Truthful Communication Prior to a Sale (preliminary and incomplete)

Jeremy Bertomeu*

Abstract

This paper offers a unified treatment of truthful communication prior to a sale under incomplete information. Among other things, the framework accommodates general reporting costs, information endowments, message spaces or possible uses of information. I show that this class of problems admits at most a unique perfect sequential rational expectations equilibrium (PRE) and provide a tractable methodology to characterize it. I provide conditions for the existence of the PRE which nest common settings and a class of problems involving selective disclosure over multiple pieces of information. Necessary and sufficient conditions are provided for unravelling to a fully-revealing equilibrium and I show that this unravelling property has a representation that is equivalent to a breakdown in a lemon's market. Then, when unravelling does not hold, I examine the social desirability of rules that reduce the level of discretion in communication. Lastly, I analyze an application in which the seller observes, and can disclose, ranges over the true state; the model casts doubt on several commonly-accepted conjectures in the area of truthful disclosure, e.g., that unfavorable information should be withheld or that an increase in costs or a reduction in information endowment would necessarily reduce communication.

.

Keywords: Voluntary, disclosure, truthful, persuasion, unravelling, market.

Models of communication generally fall into one of three broad families, which emphasize one channel through which a self-interested agent (sender) making the communication may transmit information to another party. In the *signalling* family, holding the response to the information fixed, the sender only has private information about her preferences over possible communications. In the *cheap talk* family, holding the communication fixed, the sender has private information about her preference over the possible responses to the information. In the *truthful disclosure* family, the sender is privately informed about the feasibility of making certain communications.

*Stan Ross Department of Accountancy. Baruch College, City University of New York, One Bernard Baruch Way, New York, NY 10010. I acknowledge financial support from the PSC-CUNY research fund. I thank Judson Caskey, Edwige Cheynel, Masako Darrough, Ron Dye, Prasart Jongjaroenkamol, Steve Kachelmeier, Bill Kinney, Volker Laux, Chen Li, Bob Magee, Ivan Marinovic, Lillian Mills, Igor Vaysman as well as other seminar participants at the University of Texas at Austin and at the internal workshop at Baruch College. Contact author: J. Bertomeu, email: jeremy.bertomeu@gmail.com.

While signalling and cheap talk have been the object of significant theoretical research, prior literature in the area of disclosure has primarily focused on models within specific applied settings. Existing theories of disclosure require a detailed description of the information known to the players, their preferences and the consequences of their disclosures and there is no general theoretical framework for models involving truthful disclosure.

This paper offers to fill this gap and presents a tractable methodology to analyze pure truthful disclosure in the context of a sale. By pure truthful disclosure, I assume here that the sender cannot signal her information as a result of her differential preference over messages or transmit information via unverifiable cheap talk. Throughout the paper, for expositional reasons, I use the simplifying analogy of a seller putting an item of unknown quality for sale in a market. This analogy is helpful to make the analysis more directly comparable to the rest of the literature in the area and captures well the underlying conflict of interest. It is possible to obtain the results more abstractly by reinterpreting the market price in this analogy as a more general decision that is made by another party and over which the preferences of the communicating and receiving party are completely misaligned.

My analysis offers several novel contributions to the area of truthful disclosure. First, I introduce a refinement that selects a unique equilibrium in all truthful disclosure problems and show that the unique equilibrium can always be recovered constructively from a simple algorithm. I provide conditions under which the equilibrium exists and apply the theory to prove existence and uniqueness in generalized versions of classic truthful disclosure models. Second, I develop necessary and sufficient conditions for unravelling to a fully-revealing set of disclosures. These conditions clarify and extend some examples of breakdown of unravelling described in specialized settings. In particular, I show that the existence of disclosure costs or imperfect endowment of information are neither necessary nor sufficient conditions to prevent unravelling but are instead special cases of the more general economic principles that might prevent unravelling. Third, I examine the economic consequences of regulations that restrict discretion over what sellers can communicate. While, in general, the consequences of such regulations are redistributive across sellers, eliminating discretion to withhold or forcing sellers with the most favorable information to perfectly reveal their information will be undesirable to all sellers.

Lastly, I revisit three classic predictions in this area within models that involve richer forms of truthful communication. Prior literature generally assumes at least two of the following three conditions: (a) the cost of disclosure is fixed and does not depend on what information is disclosed, (b) the seller must be either fully informed or fully uninformed, and (c) the seller must can either keep silent or fully report what she knows. When any two of assumptions (a)-(c) hold, it is known that sellers with unfavorable information withhold and both increases in disclosure costs or lower endowment of information reduce communication. When only one of these assumptions hold, however, these insights are shown not to hold in general and I provide a complete analysis of the type of communication in these environments.

A more detailed elaboration on the results follows. Truthful disclosure problems are a subset of persuasion games in which some messages are defined as "truthful" and, by assumption, a sender cannot make any untruthful report or be directly penalized for being truthful as a function of her own observed information. This class of problems has found many natural applications in law and economics. As one example, in a court of law, a person can show a piece of evidence, or strategically select which evidence to show or withhold but, plausibly, cetain forms of evidence might be difficult to fabricate or tamper with. As a second example, the communication might be about a verifiable fact to be released at some date in the future; provided the law has the ability to (ex-post) enforce sufficient fines against any misreporting, the communication will be truthful.

I introduce a simple equilibrium concept to construct the equilibrium in truthful disclosure problems which is adapted from the perfect sequential equilibrium of Grossman and Perry (1986) and denoted *perfect sequential rational expectations equilibrium* (PRE). I show that the PRE is generically unique and can be constructed from an algorithm, *the priority algorithm*, that applies to any truthful communication problem. Under the priority algorithm, the PRE is constructed iteratively in a sequence of steps: in the first step, the priority algorithm selects the report that maximizes sellers' expected utility under the belief that all sellers for whom the report is feasible make this report. Then, all types that can feasibly send this report are assigned the report as their equilibrium strategy, and the type space is updated by iterated elimination in this manner until all types have been exhausted.

I give several economic conditions that guarantee the existence of the PRE and use the theory to show the existence of a unique PRE in a class of multi-dimensional disclosure problems. Using the priority algorithm, I examine the welfare consequences of certain regulations in this environment, in which the regulator can reduce or increase discretion over the reporting space. Removing the option to withhold information or mandating full-disclosure of the most favorable information will hurt all sellers, regardless of the information they intend to communicate. The result suggests that the social value of mandatory disclosure may be caused by other externalities of disclosure (e.g., product market competition) or its effect on the bargaining power of buyers in the market.

The PRE can be applied to analyze rich models of truthful communication. A procedure is obtained to characterize the PRE in problems with a continuous type space and I apply this result along three applications that improve over weaknesses of the classic models in this area. Within a richer setting, I re-examine three well-accepted results in this literature, i.e., whether more favorable events are more tightly disclosed and whether costs or lack of information endowment reduce communication. None of these properties holds when assuming only truthful communication, but the model offers simple insights as to when and why these properties will hold.

The plan of the analysis is as follows. Section 1 contains a definition of truthful communication problems, a discussion of the literature with several examples, and the equilibrium concept of PRE. Section 2 offers the main theoretical contribution of the paper, and introduces the priority algorithm, the proof of uniqueness and conditions under which the PRE exists. Section 3 uses of the properties of the priority algorithm to revisit general conditions under which unravelling to full-disclosure will occur, and then examines the effect of particular mandatory disclosure rules. Section 4 extends the algorithm to a continuous type space and further develops three applications of the model in which some of the comparative statics discussed in the broader literature can be assessed.

1. The model

1.1. Communication and beliefs

This is a model in which a good is placed for sale by a privately informed seller after a truthful disclosure has been made. There is a finite number of possible states of the world, where the state is a random variable $\tilde{s} \in S = \{s_1, \ldots, s_n\}$ and the probability of state s being realized is denoted q_s. Hereafter, I use the interpretation of the sale of a firm's asset but, as for most models in this literature, the good for sale may also be interpreted as the sale of a product by a company, the supply of labor services, etc.

The seller has private information about the state of the world which I represent as a signal $\tilde{x}$. I refer to each realization $\tilde{x} = x$ as the seller's type and denote the set of types as a finite set

$X = \{x_1, \ldots, x_m\}$.[1] Conditional on state s, each type has probability $t_x(s)$. Prior to the sale, the seller must issue a public report r which may convey information about her type. A seller with type s can choose a report $r \in M(x)$, where $M(x)$ is a finite non-empty that contains all reports that can be truthfully made by type x. Put differently, the set $M(x)$ defines for the problem under consideration what a truthful report is. A report $r \notin M(x)$ is categorized as *untruthful* and, by assumption, cannot be made by type x.

Conditional on price p and report r, the seller achieves a utility $U(p, r)$, where U is strictly increasing in p. Implicit in this assumption, the seller has no alternative use for the asset and must sell.[2] The selling price is a function of buyers' expectations about the underlying state. For any distribution F over the set of states S, let $\mathcal{P}(F)$ denote the market pricing function.[3] When observing a report, buyers form a belief $b \subseteq X$ that the type of the seller making this report is such that $\tilde{x} \in b$. To map beliefs about types into prices, I define the function $\phi(.)$ that associates to any b the induced probability distribution over the set of states, i.e., $\phi(b)$ is the c.d.f. of the random variable $\tilde{s}|\tilde{x} \in b$. Then, $\mathcal{P}(\phi(b))$ is the market price that is offered by buyers with belief b.

The timing of the model is as follows. Nature draws the state of the world $\tilde{s} = s$ and the seller draws her information $\tilde{x} = x$. Then, she issues a report $r \in M(x)$. This report is publicly observed and the firm is priced by buyers at price p.

I make two technical assumptions, both of which are very natural for my setting. First, for any two beliefs b and b', the price function satisfies that $\mathcal{P}(\phi(b \cup b')) \geq \min(\mathcal{P}(\phi(b)), \mathcal{P}(\phi(b')))$. That is, even if there is uncertainty about the state, the resulting price should not be lower than the most pessimistic belief. Second, I assume that the problem is generic, i.e., specifically, any feasible utility level $\hat{u}$ has a unique antecedent (b, r) where $U(\mathcal{P}(\phi(b)), r) = \hat{u}$. Genericity will be implied by any small perturbation to the payoff structure and is commonly-used in finite games.[4]

[1] As I will show later on, there are several, mostly technical, difficulties when considering a continuous type space. The assumption of discrete types allows me to state the economic aspects of the model with more generality and has been used in the voluntary disclosure literature (e.g., Grossman and Hart (1980), Grossman (1981), Shin (1994)).

[2] Without loss of generality, I adopt the convention of modelling the possible cost of sending a report r directly in the utility function rather than including it in the price.

[3] At this point, I make no assumption about how buyers value risk or whether further operating decisions are made by buyers. As an example with real effects, the model accommodates a pricing function $\mathcal{P}(F) = \max_k \int H(k, s) dF(s)$ (Shavell (1994)) which can be interpreted as a production economy with a post-disclosure investment k.

[4] As for the equilibrium concept, I adapt genericity, usually defined in the context of multi-player games, to a rational expectations setting. In other words, the belief b and implied price is, in a rational expectations equilibrium, the analogue to the action of the responder in a traditional communication game. For other examples

Definition 1.1 *A disclosure problem* (Q) *is given by a state space* (S, q_s)*, a type space* $(X, t_x(s))$*, the pricing function* $\mathcal{P}(F)$ *and the utility function* $U(p, r)$*.*

Example 1 (Unravelling) In Grossman (1981) and Milgrom (1981), the seller observes information about the quality of an item placed for sale $\tilde{s} \in S$, so that the type of a seller is equal to the state $\tilde{x} = \tilde{s}$. There is a prospective buyer for the item who values quality according to a VnM utility function $V(s, p)$ where s is quality and p is price paid. For a given probability assessment F about quality, the buyer is willing to pay up to $\mathcal{P}(F)$, defined by $\int V(\tilde{s}, \mathcal{P}(F))dF = \underline{v}$, where $\underline{v}$ is the utility obtained when not buying. The seller can make any report r that indicates that the true quality lies in the set r, i.e., $M(x) = \{r \subseteq S : x \in r\}$. There are no reporting costs and the seller achieves a utility $U(p, r) = p$. In Grossman and Hart (1980), the pricing function $\mathcal{P}(F)$ can be a linear or convex function of F and the reporting space is restricted to disclosure or withholding, i.e., $M(x) = \{x, r_{nd}\}$. In Viscusi (1978), a seller can certify that she is above a certain quality, i.e.., $M(x) = \{r : x \geq r\}$ and the analogue to withholding is $r = \underline{s}$.

Example 2 (Disclosure costs) In Verrecchia (1983), a seller observes a noisy signal $\tilde{x}$ about the liquidating cash flow $\tilde{s}$ of a traded asset, i.e., $\tilde{x} = \tilde{s} + \epsilon$ where ϵ is white noise. The market price for the asset is $\mathcal{P}(\phi(b)) = \mathbb{E}(\tilde{s}|\tilde{x} \in b) - \beta Var(\tilde{s}|\tilde{x} \in b)$. A seller can truthfully disclose or withhold information and makes a report $r \in M(x) = \{x, r_{nd}\}$. There is a proprietary cost $c > 0$ that reduces the firm's cash flows conditional on a disclosure $r = x$. In Verrecchia (1983), this cost is modelled directly as part of the market price but, since the manager owns the firm, the cost can be equivalently be represented as part of the seller's utility function $U(p, r) = p - 1_{r=x}c$. In Jorgensen and Kirschenheiter (2003), $\tilde{x}$ is the variance of the asset which she can disclose for a cost. The first model in Jovanovic (1982) (p.37) also satisfies these assumptions, with the special case of risk-neutral pricing. The model of Fishman and Hagerty (1990) has more trade-offs (e.g., moral hazard and liquidity trading) but the disclosure framework follows these assumptions, namely, the manager knows his effort $\tilde{x}$ which increases the final cash flow $\tilde{s}$, and chooses a precision $r \in M(x) = \{\tau_H, \tau_L\}$, which is publicly observed, and such that the cost of a high precision is higher than the cost of a low precision. The price then forms according to a price-setting mechanism which is a function of r and is increasing in the realized disclosure.

of uses of genericity in games, see Kreps and Wilson (1982), Banks and Sobel (1987) or Blume and Zame (1994).

Example 3 (Imperfect information endowment) Dye (1985) and Jung and Kwon (1988) develop a model in which the seller can be informed about the true liquidation cash flow, denoted $\tilde{x} = \tilde{s}$, or uninformed, denoted $\tilde{x} = NI$. An uninformed seller cannot disclose that she is uninformed, and must report $M(NI) = \{r_{nd}\}$ while, by contrast, an informed seller may report her information $M(x) = \{x, r_{nd}\}$. There are no disclosure costs and all players are risk-neutral. Hence, $\mathcal{P}(\phi(b)) = \mathbb{E}(\tilde{s}|\tilde{x} \in b)$ and $U(p,r) = p$. In Shavell (1994), the information can be used for productive purposes, i.e., $\mathcal{P}(\phi(b)) = \max_k \mathbb{E}(\tilde{s}r(k) - k|\tilde{x} \in b)$ where k represents the buyer's operating choice (Shavell also models a pre-disclosure information acquisition stage). In Hughes and Pae (2004), a mandatory disclosure z is publicly observed and is a noisy unbiased signal on the terminal cash flow with mean z_0. Then, the manager might observe its precision $\tilde{x}$ which she can disclose or withhold when informed. The market prices the firm as a weighted average of the initial prior and the mandatory disclosure, $\mathcal{P}(\phi(b)) = \max(l, \mathbb{E}(\tilde{s}z + (1-\tilde{s})z_0|\tilde{x} \in b))$ where l can represent the payoff from an early liquidation.

Example 4 (Multi-dimensional information) In Shin (1994), the seller observes a random vector of signals $x = (y_i, z_i)_{i=1}^k$ where, conditional on true state $\tilde{s}$, for each i, y_i and/or z_i can be equal to NI (the seller did not observe the signal) or, otherwise, satisfy $y_i \leq \tilde{s}$ and $z_i \geq \tilde{s}$. The seller can truthfully disclose or withhold information, i.e., she can make any report $r = \{\{r_i^1, r_i^2\}_{i=1}^k\}$ where $r_i^1 \in \{y_i, NI\}$ and $r_i^2 = \{z_i, NI\}$. There are no costs and buyers and sellers are risk-neutral. Kirschenheiter (1997) develops a two-dimensional version of Verrecchia (1983) in which the seller observes $x = (x^1, x^2)$, can report $M(x) = \{(x^1, r_{nd}), (r_{nd}, x^2), (r_{nd}, r_{nd}), (x^1, x^2)\}$, where r_{nd} indicates that the signal is withheld. The seller is risk-neutral and obtains a utility $U(p,r) = p - 1_{x^1 \neq r_{nd}^1} c_1 - 1_{x^2 \neq r_{nd}^2} c_2 - 1_{x^1, x^2 \neq r_{nd}} c_{12}$. Pae (2005) develops a two-dimensional version of Dye (1985) in which the seller might receive between zero and two signals. Dye and Finn (2007) consider a n-dimensional version of this model where the seller may stochastically receive between 1 and n signals.

