Optimal News Management*

Boris Ginzburg†

December 2013

Abstract

An informed Sender can communicate news to a Receiver, who then decides on her action. The Sender can send a different signal for every realisation of news. I show that the Sender's optimal strategy is *simple* - all signals are degenerate, news that are not revealed precisely are pooled into only one set, and this set will normally consist of only a small number of disjoint intervals. Greater transparency is optimal when the Receiver is more likely to be predisposed against the Sender. These results shed some light on phenomena such as political censorship, restrictions on hate speech, central bank transparency, and disclosure of information by firms.

Keywords: censorship, political communication, hate speech, persuasion games, information disclosure, transparency.

JEL codes: D82, D72, C72

1 Introduction

There are many contexts in which an informed agent chooses how to reveal information to a decision-maker. For example, governments can often choose how much information to disclose to citizens in order to maximise public support. Central banks can decide how much transparency they want to permit - for instance, whether to make their internal forecasts public. Firms select the amount of information they reveal about their product quality, financial status, and environmental performance. How does the optimal communication policy look?

To answer this question, this paper proposes a model centred on a game between a Sender (e.g. a government) and a Receiver (e.g. a voter). The Sender chooses how to communicate news to the Receiver. He does it by mapping each realisation of the news to a probability distribution over a set of messages - in effect, selecting a signal for any realisation of news. In turn, the Reciever, upon

*I am grateful to V. Bhaskar for guidance and helpful advice. I also thank Nageeb Ali, Roland Bénabou, Alessandro Ispano, Philippe Jehiel, Emir Kamenica, Christian Krestel, Stephen Morris, John Quah, Francesco Squintani, Felix Várdy, Jörgen Weibull, Andriy Zapechelnyuk, and audiences in Barcelona, Brussels, London, Madrid, and Riga, for valuable comments.

†Department of Economics, University College London

seeing a message, forms a posterior belief about the news. She can then choose an action - a real number. The Sender would like that number to be high. The Receiver, on the other hand, wants to select a higher action (e.g. to vote for the government) only when the news are good enough - namely, when the value of the news is above the Receiver's preference parameter, which is her private information. The distribution of these preference parameters corresponds, for example, to a distribution of attitudes towards the government. The Sender thus wants to select a news management policy that, in expectation, maximises the probability that the Receiver's posterior belief about the news is above her preference parameter. In general, the Sender has a very large set of strategies at his disposal - he can pool news together into many different sets, these sets can consist of many disjoint intervals, and he can complicate the strategy further by mapping each of these sets set to a probability distribution over a large set of messages.

The key contribution of this paper is to show that the equilibrium news management policy is *simple*. First, there will always exist an optimal strategy under which the Sender will only send degenerate signals (i.e. signals that produce a single message with probability one). The Sender will never want to randomise. Thus, Sender's strategy can be reduced to a partition of the set of news, with every element of the partition containing news for which the same message is sent.

Second, it will be shown that the optimal partition will contain at most one non-singleton element. In other words, the Sender will pool some news together into one set, while disclosing the other news precisely. This corresponds to a strategy of *censorship* - the Receiver will be allowed to learn some of the news exactly, while the rest will be hidden. The strategy of disclosing news *approximately* - pooling them into a number of sets, so that the Receiver could learn that the news belong to one of them - will not be pursued, although it is potentially available.

Third, the optimal censorship policy will be simple as well. In principle it is still possible for the Sender to make the set of censored news complicated. For example, he can pool together very bad news and very good news, while disclosing "average" news. In reality, this is rarely observed, and indeed this paper will show that such a policy can rarely be optimal. More specifically, it will be shown that the set of censored news cannot consist of many disjoint intervals, unless the distribution of the Receiver's preferences has a complicated shape.

The actual news management policy will largely depend, at an equilibrium, on the shape of the Receiver's preferences. Both full revelation and complete hiding of information emerge as special cases under certain distributions of the Receiver's preference parameter. In general, it will be shown that it is optimal to reveal more information if the Receiver is predisposed against the Sender. On the positive side, this suggests that regimes that are more popular[1] are likely to restrict freedom of the press to a greater extent. Similarly, firms facing more

[1] For reasons unrelated to the news being disclosed.

skeptical consumers will reveal more information about their product quality. On the normative side, measures such as restrictions on racist hate speech are useful only when the public is largely anti-racist to start with. Similarly, central banks are better off with more openness when facing skeptical market players. More generally, the paper suggests that the optimal amount of transparency varies depending on the audience's preferences. This helps explain why, for example governments can both gain and lose from censoring bad news[2], central banks differ in their level of transparency[3], commercial banks tend to reveal varying amounts of financial information[4] , and firms differ in the amount of information they make public[5]

