Multidimensional Disclosure

Giorgio Martini*

December 2018

I study a model of verifiable information disclosure in which the Sender's type is multidimensional. I propose a new explanation for why complete unraveling does not occur in practice, in situations such as firms disclosing financial information, politicians revealing conflicts of interest, and participants in two-sided matching markets choosing which attributes to show. When the Sender's preferences over market beliefs are sufficiently convex, equilibria feature partial disclosure. Types that are low along any dimension pool together and reveal nothing. Thus, unlike traditional one-dimensional models, complete unraveling does not occur. Convex preferences arise when the Sender is uncertain over which dimensions matter, or is disclosing to a committee or jury with heterogenous preferences.

1 Introduction

Voluntary disclosure arises in many economic contexts: a CEO disclosing information on firm performance to shareholders; a job applicant composing their resume; a prosecutor selectively presenting evidence to a jury; a politician disclosing potential conflicts of interest and tax returns. In many such contexts, the Sender's information is multidimensional. For a CEO, information on revenue and profitability can be disaggregated by geography, division, or product-line; similarly, there is substantial discretion in how to break down balance-sheet information. Job applicants choose whether to reveal grades, and choose which experiences, certifications, and references to include in their resume.

My model explores verifiable disclosure in environments where information is multidimensional. I show that the standard unraveling prediction from the vertical, one-dimensional model – that all information will be voluntarily revealed, because the

*Microsoft; email: giorgio.martini@gmail.com; website: giorgiomartini.com.

inference from not disclosing would be too negative – breaks down when information is multidimensional. Instead, "good news" (types that are high along all dimensions) is fully revealed, whereas all other types pool together ("vague news") and reveal nothing. Unraveling stops midway.

The core prediction of less-than-full unraveling relies on the Sender's preferences over Receiver's beliefs being sufficiently convex. I show that this assumption is satisfied in several natural environments. First, the Sender may be uncertain about which dimensions the Receiver cares about, and is either ambiguity-averse or risk-averse. A second, independent foundation arises when the Sender is disclosing to multiple and heterogenous stakeholders, who care about a different mix of dimensions.

A leading application of my model is to financial disclosure. The full unraveling prediction is seen as counterfactual in this environment: leaving aside mandatory disclosure, there is great heterogeneity among firms in voluntary disclosure policies and large levels of apparent non-disclosure. The accounting literature has identified two disclosure frictions that break the paradoxical unraveling prediction: costly disclosure costs and lack of common knowledge that the Sender is informed. However, these frictions need to be large to meaningfully alter the equilibrium prediction. To give a recent example, Apple has announced that starting from Q1 2019 it will no longer break apart its hardware revenue by product line.[1] It is certainly common knowledge among investors that Apple possesses such a breakdown, and direct disclosure costs are negligible.[2] My model suggests an additional force that stops unraveling when information is multidimensional.

1.1 Literature review

The literature on voluntary disclosure was initiated by Grossman and Hart (1980), Grossman (1981), and Milgrom (1981), who established the seminal unraveling result. Two strands of the accounting literature then showed that the unraveling result breaks down under disclosure frictions: in Dye (1985) and Jung and Kwon (1988), because the market is uncertain whether the Sender is informed; in Jovanovic (1982) and Verrecchia (1983), because disclosure is costly.

[1] For recent news coverage, see e.g. https://www.cnbc.com/2018/11/02/heres-what-every-major-analyst-had-to-say-about-apples-earnings.html and https://www.forbes.com/sites/petercohan/2018/11/17/what-are-apple-ge-and-google-trying-to-cover-up/.

[2] However, there may be indirect disclosure costs by revealing which products are more profitable. For example, Amazon's decision to not disaggregate AWS revenue may have been driven by a desire to hide from (potential) competitors how valuable cloud computing was becoming. I thank Michael Ostrovsky for making this point.

The disclosure literature is still very active: more general models of verifiable disclosure are studied by Bertomeu and Cianciaruso (2018) and Rappoport (2017), among many others. These models consider complex evidence structures, and show that full unraveling does not in general obtain; in my model, the evidence structure is simple and unraveling stops because of the geometry of the problem.

There are other papers that study voluntary disclosure with multidimensional information. In Fishman and Hagerty (1990) the Sender can only disclose one of many dimensions. In Shin (2003), Pae (2005), and Dziuda (2011), the Sender observes multiple dimensions and chooses which to disclose, but the Receiver is uncertain over the Sender's information endowment as in Dye (1985). Acharya, DeMarzo, and Kremer (2011) and Guttman, Kremer, and Skrzypacz (2014) add dynamics and show that the timing of disclosure plays a crucial role.

To the best of my knowledge, my paper is the first to show that multidimensionality alone – without additional disclosure frictions, and without a complex evidence structure – is enough to stop unraveling.

2 Leading example

To highlight the main features of my model and give intuition for its mechanics, I first present a special case. The model is described formally and in greater generality in Section 4.