Example 5 (Untruthful disclosure) Benabou and Laroque (1992) and Marinovic (2013) offer examples of models with uncertainty about the sender's preference. In their model, a type $\tilde{x} = (\tau, s)$ has two characteristics. The parameter $\tau \in \{0,1\}$ represents whether the seller constrained and, when $\tau = 1$, the seller must report $M(x) = \{s\}$. When $\tau = 0$, the seller is unconstrained and can report any s, i.e., $M(x) = \{s : s \in S\}$. There is no credible mechanism to perfectly report τ. There are no costs and buyers price the firm as $\mathcal{P}(\phi(b)) = \mathbb{E}(\tilde{s}|\tilde{x} \in b)$.

Note that the unconstrained type can, in effect, make disclosures that do not represent her observed s and mimic the disclosures made by the constrained type. In the version developed by da Silva Pinheiro (2013), the unconstrained type must report within certain bounds that determine the probability that a favorable report is publicly released.

Example 6 (Lemon's market) It might seem that a market for lemon (Akerlof (1970)) does not fit in the model, because if the seller opts to retain and consume the asset, her utility will depend on her true type. However, the model can be easily reinpreted to accommodate lemons as a special case. As in prior examples, let $M(x) = \{x, r_{nd}\}$ but letting now $r = x$ denote the (binding) report that the seller has an asset with value x that she has opted not to sell; hence, she has no longer a reason to make a false disclosure. Then, denote $U(p, r) = (1 - 1_{r=x}\beta)p$ so that the "cost" of disclosure is reinterpreted as a proportional discount for having the seller consume her asset instead of transferring it to a buyer with a higher willingness to pay. Vice-versa, let $r = r_{nd}$ be the report in which the seller does not consumer and puts the asset for sale. As an richer version of this example, Jovanovic (1982) allows the seller to either retain the asset or disclose and sell, which can be represented as $U(p, r) = \max(1 - 1_{r=x}\beta)p, p - 1_{r=x}c)$, where $\beta \in (0, 1)$, since the seller implicitly has the option between two mechanisms to make a truthful disclosure: (a) keep the discounted asset, or (b) pay the disclosure cost but sell it at its full value.

Example 7 (Dynamic disclosure) Several recent studies focus on the issue of dynamic disclosure, when the disclosure process occurs over multiple periods with, possibly, new arrival of information (see, e.g., Einhorn and Ziv (2008), Beyer and Dye (2011), Guttman, Kremer and Skrzypacz (2012), Marinovic and Varas (2013)). Many of these models are solved by backward induction or dynamic programming where the disclosure problem is considered over a single period looking ahead on the continuation payoff of future periods. This, in turn, implies that (in each period) the utility function has the form $U_t(p, r)$ which is to be recovered endogenously and thus cannot necessarily be characterized in closed-form on which standard methods can be applied. Although the methods discussed here do not cover the issue of solving for the proper characteristics of $U_t(.)$, they allow modelers to establish the existence and uniqueness of an optimal disclosure policy in each period with little required knowledge of continuation payoffs.

1.2. Equilibrium concept

I introduce next the equilibrium concept for the model. As is standard in the truthful disclosure literature, I use a rational expectations equilibrium concept and model the receiver in reduced-form as a price-setting mechanism.[5]

Definition 1.2 *For a disclosure problem (Q), a rational expectations equilibrium (RE) Γ is defined as a price function $P(r)$, a belief structure $B(r)$ and a reporting strategy $R(x)$, such that:*

(i) All sellers maximize their utility, i.e., for any type x, $R(x) \in argmax_{r \in M(x)} U(P(r), r)$.

(ii) Prices and beliefs form in a Bayesian manner, i.e., $B(r) = \{x : R(x) = r\}$ and, for any $B(r) \neq \emptyset$, $P(r) = \mathcal{P}(\phi(B(r)))$.

(iii) For any $B(r) = \emptyset$, there exists $b_r \subseteq M^{-1}(r)$ such that $P(r) = \mathcal{P}(\phi(b_r))$.

Definition 1.3 *For a triple $\Gamma = (P(.), B(.), R(.))$, let $u^{\Gamma}(x) = U(P(R(x)), R(x))$ be defined as the utility obtained by type x conditional on prices, beliefs and reporting strategies Γ.*

Like many persuasion games, disclosure problems can have many equilibria which can be sustained by various pricing functions $P(.)$. I introduce a slightly modified version of perfect sequential equilibrium (Grossman and Perry (1986)), which I label perfect rational expectations equilibrium (PRE). This criterion cannot rule out multiplicity in all persuasion games but I will show that it predicts at most a unique equilibrium in (generic) truthful disclosure problems.

I give a general intuition for the PRE and then present a formal definition. In a PRE, the problem described above is viewed as a simplified representation of a richer interaction in which the seller chooses a report r_0 and can either take the quoted market price *or* make an alternative binding offer p_0. If this offer p_0 is rejected, the asset is not sold, an eventuality I am assuming to be a least-preferred outcome for the seller.[6]

[5] One potential limitation of this formulation is that it may seem counter-intuitive to use a refinement that applies to the price-setting mechanism and not to a real player. However, one should note that the price-setting mechanism can always be expanded here as a model in which the seller makes a take-it-or-leave-it offer to a buyer or if multiple buyers compete. The use of the price-setting mechanism is used here to save space when exposing the model but is, like in most of this literature, not critical for the main insights.

[6] Note that for this rationale to hold, the seller should not be able to costlessly revert to the market price function $P(.)$ since, in this case, any offer would be cheap talk. This assumption is consistent with most non-cooperative bargaining games, in which modifying an offer should involve non-zero costs to the party making the offer (Rubinstein (1982), Crawford (1982)). Naturally, since my purpose here is not to develop a complete bargaining game but to motivate the PRE criterion, I make the simplifying assumption that an offer is completely binding.

The PRE relies on a forward-induction logic where players assume that an off-equilibrium action must be rationalized from a particular belief and, therefore, buyers should learn from what this action means. To be specific, consider an off-equilibrium proposal (r_0, p_0). For the seller to rationally make an offer different from the posted price, she should be anticipating her offer to be accepted. Hence (if possible), buyers should assume that the seller expects the offer to be accepted.[7] Further, all types of sellers should have the same expectation and, therefore, expect this particular offer to be accepted. Therefore, all sellers who can truthfully offer (r_0, p_0) and are better-off doing so should be expected to make this offer. In turn, buyers can calculate the set of types b_0 accordingly, implying a maximal willingness to pay for the asset $\mathcal{P}(\phi(b_0))$. If the offered price p_0 is lower than the willingness to pay, the forward-induction argument stated above can be rationalized, in turn disqualifying the conjectured equilibrium.[8]

Definition 1.4 *A perfect rational expectations equilibrium (PRE) Γ is a RE such that there exists no triple (b_0, r_0, p_0) such that all of the following conditions hold:*

(i) There exists $b_0 \neq \emptyset$ such that types $x \in b_0$ are strictly better-off with an off-equilibrium price and report pair (r_0, p_0).

(ii) The offered price p_0 is below buyers' willingness to pay conditional on the belief b_0, i.e., $p_0 \leq \mathcal{P}(\phi(b_0))$.

(iii) The belief b_0 is consistent with the set of types strictly better-off sending r_0, i.e.,

$$b_0 = M^{-1}(r_0) \cap \{x : U(p_0, r_0) > u^{\Gamma}(x)\}. \tag{1.1}$$

Note that there are a few small differences with Grossman and Perry's perfect sequential equilibrium, the most obvious one being that the concept is applied in a rational expectations environment rather than a two-player game. In addition, I impose the conditions of PRE over all possible reports r_0 including reports that could be made on the equilibrium path. Conceptually, this is natural given that the logic of the pre-sale binding offer that underlies the PRE in this problem is that the seller can make a new proposal that takes the form of an off-equilibrium *pair* (r_0, p_0). Further, I allow the price p_0 to be below buyers' maximal willingness to pay because

[7] For a more general discussion of forward-induction and strategic stability, see Kohlberg and Mertens (1986) and Carlsson and Van Damme (1993).

[8] If the price p_0 is greater than the willingness to pay, buyers cannot reconcile this offer as a rational move and the offer cannot be accepted.

this pre-sale offer need not be equal to buyers' willingness to pay to be accepted.[9] Lastly, I select the set of types b_0 in terms of types that would be *strictly* better-off making the offer. I view this assumption more intuitively appealing but all results carry over using a stronger version of PRE in which indifferent types may or may not be part of b_0.

1.3. Discussion and links to the literature

As my intent is to provide a general formulation for a particular class of models usually referred to as truthful disclosure, I provide here some further discussion of the main restriction that underlies these models.[10] In the literature, truthful disclosure problems are a subset of the general class of persuasion games, in which the utility of the seller *depends only on beliefs and the signal sent, not on her original type* ("type-independent" preferences). Examples of this approach include, among many others, Grossman and Hart (1980), Grossman (1981), Milgrom (1981), Jovanovic (1982), Verrecchia (1983), Dye (1986), Teoh and Hwang (1991), Penno (1997), Jorgensen and Kirschenheiter (2003), Shin (2003), Suijs (2007), Bagnoli and Watts (2007), Acharya, DeMarzo and Kremer (2011), Guttman et al. (2012), Kumar, Langberg and Sivaramakrishnan (2012) and Marinovic (2013).[11]

Like persuasion games, disclosure problems tend to have many equilibria. In fact, various workarounds have been considered in the prior literature to address the multiplicity and I offer here some discussion of the benefits of the PRE over other concepts. In the early literature, one approach has been to reduce the message space to a seller indicating her own type ("disclose") or a single pooling signal that is always available ("withhold"), e.g., Grossman and Hart (1980), Jovanovic (1982), Verrecchia (1983), Dye (1985), Jung and Kwon (1988). This approach generally addresses the multiplicity that might emerge from the off-equilibrium but does not readily accommodate coarse disclosures or partial withholding of information; for example, an informed seller might be willing to disclose only a piece of what she knows. Furthermore, this approach

[9] One limitation of this restriction is that buyers might wonder why the seller does not make the best offer that would maximize her utility - for example, filtering out types that should have made a different offer pair. However, this is a more fundamental problem in that even if we were to set p_0 equal to the willingness to pay, there could be other report-price pairs $(r, p) \neq (r_0, p_0)$ that would further raise the utility of certain types in b_0 if accepted by the above logic. This is a limitation that we share with perfect sequential equilibrium, as well as other common refinements (see also Mailath, Okuno-Fujiwara and Postlewaite (1993) for a more complete discussion of this problem).

[10] For recent surveys of this literature, see Verrecchia (2001), Dye (2001), Dranove and Jin (2010) and Beyer, Cohen, Lys and Walther (2010).

[11] Note that the assumption rules out many other forms of communication as in reporting models with cheap talk (Stocken (2000), Baldenius, Melumad and Meng (2011)), costly misreporting (Guttman, Kadan and Kandel (2006), Kartik, Ottaviani and Squintani (2007), Beyer (2009), Caskey, Nagar and Petacchi (2010), Laux and Stocken (2012)) or if sellers remain exposed to the residual value of the asset (Bushman and Indjejikian (1995), Huddart, Hughes and Levine (2001)).

does not always guarantee a unique RE.

A workaround, related to the PRE, is Dye's optimal policy (Definition 3, p. Dye (1986)), which is also used in Kirschenheiter (1997). A RE is optimal if it is not preferred by all types (strictly by some) by another RE. However, while optimality is a compelling criterion, it is often too demanding to eliminate most REs because its application requires (i) *all* types to favor one RE and (ii) compare payoffs in a RE to another specified RE. The PRE always satisfies Dye's optimality criterion since, if this were not the case, some subset of sellers would be able to offer a new pair of report and price consistent with the other RE. But the PRE is a more demanding criterion than optimality given that it needs only be applied to a subset of seller types.

2. Motivational examples

In the following examples, I illustrate how multiplicity of equilibria is a prevalent feature of voluntary disclosure models and how, in certain simple models, a unique PRE can be derived with minimal formalism. In later sections, I will make the logic more systematic to apply it to the general disclosure model.

2.1. A binary disclosure problem

In order to show how to apply PRE, I begin with a simple introductory example with only two types $X = \{l, h\}$, where l is a low type and h is a high type, and the price function is such that:

$$\underbrace{\mathcal{P}(\phi(\{h\}))}_{=P_h} > \underbrace{\mathcal{P}(\phi(\{h,l\}))}_{=P_l} > \underbrace{\mathcal{P}(\phi(\{l\}))}_{=0}. \tag{2.1}$$

The low type has only one message $M(l) = \{0\}$ and the high type has two messages $M(h) = \{0, 1\}$; for example, "0" might represent 'withhold" and "1" might represent "disclose to be a h type". Sellers are risk-neutral but there is a cost rc when sending report $r \in \{0, 1\}$, where $c \in (P_m, P_h)$.

It is immediately seen that the model has two REs. In the fully-separating equilibrium (RE-1), type h reports $r = 1$ and type l reports $r = 0$. In the pooling equilibrium (RE-2), type h and type l reports $r = 0$ and obtains a price P_m.[12] Many reporting games outside of the voluntary disclosure literature tend to favor fully-separating equilibria when they exist and, in such games, full-separation is sometimes used as a selection criterion, see, e.g., Dye (1988), Kanodia and Lee

[12]The problem also has a mixed strategy equilibrium which I do not discuss here.

(1998) or Fischer and Verrecchia (2000). On the other hand, other studies argue that pooling equilibria can lead to lower deadweight cost of separation on welfare grounds (Guttman et al. (2006), Guttman, Kadan and Kandel (2010)), so it is an open question as to whether RE-1 or RE-2 is the proper equilibrium for this problem; however, RE-1 is not a PRE. To see this, note that when playing RE-1, a type could make an alternative offer $(r_0, p_0) = (1, P_m)$, i.e., *"I am sending message r=0 but you should not give me the current quoted price* $P(0) = 0$ *and, instead, trade at my offer* P_m*."* Types $b_0 = \{h, l\}$ would have made this offer, thus validating the price $p_0 = P_m$ and making it irrational for type h to issue $r = 1$ or type l to accept a price $P(0) = 0$.

2.2. Costly disclosure

Consider the following model adapted from Jovanovic (1982) and Verrecchia (1983). For expositional purposes, I use here the continuous type space in these models to better fit this literature (the definitions of RE and PRE carry over to the continuous setting). As in the previous example, the seller is risk-neutral. The type space is X is an interval and $\mathcal{P}(\phi(b)) = \mathbb{E}(\tilde{x}|\tilde{x} \in b)$; the message space is $M(x) = \{x, r_{nd}\}$ where r_{nd} indicates "withhold information." There is a cost $c \in (0, \sup X - \mathbb{E}(\tilde{x}))$ when sending $r = x$ and no cost when sending $r = r_{nd}$.

This model admits REs that take the form of a threshold τ and such that $R(x) = x$ if $x \geq \tau$ and $R(x) = r_{nd}$ if $x < \tau$. A commonly-used solution technique is to note that τ corresponds to the marginal type indifferent between disclosing and withholding.

$$\mathbb{E}(\tilde{x}|\tilde{x} < \tau) = \tau - c \tag{2.2}$$

However, Equation (2.2) does not necessarily have a unique solution. For example, it is known that a sufficient condition for uniqueness is logconcavity and a sub-exponential lower tail (see Bertomeu (2012)). Logconcavity and sub-exponential tails are satisfied by the Normal distribution, as shown in Verrecchia (1983). However, putting aside the issue of negative prices with normal distributions, logconcavity is a technical condition that is not easily derived from economic behavior. Furthermore, logconcavity is only sufficient if the price under risk-neutral pricing (e.g., $\mathcal{P}(\phi(b)) = \mathbb{E}(\tilde{x}|\tilde{x} \in b)$) and might not guarantee uniqueness in problems that involve real operating decisions.

I will illustrate a reasonable environment in which logconcavity would not hold and develop further what economic problems this may create. Borrowing from Subramanyam (1996), herafter

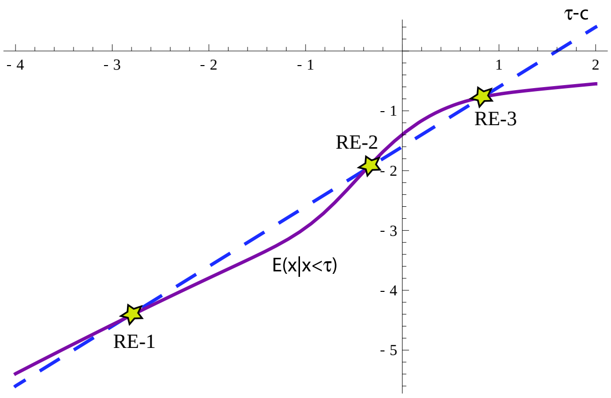

Figure 1: Costly voluntary disclosure with uncertain variance

KRS, assume that firms have normally-distributed types with mean normalized to zero but whose variance might be different. Suppose that the variance can be either $\sigma^2 = 1/2$ or $\sigma^2 = 3$ (KRS uses a continuous distribution for the precision). In Figure 1, I plot $\mathbb{E}(\tilde{x}|\tilde{x} < \tau)$ against $\tau - c$ at $c = 1.6$. An intersection between these two curves is a solution to Equation (2.2).

