1}{2} - 1\right)$. Denoting the probability with which voter v votes for $m = 1$ by q_v, the total probability of a draw among these $V - 1$ voters is thus $q_v X + (1 - q_v) Y$. From $\widehat{q} > q \geq 1/2$ it follows that $X < Y$[48] and that $q_v X + (1 - q_v) Y$ is larger for $q_v = q$ than for $q_v = \widehat{q} > q$. **Q.E.D.**

[48] Note that for each possible outcome with $V_1 = (V-1)/2 - 1$ and $V_2 = (V-1)/2$ there exists a respective outcome with $V_1 = (V-1)/2$ and $V_2 = (V-1)/2 - 1$ where all but one voter (call this voter v'') take the same election decision. But as v'' is more likely to vote for $m = 1$, the result follows.

Appendix B: Selective *Targeted* Disclosure

This appendix provides additional details for the case of selective targeted disclosure. It is convenient to suppress the subscript m. When disclosure is not selective (as the sender did not learn receiver preferences), the receiver's perceived utility will be $U = u_d + E[u^i]$, so that the ex-ante distribution $G(U)$ equals $N(U) = F(U - E[u^i])$. Instead, an informed sender will disclose the attribute with the highest "fit", $d = \arg\max_{i=1,2} u^i$.[49] The receiver, when aware of this, will discount the value of the non-disclosed attribute and obtain perceived utility

$$U = u_d + E[u^i \mid u^i \leq u_d].$$

This gives rise to a unique and increasing value $u_d(U)$, which retrieves the value that the receiver must have learned for a given perceived utility U to arise. The perceived utility of a receiver who is aware that the sender discloses selectively the attribute with the *highest* value is, thus, distributed according to $I(U) = F^2(u_d(U))$. Instead, the perception about the non-disclosed attribute remains unchanged if the receiver is unaware that the sender discloses selectively, so that $U = u_d + E[u^i]$ with distribution $\widehat{I}(U) = F^2(U - E[u^i])$.

Take now the case of a uniform distribution (as in Figure 1). With support $[\underline{u}, \overline{u}]$, $N(U)$ and $I(U)$ are given by

$$N(U) = \frac{2U - 3\underline{u} - \overline{u}}{2(\overline{u} - \underline{u})} \text{ and } I(U) = \left(\frac{2(U - 2\underline{u})}{3(\overline{u} - \underline{u})}\right)^2, \tag{20}$$

which clearly can have at most two intersections. It is easily verified that these occur at the shared upper bound of the support $\overline{U} = \overline{u} + E[u^i] = (3\overline{u} + \underline{u})/2$ and at $\widetilde{U} = (3\overline{u} + 5\underline{u})/4$, with $\underline{u} + E[u^i] < \widetilde{U} < \overline{u} + E[u^i]$. For completeness, note also that

$$\widehat{I}(U) = \left(\frac{2U - 3\underline{u} - \overline{u}}{2(\overline{u} - \underline{u})}\right)^2. \tag{21}$$

Claim 1 (Appendix B). *When attribute values u^i are uniformly distributed, then the respective distributions of perceived utility (as given explicitly in (20) and (21)) are related by a SCMPS ($N(U)$ and $I(U)$) and FOSD ($N(U)$ and $\widehat{I}(U)$).*

Next, we formally support the claim that logconcavity of $F(u^i)$ is generally sufficient to ensure that $I(U)$ has more mass than $N(U)$ also in the upper tail, i.e., that there exists $U' < \overline{U}$ such that $N(U) > I(U)$ for all $U \geq U'$.[50] With a slight abuse of notation, denote

[49] A strategy where the sender does not disclose any attribute, but where the receiver knows that he is informed, would also not arise in equilibrium due to a standard unraveling argument.

[50] For the respective comparison in the lower tail, recall from the main text that the lower bound of the support of $I(U)$ equals $2\underline{u}$ which is strictly smaller than the respective value of $2\underline{u} + E[u^i]$ in case of $N(U)$.

for given realized value U the corresponding disclosed values by $u_N = U - E[u^i]$ and $u_I = u_d(U)$ respectively, where at $U = \overline{U}$ we have $u_I = u_N = \overline{u}$. The densities of the two distributions are then, generally,

$$\begin{aligned} n(U) &= f(u_N), \\ i(U) &= f(u_I)\frac{2F(u_I)}{1+\frac{dE[u^i|u^i\leq u_I]}{du_I}}. \end{aligned}$$

Together with $u_I = u_N$ at $U = \overline{U}$, next to $F(u_I = \overline{u}) = F(u_N = \overline{u}) = 1$, this yields $n(\overline{U}) < i(\overline{U})$ *if and only if*, at this point, condition (7) holds (with $u' = u_I$). This is implied by logconcavity of $F(u^i)$.

Claim 2 (Appendix B). *$I(U)$ has more mass also in the upper tail, i.e., $I(U) < N(U)$ holds for all sufficiently high $U < \overline{U}$.*

As noted in the main text, we can also show the SCMPS for a case where the support of the distribution is unbounded, namely for that where u^i is distributed exponentially with parameter λ.