Let Sender's private information $\theta = (\theta_1, \theta_2)$ be two-dimensional, distributed according to the bivariate uniform distribution on the unit square $[0, 1] \times [0, 1]$.[3] Sender decides whether to reveal θ (full disclosure), reveal nothing (no disclosure) or possibly reveal only one or the other component (selective disclosure). Disclosure is verifiable (only true information can be revealed) and costless, and it is common knowledge that the Sender is informed.

Along each of the two dimensions, the Receiver (or market) forms rational expectations given the Sender's disclosure strategy. The market belief is denoted by $x = (x_1, x_2)$. When the Sender discloses a component i, market beliefs are correct: $x_i = \theta_i$. When the Sender does not disclose θ_i, the market belief x_i is the expected value of θ_i given non-disclosure (whenever this is well-defined). Sender's payoff is $u(x_1, x_2) = \min\{x_1, x_2\}$.

An *equilibrium* is a disclosure strategy and market belief such that the Sender's choice

[3] In particular, the two components of the private information are symmetric and independent. Neither of these assumptions are required for my results; in the full model, arbitrary distributions are allowed.

of disclosure is optimal given the market beliefs, and market beliefs are correct given the Sender's disclosure strategy.

To build intuition, suppose that the Sender never discloses their type: all (θ_1, θ_2) pool together. (Since Sender's utility is concave, this disclosure strategy would indeed be optimal if it could be sustained, as can be shown formally using the results of Chapter 2.) After observing no disclosure, market beliefs are set by taking conditional expectations: $x_i = E[\theta_i| \text{ no disclosure }] = 1/2$, for $i = 1, 2$. Therefore, types with $\theta_1 > 1/2$ and $\theta_2 > 1/2$ — that is, the upper right quadrant of the unit square — have a strict incentive to deviate to full disclosure: market beliefs would then be set to $x_1 = \theta_1$ and $x_2 = \theta_2$ (despite being off-path, disclosure is always interpreted correctly since it must be truthful), and those Sender types are better off.

We can continue this logic one step further. Suppose that Senders with $(\theta_1, \theta_2) \gg (1/2, 1/2)$ disclose their type, and all other types pool together and disclose nothing. Then market beliefs after non-disclosure can be verified to be $(x_1, x_2) = (5/12, 5/12)$. Again, there is a set of types — the L-shaped region with $5/12 < \theta_i \leq 1/2$, for $i = 1, 2$ — where deviating to full disclosure is strictly profitable.

This logic is, of course, the same as that of the classic unraveling results of Grossman (1981) and Milgrom (1981): for every disclosure strategy where types at the bottom pool (any other disclosure strategy clearly can not be sustained in equilibrium), there are above-average types that benefit from deviating to disclosure. In the one-dimensional case, the logic inexorably leads to full unraveling: all types (except possibly the lowest type) disclose their private information, and after non-disclosure the market believes that the Sender is the worst type possible.

However, this logic does not lead to full unraveling in two or more dimensions. Suppose the unraveling logic had brought us to consider the disclosure strategy where almost all types disclose; only those with $\theta_1 < \varepsilon$ or $\theta_2 < \varepsilon$, for some small ε, do not disclose their type. What are market beliefs after observing non-disclosure? Clearly, at least one of the two dimensions is below ε, by construction. Market beliefs, however, are determined dimension-by-dimension, and the adverse inference along any one dimension is not as extreme. Take belief x_1. With probability $\sim 1/2$, non-disclosure was caused by a low θ_1, in which case the inference is extremely low: $\varepsilon/2 \simeq 0$. However, with probability $\sim 1/2$, non-disclosure was caused by a low θ_2; this is uninformative about θ_1, which on average is $1/2$. Putting things together, $x_1 \simeq \frac{1}{2} \cdot 0 + \frac{1}{2} \cdot \frac{1}{2} = 1/4$. (All quantities are approximate because this calculation double-counts the second-order event in which both θ_1 and θ_2 are below the threshold ε.) But $1/4 > \varepsilon$: thus, there is a region of types that has a profitable deviation to non-disclosure.

To summarize, there are two contrasting phenomena. When non-disclosure is too common, the marginal non-disclosing types strictly prefer to distinguish themselves from the (average of the) pooling types, and deviate to disclosure. When non-disclosure is too rare, the marginal disclosing types benefit by deviating to non-disclosure, where they enjoy the benefit of "cross-subsidization" across the two dimensions. There is an intermediate level of non-disclosure where these two forces cancel out, which constitutes an equilibrium of this example.