There are three possible voluntary disclosure equilibria once one considers the (plausible) scenario in which not all firms have the same volatility. In fact, the number of equilibria can be even greater with more than two possible volatilities, implying that a single indifference condition like Equation (2.2) is far from sufficient to pin down an equilibrium. Worse still, one of these equilibria, RE-2, has comparative statics that are the opposite of the standard models: a small increase in c (raising the dotted curve) will *decrease* the voluntary disclosure threshold, implying more voluntary disclosure.

Fortunately, the PRE performs again to select which equilibrium is reasonable and the counter-intuitive RE-2 fails to be a PRE. This can be again shown with minimal need for notation, but the reader may stop at this point and throw in an educated guess as to whether RE-1 or RE-3 (or both) will be the PRE in this model. For each RE-i, I label the disclosure threshold τ_i where, by definition, $\tau_1 < \tau_2 < \tau_3$.

Let me first dismiss RE-2 as a reasonable equilibrium. Suppose that some types send an alternative pair (p_0, r_0) where $p_0 = \mathbb{E}(\tilde{x}|\tilde{x} \leq \tau_3)$ and $r_0 = \emptyset$, i.e., this message would have the form: *"I am not disclosing but consider pricing the asset as if we were playing according to RE-3".* Naturally, because RE-3 is itself an RE, this price p_0 can be rationally sustained and, further, it would strictly benefit all sellers with $c < \tau_2$ which were obtaining $\mathbb{E}(\tilde{x}|\tilde{x} \leq \tau_2)$ in RE-2. Because RE-3 is an RE, sellers with $x \in (\tau_2, \tau_3)$ must be better-off obtaining $\mathbb{E}(\tilde{x}|\tilde{x} \leq \tau_3)$ over their disclosure utility $x - c$ when playing RE-2. This confirms that RE-2 is not a PRE.

By the same argument, RE-1 is also a PRE either and the PRE must be the *maximal* solution to Equation (2.2).

The equilibrium RE-3 has the intuitive comparative statics that higher cost reduces disclosure. This comparative statics even holds when the implicit function theorem does not apply. If the cost c decreases sufficiently so that RE-3 ceases to exist, then the PRE will shift to RE-1 leading to a non-marginal increase in voluntary disclosure.

2.3. Many dimensions

A well-known limitation of voluntary disclosure models is that the standard tools used when a single piece of news is being voluntarily disclosed do not generalize well to two or more dimensions. As I will argue later on, the PRE will address this problem in a very general sense but I illustrate at this point the fundamental problem of multiplicity under multi-dimensional disclosure.

In the simplest of such multidimensional problems, consider a continuous two-dimensional state space $(s_1, s_2) \in S_1 \times S_2$, where the price is $\mathcal{P}(G) = \mathbb{E}_G(\tilde{s}_1 + \tilde{s}_2)$ such that $\mathbb{E}_G$ indicates the conditional of expectation according to c.d.f. $G(.)$. The type space is $(\{r_{nd}\} \cup S) \times (\{r_{nd}\} \cup S')$. As in Dye (1985) and Jung and Kwon (1988), each type may not receive information and assume that $\tilde{s}_1$ and $\tilde{s}_2$ are observed independently with the same probability $p \in (0, 1)$. To make this example entirely straightforward, assume that buyers can identify whether the first or second dimension is disclosed. There are no costs and the seller achieves a utility equal to the selling price.

Obviously, this model has a very apparent RE. Because the two dimensions are fully separable (additive marginal effects and uncorrelated endowments/signals), one can appeal to the solution to the one-dimensional case to construct an equilibrium. Denoting τ_1 and τ_2, the (unique) disclosure thresholds in Jung and Kwon (1988) for each dimension, an RE exists in which each $\tilde{s}_i$ is disclosed whenever information is received and $\tilde{s}_i \geq \tau_i$ is greater than the designated threshold. Let this RE be denoted as RE-1 and represented in Figure 2 where $\tilde{s} = (\tilde{s}_1, \tilde{s}_2)$ in the upper-right (lower-left) quadrant is fully disclosed (withheld) and only one dimension of $\tilde{s}$ is disclosed along the off-diagonal quadrants.

However, this intuitive RE-1 is not the unique RE in the two-dimensional problem and it is relatively easy to create many other equilibria by small changes over off-equilibrium beliefs. As an example, I carved out zone A from the region in which $\tilde{s}_2$ should be disclosed in the original

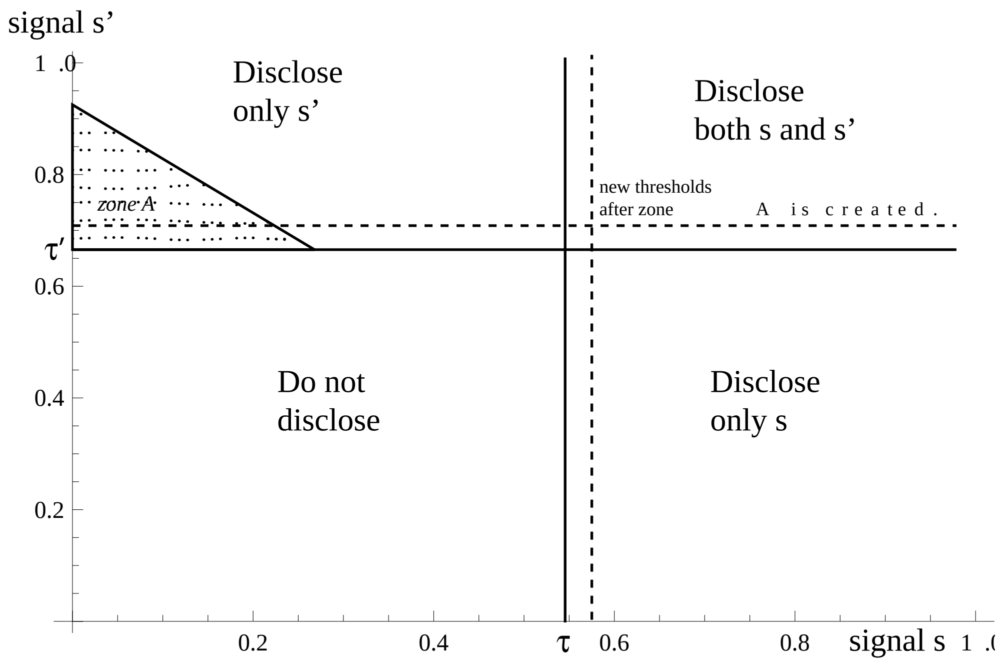

Figure 2: Two-dimensional Dye model

RE and, for now, let me define a new candidate RE in which the seller does not disclose when lying in zone A. To justify this conjecture, assume that if a single dimension is disclosed in zone A, buyers assign a (sufficiently) negative belief about the other piece of information that is withheld, thus confirming that no disclosure in zone A. If the support of $(\tilde{s}_1, \tilde{s}_2)$ is unbounded from below, any zone can be carved out of the regions where only one signal is disclosed and reclassified into the non-disclosure region.

To make things worse, it is *not* the case that such manipulations would simply change the equilibrium in places where they occur. Creating zone A would imply an increase in the non-disclosure price, thus shifting the entire regions defined earlier and globally changing the nature of the equilibrium. Furthermore, this type of equilibrium cannot be removed with Dye's optimality criterion because, evidently, reclassifying some high disclosed events into the non-disclosure region is desirable to uninformed types.

It is clear that RE-1 is the most intuitive equilibrium in this problem, but recall that it can only be obtained as such in the very simple setting described above. How does one identify which off-equilibrium beliefs are reasonable or unreasonable in more complex problems that involve correlated types or information endowments? To avoid these nagging problems, a few existing multi-dimensional studies specifically rule out a situation in which no signal is disclosed (Kirschenheiter (1997), Dye and Finn (2007)) which, however, severely restricts what problems can be analyzed using the standard tools.[13]

Fortunately, the PRE selects the "right" RE for this example and, as I will show later on, guarantees a unique prediction in fairly general multi-dimensional disclosure problems. The interested reader may refer to the general proof later on, but I can give here a heuristic explanation. In any RE in which no-disclosure is possible, RE-1 is the unique RE that minimizes the non-disclosure price across all REs. So, if an RE does not coincide with RE-1, some types that were disclosing under RE-1 do not disclose under, say, another equilibrium RE-2. Therefore, some types are worse-off under RE-2. The concept of PRE follows immediately. These types would adopt their disclosures of RE-1 *and* suggest the price that would have occurred under RE-1, thus moving away from the price function in RE-2 and ruling it out as a PRE.

3. Le main result

Although PRE has performed well in a few examples (in these cases, it can be applied with almost no need for formalism), the question remains as to whether the PRE will be effective in finding equilibria in the general class of disclosure problems. Within this Section, I establish that this class of models has (generically) at most one PRE, provide a simple constructive characterization of the PRE and, lastly, derive a necessary and sufficient condition for a PRE to exist.

3.1. Uniqueness

As a first step for my analysis, I will construct a set of strategies, prices and beliefs from a particular algorithm that, as it turns out, is closely related to the concept of PRE.

For any type space $X' \subseteq X$, define $V(X')$ as the report that maximizes the seller's utility provided that the market believes that a type reports r if she can do so. Genericity and the finiteness of the reporting space guarantee that the maximizer exists and is unique.[14] In formal

[13] Other studies have fallen into the precipice of inadvertently selecting one among the many possible REs, without much consideration as to why this equilibrium would be more reasonable. A good example is Pae (2005) who describes the equilibrium of a two-dimensional Dye model, in contradiction to the multiplicity of RE made apparent in my example. His study uses the same restriction as Kirschenheiter (1997) and Dye and Finn (2007) by first focusing on a model in which at least one signal is received (p. 388-395) and for which the RE can be examined. Then, considering the case in which no firm may receive information, the study argues that the same functional form for the updating function with one signal derived earlier can be used. It *can* be used but does not *have* to be, because for any off-equilibrium message (and there are many of them) this particular manner of updating need not hold. The results that follow p. 396-401 are only true for one among a continuum of other REs in this game.

[14] For non-generic models, non-uniqueness might (occasionally) imply that the algorithm below would branch out, leading to more than one PRE. Otherwise, the existence and construction would still hold.

terms,

$$V(X') = argmax_r \;\; U(\mathcal{P}(\phi(\{x \in X' : r \in M(x)\})), r) \tag{3.1}$$

I define the algorithm presented below as the *priority* algorithm. The triplet Γ^a given by $(P^a(r), B^a(r), R^a(x))$ is constructed iteratively as follows:

1. Initialize the algorithm at $i = 1$ and $X_1 = X$,
2. Calculate $r_i = V(X_i)$ and set $b_i = B^a(r_i) = \{x \in X_i : r_i \in M(x)\}$, $R^a(x) = r_i$ for all $x \in b_i$ and $P^a(r_i) = \mathcal{P}(\phi(b_i))$,
3. Set $X_{i+1} = X_i \backslash b_i$,
4. Stop if $X_{i+1} = \emptyset$, otherwise update to i+1 and return to step 2.
5. Complete the process for any off-equilibrium r, with $B^a(r) = argmin_{b \in K(r)} U(\mathcal{P}(\phi(b)), r)$ s.t. $K(r) = \{b : \text{ if } x \in b, r \in M(x)\}$ and $P(r) = \mathcal{P}(\phi(B^a(r)))$.

The algorithm is initialized with the complete type space (step 1.), at which point, it selects the best attainable report, i.e., the report that maximizes the utility if it is issued by all sellers that can send the report (step 2.). Then, these types are removed from the type space (step 3.) and the procedure repeats until all types have been exhausted (steps 2-4.). Lastly, all off-equilibrium beliefs are set to be sufficiently pessimistic (step 5.).[15]

Definition 3.1 *I say that type x has priority over type x', denoted $x \Vdash x'$, if x is selected by the algorithm at the same step as x' or earlier. Denote $x||x'$ if x and x' are selected in the same step.*

Note that the priority order is a complete order over the type space X. I show next that being recovered from the priority algorithm is a necessary condition for an equilibrium to be a PRE.

Lemma 3.1 *When it exists, the PRE is unique and must coincide with Γ^a (except, possibly, for off-equilibrium beliefs and prices).*

[15] In theory, in a RE one might set any arbitrarily small value for the off-equilibrium price (Fudenberg and Tirole (1991)). I use here a more demanding specification for the off-equilibrium so that the specification would also qualify as a sequential equilibrium in the sense of Kreps and Wilson (1982).

When a PRE exists, it must be recovered from the priority algorithm. At step $i = 1$, for any RE that differs from Γ^a, there is a subset of types that could have achieved a higher utility if they had sent a different message, contradicting the requirements of a PRE. This argument can be repeated at steps $i > 1$, implying that after accounting for those types that have already selected their utility-maximizing report, the algorithm must keep selecting the utility-maximizing report across all remaining types.

3.2. Existence of the PRE

Lemma 3.1 does not guarantee that a PRE exists. Indeed, there are examples of disclosure problems that do not always admit a pure-strategy RE (Benabou and Laroque (1992), Marinovic (2013), da Silva Pinheiro (2013)). I formally examine next when the candidate equilibrium obtained from the priority algorithm qualifies as a PRE.

Lemma 3.2 Γ^a *is a PRE if and only if* $u^\Gamma(x) \geq u^\Gamma(x')$ *for any* $x \Vdash x'$.

Lemma 3.2 offers a necessary and sufficient characterization of the existence of a PRE. However, verifying this property requires a direct computation of the algorithm and, therefore, whether existence holds is not evident from the description of the game. I develop next simpler conditions that can guarantee the existence of the PRE in various disclosure problems.

Condition (A). The message space is complete if, for any r and r', there exists r'' such that $M^{-1}(r'') = M^{-1}(r) \cup M^{-1}(r')$ and $U(p, r'') \geq \max(U(p, r), U(p, r'))$ for any p.

Condition (A) has an interpretation in terms of a semantical connector "or"; that is, if a seller can make a truthful sentence that she has a certain piece of information with certainty, then, she should be able to truthfully make a report that she has this information or has another piece information consistent with a different report. Because this new statement is coarser and conveys less information, it should be weakly less costly to the seller (e.g., an external party that can verify that r or r', might require more verification effort to verify only r and only r'). Consider, as an example, an interval forecast made by a manager: even if the manager knows that the true expected cash flow will be a certain value, she can truthfully report that it will lie within an interval, plausibly letting out less proprietary information to the firms' competitors. Another example is Grossman (1981) and Milgrom (1981) who allow the seller to make any coarse representation of her information.

Theorem 3.1 *If condition (A) holds, there exists a unique PRE, and it is given by* Γ^a.

A limitation of condition (A) is that it is rarely assumed in voluntary disclosure models, in part for its inherent economic limitations but also for technical reasons. Condition (A) significantly increases the size of the reporting space (and the implied off-equilibrium path) which since most of these models do not use refinements, creates many REs. As an economic assumption, also, it is possible that the reporting space might be constrained by feasibility considerations (a suspect cannot simultenously report two disjoint alibis).

I consider here an alternative to condition (A) that does not require a complete message space and can be directly applied to existing voluntary disclosure models as-is.

Definition 3.2 *A type x is higher than type x', denoted $x \succeq x'$, if, for any b s.t. $x, x' \notin b$, $\mathcal{P}(b \cup \{x\}) \geq \mathcal{P}(b \cup \{x'\})$.*

I define this order as the *value* order given that it represents the impact of a type on price; I further denote the implied strict value order as $\succ$ and equivalence relation $\sim$. Using the value order, I introduce the following three conditions.

Condition (B1). X is fully ordered according to the value order $\succeq$, with $x \sim x'$ if and only if $x = x'$.

Condition (B2). For any r, $M^{-1}(r)$ are intervals in the sense of $\succeq$.

Condition (B3). For any two reports r and r' such that (a) $M^{-1}(r) \cap M^{-1}(r') \neq \emptyset$ and (b) all maximal elements $M^{-1}(r')$ are higher than all maximal elements of $M^{-1}(r)$, then the following holds: $U(p, r') \geq U(p, r)$ for any p.