Claim 3 (Appendix B). *When attribute values u^i are exponentially distributed (and thus have unbounded support), the distributions $N(U)$, $I(U)$, and $\widehat{I}(U)$ are still related by a SCMPS or FOSD, respectively.*

Proof. As $\widehat{I}(U)$ always FOSD dominates $N(U)$, we can restrict ourselves to a comparison of $N(U)$ and $I(U)$, which are given by $N(U) = 1 - e^{1-\lambda U}$ and $I(U) = \left(1 - e^{-\lambda u_d(U)}\right)^2$, where $u_d(U)$ solves

$$U = u_d + \frac{1}{\lambda} - \frac{e^{-\lambda u_d}}{1-e^{-\lambda u_d}}u_d. \tag{22}$$

For $Z(U) = N(U) - I(U)$ we have

$$Z(U) = -\exp\{-\lambda u_d(U)\}\left[\exp\{-\lambda u_d(U)\} + \exp\left\{\lambda\frac{e^{-\lambda u_d(U)}u_d(U)}{1-e^{-\lambda u_d(U)}}\right\} - 2\right].$$

As $u_d(U)$ is strictly monotonic in U, $Z(U) \to 0$ as $U \to \infty$, while for bounded U, $Z(U) = 0$ holds if and only if the term in brackets is equal to zero, i.e., in case u_d solves

$$\exp\{-\lambda u_d\} + \exp\left\{\lambda\frac{e^{-\lambda u_d}u_d}{1-e^{-\lambda u_d}}\right\} = 2. \tag{23}$$

To see that a solution to this must be unique, we have for the derivative w.r.t. u_d

$$-\lambda\exp\{-\lambda u_d\} - \lambda e^{-\lambda u_d}\frac{e^{-\lambda u_d} + \lambda u_d - 1}{\left(1-e^{-\lambda u_d}\right)^2}\exp\left\{\lambda u_d\frac{e^{-\lambda u_d}}{1-e^{-\lambda u_d}}\right\},$$

which is clearly negative if $\upsilon(u_d) = e^{-\lambda u_d} + \lambda u_d \geq 1$. This follows from $\upsilon'(u_d) = \lambda\left(1 - e^{-\lambda u_d}\right) > 0$ together with $\upsilon(u_d = 0) = 1$. (Incidentally, from $Z(U = 1/\lambda) = -I(1/\lambda) < 0$ we can also establish that the point of intersection satisfies $\frac{1}{\lambda} < \widetilde{U} < \infty$.) **Q.E.D.**

We conclude the discussion of selective targeted disclosure with an analysis of the case where attributes have different weights:

$$u = \alpha^1 u^1 + \alpha^2 u^2.$$

Without loss of generality we stipulate $\alpha^1 > \alpha^2 > 0$. We focus on the tractable case where, as in Claim 1, u^1 and u^2 are uniformly distributed on $[\underline{u}, \overline{u}]$. When the considered sender observes the receiver's preferences and can thus practice targeted selective disclosure, we restrict attention to the characterization of a rational expectations equilibrium where the disclosure rule is linear: $d = 1$ whenever $u^1 \geq a + bu^2$. If this rule is rationally anticipated by the receiver, then choosing $d = 1$ is indeed optimal if and only if

$$\alpha^1 u^1 + \alpha^2 E\left[u^2 | u^2 \leq \frac{u^1 - a}{b}\right] \geq \alpha^2 u^2 + \alpha^1 E\left[u^1 | u^1 \leq a + bu^2\right],$$

which can be transformed to obtain

$$u^1 \geq \frac{(\alpha^1 - \alpha^2)}{\alpha^1}\underline{u} + \frac{\alpha^2}{\alpha^1}u^2. \tag{24}$$

If (24) does not hold, $d = 2$ is disclosed. With this rule at hand, after disclosing $d = 1$ the expected utility equals

$$U = \frac{3}{2}\alpha^1 u^1 - \frac{1}{2}\left(\alpha^1 - 2\alpha^2\right)\underline{u},$$

so that $U \in \left[(\alpha^1 + \alpha^2)\,\underline{u}, \frac{3}{2}\alpha^1\overline{u} - \frac{1}{2}\left(\alpha^1 - 2\alpha^2\right)\underline{u}\right]$. With $d = 2$ we obtain

$$U = \frac{3}{2}\alpha^2 u^2 + \frac{1}{2}\left(2\alpha^1 - \alpha^2\right)\underline{u},$$

so that now $U \in \left[(\alpha^1 + \alpha^2)\,\underline{u}, \frac{3}{2}\alpha^2\overline{u} + \frac{1}{2}\left(2\alpha^1 - \alpha^2\right)\underline{u}\right]$. Note that from $\alpha^1 \geq \alpha^2$, which we stipulated without loss of generality, the highest value of U is attained when disclosing $u^1 = \overline{u}$, $\overline{U} = \frac{3}{2}\alpha^1\overline{u} - \frac{1}{2}\left(\alpha^1 - 2\alpha^2\right)\underline{u}$, while $\underline{U} = (\alpha^1 + \alpha^2)\,\underline{u}$. The following characterization can now be obtained after some calculations:

$$I(U) = \begin{cases} \frac{1}{\alpha^1\alpha^2}\left(\frac{2\left(U - \left(\alpha^1 + \alpha^2\right)\underline{u}\right)}{3(\overline{u} - \underline{u})}\right)^2 & \text{for } \underline{U} \leq U \leq U' \\ \frac{1}{\alpha^1}\frac{2\left(U - \left(\alpha^1 + \alpha^2\right)\underline{u}\right)}{3(\overline{u} - \underline{u})} & \text{for } U' < U \leq \overline{U} \end{cases} \tag{25}$$

where

$$U' = \frac{3}{2}\alpha^2\overline{u} + \frac{1}{2}\left(2\alpha^1 - \alpha^2\right)\underline{u}, \tag{26}$$

with $U' \in \left(\underline{U}, \overline{U}\right)$ for $\alpha^1 > \alpha^2$.

We next derive $\widehat{I}(U)$ (where the receiver is not wary of the fact that the sender observes her preferences before disclosure). In this case, for the sender it is optimal to choose $d = 1$, so as to maximize the perceived valuation, when

$$\alpha^1 u^1 + \alpha^2 E[u^2] \geq \alpha^1 E[u^1] + \alpha^2 u^2,$$

which transforms to

$$u^1 \geq \frac{(\alpha^1 - \alpha^2)}{\alpha^1}\frac{(\overline{u} + \underline{u})}{2} + \frac{\alpha^2}{\alpha^1}u^2,$$

and otherwise to disclose $d = 2$. Again after some tedious calculations we obtain:

$$\widehat{I}(U) = \begin{cases} \frac{\left[2U - \left(\alpha^1 + \alpha^2\right)\left(\frac{(\overline{u}+\underline{u})}{2} + \underline{u}\right)\right]^2 - \left[\left(\alpha^1 - \alpha^2\right)\left(\frac{(\overline{u}+\underline{u})}{2} - \underline{u}\right)\right]^2}{4\alpha^1\alpha^2(\overline{u}-\underline{u})^2} & \text{for } \underline{\widehat{U}} \leq U \leq \widehat{U}' \\ \frac{2U - \alpha^2(\overline{u}+\underline{u}) - 2\alpha^1\underline{u}}{2\alpha^1(\overline{u}-\underline{u})} & \text{for } \widehat{U}' < U \leq \overline{\widehat{U}} \end{cases} \tag{27}$$

where

$$\underline{\widehat{U}} = \alpha^2\underline{u} + \alpha^1\frac{(\overline{u} + \underline{u})}{2}, \tag{28}$$

$$\overline{\widehat{U}} = \alpha^1\overline{u} + \alpha^2\frac{(\overline{u} + \underline{u})}{2}, \tag{29}$$

$$\widehat{U}' = \alpha^2\overline{u} + \alpha^1\frac{(\overline{u} + \underline{u})}{2}, \tag{30}$$

with $\widehat{U}' \in \left(\underline{\widehat{U}}, \overline{\widehat{U}}\right)$ for $\alpha^1 > \alpha^2$.

What now complicates the analysis is that with unequal weights $\alpha^1 \neq \alpha^2$ the sender's strategy is no longer immediate even when he does not observe the receiver's preferences. This is despite the fact that the sender arguably still applies the same disclosure rule to each receiver. Without loss of generality, we can limit the sender's strategies to always disclosing the first attribute *or* to always disclosing the second attribute. For our subsequent derivations we need not determine which one is optimal. When the sender discloses $d = 1$, then

$$N(U) = \frac{2U - \alpha^2\left(\overline{u} + \underline{u}\right) - 2\alpha^1\underline{u}}{2\alpha^1\left(\overline{u} - \underline{u}\right)} \tag{31}$$

for $U \in \left[\alpha^1\underline{u} + \alpha^2\frac{(\overline{u}+\underline{u})}{2}, \alpha^1\overline{u} + \alpha^2\frac{(\overline{u}+\underline{u})}{2}\right]$. When he discloses $d = 2$, then

$$N(U) = \frac{2U - \alpha^1\left(\overline{u} + \underline{u}\right) - 2\alpha^2\underline{u}}{2\alpha^2\left(\overline{u} - \underline{u}\right)} \tag{32}$$

for $U \in \left[\alpha^2 \underline{u} + \alpha^1 \frac{(\overline{u}+\underline{u})}{2}, \alpha^2 \overline{u} + \alpha^1 \frac{(\overline{u}+\underline{u})}{2}\right]$.

Based on the derived expressions, one can show the following result. (The proof is available from the authors upon request.)

Claim 4 (Appendix B). *For the case with targeted selective disclosure where u^i is uniformly distributed but where the receiver applies different weights, $\alpha^1 \neq \alpha^2$, the distributions for the receiver's perceived utility are still related by a SCMPS or FOSD, respectively.*