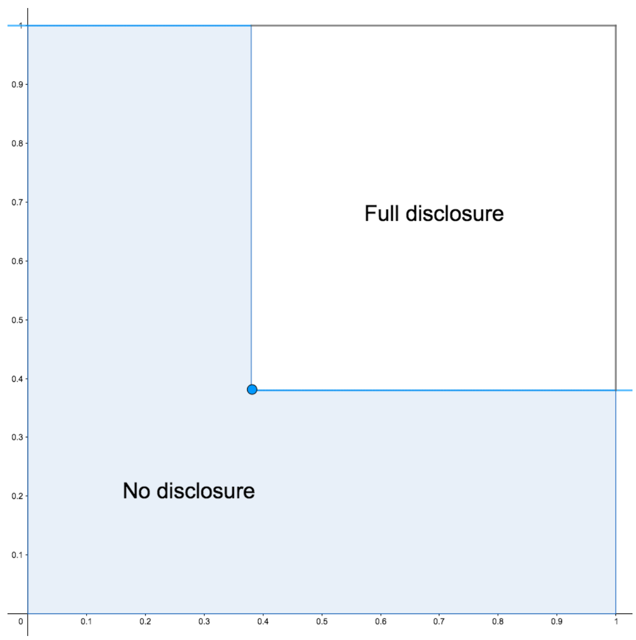

Figure 1: Partially-unraveled equilibrium

Figure 1 shows an equilibrium of the model. All types in the shaded L-shaped region – where either θ_1 or θ_2 or both are below a critical threshold – disclose neither dimension. Types in the upper-right square disclose both dimensions. Unraveling stops here because, conditional on no disclosure, market beliefs (represented by the blue dot) lie on the threshold of the non-disclosure region. Since market beliefs are not interior to the non-disclosure region – as they always would be in a one-dimensional model – there are no types that want to deviate to disclosing. This equilibrium is prototypical, and the main focus on the paper is on equilibria of the same form:

Definition 1. An equilibrium is a *partially-unraveled equilibrium* if all on-path actions are either full-disclosure (all dimensions are revealed) or no-disclosure, and no-disclosure happens with positive probability.

2.1 Full disclosure equilibrium

The basic example features a second, less interesting, equilibrium (and it can be shown that there are no other equilibria). Suppose that the market expects the Sender to always disclose both dimensions, and believes that the Sender's type is 0 along any dimension that is not disclosed (the same skeptical beliefs that sustain the classic one-dimensional equilibrium). Then for all Sender types the best response is to disclose both dimensions (types that have one dimension equal to exactly 0 are indifferent). In turn, this validates the market beliefs; in particular, since non-disclosure is off path, Bayes rule does not restrict market beliefs.

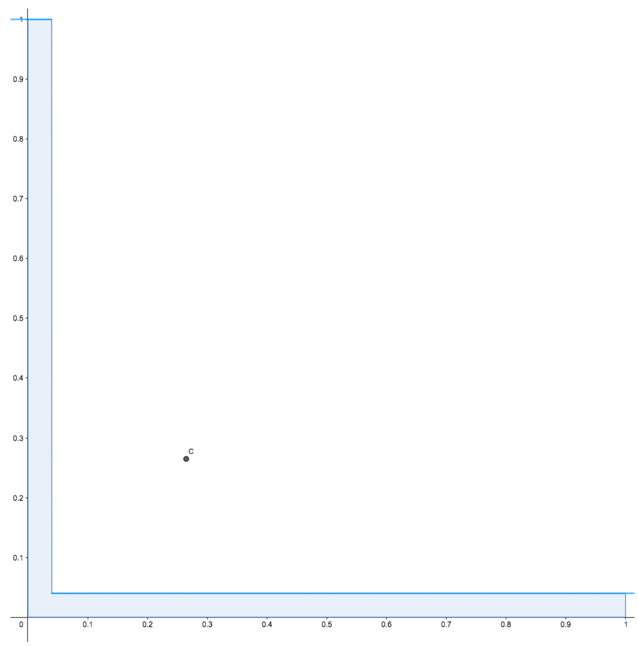

Figure 2: Non-disclosure set with small disclosure costs

I will now show that the full-disclosure equilibrium does not survive the addition of small disclosure costs, à la Verrecchia (1983). (Similar arguments work for a Dye (1985)-like perturbation.) Suppose that the Sender incurs a small disclosure cost $\varepsilon > 0$ if they choose to disclose. Then Sender types for whom $\theta_1 < \varepsilon$ or $\theta_2 < \varepsilon$ will prefer not to disclose, regardless of market beliefs: by disclosing their payoff would be negative. Therefore the non-disclosure region includes at least the L-shaped strip depicted in Figure 2. Now Bayes rule pins down market beliefs after non-disclosure, and those beliefs are far from the threshold non-disclosing type, regardless of how small ε is. (In the uniform case, they are always ≥ 0.25.) It can be shown that the only equilibrium that survives this perturbation is the partially-unraveled equilibrium.

This is in contrast the one-dimensional model, where the full disclosure equilibrium is robust to *small* frictions of this kind (in that the equilibrium prediction moves continuously with the size of the friction).

2.2 Equilibria with selective disclosure

The equilibria described above feature only full disclosure and no disclosure on the equilibrium path. A natural question is then: are there other equilibria where information is disclosed selectively? For example, it would be reasonable to conjecture that there are equilibria in which agents disclose their best feature only, provided it is sufficiently high. Indeed, this is the driving force in other models of multidimensional disclosure (eg, Fishman and Hagerty (1990)), where selective disclosure is the only possibility (because the Sender is not allowed to disclose all information, or because the Receiver has limited information processing capacity). In contrast, in the leading example (and more generally when there are only two dimensions) there is no equilibrium with selective disclosure. The conjectured equilibrium in which the Sender discloses their best attribute cannot be sustained because their utility will then depend (also) on the other attribute, on which the market forms rational expectations; but since now there is only one dimension of uncertainty left, the standard unraveling argument applies. When there are three or more dimensions, the results are not as sharp: see Section 5.