Condition (B1) states that the set of types can be ordered from the type that most increases prices to the type that most decreases prices. Condition (B2) states that unconditional beliefs after observing a report must be convex (interval) sets in the sense of the value order. Condition (B3) is generally satisfied in many models of voluntary disclosure. This condition is a technical restriction that states that if there are two reports available to some moderate types but one of these reports is only available to higher types (the "good" report) while the other is only available to lower types (the "bad" report), then issuing the bad report should be weakly more costly. One interpretation of this condition is that it may expensive to verify that a report is inconsistent with some events being very favorable (i.e., good news is hard to objectively verify). Importantly, I do not mean here that it should be necessarily cheaper to make disclosures consistent with higher types in that this property need only hold across reports that span over the same types. This condition is always satisfied in models in which the only pooling report is "withhold" which is both costless to send with a maximal upper bound.

Theorem 3.2 *If conditions (B1), (B2) and (B3) hold, there exists a unique PRE, and it is given by* Γ^a.

3.3. The priority order: a simplified approach

A notable difficulty in applying the algorithm Γ^a is that, at each step i, the report $r_i = V(X_i)$ must be evaluated by considering every possible report that can be sent by types in the set X_i. This presents two notable challenges. First, the set X_i is not analytically characterized unless the algorithm is applied. Second, the search for a PRE can present computational challenges if the set of types or the message space is large. Technically, I am thus interested in considering when a simplified algorithm can be used in which, at each step, one may restrict the attention to reports that could be sent by certain types that are easy to identify in a subset of X_i.

There is a subclass of disclosure problems that admits a simplified computation of the priority order. To show this, I introduce a new order that can be directly computed from knowledge of the type and reporting spaces. I refer to this order as the *dominance* order.

Definition 3.3 *Type x dominates type x', denoted $x \trianglerighteq x'$, if for any $r' \in M(x')$, there exists $r \in M(x)$ such that (a) $U(p,r) \geq U(p,r')$ for all p, and (b) A type in $M^{-1}(r) \backslash M^{-1}(r')$ is always higher than a type in $M^{-1}(r') \backslash M^{-1}(r)$.*

In plain language, a type is dominant when she has access to cheaper message that pool with higher types in the sense of the value order. As before, I will introduce an algorithm, which I call the dominance algorithm, and then examine its properties. The triplet $\Gamma^b = (P^b(r), B^b(r), R^b(x))$ is constructed iteratively in the following manner:

1. Initialize the algorithm at $i = 1$ and $Y_1 = Y$,

2. *Select the maximal set of $Z_i \subseteq Y_i$, in the sense of the dominance order $\trianglerighteq$.* Calculate $r_i = argmax_{r \in \bigcup_{x \in Z_i} M(x)} U(\mathcal{P}(\phi(M^{-1}(r))), r)$ and set $b_i = B^b(r_i) = \{x \in Y_i : r_i \in M(x)\}$, $R^b(x) = r_i$ for all $x \in b_i$ and $P^b(r_i) = \mathcal{P}(\phi(b_i))$.

3. Set $Y_{i+1} = Y_i \backslash b_i$,

4. Stop if $Y_{i+1} = \emptyset$, otherwise update to i+1 and return to step 2.

5. Complete the off-equilibrium r, with $B^b(r) = argmin_{b \in K(r)} U(\mathcal{P}(\phi(b)), r)$ s.t. $K(r) = \{b :$ if $x \in b, r \in M(x)\}$ and $P(r) = \mathcal{P}(\phi(B^b(r)))$.

The main simplification obtained in algorithm b is that the types that will send the report at step i are a maximal element according to the dominance order $\trianglerighteq$.

Theorem 3.3 *The reporting strategies, and beliefs and prices following any report that may be made with positive probability under* Γ^a *and* Γ^b *coincide. Further, if* $x \trianglerighteq x'$, *then* $x \Vdash x'$ *and* $u^{\Gamma^a}(x) \geq u^{\Gamma^a}(x')$.[16]

One remaining difficulty is that the set Z_i may typically include multiple types which might still cause the search for the report to be selected in step i to be cumbersome. Under a stronger condition, as I show next, the simplified approach can be used to directly remove the search for the "right" type by selecting any maximal type.

Condition (C) X is fully-ordered according to the dominance order $\trianglerighteq$.

Theorem 3.4 *Suppose that conditions (B1) and (C) hold. If* Γ *is a PRE, all types that are maximal at step* i *and have a report in common must send the same report.*

When the conditions of Theorem 3.4 hold, step 2 can be simplified to selecting *any* $\trianglerighteq-$maximal type, instead of the maximal type that maximizes the price. The search for an optimal price remains only present across reports available to that particular maximal type. There are well-known examples in which the type space is fully ordered according to $\trianglerighteq$ and therefore algorithm b yields a very simple procedure to compute equilibria. Under the canonical costly disclosure models of Jovanovic (1982) and Verrecchia (1983), each type $x \in \mathbb{R}$ is such that $M(x) = \{r_{nd}, x\}$ and such that $r = x$ involves an additive cost $c > 0$ while $r = r_{nd}$ does not. If $x > x'$, type x has access to the same report as x' as well as report $r = x$ which involves the same cost but a higher type than the message available to x', i.e., $r = x'$. It then follows that $x \trianglerighteq x'$ if and only if $x \geq x'$ and, in this case, the dominance order coincides with the value order.

[16] Algorithm b need not pick types in the same order as algorithm a. As an example, if a high type has a single message to send which she can only send and this message is very costly, then she will obtain an equilibrium that is very low, and thus will be captured by algorithm a at a later step but would be captured by algorithm b much earlier.

Yet, even when (B1) and (C) are satisfied, the dominance order need not exactly coincide with the value order, as shown in the uncertain information endowment models of Dye (1985) and Jung and Kwon (1988). These models have a set of types where either $x \in S$ (the seller is informed and anticipates value x) or $x = NI$ (the seller is not informed and anticipates value $\mathbb{E}(\tilde{x})$). The message space is then given by $M(x) = \{r_{nd}, x\}$ where, by definition, the uninformed can only report r_{nd}. Consider the orders $\succeq$ and $\trianglerighteq$. It is clear that $NI \succeq x$ if $x \leq \mathbb{E}(\tilde{x})$ (below-average types are lower types than the uninformed). Yet, these types have more messages than type NI and therefore satisfy $x \trianglerighteq NI$, with NI being the type that is (indeed) dominated by all other types. In fact, within this class of models, some informed below-average types have priority over uninformed types and, as is well-known, achieve a greater equilibrium surplus.[17]

3.4. Existence with multi-dimensional information

I develop here a basic application of the theory to the case of a seller informed about several characteristics and who may disclose/withhold each characteristic separately. In what follows, I use upperscript for dimensions to distinguish x^i, the i^{th} component of vector x from x_i, the element of the type space X.

Let (Q) be the disclosure problem defined as follows. A type is given by a vector $x = (x^1, \ldots, x^d)$ where $x^i \in \{NI, s^i\}$, NI represents the eventuality that the seller is not informed about component i, and d is the number of dimensions. Note that I make no assumption about distributions or the pricing function, i.e., (a) whether components of the state s^i are correlated, (b) whether the probabilities of receiving information on one state are correlated or depend on the state, (c) whether the economy is pure-exchange or there are uses for information. There is a cost to disclose so that, for each p, $U(p, r)$ is a non-decreasing function of the number of reports disclosed (not their value). For any report, the reporting space is $M(x) = \{\{r^1, \ldots, r^d\} : r^i \in \{r_{nd}, x^i\}\}$, i.e., a seller can withhold information. The results apply more generally to ordered reports with the form $r = (r_1, \ldots, r_d)$, when buyers do know which signal is disclosed and, since the proofs are unchanged, I will focus here on unordered reports.[18]

[17] Note that the condition that the type space should be fully ordered according to dominance is demanding. For example, this condition is not verified in multi-dimensional disclosure models of the kind described in Section 2.3. In this example, a type $x = (1, 2)$ neither dominates nor is dominated by a type $x = (1.5, 1.5)$ or by type $(2.5, NI)$ for that matter. For these problems, only Theorem 3.3 applies.

[18] One issue with this multi-dimensional model is that it does not fit (A) or (B2). The latter is because, as an example, the report $\{x^1, \ldots, x^{d-1}, NI\}$ can be sent by all types with $\{x^1, \ldots, x^{d-1}, y\}$ implying states in the range, say, of the price when y is minimal or y is maximal. But this range might include the price of types with $\{y', \ldots, x^{d-1}, y\}$ with $y' \neq x^1$ and which cannot send this particular report, violating the convexity requirement in (B2). Fortunately, the results from Lemmas 3.1 and 3.2 are sufficient to obtain a PRE even for this case.

Proposition 3.1 *(Q) has a unique PRE.*

One notable feature of multi-dimensional information is that the signal space available to each type is ranked from most precise to least precise (more non-disclosure). This property guarantees that under fairly general conditions, there is a unique PRE in many multi-dimensional disclosure problems. Of note, while I focus here on existence and uniqueness, a more detailed analysis of such models is not trivial and goes beyond the scope of my study.

4. Properties

4.1. The unravelling property

I first analyze conditions under which an unravelling property holds. By "unravelling," I mean here that the type is perfectly revealed from the report, i.e., in my notation, each set b_i obtained from the priority algorithm contains a single type.

Definition 4.1 *The unravelling property holds if a PRE exists and, for any type (x, x') such that $x \Vdash x'$ and $x' \Vdash x$, $x = x'$. That is, the unravelling property is equivalent to the priority order $\Vdash$ being a complete order.*

I present next an equivalent characterization of unravelling. First, each type x must have access to (at least) one report $r \in M(x)$ *that cannot be sent by any lower priority type.*

$$\exists r \in M(x), \quad r \notin \bigcup_{x \Vdash x', x' \neq x} M(x') \tag{4.1}$$

This is the key restriction that tends to be violated in models with uncertainty about information endowment, such as Dye (1985) or Jung and Kwon (1988), because types that did not receive information cannot separate. Note that this condition does not require each type to have a report that no other type can send.

Second, each type must be unwilling to pool with lower types. To be more precise, whenever Equation (4.1) is satisfied, there exists a report $r_0(x)$ that maximizes $U(\mathcal{P}(\phi(\{x\})), r)$ over the set of reports that satisfy (4.1). Then, the second condition for unravelling to occur is:

$$U(\mathcal{P}(\phi(\{x\})), r_0(x)) \geq \max U(\mathcal{P}(\phi(M^{-1}(r) \cap \{x' : x \Vdash x'\})), r) \tag{4.2}$$

This second condition is generally violated in models that feature disclosure costs, as in Jovanovic (1982) and Verrecchia (1983), and in which less informative reports tend to involve lower costs The next Lemma follows immediately.

Lemma 4.1 *Suppose that a PRE exists. The unravelling property holds if and only if Equations* (4.1) *and* (4.2) *hold.*

I present next sufficient conditions under which the unravelling property does not hold when the simplified approach of Section 3.3 can be applied. One benefit of the simplified approach is that the absence of unravelling can be demonstrated from properties of the dominance order.

Theorem 4.1 *Suppose that there exists two distinct types x and x' such that (i) $x \trianglerighteq x'$, (ii) $x' \succeq x$ and (iii) either condition (A) or conditions (B1)-(B3) hold. Then, the unravelling property does not hold.*

Theorem 4.1 illustrates that a failure of unravelling can be summarized as a mismatch between the ability to issue reports and the underlying value to buyers. When lower types (according to the value order) are less constrained to send reports, they will strategically choose to pool with higher-value constrained types and prevent the unravelling property.,

As an interesting aside, note that the unravelling property leads to a situation in which all sellers receive a price that is entirely based on their observed information $\tilde{x}$. But, as noted in example 6 earlier, that a seller discloses her information for a cost (i.e., a reduction in utility) can be reinterpreted, in a lemon's market, as a situation in which the seller retains her asset and incurs a discount for the trade inefficiency. In other words, any environment in which the unravelling property leads to an analogous model in which the market breakdown due to a lemon's problem. To put this differently, while disclosure is commonly-viewed as a solution to the lemon's problems (Viscusi (1978), Jovanovic (1982)), the mathematical representation of the unravelling is equivalent to that of the market breakdown. One implication of this is that any result about disclosure has a close analogue in the context of a lemon's market.[19]

4.2. The social value of withholding

In this Section, I re-explore an old question of the voluntary disclosure literature. A number of unravelling results have shown that having the option *not* to release information might not

[19] This similarity has been somewhat clouded in the literature because a large portion of the disclosure literature has used an assumption of additive costs while the lemon's literature favors a multiplicative cost; there is, however, no fundamental reason in either type of model why the cost should be specified in a particular manner.

affect communication; nevertheless, this unravelling property requires restrictive assumptions that are typically not satisfied in most voluntary disclosure models (of which I have shown a few examples). This leaves open the following more general question: does the option to withhold information implies a negative informational externality that is detrimental to at least some of the sellers? This question is of particular interest in the context of financial reporting given that, over the last century, many regulatory actions have removed the option to withhold in the US, e.g., a publicly traded firm cannot withhold an annual report or its mandatory Securities and Exchange Commission (SEC) filings, cannot withhold an audit opinion or, in theory at least, cannot withhold material information. This evolution is in sharp contrast with the more "optional" disclosures that were the norm pre-SEC and remains the object of much unresolved controversy (Stigler (1964), Dye and Sunder (2001), Bushee and Leuz (2005), Greenstone, Oyer and Vissing-Jorgensen (2006)).

A seller who strategically withholds information does not make a false statement, regardless of the underlying information. Hence, I model the option to withhold as a special report r_{nd} which is available to all types and involves no cost to the seller.

Definition 4.2 *The disclosure problem (Q) has a non-disclosure report* r_{nd} *if (i)* $r_{nd} \in M(x)$ *for any* $x \in X$*, (ii) for any* p, r*,* $U(p, r_{nd}) \geq U(p, r)$*.*

I am equipped next to answer whether the option to withhold can be socially costly. This question is of importance because, as long as the reporting space does not solely include the non-disclosure message for any type, a regulator could potentially prohibit non-disclosure at, plausibly, minimal enforcement cost (since this would only involve observing the publicly released report). In fact, even if some types had no other information to disclose (e.g., the uninformed type in Dye (1985)), one could fathom that some alternative disclosure would involve collecting additional information for a possibly very high cost and which might not have been optimal absent regulation - see Shavell (1994) for an example.

Theorem 4.2 *Let* (Q) *be a disclosure problem with a non-disclosure report* r_{nd} *and let* (Q') *be an associated disclosure problem with* $M'(x) = M(x)\backslash\{r_{nd}\}$ *for all* x*. Assume that* (Q) *and* (Q') *admit a PRE* Γ *and* Γ'*, respectively. Then, the allocation* $(u^{\Gamma}(x))_{x \in X}$ *Pareto dominates the allocation* $(u^{\Gamma'}(x))_{x \in X}$*.*

I show in Theorem 4.2 that, regardless of the uses of information or the specification of the reporting space, non-disclosure has a social value in a very strong sense: it makes all types better-off and, generically, at least one type strictly better-off. The reason for this is as follows. The PRE maximizes the welfare of each type sequentially, according to the priority order. Types that choose to withhold always rank last according to the priority order so that any regulation that solely affects their behavior can have no effect on types with greater priority. In fact, removing the option to withhold can only hurt those that would have preferred to withhold.[20]

4.3. Mandatory disclosure

I discuss next whether requiring certain types to make a fully-revealing disclosure would have social value. I introduce first the following characteristic for problems in which mandating disclosure is feasible.

Definition 4.3 *A disclosure problem (Q) has the disclosure property if for any type* $x \in X$ *except at most for one type* x_0*, there exists a "disclosure" report* $r_D(x)$ *defined by* $M^{-1}(r_D(x)) = \{x\}$.

The problem (Q) has the disclosure property if each type can send a separating report, except for one type that would then be fully revealed by sending any of her feasible reports. Note that the disclosure property is (almost) without loss of generality given that it is always possible to assume that this property holds but such that $U(p, r_D(x))$ is sufficiently low for certain types so that it would never be in the best interest of the seller to send the disclosure report.

I am interested next whether, in general, imposing that certain types should send their disclosure report, even if this involved no public monitoring cost, would be desirable. To motivate this idea, note that (like in most persuasion games) the PRE need not attain a first-best allocation if a planner could select which signals are sent by each type.

Theorem 4.3 *Let (Q) be a disclosure problem with the disclosure property and (Q′) be an associated disclosure problem in which* $M(x) = \{r_D(x)\}$ *for any* $x \in \Omega$ *for which a disclosure report exists. Assume that both problems admit a PRE* Γ *and* Γ'*, respectively, and one of these reports* $r_D(x)$ *is used in* Γ*. Then, the following holds:*

(i) The allocation $(u^{\Gamma'}(x))_{x\in X}$ *does not Pareto dominate the allocation* $(u^{\Gamma}(x))_{x\in X}$.

(ii) If $\Omega = \{x : x \Vdash \hat{x}\}$ *for some* $\hat{x}$*, the allocation* $(u^{\Gamma}(x))_{x\in X}$ *Pareto dominates the allocation* $(u^{\Gamma'}(x))_{x\in X}$.

[20] Note that this property would not necessarily hold over the entire set of REs and relies on the characterization of the PRE.

What forms of mandatory disclosure requirements could have social value? Imposing unambiguous disclosures does not. Part (i) establishes that, in the sense of Pareto-efficiency, no disclosure requirement can be beneficial to all sellers if it mandates some firms to use their disclosure signal (recall that, more generally, some mandatory disclosure schemes could not meet condition (i) if they involved some form of mandatory pooling). Part (ii) further reveals that imposing a mandatory disclosure for types with higher priority is necessarily welfare-decreasing. In many commonly-used disclosure models, the notion of priority fully or partly coincides with the value order, implying that (in these settings) imposing disclosures "at the top" is not desirable. By contrast, imposing disclosures over types with low or intermediate priority might be beneficial to some types since it might allow higher priority types to filter out otherwise undesirable low-value types. For problems in which the dominance order coincides with the value order (e.g., Jovanovic (1982), Verrecchia (1983)), "bad news" mandatory disclosure can benefit some sellers while "good news" mandatory disclosure hurts all.

5. Application to continuous problems

5.1. The priority algorithm

I develop in this Section several extensions of the theory for models in which the type space is continuous. Unlike with finite types, the continuous model does not admit an algorithm which reveals the priority order and the order must be conjectured from the nature of the problem under consideration. Nevertheless, the previous analysis suggests the following observation: the priority algorithm closely tracks the ranking of equilibrium utilities so that, if the ranking of utilities can be correctly conjectured, this ranking should also point to the priority order, in which case an algorithm can be properly defined to compute the reporting strategies and actual utilities achieved. Here, I will show this result formally, and then present three examples in which this logic can be applied.

Let X be a measurable type space and, for simplicity, I assume that X has the same cardinality as the real line and, without loss of generality, set $X = [\underline{x}, \overline{x}]$ equal to a proper bounded real interval. To ensure that Bayesian expectations are well-defined, I restrict the attention to REs in which if a report is with measure zero, it is sent by at most one type. Uniqueness is now defined up to differences over sets that have measure zero. In addition, I need to eliminate some pathological equilibria that could potentially arise: formally, I restrict the attention to

REs in which there is no sequence $\{R^{-1}(x_i)\}_{i=1}^{+\infty}$ such that $R^{-1}(x_i)$ has non-zero probability and $\lim Pr(\tilde{x} \in R^{-1}(x_i)) = 0$. The rest of the definitions of a RE and a PRE apply to the continuous setting presented above so I do not repeat them here.

The general idea to solve continuous problems is to find a manner to conjecture the priority algorithm. To do so, assume that the modeler can identify from the problem a unique ranking of utilities (I will show examples later). To be more precise, let $\hat{\geq}$ indicate a preorder that has the property that for any two elements x, x' of X, only one of the following is true: (a) $x \hat{\geq} x'$, or (b) $x' \hat{\geq} x$. Note that that there is no need, at this stage, to know whether a type will achieve a strictly greater equilibrium utility. For any such preorder, I define a "disclosure" utility $u_D(x)$ and the "pooling" utility $u_P(x)$ as:

$$u_D(x) = \sup_{r \in M(x) \bigcap_{x \hat{\geq} x' \neq x} \overline{M}(x')} U(\mathcal{P}(\phi(\{x\})), r) \tag{5.1}$$

$$u_P(x, Y) = \sup_{r \in M(x) \bigcup_{x \hat{\geq} x'} M(x')} U(\mathcal{P}(\phi(M^{-1}(r) \cap Y)), r) \tag{5.2}$$

The disclosure utility $u_D(x)$ can be interpreted as the utility that type x would achieve which would ensure separation from all types that are lower in the sense of the preorder. The pooling utility $u_P(x, Y)$ corresponds to the maximal feasible payoff when type x chooses to pool with some types in a set Y.

Next, let the implied priority algorithm be defined as follows (I omit the dependence on the preorder to save space).

1. Initialize the algorithm at $i = 1$ and $X_1 = X$,

2. Select the maximal element x of X_i according to $\hat{\geq}$. Calculate $u_D(x)$ and $u_P(x, X_i)$. Set $R^a(x) = argmax_{r \in M(x)} U(\mathcal{P}(\phi(M^{-1}(r) \cap X_i)), r)$; if this report does not exist, go to step 7. Otherwise, if $u_P(x) \geq u_D(x)$, go to step 3 and, in the remaining case, go to step 4.

3. If $R^a(x) \notin \bigcup_{x \hat{\geq} x' \neq x} M(x')$, go to step 7. Otherwise, set $B^a(R^a(x)) = \{x \in X_i : R^a(x) \in M(x)\}$, and for any $x' \in M^{-1}(R^a(x)) \cap X_i$, $R^a(x') = R^a(x)$ and $P^a(R^a(x)) = \mathcal{P}(\phi(B^a(R^a(x))))$. Set $Z = B^a(x)$. Go to step 5.

4. Determine the set $A \subseteq X_i$ defined by $A = \{x' : u_D(x') > u_P(x', X_i \cap \{x'' : x' \hat{\geq} x'' \neq x'\})$. If this set has measure zero, go to step 7. Otherwise, for any $x' \in A$, set $R^a(x) = argmax_{r \in M(x) \bigcap_{x \hat{\geq} x' \neq x} \overline{M}(x')} U(\mathcal{P}(\phi(\{x\})), r)$. If any of these is not defined, go to step 6. Then, set $B^a(R^a(x)) = \{x\}$ and $P^a(R^a(x)) = \mathcal{P}(\phi(B^a(x)))$. Set $Z = B^a(x)$. Go to step 5.

5. Set $X_{i+1} = X_i \backslash Z$. If $X_{i+1} = \emptyset$, go to step 2. Otherwise, go to step 6.

6. Complete the off-equilibrium r, with $B^a(r) = argmin_{b \in K(r)} U(\mathcal{P}(\phi(b)), r)$ s.t. $K(r) = \{b :$ if $x \in b, r \in M(x)\}$ and $P^a(r) = \mathcal{P}(\phi(B^a(r)))$. If this is not well-defined, go to step 7 and, otherwise, end the algorithm.

7. End the algorithm: the algorithm crashed and was unable to identify a tentative PRE.

Note that Γ^a is a function of the chosen preorder. Further, by contrast to the finite case, the algorithm is not guaranteed to deliver a proper solution and can "crash" for two reasons. First, it is possible that a most-preferred report is ill-defined in a problem whenever the set X_i is not closed or payoffs are discontinuous. Second, the algorithm may run an infinite loop and never fully exhaust all types. Unfortunately, it is difficult to provide conditions under which the algorithm works in general but, as I will show later on, specialized problems will immediately indicate whether the algorithm crashes. As before, let $\hat{\Vdash}$ denote the priority algorithm implied by the preorder.

Theorem 5.1 *Let $\hat{\geq}$ be a preorder that satisfies the previous conditions and Γ^a be the triplet implied by the priority algorithm. Assume that the algorithm does not crash (reach step 7), includes a finite number of steps and step 2 always has a unique maximizer. Then, a PRE that is consistent with preorder $\hat{\geq}$ must be equal to Γ^a. If $u^{\Gamma^a}(x) \geq u^{\Gamma^a}(x')$ when $x \hat{\Vdash} x'$, γ^a is a PRE.*

Theorem 5.1 offers a helpful constructive argument to obtain PREs in continuous models. Specifically, as long as the weak ranking of equilibrium utilities in any PRE can be determined (generally, from the nature of the problem), then a tentative PRE can be obtained from Γ^a. To verify this is indeed the unique PRE, it remains to be shown that (a) there is only one weak ranking, (b) the algorithm completes in a finite number of steps, and (c) the resulting utilities are monotonic in the steps of the algorithm (and must then confirm the original weak ranking). I will apply next this principle to find PREs in three continuous models.

5.2. Three Applications

Although Theorem 5.1 does not offer a complete characterization of the equilibrium and some further analysis will be required to determine a PRE, some straightforward analysis can often pin down the correct pre-order for the problem. I illustrate this principle with three examples which also offer some interesting stand-alone results.

Assume that the state space is $S = [\underline{s}, \overline{s}]$, a bounded interval of $\mathbb{R}$, and the true state $\tilde{s}$ is a continuous random variable with p.d.f. $f_s(.)$ and c.d.f. $F_s(.)$.[21] The type space is defined as a collection of subsets of S and always satisfies that $\tilde{s} \subseteq \tilde{x}$. I make an additional restrictions to avoid technical issues that can arise in continuous variables and which are common to all the settings. I specialize the market pricing function to have the form, for any measurable subset b of X,

$$\mathcal{P}(\phi(b)) = \mathbb{E}(g(\tilde{s}|\tilde{x} \in b)) \tag{5.3}$$

where $g(.)$ is a continuous increasing function.

This restriction is with loss of generality but allows me to write prices as a continuous function of the underlying distribution about the states of the world.[22] Note that the restriction can accommodate a pure-exchange setting (if $g(.)$ is the identity) as well as settings in which information has social value (if $g(.)$ is convex) or buyers are risk-averse (if $g(.)$ is concave).

The utility function is additively-separable $U(p, r) = u(p) - c(r)$ where $u(.)$ is a continuous and increasing function. As in Grossman (1981) and Milgrom (1981), assume for now that r is a measurable closed subset of S and, when reporting r, a truthful report must be such that $x \subseteq r$. That is, the disclosure can be less precise than the observed information but may not contradict it. There is a cost function $c(r)$, which is non-decreasing when a report is unconditionally more precise (in the sense of Blackwell), i.e., if $r \subseteq r'$, $c(r) \geq c(r')$. The assumption is intuitively appealing because if a measurement system can verify that the observed information is compatible with r, then it will imply that the observed information is compatible with r'. There is no-disclosure report $r_{nd} = S$ whose cost I normalize to $c(r_{nd}) = 0$.

There could be some minor multiplicity of equilibria that could occur if various equilibrium reports are interchangeable. To address this, I view PREs in which the reports imply exactly the same utilities for all types as part of the same equivalence class and define uniqueness up to this equivalence class.[23]

Lastly, I make a (innocuous) regularity condition. For any x, y, y' such that $y \leq y'$, assume

[21] The assumption of a bounded interval is for expositional purposes and does not affect the results.

[22] Without Equation (5.3), because $\mathcal{P}(\phi(.))$ is defined over the measurable subsets of X and thus a topological presentation of continuity would be required. This makes the analysis more cumbersome but does not add much economics to the problem. The results are very similar with other formulations of the production problem, such as the decision problem used in Shavell (1994), i.e., $\mathcal{P}(\phi(b)) = \max_k \mathbb{E}(\tilde{s}r(k) - k|\tilde{x} \in b)$.

[23] One manner to rule out this multiplicity is to assume that the cost function is strictly decreasing in y', i.e., a coarser report is strictly cheaper to send. However, the multiplicity can only occur when there are types in $[y', \overline{s}]$ that make a different equilibrium report; therefore, there is question as to whether making a report that is coarser but in a sense that includes higher types that do not want to send this message in the first place should feature *strictly* lower costs since, plausibly, it should involve no additional verification and does not reveal less information.

that the function $U(\mathcal{P}(\phi([x,\max(x,y)])),[x,\max(x,y')])$ is not constant over any interval. This regularity condition is a translation of condition (G); if it is violated, multiple PRE could exist if the disclosure threshold were to lie exactly at locations where this utility is constant.

5.2.1 Application 1: The canonical disclosure model

I will develop here a solution to the "canonical" disclosure model, in which a seller has a single piece of information that she may decide to disclose or withhold (Grossman and Hart (1980), Jovanovic (1982), Verrecchia (1983), Dye (1985), Jung and Kwon (1988)). To do so, I add to the basic model the following assumptions. First, I set the type space $X = \{S, \{s\}_{s\in S}\}$ where $x = S$ means that the seller does not receive information and $x = s$ means that the seller knows the true state. Conditional on state s, the seller has probability $h(s) \in [0,1)$ not to receive information. Note that this formulation is slightly more general than Dye (1985) and Jung and Kwon (1988) to the extent that, if $h(s)$ is not a constant, a seller who does not receive information is partially informed and can update her prior to $\mathbb{E}(\tilde{s}|h(\tilde{s}) = 1)$. If a seller does not receive information, she has to report $r = r_{nd}$, i.e., $M(S) = \{r_{nd}\}$. If a seller does receive information, she can either report her information or withhold, i.e., for any $x \in S$, $M(x) = \{r_{nd}, s\}$. In this problem, $u_D(x) = U(\mathcal{P}(\phi(\{x\}), x)$ is the utility obtained by informed type $x \in S$ when disclosing the information.

First, note that a no-disclosure REs exists if and only if the following condition holds.

$$U(\mathcal{P}(\phi(S)), r_{nd}) \geq u_D(\overline{s}) \tag{5.4}$$

Second, I consider whether there are REs that have the unravelling property, namely, all types perfectly reveal their information. To begin with, consider the Dye (1985) and Jung and Kwon (1988) setting, in which $h(s) \notin \{0,1\}$ over a set with non-zero measure. Then,

$$U(\mathcal{P}(\phi(S)), r_{nd}) > U(\mathcal{P}(\phi(\{\underline{s}\}), r_{nd}) \geq u_D(\underline{s}) \tag{5.5}$$

Equation (5.5) implies that there is no RE in which the unravelling property holds.

Next, assume that $h(s)$ is always zero for any s, which corresponds to the Jovanovic (1982) and Verrecchia (1983) environments. Consider the case in which $c(\{\underline{s}\}) > 0$. If the unravelling property holds, it must be the case that only type $x = \underline{s}$ reports $r = r_{nd}$ and strictly prefers to do so. But, by continuity, this should also be the case for types x close enough to $\underline{s}$, implying a failure of unravelling. Suppose next that $c(\{\underline{s}\}) = 0$. Then, an unravelling RE exists if and only if, for any $x \in S$,

$$u_D(x) \geq u_D(\underline{s}) \tag{5.6}$$

Proposition 5.1 *The following holds:*

(i) A non-disclosure RE exits if and only if Equation (5.4) *holds.*

(ii) A RE with the unravelling property exists if and only if (a) $Prob(h(\tilde{s}) = 1) = 1$ for all s, or (b) $Prob(h(\tilde{s}) = 1) = 0$, $c(\{\underline{s}\}) = 0$ and Equation (5.6) *holds.*

Note that it is entirely possible that both an unravelling and a no-disclosure RE co-exist, making the analysis of such models without the concept of PRE problematic.

I move next to the characterization of partial disclosure REs in which the report r_{nd} occurs with probability between zero and one. Define an auxiliary function $L(z)$ as follows:

$$L(z) = \{z : x \in K, U(\mathcal{P}(\phi(\{x\})), x) \leq z\}. \tag{5.7}$$

By continuity of the pricing functional, any RE must be such that the set of informed types that do not disclose, b_{nd} has the following form:

$$u_{nd} = U(\mathcal{P}(\phi((L(u_{nd}) \cup \{S\})), r_{nd}) \tag{5.8}$$

Equation (5.8) states that a partial disclosure RE must be such that the utility received conditional on withholding is consistent with the price that would be paid if all sellers better-off withholding were to do so. Note that u_{nd} is defined implicitly and is present on both sides of the Equation; therefore, it needs not have a unique solution. Fortunately, the next Proposition can make use of Theorem 5.1 to characterize the unique PRE of the problem.

Proposition 5.2 *The PRE exists and is unique. Whenever the no-disclosure RE exists, it is the PRE. Otherwise, in the PRE, all non-disclosers achieve an expected utility u^*_{nd} given by the maximal solution of Equation* (5.8) *and an informed type $x \in S$ withholds if and only if $u_D(x) \leq u^*_{nd}$.*

The proof of Proposition 5.2 is immediate as a Pareto-dominated RE cannot be a PRE and a RE with a higher non-disclosure price is preferred by all, strictly so by non-disclosers.[24] Note

[24] Continuity implies that the supremum of all solutions to Equation (5.8) satisfies this Equation as well.

that the argument also offers an immediate proof why the RE is unique in a pure Jung and Kwon (1988) model, without any need to compute the non-disclosure price (or to integrate by parts). If the price functional is the expected state and there are no disclosure cots (a pure-exchange setting), no feasible utility vector is Pareto-dominated, implying that there must be a unique solution to Equation (5.8). The argument further extends to situations in which $g(.)$ is concave (e.g., risk-averse buyers). For the more general problems in which $g(.)$ is neither linear nor concave (e.g., real effects) or costs are non-zero (e.g., proprietary costs), there is no guarantee that the RE is unique; yet, even then, the PRE is unique and selects the RE with the *least* amount of voluntary disclosure.