2.3 Disclosure in many dimensions

I now consider equilibrium behavior (in partially-unraveled equilibria) as the number of dimensions k grows. I continue to restrict attention to the multivariate uniform prior and min preferences, where the partially-unraveled equilibrium can be characterized in closed form.

Proposition 1. *In dimension k, the unique partially-unraveled equilibrium is given by the disclosure strategy $\sigma = \mathbb{1}_{[a,1]^k}(\theta)$, where disclosure threshold a is the unique solution that satisfies $a \in (0,1)$ to*

$$(1-a)^{(k+1)} + 2a - 1 = 0.$$

Proof. See Section 7.1. □

Corollary 2. *In the limit as $k \to \infty$, the probability of full disclosure in the partially-unraveled equilibrium goes to 0.*

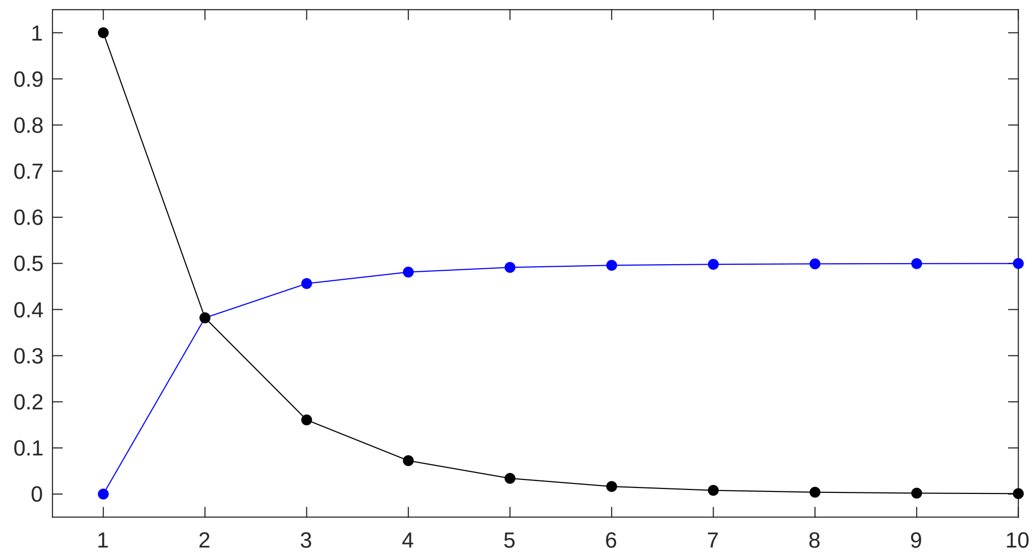

Figure 3: Disclosure threshold a (in blue) and probability of disclosure (in black) as a function of the number of dimensions k. Also shown is the classical unraveled equilibrium for $k = 1$.

Figure 3 shows the threshold $a(k)$ and the probability of disclosure $P(\theta \in [a(k), 1]^k)$ for different values of k. As k grows, the probability of disclosure goes to 0. The larger k is the likelier it is that at least one dimension is low; hence, there are less and less types that gain from separating (which in turn makes pooling less informative about any one dimension).

2.4 More general preferences

The examples discussed so far all used variations on the $u(x, y) = \min\{x, y\}$ preferences. The driving force of the model – that non-convexities in the lower contour sets can stop unraveling – applies much more widely. Figure 4 shows an example of another utility function that also has a unique partially-unraveled equilibrium (in two dimensions with uniform prior).

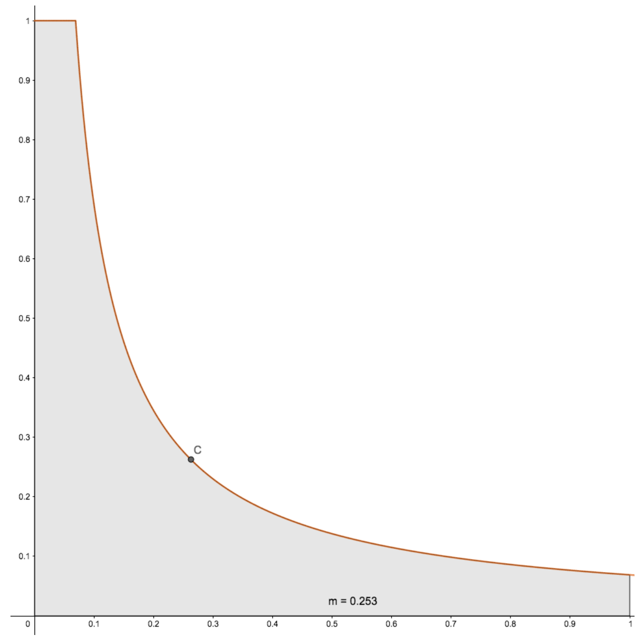

Figure 4: Partially-unraveled equilibrium with $u(x, y) = xy$

3 Sources of convex preferences

For the model to make predictions that are different from those in the classic one-dimensional model, the Sender's preferences must be sufficiently convex, and the components of θ not too correlated. The one-dimensional model can be seen as a special case of my model when either preferences are linear (so that the components of θ can be summarized by their weighted average) or when the prior belief is concentrated along a ray. In either case, full unraveling would be the unique equilibrium prediction.