Note that this result includes unravelling as a special case. With unravelling, the lowest value type does not disclose, implying $u^*_{nd} = u_D(\underline{s})$ (note that unravelling only occurs when the cost of disclosure for the lowest type is zero). This implies that unravelling is a PRE only if there are no other REs. By contrast, if a no-disclosure RE exists, it is always the PRE of the model.[25]

An important property of the model is that, if there are no costs to disclose, the higher informed types are the most willing to disclose, i.e., $L(u_{nd})$ is an interval with upper bound $\overline{s}$. This property does not generalize to the case with disclosure costs, as is made evident from the functional $L(.)$. Specifically, sellers with the highest disclosure utility net of costs are the most willing to disclose. If costs are low for "medium" types, sellers with medium value might disclose while those with high value might not. Nevertheless, the following weaker property can be shown.

Corollary 5.1 *In any PRE that does not have the unravelling property, there exists $s_0 > \underline{s}$ such all types with $s \in [\underline{s}, s_0]$ choose to withhold information.*

The idea of this Corollary is straightforward. Types with the lowest value are those that lose the most from revealing themselves. Therefore, these firms are necessarily the most willing to withhold.

I move next to the comparative statics of the model. To do so, let (Q) indicate a disclosure problem and, to obtain strict comparative statics, I assume that the PRE of (Q) is neither the

[25]Note, lastly, that this PRE has been effectively computed from the priority algorithm obtained earlier. In this problem, it is easy to see that the ranking of equilibrium utilities must be such that the utility of the uninformed is (weakly) lowest and, then, the utility of an informed type is increasing in her utility from disclosure $u_D(x)$. Thus, the PRE must maximize the utility of those informed sellers with highest $u_D(x)$ first which, in turn, involves trying to maximize the non-disclosure price but, if it is below $u_D(x)$, classify type x as a discloser.

(corner) solution of no-disclosure or unravelling. Define $(\hat{Q})$ as a modified problem along one dimension: $(\hat{Q})$ is identical to (Q) except that the cost $\hat{c}(x) > c(x)$ for $x = s$ is greater for all types (except possibly over a set with measure zero). To make the dependence explicit, I index the PRE Γ by the problem under consideration.

Corollary 5.2 *The PRE $\Gamma_{(Q)}$ features more disclosure than the PRE $\Gamma_{(\hat{Q})}$ and a higher non-disclosure price.*

An increase in the cost of disclosure induces more type to withhold and, given that disclosure is costly, these types tend to be higher types that positively contribute to the non-disclosure price. An implication of this feature is that greater disclosure costs always benefits sellers that were previously non-disclosers.

The comparative static with respect to the information endowment is more ambiguous, as I will illustrate. Let $(\hat{Q})$ be defined as the problem with the same cost but a different probability of information endowment $\hat{h}(s)$.

Corollary 5.3 *If $\mathbb{E}(\tilde{s}|\tilde{s} \in L(u^*_{nd})) > \int \hat{h}(s) s f_s(s) ds$ (resp., $<$), the PRE $\Gamma_{(Q)}$ features less disclosure (resp., more) disclosure than the PRE $\Gamma_{(\hat{Q})}$.*

As a special case, if there are no costs of disclosure (thus $L(u^*_{nd})$) is an upper interval and $\hat{h}(s) = h(s) + k(s)$, where $k(s)$ is non-decreasing, i.e., the probability of not receiving information is increased either irrespective of the state or more for good state), the level of strategic disclosure will be lower. Note also that the opposite effect can occur if there are costs of disclosure, because the uninformed types might be types that decrease the non-disclosure price, or if low-value events are more likely to generate no information. The latter effect is, in fact, consistent with the theory of Dye (1985) and Jung and Kwon (1988) given that a function $\hat{h}(s)$ that is steeper may imply that types that do not receive the information are, effectively, better informed about the true state.

5.2.2 Application 2: Costly communication

A significant limitation of the canonical model is that it limits the forms of communication that might be made by sellers. As an example, while it may be possible to truthfully disclose that information $x = s$ has been received, doing so might involve significant proprietary and/or certification costs. As an alternative to this, the seller might prefer to divulge a less precise report r where $\{x\} \subseteq r$. In the previous application, this can only be achieved by making no disclosure but no other coarse disclosure can be made.

I will first develop the idea of coarse disclosures in the context of costly communication and, mainly for expositional purposes, assume that $h(s) = 0$ so that all sellers are informed. I will develop the case of $h(s) > 0$ separately in the model with no costs, although the two approaches can be merged together with a little extra notation.

In this model, I restrict the attention to reports that take the form of closed intervals, so that the cost function is $c([y, y'])$. This restriction is with loss of generality given that, if it were extremely cheap to verify non-interval disclosures, then reports that are not intervals could be sent. Later on, I will provide a sufficient condition on the cost function such that only intervals reports would be used.

The next Lemma formally establishes that one can characterize a PRE in which sellers report a lower bound on their observed information.

Lemma 5.1 *If Γ is a PRE, it has an equivalent PRE in which sellers report $R(x) = [r_1(x), \overline{s}]$, where $r_1(x)$ is an injective mapping defined over S. In particular, the surplus of a seller with type x, $u^\Gamma(x)$, is non-decreasing in x.*

Lemma 5.3 is closely related to Shin (1994)'s sanitization strategies in which a seller removes all potentially adverse disclosures from the feasible set. This principle is generalized here to a problem in which costs restrict how much of this sanitization can occur. Following this general principle, it is never optimal to disclose an upper bound on the true state but, because of the cost function, it may not be optimal to simply set $r_1(x) = x$. I also emphasize here that this property is specific the PRE and would not necessarily hold in all REs of this problem.[26]

Lemma 5.2 *In a PRE, if $r_1(x) < x$, then $r_1(x') = r_1(x)$ for any $x' \in [r_1(x), x)$.*

An implication of Lemma 5.2 is that a seller reporting a lower bound less than her true observed value will pool with other sellers that have value higher than this lower bound and issue the same report. In this respect, a PRE may feature several of such pooling regions in which the issued report imperfectly reveals the observed $\tilde{s}$. In formal terms, a PRE can be

[26] The reason for this is because of some very counter-intuitive REs in this type of model, in which intermediate type t sends a report r where $\max r < \overline{s}$ and higher types with $x > \max r$ would receive an extremely unfavorable price for any (off-equilibrium) report with $\min r \geq t$. This higher type cannot send the report sent by t and thus will send a much coarser report in the equilibrium with $\min r < x$ which in turn might involve an equilibrium price and utility lower than type t.

written as a partition of $S = \{b_i\}$ composed *only* of singletons and intervals, and such that there is one-to-one mapping between an element of the partition and the achieved utility and all sellers whose observation is in the same element of the partition make the same report.

The function $u_D(x)$ is now introduced, noting that if a seller to discloses some of its information, it should always to do by making a disclosure of the form $r = [x_1, \overline{s}]$ in a PRE. As I have shown, setting $r(x) < x$ would necessarily trigger pooling by lower types while setting $\max r < \overline{s}$ is potentially costly and, at best, would imply an equivalent equilibrium. I thus define the partial disclosure utility as:

$$u_P(x, y) = u(\mathbb{E}(g(\tilde{s})|\tilde{s} \in [x, y])) - c([x, \overline{s}]) \tag{5.9}$$

This is the utility obtained when a type sends the equilibrium report $[x, \overline{s}]$ and the market believes that all types with $\tilde{s} > x$ would have chosen a different report.

The fact that utilities are monotonic in the observed information strongly suggests to adapt the priority algorithm to proceed from the highest $\tilde{x}$, sequentially to the lowest $\tilde{x}$. From Theorem 5.1, this can be obtained by applying the priority algorithm and, if it completes in a finite number of steps, the analysis is then complete. I focus first on situations in which a PRE exists such that no information is conveyed. That is, all firms report $r_1(x) = \underline{s}$ and issue the uninformative report.

Proposition 5.3 *There exists a no-disclosure PRE if and only if $u_P(\underline{s}, \overline{s}) \geq u_P(x, \overline{s})$ for any $x \in S$. Then, it is unique PRE.*

I turn next to the polar opposite in which all firms fully disclose their information.

Proposition 5.4 *There exists a PRE with the unravelling property if and only if, for any $x > \underline{s}$, $u_P(x, x) > u_P(x', x)$ as long as $x' < x$. Then, it is unique PRE.*

The cases of no-disclosure and unravelling are special cases in which the priority algorithm completes in a single step. Proposition 5.4 develops an extension of the disclosure unravelling result (Grossman and Hart (1980), Milgrom (1981) and Grossman (1981)) in the presence of disclosure costs. Unravelling may occur even when certification is costly provided costs increase sufficiently slowly relative to the benefit of certification. The Proposition has two well-known special cases. If $g(.)$ and $u(.)$ are linear and $c([x, y])$ is set equal to zero, the condition for the existence and uniqueness of the full-disclosure equilibrium simplifies to $x' > \mathbb{E}(\tilde{s}|\tilde{s} \in [x, x'])$ for all $x < x'$. This condition is always true and corresponds to the standard unravelling result. Vice-versa, if $c([x, y])$ is equal to $c > 0$ for all $x > \underline{s}$, $\underline{s} + \epsilon - c([\underline{s} + \epsilon, \overline{s}]) < \mathbb{E}(\tilde{s})$ so that the condition for unravelling is always violated for ϵ small enough. This is the standard failure of unravelling with a fixed disclosure cost in Jovanovic (1982) and Verrecchia (1983).[27]

The analysis is extended to PREs in which the seller chooses a reporting strategy that conveys some information to buyers. Applying the priority algorithm implies the following description of the PRE. First, I define the function $\zeta(x)$ in the following manner: $\zeta(x)$ is the minimal x' such that $u_P(\zeta(x), z)$ maximizes $u_P(x', z)$. I then run the priority algorithm as follows.

1. Initialize the algorithm at $z_1 = \overline{s}$.

2. Calculate $\zeta(z_i)$ and consider each of the following cases. If $\zeta(z_i) < z_i$, go to step 3 and, otherwise, go to step 4.

3. Set $r_1(x) = \zeta(z_i)$ for any $x \in (\zeta(z_i), z_i]$, $R^a(x) = [r_1(x), \overline{s}]$, $B^a(R^a(x)) = (\zeta(z_i), z_i]$ and $P^a(R^a(x)) = \mathcal{P}(\phi(B^a(R^a(x))))$. Define $Z = B^a(R^a(x))$. Go to step 5.

4. Set $r_1(x) = x$ for all x such that $\zeta(x) = x$, which is denoted as a set A. Then, set $R^a(x) = [x, \overline{s}]$, $B^a(R^a(x)) = \{x\}$ and $P^a(R^a(x)) = \mathcal{P}(\phi(\{x\}))$. Define $Z = A$.

5. Set $z_{i+1} = \min Z$. Whenever $z_{i+1} = \underline{s}$, go to step 6, otherwise, go to step 2.

6. End of the algorithm and complete with off-equilibrium beliefs $B^a(r) = \{\min r\}$ and implied price for any off-equilibrium report.

Proposition 5.5 *There exists a unique PRE (up to off-equilibrium beliefs) which is given by* Γ^a.

A partially-pooling equilibrium may feature one or more pooling regions in which firms with different qualities submit the same report. Note that these pooling regions are obtained sequentially starting with sellers with the highest value. Once it is determined that high-value sellers wish to pool or separate, these sellers are removed from the distribution and the same argument is invoked over the truncated distribution. As a special case, the Proposition admits as a special case the Jovanovic (1982) equilibrium in which there is (at most) a single pooling region that

[27]This is not sufficient in Verrecchia (1983) given that his study uses a support unbounded from below, in which case a condition on the tails of the distribution (exponential decrease) is required but is satisfied by normal distributions.

lies at the lower-tail (low quality sellers do not disclose). More generally, the characterization reveals that higher value sellers may choose to give some imperfect information.

As an application of this characterization, assume that $g(z) = u(z) = z$, and $\tilde{s}$ is logconcave with a truncated expectation $\partial \mathbb{E}(\tilde{s}|\tilde{s} \in [x, y])/\partial x$ decreasing in y (this is satisfied by many common distributions). Then, if $c([x, \overline{s}])$ is concave, the priority algorithm implies that there is a single disclosure region such that types with $x \geq y_0$ disclose. This is the classic result that disclosure occurs for high types generalized to a particular class of distributions and utility functions. Note that, in this case, an increase in disclosures costs expands the non-disclosure region and leads to a coarser communication. On the other hand, if the cost function $c([x, \overline{s}])$ is convex, then any x such that $c_x > 1$ must be part of a pooling region "at the top." More generally, the case of a convex cost function could feature multiple pooling regions.

Example: Suppose that $\tilde{x} \sim U[0, 1]$ and $c(x, y) = \gamma(1 - (y - x)^\alpha)$, then the following holds:

(i) If $\alpha > 1$, all sellers with $x > x_1$ report $r_1(x) = x$ and all sellers with $x < x_1$ report $r_1(x) = \underline{s}$ where:

$$x_1 - \gamma(1 - (1 - x_1)^\alpha) - \frac{x_1}{2} = 0 \tag{5.10}$$

(ii) If $\alpha < 1$, all sellers with $x < x_1$ report $r_1(x) = x$ and all sellers with $x > x_1$ report $r_1(x) = x_1$ where:

$$(x_1 + 1)/2 - \gamma(1 - (1 - x_1)^\alpha) - (1 - \gamma) = 0 \tag{5.11}$$

I conclude this Section application by noting that, if sellers can send reports that are not intervals, the reporting strategy need not be with the form $[r_1(x), \overline{s}]$ because, for high enough types, it might be cheap to pool only with low types and avoid the adverse price effect of pooling with too many intermediate types. Even in this scenario, however, a tentative PRE can be examined by redefining the pooling utility $u_P(.)$ in terms of all possible reports and apply step 3 across all such messages. When applied once using the original preoder, this procedure will yield a new preorder that is reverse-engineered from the priority algorithm. Then, Theorem 5.1 will apply to this new preorder and, given that this order can only deliver the priority algorithm discussed earlier, implies that a tentative PRE for the game. As I note below, however, the possibility of non-interval reports can be also ruled out in a class of cost functions.

Theorem 5.2 *Suppose that in problem (Q), a seller can report any r that contains $\{x\}$ and is* *a finite union of closed intervals. Then, if* $c(.)$ *is only a function of the unconditional probability of* $\tilde{x} \in r$*, an equivalent PRE can be obtained with strategies in which* r *is a closed interval.*

The intuition for this Theorem is simple. If the cost function is not biased to make disclosure of certain realizations of the events more or less costly, then, the utility-maximizing report is always the report that features pooling with higher types. Therefore, the report features a strategy in which sellers report that they are above a certain threshold.

5.2.3 Application 3: Uncertain information endowment

I conclude with a development of a version of Dye (1985) and Jung and Kwon (1988) in which sellers need not be either informed or completely uninformed, and instead various sellers receive different pieces of information (á la Shin (1994), but with a larger reporting space). Let a type be defined as an interval $\tilde{x} = [\tilde{x}_l, \tilde{x}_h]$ where the bounds of this interval are drawn from a continuous bivariate distribution $f_x(.,.|s)$ such that $\tilde{x}_l \leq s \leq \tilde{x}_h$ with probability one so one may interpret each sellers has knowing at least a range about $\tilde{s}$. Assume that $\tilde{x}_l|\tilde{x}_h = x_h, \tilde{s} = s$ has full support over $[\underline{s}, s]$ and $\tilde{x}_h|\tilde{x}_l = x_l, \tilde{s} = s$ has full support over $[s, \overline{s}]$. Note that no one is perfectly informed but there could certain types that have excellent information. The seller can make a report r that is a closed intervals as long as $x \subseteq r$ and there is no cost to make a particular report.

Many of the steps and intuitions of the model with costly disclosure apply to this setting, so that I will mainly focus here on the new facets of this particular model.

Lemma 5.3 *If* Γ *is a PRE, it has an equivalent PRE in which sellers report* $R(x) = [r_1(x_l), \overline{s}]$*, where* $r_1(.)$ *is an injective mapping defined over* S*. In particular, the surplus of a seller with type* x*,* $u^\Gamma(x)$*, is non-decreasing in* x_l*.*

When sellers are imperfectly informed, they never report any information about their upper bound, so that the realization of $\tilde{x}_h$ has no role to play. Then, sellers reporting strategies and utility depends only on the lower bound of what they observe, which effectively constrains their reporting choices. This property, in turn, suggests that the pre-order to be used is one in which a type is higher in the order if x_l is higher, which may introduce a distinction between value and the priority of a type. Interestingly, an algorithm that is very nearly identical to that developped in the case of costly disclosure, but replacing $\tilde{x}$ by $\tilde{x}_l$ applies here. To be more precise, one can redefine $u_D(.)$ and $u_P(.)$ as in Section 5.2.2 with a cost set to zero and $u_P(.)$ as:

$$u_D(x) = u(\mathbb{E}(g(\tilde{s})|\tilde{x}_l = x) \tag{5.12}$$

$$u_P(x, y) = u(\mathbb{E}(g(\tilde{s})|\tilde{x}_l \in [x, y])) \tag{5.13}$$

The function $\zeta(.)$ is then redefined accordingly.

The priority algorithm obtained in Section 5.2.2 applies to the case and delivers the PRE for the model. Under certain conditions, the PRE can be derived explicitly. Suppose that $u_D(x)$ has a unique maximum or a unique minimum. Then, there are four possible types of PRE that may occur in the model.

Proposition 5.6 *Suppose that $g(x) = x$ (risk-neutral pricing). Then:*

(i) If $u_D(x)$ is U-shaped with a unique interior minimum, there is a unique PRE and it features withholding $r_1(x) = S$ for $x_l \leq x^0$ and $r_1(x) = x$ (full disclosure of x_l) for $x_l > x^0$.

(ii) If $u_D(x)$ is inverse U-shaped with a unique interior maximum, there is a unique PRE and it features $r_1(x) = x^0$ for $x_l \geq x_0$ (partial disclosure) and $r_1(x) = x$ for $x < x^0$.

(iii) If $u_D(x)$, there is a unique PRE and it features full disclosure of x_l, i.e., $r_1(x) = x$ for all x.

Within Proposition 5.6, part (i) is consistent with the primary insight from Dye (1985) and Jung and Kwon (1988) that more disclosure should occur for higher outcomes. Naturally, this is a generalization of their analysis since these models restrict the attention to either $\tilde{x}_l = \tilde{s} = \tilde{x}_h$ or $\tilde{x}_l = \underline{s}$ which, in turn, will cause $u_D(x)$ to be decreasing for $x \in (\underline{s}, \overline{s}]$ and with a discontinuity at $x = \underline{s}$ (a degenerate U-shape in which the left branch of the U is vertical). I show in part (ii) that this insight does not entirely generalize to arbitrary information endowments. In fact, the communication can be more precise for lower realizations of $\tilde{x}_l$ if $u_D(x)$ is inverse-shaped; this can occur if (sensibly) sellers that observe a low $\tilde{x}_l$ are also likely to observe a low $\tilde{x}_h$. Part (iii), lastly, develop a weaker notion of unravelling that may nevertheless be present in this model. Although perfect unravelling cannot occur because sellers do not reveal any information about $\tilde{x}_h$, there is a possibility that the lower bound unravels if the effects present in part (ii) are strong enough.

Corollary 5.4 *Under the conditions of Proposition 5.6, a decrease in the quality of information by sellers, defined by* $\tilde{x}_h|\tilde{s}, \tilde{x}_l$ *increasing in the sense of the first-order stochastic dominance, implies a decrease in communication (in the Blackwell sense).*

Corollary 5.4 extends the idea that lower information endowments induces types to be weakly more strategic and pool more, causing total information to decrease as well. This property holds in the type of equilibria described in part (ii) that do not have the form in Dye (1985) and Jung and Kwon (1988). Lastly, note that this property is generally not true if there is more than a single pooling region for $\tilde{x}_l$ since, in these cases, a change in information endowment will tend to change which types are pooled together.

6. Concluding remarks

This paper offers a general methodology to solve and analyze a class of games in which a piece of information can be truthfully disclosed but, beyond the report that is made, there are no other sources of communication that may reveals the seller's type. I show that this class of model has a single reasonable equilibrium that can be explicitly computed in finite games and, for many classic problems, extends to continuous games. The solution concept addresses the large multiplicity of rational expectations equilibria that arises in any voluntary disclosure game in which not all feasible reports are on the equilibrium path.

I offer several theoretical contributions above and beyond the characterization of the perfect sequential rational expectations equilibrium. First, I provide a set of conditions under which the equilibrium exists which include, among others, generalized multi-dimensional versions of the Grossman-Hart-Jovanic-Dye-Verrecchia-Jung-Kwon models, as well as problems with large reporting spaces. Second, I develop several fundamental properties of such voluntary disclosure frameworks, in particular I analyze conditions under which unravelling occurs and, then, show that whenever unravelling does not occur, setting a regulation that removes the option to withhold information leads to a Pareto-decrease in welare. Under certain conditions, requiring disclosure for higher types is also Pareto-decreasing which may partially rationalize why regulations tend to require more demanding disclosures for low-value types. Third, I illustrate the predictions of the model within three classic continuous models: (a) a generalization of the Dye-Verrecchia model with arbitrary costs, real effects and information endowment that correlates to the state, (b) a generalization of the Verrecchia model in which sellers can make interval disclosures, and (c) a generalization of the Dye model in which sellers can be partially informed.

Although the study is not intended to provide all properties for all specialization of voluntary disclosure models, it does suggest a general methodology to approach these problems. This methodology might allow future research to examine previously poorly understood problems, such as multi-dimensional information or when some types are unconstrained about their reporting strategies.

Technical Appendix

Proof of Lemma 3.1: Suppose Γ is a PRE that does not coincide with Γ^a. Let i denote the first step in the algorithm in which at least one type makes a different report under Γ^a. With generic payoffs, this implies that any type $x \in X_i$ must achieve $U(P(R(x)), R(x)) \neq U(P^a(r_i), r_i)$.

I first show that, for all types $x \in X_i$, $U(P(R(x)), R(x)) < U(P^a(r_i), r_i)$. Suppose, by contradiction, that a type $x \in X_i$ exists such that:

$$U(P(R(x)), R(x)) > U(P^a(r_i), r_i). \tag{A.1}$$

By construction of r_i, there must be a type $x' \in X_i$ such that $R(x) \in M(x')$ and $R(x') \neq R(x)$. Then:

$$U(P(R(x')), R(x')) > U(P(R(x)), R(x)).$$

This last inequality implies that x' also verifies Equation (A.1) and one can repeat this argument to find an additional type x'' such that $U(P(R(x'')), R(x'')) > U(P(R(x')), R(x'))$. Iteratively, this implies an infinite sequence of such types, a contradiction.[28]

The claim establishes that pooling on the message r_i must be strictly preferred by all sellers that can send r_i. Therefore, the triple $(r_0, p_0, b_0) = (r_i, P^a(r_i), B^a(r_i))$ violates PRE.□

Proof of Lemma 3.2: I first show that the existence of a PRE implies that the sequence of utilities is decreasing. Assume that Γ^a is the PRE and, by contradiction, suppose that i is the first step of the algorithm such that $U(P^a(r_i), r_i)) < U(P^a(r_{i+1}), r_{i+1}))$. For Γ to be a RE, no

[28]Of note, this proof relies on a *finite* type space, in part because the algorithm itself is only guaranteed to exhaust all types for the case in which the type space is not countably infinite. If the type space is continuous (of which some examples will be given later on), the algorithm must be suitably adapted and additional conditions must hold for its output to be well-defined.

type in X_i that can send r_i should be able to send r_{i+1}, i.e. $X_i \cap M^{-1}(r_i) \cap M^{-1}(r_{i+1}) = \emptyset$. But, then, at step i,

$$\mathcal{P}(\phi(X_i \cap M^{-1}(r_{i+1}))) = \mathcal{P}(\phi(X_{i+1} \cap M^{-1}(r_{i+1}))) = P^a(r_{i+1}). \tag{A.2}$$

Conversely, I show that Γ^a is a PRE when the sequence of utilities is decreasing. Assume that $(U(P^a(r_i), r_i)))$ is decreasing in i. Note first that if $x \in B^a(r_i)$ then (a) type x could not send any of the reports made prior to step i, and (b) type x would be worse-off by sending a report that is made after step i (or any off-equilibrium). It follows that Γ^a is a RE. Next, suppose by contradiction that Γ^a is not a PRE, in which case there exists (r_0, p_0, b_0) that violates PRE. Let i be the last step such that $b_0 \subseteq X_i$, i.e., $b_0 \cap B^a(r_i) \neq \emptyset$. The violation of PRE must make all types in $x \in b_0 \cap B^a(r_i)$ better-off than under Γ^a, that is,

$$U(P^a(r_i), r_i)) < U(\mathcal{P}(\phi(b_0)), r_0). \tag{A.3}$$

Further, by definition of r_i, b_0 cannot contain all types in X_i that can send r_0 and, therefore, there exists x such that $r_0 \in M(x)$ and:

$$U(P^a(R^a(x)), R^a(x)) > U(\mathcal{P}(\phi(b_0)), r_0) \tag{A.4}$$

However, because $(U(P^a(r_i), r_i)))$ is decreasing and $x \in X_i$,

$$U(P^a(R^a(x)), R^a(x)) < U(P^a(r_i), r_i)) \tag{A.5}$$

Equations (A.3), (A.4) and (A.5) imply a contradiction.$\Box$

Proof of Theorem 3.1: By Lemma 3.2, I need to verify that the sequence $(U(P^a(r_i), r_i)))$ is decreasing. By contradiction, suppose that i is the first step of the algorithm such that $U(P^a(r_i), r_i)) < U(P^a(r_{i+1}), r_{i+1}))$. By condition (A), there exists r such that $M^{-1}(r) = M^{-1}(r_i) \cup M^{-1}(r_{i+1})$ and $U(p, r) \geq \min(U(p, r_i), U(p, r_{i+1}))$.

The following holds:

$$X_{i+1} \cap M^{-1}(r) = (X_i \backslash M^{-1}(r_i)) \cap (M^{-1}(r_i) \cup M^{-1}(r_{i+1})) = X_{i+1} \cap M^{-1}(r_{i+1}). \tag{A.6}$$

Substituting in the price in the utility functions,

$$U(\mathcal{P}(\phi(X_{i+1} \cap M^{-1}(r))), r) \geq U(\mathcal{P}(\phi(X_{i+1} \cap M^{-1}(r_{i+1}))), r_{i+1}). \tag{A.7}$$

Since r_{i+1} is selected at step $i+1$, Equation (A.7) must hold at equality, implying that $r_{i+1} = r$. Therefore, $M^{-1}(r_i) \subset M^{-1}(r_{i+1})$.

Next, I examine the price that would be attained when sending r_{i+1} at step i.

$$\mathcal{P}(\phi(X_i \cap M^{-1}(r_{i+1}))) \geq \min(\mathcal{P}(\phi(X_i \cap M^{-1}(r_i))), \mathcal{P}(\phi(X_{i+1} \cap M^{-1}(r_{i+1})))). \tag{A.8}$$

It then follows that:

$$U(\mathcal{P}(\phi(X_i \cap M^{-1}(r_{i+1}))), r_{i+1}) \geq U(\mathcal{P}(\phi(X_i \cap M^{-1}(r_i))), r_i). \tag{A.9}$$

This implies that $r_i = r_{i+1}$, a contradiction.$\square$

Proof of Theorem 3.2: By Lemma 3.1, I need to verify that the sequence $(U(P^a(r_i), r_i)))$ is decreasing. By contradiction, suppose that i is the first step of the algorithm such that $U(P^a(r_i), r_i) < U(P^a(r_{i+1}), r_{i+1})$. Note that $X_i \cap M^{-1}(r_i) \cap M^{-1}(r_{i+1}) \neq \emptyset$ (or else, $r_i = r_{i+1}$). By condition (B2), $M^{-1}(r_i) = [x_i, y_i]$ and $M^{-1}(r_{i+1}) = [x_{i+1}, y_{i+1}]$.

Case 1. Suppose that $y_{i+1} \leq y_i$. Then, $x_i \geq x_{i+1}$.

$$U(P^a(r_i), r_i) < U(\mathcal{P}(\phi(b_{i+1})), r_{i+1}) \leq U(\mathcal{P}(\phi(M^{-1}(r_{i+1}) \cap X_i)), r_{i+1}) \tag{A.10}$$

This is a contradiction to r_i being selected at step i.

Case 2. Suppose that $y_{i+1} > y_i$ and $x_{i+1} \geq x_i$.

$$\mathcal{P}(\phi(M^{-1}(r_{i+1}) \cap X_i)) \geq \underbrace{\min(\mathcal{P}(\phi((x_i, y_i) \cap X_i), \mathcal{P}(\phi((y_i, y_{i+1}) \cap X_i)))}_{=\mathcal{P}(\phi(b_i))} \tag{A.11}$$

By condition (B3),

$$U(\mathcal{P}(\phi(M^{-1}(r_{i+1}) \cap X_i)), r_{i+1}) \geq U(\mathcal{P}(\phi(M^{-1}(r_{i+1}) \cap X_i)), r_i) \geq U(\mathcal{P}(\phi(b_i)), r_i) \tag{A.12}$$

This would imply that $r_i = r_{i+1}$, a contradiction.

Case 3. Suppose that $y_{i+1} > y_i$ and $x_{i+1} < x_i$. By condition (B),

$$\mathcal{P}(\phi(M^{-1}(r_{i+1}) \cap X_i)) \geq \min(\mathcal{P}(\phi(b_i)), \mathcal{P}(\phi(b_{i+1}))) \tag{A.13}$$

By condition (D),

$$U(\mathcal{P}(\phi(M^{-1}(r_{i+1}) \cap X_i)), r_{i+1}) \geq \underbrace{\min(U(\mathcal{P}(\phi(b_i)), r_i), U(\mathcal{P}(\phi(b_{i+1})), r_{i+1}))}_{=U(\mathcal{P}(\phi(b_i)), r_i)}$$

This implies that $r_i = r_{i+1}$, a contradiction.□

Proof of Theorem 3.3: By construction of algorithm b, I need to show that at step i of algorithm a, the set b_i must include a $\trianglerighteq-$maximal type in X_i. By contradiction, if this is not the case, there exists $x \in X_i$ such that $x \notin b_i$ and $x \trianglerighteq x'$ where $x' \in b_i$. Then, there exists a report r such that (a) $U(p, r) \geq U(p, r_i)$ for any p and (b) $M^{-1}(r) \succeq_s M^{-1}(r_i)$. Then,

$$U(p, \mathcal{P}(\phi(X_i \cap M^{-1}(r)))) \geq U(p, \mathcal{P}(\phi(X_i \cap M^{-1}(r)))). \tag{A.14}$$

This would imply that $r_i = r$, a contradiction to $x \notin b_i$.□

Proof of Theorem 3.4: In this proof, I use upperscript a to refer to algorithm a and upperscript b to refer to algorithm b.

Define for any report $r \in M(x)$ (resp., $r \in M(y)$), a function $H^y(r)$ (resp., $H^x(r)$) that maps to a report in $M(y)$ (resp., $M(x)$) that satisfies (a) and (b) in Definition 3.3. Then, define the function $H^n(r) = (H^x \circ H^y)^n$. By condition (a), for any p, $U(p, H^n(r))$ is increasing in r and therefore attains a maximum at $H^n(r) = \psi(r)$. By condition (b), the set $M^{-1}(H^n(r))$ is increasing in the sense of $\succeq_s$ and must also attain its maximum at $M^{-1}(\psi(r))$. It then holds that:

$$M^{-1}(\underbrace{H(\psi(r))}_{=\psi(r)}) \succeq_s M^{-1}(H^y(\psi(r))) \succeq_s M^{-1}(\psi(r)).$$

By condition (B1), this implies that $M^{-1}(H^y(\psi(r))) = M^{-1}(\psi(r))$. From the same argument,

$$U(p, \underbrace{H(\psi(r))}_{=\psi(r)}) \geq U(p, H^y(\psi(r))) \geq U(p, \psi(r)).$$

It then follows that $U(p, H^y(\psi(r))) = U(p, \psi(r))$. In a generic problem, it must then hold that $\psi(r) = H^y(\psi(r))$. Let $M_2(x) \subseteq M(x) \cap M(y)$ be the set of reports that can be written as $\psi(r)$, then, by construction of algorithms a and b, a type picked at step i would always choose a report in $M_2(x)$, at which point both type x and y would be included in b_i. This establishes the claim.□

Proof of Theorem 4.1: Assume that (i)-(iii) hold. From Theorem 3.3, $x \Vdash x'$, so that let r_i be the report made by x and r_j be the report made by x', where $j > i$. Suppose by contradiction that the unravelling property holds.

Assume that condition (A) holds. I show first that $M^{-1}(r_i) \subset M^{-1}(r_j)$. From the construction of the algorithm, $x' \notin M^{-1}(r_i)$ so that the two sets cannot be equal. By condition (A), there exists a report r such that $M^{-1}(r) = M^{-1}(r_i) \cup M^{-1}(r_j)$ and:

$$U(\mathcal{P}(\phi(X_j \cap M^{-1}(r))), r) \geq U(\mathcal{P}(\phi(X_j \cap M^{-1}(r_j))), r_j) \tag{A.15}$$

It follows that $r = r_j$ which, in turn, implies the claim that $M^{-1}(r_i) \subset M^{-1}(r_j)$. Then,

$$U(\mathcal{P}(\phi(X_j \cap M^{-1}(r_j))), r_j) = U(\mathcal{P}(\phi(\{x'\}, r_j) \geq \underbrace{U(\mathcal{P}(\phi(X_i \cap M^{-1}(r_i))), r_i)}_{=U(\mathcal{P}(\phi(\{x\})), r_i)} \tag{A.16}$$

By Lemma 3.2, this implies that $r_i = r_j$, a contradiction.