A ***non***-example of a model that satisfies my assumptions is complementarity among dimensions in the market's evaluation of the Sender's type. For example, suppose that Sender is CEO of a firm whose true market value depends on the weakest of two crucial features, à la O-ring model. The CEO observes the two features and makes a disclosure decision to maximize expected market value. After no disclosure, the Sender's payoff is

$$u(\text{no disclosure}) = E[\min\{\theta_1, \theta_2\} \mid \text{region of no disclosure}]. \tag{1}$$

Note how the minimum function is *inside* the expectation. Suppose the no-disclosure region were L-shaped like in Section 2. Then, after observing no disclosure, the market cannot infer which dimension was low – but they *can* infer that the minimum of the two is low, and hence so is the firm's value. As a result, full unraveling is the unique equilibrium.

For my model to have bite, convexity must arise from the Sender's preferences only. In contrast to (1), the Sender's preferences must be of the form

$$u(\text{no disclosure}) = \min\{E[\theta_1|\text{ no disclosure}], E[\theta_2|\text{ no disclosure}]\}, \tag{2}$$

where the minimum (or more generally any sufficiently convex function) is taken *outside* the expectations. Below I present three applications where this required structure arises naturally.

3.1 Risk aversion

Suppose that the Sender is uncertain about which dimensions of θ are considered important by the Receiver. For example, a CEO may be uncertain about which of two features (say, performance of two different divisions) will matter more in determining the share price. Furthermore, the CEO is risk-averse: a non-trivial fraction of the CEO's wealth is tied to the share price.

For concreteness, let $k = 2$ and suppose that there are two states of the world (equally likely and independent of θ): θ_1 1 matters (and the share price is proportional to market belief x_1), or θ_2 matters. Sender's utility is a concave function of share price, say the logarithm. The Sender maximizes expected utility. Then, the Sender's utility function is

$$u(x_1, x_2) = \frac{1}{2}\log(x_1) + \frac{1}{2}\log(x_2)$$

which can be monotonically transformed into $u(x_1, x_2) = x_1 x_2$. As shown in Figure 4, under such preferences there is a unique partially-unraveled equilibrium.

3.2 Knightian uncertainty

Convexity of the Sender's preferences can also arise when the Sender is uncertain over which dimension matters, and faces Knightian uncertainty or is ambiguity averse. The starkest example is when the Sender is extremely ambiguity averse, and believes that the dimension that matters will always be the one on which the market has the most negative beliefs. This gives rise to the min preferences discussed throughout Section 2.

Such non-Bayesian uncertainty aversion can arise for behavioral reasons (for example, a job applicant that fears they will always be judged on their worst characteristic) or out of a desire for robustness (for example, a marketer deciding not to release a mixed set of signals out of fear that it may backfire).

3.3 Multiple stakeholders

A third setting where convexity may arise is when the Sender is disclosing to multiple stakeholders simultaneously. Examples include: a prosecutor choosing what to disclose to a jury, who will then vote to convict or acquit; a job applicant choosing what to include on their resume, which will be examined by multiple interviewers; finally, a politician disclosing information to different segments of the electorate. The key features for my purposes are that stakeholders put different weights on different dimensions of the Sender's information, and that the Sender's preferred decision requires unanimity. In the job applicant example, one interviewer may put more value on the signal given by a detailed GPA breakdown, whereas another interviewer may value precise descriptions (job title, responsibilities) of previous roles instead; and both have to approve the candidate for hiring to take place.

For concreteness, suppose that there are two stakeholders $i = 1, 2$, and that stakeholder i only cares about dimension i. If the Sender's payoff is determined by the stakeholder with the more negative belief (because both stakeholders have veto power, for example), then we are in the world of Section 2. If stakeholders take a binary decision (hire/no-hire, convict/acquit), unanimity is required, and θ_i determines the probability (all else held constant) that stakeholder i takes the preferred action, then the example in Section 2.4 applies.

4 General model

I now generalize the example of Section 2 to an arbitrary number of dimensions and possibly asymmetric utility function and prior belief. To make the analysis more stark, the Sender's decision is now restricted: they either reveal all information (truthfully) or reveal none. In some situations, this is indeed the Sender's decision problem. For example, a job candidate may have discretion over submitting a detailed transcript or not; but a partially-redacted transcript cannot be produced. In Section 5 I discuss how the analysis changes when selective disclosure is allowed: the all-or-nothing equilibrium discussed here survives, but generally (when there are 3 or more dimensions) other equilibria arise as well.