Assume that conditions (B1)-(B3) hold. As before, $x' \notin M^{-1}(r_i)$ or else unravelling would not hold. By conditions (B1) and (B2), this implies that $x' \succeq \max M^{-1}(r_i)$, hence $\max M^{-1}(r_j) \succeq \max M^{-1}(r_i)$ and, by condition (B3), $U(p, r_j) \geq U(p, r_i)$ for any p.

$$U(\mathcal{P}(\phi(\{x'\})), r_j) \geq U(\mathcal{P}(\phi(\{x\})), r_i) \tag{A.17}$$

By Lemma 3.2, this implies that $r_i = r_j$, a contradiction.□

Proof of Theorem 4.2: In this proof, I denote with ' (prime) the variables that correspond to problem (Q'). Let j be the last step in the priority algorithm applied to the disclosure problem (Q). For any $i < j$, $r_i \neq r_{nd}$ or else i would have been the last step. At step j,

$$U(\mathcal{P}(\phi(b_j)), r) = U(\mathcal{P}(\phi(X_j)), r) \leq U(\mathcal{P}(\phi(\underbrace{M^{-1}(r_{nd}) \cap X_j}_{X_j})), r_{nd}) \tag{A.18}$$

Therefore, $r_j = r_{nd}$.

For any $i < j$, the priority algorithm can achieve the same utility with the smaller set of messages under (Q') as under (Q). Therefore, the $r_i = r'_i$ and all types selected at a step prior to j achieve the same utility in both problems.

Consider step j, by construction of the priority algorithm,

$$U(\mathcal{P}(\phi(X_j), r_{nd})) \geq \max_r U(\mathcal{P}(\phi(M^{-1}(r) \cap X_j), r)) \tag{A.19}$$

Therefore, all types in b'_j are better-off under (Q) than under (Q') (strictly so, in a generic problem). To conclude, note by Lemma 3.2, for any $i' > j$,

$$\max_r U(\mathcal{P}(\phi(M^{-1}(r) \cap X_j), r)) \geq U(\mathcal{P}(\phi(b_{i'}), r_{i'})) \tag{A.20}$$

It then follows that *all* types in X_j are strictly better-off under (Q) than under (Q').□

Proof of Proposition 3.1: By Lemma 3.1, there is at most one PRE and it must be given by the priority algorithm. Suppose, by contradiction, that there exists a first step i such that:

$$U(P^a(r_i), r_i) < U(P^a(r_{i+1}), r_{i+1}) \tag{A.21}$$

By construction of the algorithm, this can only occur if $X_i \cap M^{-1}(r_i) \cap M^{-1}(r_{i+1}) \neq \emptyset$ and $M^{-1}(r_{i+1}) \not\subseteq M^{-1}(r_i)$. In this problem, this implies that $M^{-1}(r_i) \subset M^{-1}(r_{i+1})$.

Suppose that r_i contains l iterations of NI and $(d-l)$-uple set of disclosures y. Then, r_{i+1} must contain $l' > l$ iterations of NI and a $(d-l')-$uple of disclosures y' where $y' \subset y$. Then,

$$U(\mathcal{P}(\phi(X_i \cap M^{-1}(r_{i+1}))), r_{i+1}) \geq U(\min(P^a(r_{i+1}), P^a(r_i)), r_{i+1}) \tag{A.22}$$

This implies that:

$$U(\mathcal{P}(\phi(X_i \cap M^{-1}(r_{i+1}))), r_{i+1}) \geq U(P^a(r_i), r_i) \tag{A.23}$$

This contradicts that $r_i \neq r_{i+1}$.□

Proof of Theorem 4.3: (i) Let i be a step in the priority algorithm in which $x \in b_i \cap \Omega$ and a disclosure report is not used. Then,

$$U(\mathcal{P}(\phi(b_i)), r_i) > U(\mathcal{P}(\phi(\{x\})), r_D(x)) \tag{A.24}$$

Therefore type x will be strictly worse-off under (Q') and $(u^{\Gamma'}(x))_{x\in X}$ cannot Pareto dominat-evthe allocation $(u^{\Gamma}(x))_{x\in X}$.

(ii) The argument used in (i) also implies that any $x \in \Omega$ is weakly worse-off under (Q'), strictly if this type did not use $r_D(x)$. Types $x \notin \Omega$ are selected after type $\hat{x}$ which implies that utilities are computed over the subset of types $X \backslash \Omega$ and therefore the PRE for (Q) and (Q') coincide.$\square$

Proof of Lemma 5.3: For expositional purposes, I state the proofs with a single type; however, to be fully rigorous, this proof needs to be applied over a subset of types with non-zero mass. Suppose that Γ admits a type x such that $\max R(x) < \overline{v}$. Define a new Γ' such that all sellers that were reporting $R(x)$ now report $r' = [\min R[x], \overline{v}]$ and buyers respond to this price by offering $P(r') = P(R(x))$. Note that if all types with $y \in [\max R(x), \overline{s}]$ achieve a utility $u^{\Gamma}(y) \geq u^{\Gamma}(x)$, these types would not send the report r' in the new RE and therefore the new RE is an equivalent RE, and therefore also an equivalent PRE.

For the new Γ' not to be an RE, it must that for a non-zero mass of types $y \in [\max R(x), \overline{s}]$, $u^{\Gamma}(y) \geq u^{\Gamma}(x)$; denote this set A_1. Then, a contradiction to PRE can be obtained by setting $p_0 = P(R(x))$, $r_0 = r'$ and $b_0 = A_1$.$\square$

Proof of Lemma 5.2: Assume that $r_1(x) < x$ for some x.

Define $K(x) = \{x' : r_1(x) \leq x' \leq x\}$. Let $x' \in K(x)$. Then, $R(x)$ is a feasible message for type x' and, vice-versa, by Lemma 5.3, $R(x')$ is feasible for type x, which then implies that $u^{\Gamma}(x) = u^{\Gamma}(x')$. Then, there exists an interval T such that $u_0 = u^{\Gamma}(x)$ if and only if $x \in T$. Let the report r_0 be defined as $r_0 = [\min T, \overline{s}]$

Suppose (by contradiction) that $r_1(x) > \min T$. For any $y \in T$,

$$u_0 = u(E(g(\tilde{s})|R(\tilde{s}) = R(y))) - c(R(y)) \tag{A.25}$$

Slightly rearranging this Equation and taking expectations over $y = \tilde{s} \in T$,

$$E(g(\tilde{s})|\tilde{s} \in T) = \mathbb{E}(u^{-1}(u_0 + c(R(\tilde{s})))|\tilde{s} \in T) \tag{A.26}$$

$$u(E(g(\tilde{s})|\tilde{s} \in T)) - c(r_0) = u(\mathbb{E}(u^{-1}(u_0 + c(R(\tilde{s})))|\tilde{s} \in T) - c(r_0)) \tag{A.27}$$

There are two cases to consider. First, suppose that there exists a positive mass of such types y such that $c(R(y)) > c(r_0)$. Then:

$$u(E(g(\tilde{s})|\tilde{s} \in T)) - c(r_0) > u(u^{-1}(u_0)) = u_0 \tag{A.28}$$

Then, setting r_0, $p_0 = E(g(\tilde{s})|\tilde{s} \in T)$ and b_0 as the set of all types in T as well as all types higher than $\max T$ that obtain a utility less than u_0 under Γ would contradict that Γ is a PRE.

Second, if $c(R(\tilde{s})) = c(r_0)$ for all $\tilde{s} \in T$, then an equivalent PRE can be constructed by substituting r_0 instead of $R(\tilde{s})$ with no effect on the equilibrium utilities.$\square$

Bibliography

Acharya, Viral V., Peter DeMarzo, and Ilan Kremer, "Endogenous Information Flows and the Clustering of Announcements," *American Economic Review*, September 2011, *101* (7), 2955–79.

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

Bagnoli, Mark and Susan G. Watts, "Financial Reporting and Supplemental Voluntary Disclosures," *Journal of Accounting Research*, 2007, *45* (5), 885913.

Baldenius, Tim, Nahum Melumad, and Xiaojing Meng, "Advising and Monitoring CEOs: The Dual Role of Boards," Technical Report 2011. Journal of Financial Economics, forth.

Banks, Jeffrey S and Joel Sobel, "Equilibrium selection in signaling games," *Econometrica*, 1987, pp. 647–661.

Benabou, Roland and Guy Laroque, "Using Privileged Information to Manipulate Markets: Insiders, Gurus, and Credibility," *Quarterly Journal of Economics*, 1992, *107* (3), 921–958.

Bertomeu, Jeremy, "Ph.D Lecture Notes, Chapter II: Voluntary disclosure," 2012.

Beyer, Anne, "Capital Market Prices, Management Forecasts, and Earnings Management," *The Accounting Review*, 2009, *84* (6).

___ **and Ronald A. Dye**, "Reputation Management and Voluntary Disclosures," 2011. Review of Accounting Studies, forth.

___**, Daniel A. Cohen, Thomas Z. Lys, and Beverly R. Walther**, "The financial reporting environment: Review of the recent literature," *Journal of Accounting and Economics*, 2010, *50* (2-3), 296–343.

Blume, Lawrence E and William R Zame, "The algebraic geometry of perfect and sequential equilibrium," *Econometrica*, 1994, pp. 783–794.

Bushee, Brian J. and Christian Leuz, "Economic consequences of SEC disclosure regulation: evidence from the OTC bulletin board," *Journal of Accounting and Economics*, 2005, *39* (2), 233 – 264.

Bushman, Robert M and Raffi J Indjejikian, "Voluntary disclosures and the trading behavior of corporate insiders," *Journal of Accounting Research*, 1995, *33* (2), 293–316.

Carlsson, Hans and Eric Van Damme, "Global games and equilibrium selection," *Econometrica*, 1993, pp. 989–1018.

Caskey, Judson, Venky Nagar, and Paolo Petacchi, "Reporting bias with an audit committee," *The Accounting Review*, 2010, *85* (2), 447–481.

Crawford, Vincent P., "A Theory of Disagreement in Bargaining," *Econometrica*, 1982, *50* (3), 607–638.

da Silva Pinheiro, Tiago, "The effect of incentive-pay and accounting discretion in a dynamic model of earnings management ," 2013. Ph.D dissertation, University of Chicago.

Dranove, David and Ginger Zhe Jin, "Quality Disclosure and Certification: Theory and Practice," *Journal of Economic Literature*, 2010, *48* (4), 935–963.

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

——, "Proprietary and Nonproprietary Disclosures," *Journal of Business*, 1986, *59* (2), 331–366.

——, "Earnings Management in an Overlapping Generations Model," *Journal of Accounting Research*, 1988, *26* (2), 195–235.

——, "An evaluation of 'essays on disclosure' and the disclosure literature in accounting," *Journal of Accounting and Economics*, 2001, *32* (1-3), 181–235.

—— **and Mark Finn**, *Equilibrium Voluntary Disclosures When Firms Possess Random Multidimensional Private Information* Essays on Accounting Theory in Honor of Joel S. Demski, edited by Rick Antle, Pierre Jinghong Liang, Frystein Gjesdal, New York, NY: Springer Publishing, 2007.

—— **and Shyam Sunder**, "Why Not Allow FASB and IASB Standards to Compete in the U.S.?," *Accounting Horizons*, 2001, *15* (3), 257–271.

Einhorn, Eti and Amir Ziv, "Intertemporal Dynamics of Corporate Voluntary Disclosures," *Journal of Accounting Research*, 2008, *46* (3), 567–589.

Fischer, Paul E. and Robert E. Verrecchia, "Reporting Bias," *The Accounting Review*, 2000, *75* (2), 229–245.

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

Fudenberg, Drew and Jean Tirole, *Game Theory*, MIT Press, 1991.

Greenstone, Michael, Paul Oyer, and Annette and Vissing-Jorgensen, "Mandated Disclosure, Stock Returns, and the 1964 Securities Acts Amendments," *Quarterly Journal of Economics*, 2006, *121* (2), 399–460.

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

——— **and Motty Perry**, "Perfect sequential equilibrium ?," *Journal of Economic Theory*, 1986, *39* (1), 97–119.

——— **and Oliver D. Hart**, "Disclosure Laws and Takeover Bids," *Journal of Finance*, 1980, *35* (2), 323–334.

Guttman, Ilan, Ilan Kremer, and Andrzej Skrzypacz, *Not Only What But also When: A Theory of Dynamic Voluntary Disclosure*, Stanford Graduate School of Business, 2012.

———**, Ohad Kadan, and Eugene Kandel**, "A Rational Expectations Theory of Kinks in Financial Reporting," *The Accounting Review*, 2006, *81* (4), 811–848.

———**,** ———**, and** ———, "Dividend stickiness and strategic pooling," *Review of Financial Studies*, 2010, *23* (12), 4455–4495.

Huddart, Steven, John S Hughes, and Carolyn B Levine, "Public disclosure and dissimulation of insider trades," *Econometrica*, 2001, *69* (3), 665–681.

Hughes, John S. and Suil Pae, "Voluntary Disclosure of Precision Information," *Journal of Accounting and Economics*, 2004, *37* (2), 261–289.

Jorgensen, Bjorn and Michael Kirschenheiter, "Discretionary Risk Disclosures," *The Accounting Review*, 2003, *78* (2), 449–469.

Jovanovic, Boyan, "Truthful Disclosure of Information," *Bell Journal of Economics*, 1982, *13* (1), 36–44.

Jung, Woon-Oh and Young K. Kwon, "Disclosure When the Market Is Unsure of Information Endowment of Managers," *Journal of Accounting Research*, 1988, *26* (1), 146–153.

Kanodia, Chandra S. and Deokheon Lee, "Investment and Disclosure: The Disciplinary Role of Periodic Performance Reports," *Journal of Accounting Research*, 1998, *36* (1), 33–55.

Kartik, Navin, Marco Ottaviani, and Francesco Squintani, "Credulity, lies, and costly talk," *Journal of Economic theory*, 2007, *134* (1), 93–116.

Kirschenheiter, Michael, "Information Quality and Correlated Signals," *Journal of Accounting Research*, 1997, *35* (1), 43–59.

Kohlberg, Elon and Jean-Francois Mertens, "On the Strategic Stability of Equilibria," *Econometrica*, 1986, *54* (5), 1003–1038.

Kreps, David M and Robert Wilson, "Sequential equilibria," *Econometrica: Journal of the Econometric Society*, 1982, pp. 863–894.

Kumar, Praveen, Nisan Langberg, and K Sivaramakrishnan, "Voluntary disclosures, corporate control, and investment," *Journal of Accounting Research*, 2012, *50* (4), 1041–1076.

Laux, Volker and Phillip C Stocken, "Managerial reporting, overoptimism, and litigation risk," *Journal of Accounting and Economics*, 2012, *53* (3), 577–591.

Mailath, George J., Masahiro Okuno-Fujiwara, and Andrew Postlewaite, "Belief-based Refinements in Signalling Games," *Journal of Economic Theory*, 1993, *60* (1), 241–276.

Marinovic, Ivan, "Internal control system, earnings quality, and the dynamics of financial reporting," *The RAND Journal of Economics*, 2013, *44* (1), 145–167.

—— **and Felipe Varas**, "Dynamic costly disclosure," 2013. Stanford GSB Working Paper.

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

Pae, Suil, "Selective disclosures in the presence of uncertainty about information endowment," *Journal of Accounting and Economics*, 2005, *39* (3), 383 – 409.

Penno, Mark C., "Information Quality and Voluntary Disclosure," *The Accounting Review*, 1997, *72* (2), 275–284.

Rubinstein, Ariel, "Perfect Equilibrium in a Bargaining Model," *Econometrica*, 1982, *50* (1), 97–110.

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

Shin, Hyun Song, "News management and the value of firms," *Rand Journal of Economics*, 1994, pp. 58–71.

—— , "Disclosures and asset returns," *Econometrica*, 2003, *71* (1), 105–133.

Stigler, George J., "Public Regulation of the Securities Markets," *Journal of Business*, 1964, *37* (2), 117–142.

Stocken, Phillip, "Credibility of Voluntary Disclosure," *Rand Journal of Economics*, 2000, *31* (2), 359–374.

Subramanyam, K. R., "Uncertain Precision and Price Reactions to Information," *The Accounting Review*, 1996, *71* (2), pp. 207–219.

Suijs, Jeroen, "Voluntary disclosure of information when firms are uncertain of investor response," *Journal of Accounting and Economics*, 2007, *43* (2-3), 391–410.

Teoh, Siew Hong and Chuan Yang Hwang, "Nondisclosure and adverse disclosure as signals of firm value," *Review of Financial Studies*, 1991, *4* (2), 283–313.

Verrecchia, Robert E., "Discretionary Disclosure," *Journal of Accounting and Economics*, 1983, *5*, 179–194.

—— , "Essays on Disclosure," *Journal of Accounting and Economics*, 2001, *32*, 97–180.

Viscusi, W Kip, "A note on" lemons" markets with quality certification," *The Bell Journal of Economics*, 1978, pp. 277–279.