The state of the world is a random vector $\theta \in \Theta$ distributed according to some distribution F on a convex and compact set $\Theta \subseteq \mathbb{R}^k$, with $k \geq 2$. There is a single player, Sender, who privately observes the realization of θ. After observing θ, Sender chooses whether to disclose the entire vector θ or to disclose nothing. Disclosure is verifiable: only the true value can be disclosed. Disclosure is costless and the Sender is always informed.

Let $\sigma : \Theta \to \{0, 1\}$ represent the Sender's (pure) disclosure strategy, where 0 represents non-disclosure. Along each dimension i, the Receiver (or market) forms rational expectations x_i given the Sender's disclosure strategy. When the Sender discloses component i, market beliefs are correct: $x_i = \theta_i$. When the Sender does not disclose θ_i, the market belief x_i is the expected value of θ_i given non-disclosure when such expectation is well-defined, and arbitrary otherwise. Let x_{ND} be the vector of market beliefs following non-disclosure.

Sender's payoff u is type-independent and only depends on the vector of market beliefs x. We assume throughout that u is continuous, weakly increasing ($u(x) \geq u(x')$ if $x_i \geq x'_i$ for all dimensions i), and strictly increasing for strict increases ($u(x) > u(x')$ if $x_i > x'_i$ for all i). Normalize $\min_{\theta\in\Theta} u(x) = 0$ and $\max_{x\theta\in\Theta} u(x) = 1$ (well-defined by compactness of Θ and continuity).

Definition 2. A (Perfect Bayesian) *equilibrium* is a pair of disclosure strategy σ and market beliefs after non-disclosure x_{ND} such that

1. σ is optimal given x_{ND}: $\sigma(\theta) = 1$ if and only if $u(\theta) \geq u(x_{ND})$;

2. If $P(\theta|\sigma(\theta) = 1) < 1$, then x_{ND} is correct given σ: $x_{ND} = E[\theta|\sigma(\theta) = 0]$.

An equilibrium is *partially-unraveled* if $P(\theta|\sigma(\theta) = 1) < 1$. For all utility levels $v \in (0, 1]$, let $\mathbf{c}(v)$ be the center of mass of the lower contour set $L(v) := \{\theta \in \Theta : u(\theta) \leq v\}$ (which is non-empty by construction):

$$\mathbf{c}(v) := \frac{\int_{L(v)} \theta dF(\theta)}{\int_{L(v)} dF(\theta)}$$

Since Θ is convex, $\mathbf{c}(v) \in \Theta$ for all v. Thus, $u(\mathbf{c}(v))$ is well-defined. Consider the $[0, 1] \to \mathbb{R}$ function

$$d(v) := u(\mathbf{c}(v)) - v.$$

Proposition 3. *A disclosure strategy σ is sustained in a partially-unraveled equilibrium if and only if it is of the form $\sigma = \mathbb{1}_{\{\theta\in\Theta\,:\,u(\theta)\geq v\}}(\theta)$ for a utility threshold v that satisfies $d(v) = 0$.*

Proof. For any disclosure strategy, all types that do not disclose receive the same utility (because utility is only a function of market beliefs, i.e., type-independent). Thus, only disclosure strategies that disclose on some upper contour set can be sustained in equilibrium. Next, if v is such that $d(v) > 0$, then threshold types (that is, θ such that $u(\theta) = v$) would be better off by deviating to non-disclosure. Similarly, if $d(v) < 0$, then types immediately below the threshold indifference curve would be better off by deviating to full disclosure. Thus, σ is part of a partially-unraveled equilibrium if and only if $d(v) = 0$. $\square$

By Proposition 3, to find equilibria of the model it is enough to find zeros of the function $d(v)$. The following lemma gives sufficient conditions under which $d(v)$ is continuous.

Lemma 1. *If u is continuous, monotone, and strictly increasing in strict increases, and F admits a bounded density f, then d is continuous.*

Continuity of d allows us to use the Intermediate Value Theorem to prove that there exists a partially-unraveled equilibrium, when further conditions can be verified. When u is sufficiently concave, for low values of v the set $L(v)$ is highly non-convex. Thus, as long as the components of θ are not too correlated, the center of mass of $L(v)$ falls outside the set. That is, there exists a $\underline{v}$ such that $d(\underline{v}) > 0$. Conversely, when $L(v) = \Theta$, by convexity the center of mass falls in the interior of the set, hence $d(\overline{v}) < 0$ is always satisfied for large enough $\overline{v}$. When such $\underline{v}$ and $\overline{v}$ can be found, then a partially-unraveled equilibrium is guaranteed to exist.

It is tempting to guess that the function $d(v)$ must be monotone: when the disclosure threshold v increases – so that increasingly "better" types are pooling with the low-types – market beliefs after non-disclosure should increase. While this is typically the case, there are pathological examples where higher disclosure thresholds change the distribution of pooling types is such a way that $\mathbf{c}(v)$ is not increasing. This can occur when the mass below $u(\theta) = v'$ is concentrated in the "corners" of Θ, and the mass in the region $v' < u(\theta) \leq v$ is both greater and concentrated along the diagonal. Monotonicity of $d(v)$, when it can be verified, is sufficient to prove that there is a unique partially-unraveled equilibrium.

5 Selective disclosure

In the general model of Section 4 the Sender is restricted to either disclosing the entire vector θ or to disclosing none of it. In many applications, this is the only type of disclosure that is feasible. For example, it is typically impossible for a graduate to produce a transcript that only shows grades for a favorably chosen subset of classes. In other applications, however, *selective* disclosure — only revealing private information along certain dimensions — is technologically feasible. For example, a CEO may disclose information about the performance of certain company divisions, products, or geographic areas, while simultaneously not disclosing any information on the rest. In this section, I ask whether such selective disclosure can be sustained in equilibrium.

Consider then an extension to the main model in which the Sender, after observing θ, chooses a subset $M \subset 2^K = 2^{\{1,\dots,k\}}$ of dimensions to disclose. (In the main model, the Sender is restricted to choosing $M \in \{\emptyset, K\}$.) The Sender's disclosure strategy is now a function $\sigma : \Theta \to 2^K$ that indicates which dimensions are disclosed. Upon observing selective disclosure, the market forms rational expectations along all undisclosed dimensions (and when the specific subset of dimensions revealed is off-path, or when the "wrong" type chooses a certain subset of dimensions to reveal, then beliefs are unrestricted). Note that the inference for one dimension will in general depend on the disclosure (or lack thereof) along other dimensions, even when the dimensions are independent.

First, note that all equilibria discussed in the main model – that feature only full or no disclosure – can be sustained as equilibria in the richer model. To do so, it is enough to specify that the Receiver holds skeptical beliefs after all (off-path) selective disclosures: when dimension i is not disclosed (but at least one other dimension is disclosed), then $x_i = \inf \Theta_i$. This makes such a deviation unprofitable: no or full disclosure (or both) give a higher payoff. Unlike the one-dimensional model, such extreme beliefs are not necessary to sustain the equilibrium: as long as x_i is low enough to always place the Sender on an indifference curve lower than that attained by no disclosure, no type will deviate to selective disclosure.

Next, we look at equilibria that feature selective disclosure. I first prove a negative result: there is no equilibrium in which all but one dimension are ever revealed on-path. Then I show that there exist equilibria that are unraveled along a subset of dimensions (meaning those dimensions are always revealed).

Proposition 4. *There is no equilibrium in which the Sender discloses exactly $k-1$ dimensions with positive probability.*

Proof. When, in addition to the maintained assumptions, we also have that u is strictly increasing ($u(\theta) > u(\theta')$ if $\theta_i \geq \theta'_i$ for all i and $\theta_j > \theta'_j$ for at least one j), the logic familiar from unraveling in the one-dimensional model applies. Suppose without loss of generality that with positive probability dimensions 1 through $k-1$ are disclosed, and dimension k is not disclosed. After such selective disclosure, market belief x_k is pinned down by Bayes rule. Since the event has positive probability, it includes types with $\theta_k > x_k$. By strict monotonicity of u, these types can profitably deviate to full disclosure. I omit the (more complex) argument needed when u is not strictly monotone, but only monotone and strictly monotone for strict increases. $\square$

In particular, when $k = 2$, all equilibria feature only full and no disclosure.

When $k \geq 3$, there exist many asymmetric equilibria in which arbitrary subset of dimensions (of size less than or equal to $k - 2$) must be disclosed. These equilibria are sustained by assigning skeptical beliefs after all disclosure choices that don't disclose the "mandated" dimensions. When the number of always-disclosed dimensions is 2 or higher, I conjecture that such skeptical beliefs do not survive the cost-based perturbation of Section 2.1. When the number of dimensions is high, the spirit of the partially-unraveled equilibria should be preserved – the only adjustment is that there exist equilibria in which one dimension and only one is fully unraveled.

6 Concluding remarks

My model studies a novel force that stops unraveling: multidimensional information coupled with convex preferences over market beliefs. This is independent of the well-known forces that come from costly disclosure and uncertainty over the Sender's information endowment. We share the same basic policy implication: since full unraveling will not occur under voluntary disclosure, making disclosure mandatory can be effective at increasing information revelation.

Because information is multidimensional, however, more nuanced implications can be drawn. Suppose that the designer only has the authority to make disclosure of one dimension mandatory. In the leading example with k dimensions (Section 2.3), making disclosure along any dimension mandatory decreases the disclosure threshold for the remaining $k - 1$ dimensions. Thus, in equilibrium, disclosure across different dimensions act as complements, unlike the standard intuition (and prediction of other models of multidimensional disclosure) that more disclosure along one dimension means less along others. A similar effect arises when the designer can *forbid* a certain dimension from being disclosed, effectively eliminating it from the game. This causes no information disclosure along the blocked dimension, but it increases equilibrium disclosure along remaining dimensions. In both cases, by making "vague news" less vague, there is less room for low-type Senders to pool together.

7 Omitted proofs

7.1 Proof of Proposition 1

Let $\mathbf{c}(S)$ denote the center of mass of subset $S \subseteq [0,1]^k$. Consider the accounting identity

$$\mathbf{c}([0,1]^k) = P(\theta \in [a,1]^k)\mathbf{c}([a,1]^k) + P(\theta \notin [a,1]^k)\mathbf{c}([0,1]^k \setminus [a,1]^k)$$

We have $P(\theta \in [a_k,1]^k) = (1-a_k)^k$. Thus, rearranging,

$$\mathbf{c}([0,1]^k \setminus [a,1]^k) = \frac{1}{1-(1-a)^k}\left(\left(\frac{1}{2},\ldots,\frac{1}{2}\right) - (1-a)^k\left(\frac{1+a}{2},\ldots,\frac{1+a}{2}\right)\right)$$

The equilibrium condition is $\mathbf{c}([0,1]^k \setminus [a,1]^k) = (a,\ldots,a)$. By symmetry, we can focus on any one component of the k-dimensional vectors:

$$a = \frac{1}{1-(1-a)^k}\left(\frac{1}{2} - (1-a)^k\frac{1+a}{2}\right) = \frac{1-(1-a)^k(1+a)}{2-2(1-a)^k}$$

which can be rearranged as

$$(1-a)^{(k+1)} + 2a - 1 = 0.$$

7.2 Proof of Lemma 1

Since u is continuous, it is enough to show that $\mathbf{c}(v)$ is continuous. I first prove that the numerator $\int_{L(v)} \theta dF(\theta) = \int_{L(v)} \theta f(\theta) d\theta$ is continuous, using continuity along sequences. Let $v_n \to v$. The sequence of indicator functions $\mathbb{1}_{L(v_n)}$ converges pointwise to $\mathbb{1}_{L(v)}$ almost everywhere (with respect to the Lebesgue measure on Θ), because the indifference curve $\{\theta : u(\theta) = v\}$ (the only place where convergence may fail) has measure zero, by the assumptions on u. Next, note that $\theta f(\theta) d\theta$ is bounded above by the constant vector with all components equal to $\max_{\theta\in\Theta} f(\theta)$, which is finite by assumption. Thus, by the dominated convergence theorem,

$$\int_{L(v_n)} \theta dF(\theta) = \int_{\Theta} \theta \mathbb{1}_{L(v_n)} dF(\theta) \to \int_{\Theta} \theta \mathbb{1}_{L(v)} dF(\theta) = \int_{L(v)} \theta dF(\theta)$$

which proves continuity. An identical argument applies to the denominator. Thus, $\mathbf{c}(v)$ is continuous as long as the denominator does not go to 0; this can only happen for $v = \inf_{\theta\in\Theta} u(\theta)$, and thus $\mathbf{c}$ can be extended by continuity there. $\square$

References

Acharya, Viral V, Peter DeMarzo, and Ilan Kremer. Endogenous information flows and the clustering of announcements. *American Economic Review*, 101(7):2955–79, 2011.

Bertomeu, Jeremy and Davide Cianciaruso. Verifiable disclosure. *Economic Theory*, 65(4): 1011–1044, 2018.

Dye, Ronald A. Disclosure of nonproprietary information. *Journal of accounting research*, pages 123–145, 1985.

Dziuda, Wioletta. Strategic argumentation. *Journal of Economic Theory*, 146(4):1362, 2011.

Fishman, Michael J and Kathleen M Hagerty. The optimal amount of discretion to allow in disclosure. *The Quarterly Journal of Economics*, 105(2):427–444, 1990.

Grossman, Sanford J. The informational role of warranties and private disclosure about product quality. *The Journal of Law and Economics*, 24(3):461–483, 1981.

Grossman, Sanford J and Oliver D Hart. Disclosure laws and takeover bids. *The Journal of Finance*, 35(2):323–334, 1980.

Guttman, Ilan, Ilan Kremer, and Andrzej Skrzypacz. Not only what but also when: A theory of dynamic voluntary disclosure. *American Economic Review*, 104(8):2400–2420, 2014.

Jovanovic, Boyan. Truthful disclosure of information. *The Bell Journal of Economics*, pages 36–44, 1982.

Jung, Woon-Oh and Young K Kwon. Disclosure when the market is unsure of information endowment of managers. *Journal of Accounting research*, pages 146–153, 1988.

Milgrom, Paul R. Good news and bad news: Representation theorems and applications. *The Bell Journal of Economics*, pages 380–391, 1981.

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

Rappoport, Daniel. Evidence and skepticism in verifiable disclosure games. 2017.

Shin, Hyun Song. Disclosures and asset returns. *Econometrica*, 71(1):105–133, 2003.

Verrecchia, Robert E. Discretionary disclosure. *Journal of accounting and economics*, 5: 179–194, 1983.