This paper is related to a large literature on the so-called persuasion games, in which a Sender ho to communicate a state of the world to a Receiver, who then selects an action that affects payoffs of both sides. Past research on persuasion games has largely concluded that in general, the Receiver will learn the state of the world[6]. This is because if the Sender chooses to pool states over some set S, then, whenever the news fall in S, he will always want to deviate to disclosing the "best" news in that set, as this will give him a higher payoff than letting the Receiver make a decision based on her posterior belief conditional on state being in S. This paper, however, differs from much of the previous research by examining what happens when the Sender has to commit to a particular disclosure strategy before learning the state. Using the commitment assumption, I am able to show that in wide range of cases, full disclosure is not an equilibrium strategy for the Sender - just like in many real-world instances, agents tend to hide some information.

There are a number of situations in which governments or similar informed agents commit to a revelation strategy beforehand. For example, hate speech or incitement of violence are restricted in many jurisdictions, and these restrictions are typically specified in laws that are approved in advance and cannot be changed even when the government feels it is advantageous to release a particular

[2]On the one hand, studies have shown that media bias has played a role in determining election outcomes in Peru and Brazil (Boas, 2005), Mexico (Lawson and McCann, 2005), and Russia (Enikolopov et al., 2011). On the other hand, Dyczok (2006) questions the effectiveness of censorship in supporting Kuchma's government in Ukraine; Kern and Hainmueller (2009) report that the East German government enjoyed greater public support in regions where the population had access to West German television; Goldstein (1989) shows that censorship of anti-government caricatures in nineteenth-century France could increase support for the message they contained.

[3]Eijffinger and Geraats (2006). See also Geraats (2002).

[4]Pérignon and Smith (2010).

[5]Including information on their product quality (Jin and Leslie, 2003), financial status (see a review by Healy and Palepu, 2001), or environmental performance (Patten (2002), Cho and Patten (2007)).

[6]Grossman (1981), Milgrom (1981), Seidmann and Winter (1997), Koessler (2003), and Mathis (2008) show that under fairly general settings, the Receiver will learn every state of the world at an equilibrium. Exceptions to the full disclosure result have largely been due to uncertainty over whether the informed agent has precise information (Shin, 1994), or due to informed agent's preferences being either uncertain (Wolinsky, 2003) or non-monotonic in decision-maker's action (Giovannoni and Seidmann, 2007).

piece of information[7]. Even in countries without rule of law, censorship of politically sensitive news is often regulated by bureaucratic instructions, rather than by decisions that are made every time the news arrive[8]. Similarly, firms can commit to revealing or concealing particular facts (such as information about their product quality, financial status, etc.) by asking for an evaluation by an independent expert (e.g. an auditor assessing a company's financial situation, or a reviewer evaluating a theatre play) who then makes her findings public. Commitment may also arise as a credible equilibrium strategy in a repeated interaction - for example, central banks can credibly commit to a particular level of transparency, since they are concerned about their reputation.

One recent work - namely, Kamenica and Gentzkow (2011), hereinafter KG - is similar in spirit to this paper in that it also posits an informed agent who commits to a disclosure policy beforehand. The crucial difference between this paper and KG is that in KG, the Receiver's preferred action is a generic function of her belief about the state. In this paper, however, the Receiver always wants to take a high action if the news are above her preference parameter. The Sender, however does not know the value of that preference parameter, but only it distribution. Hence, the Sender's expected payoff is increasing in the Receiver's posterior belief about the news. This fact drives the results summarised above and ensures that, unlike in KG, the equilibrium disclosure policy is simple - signals are deterministic, news that are not disclosed precisely are pooled into one set, and that set has a small number of disjoint intervals.

Hence, the idea that the probability that the Receiver takes an action which the Sender likes increases with the posterior belief about the news is one of the key ideas behind this paper. This framework is relevant to a large range of settings in which an informed party is dealing with a population of decision-makers who have different preferences. For example, the number of citizens willing to support the government can be higher the image of the government is better. The number of investors who choose not to sell a currency may be greater if the belief about the fundamentals of the economy is more favourable. More customers will buy a firm's product if they hold a more favourable opinion about its quality. Accordingly, this paper shows that in these situations, the informed party will opt for a simple disclosure strategy.

The rest of the paper is structured as follows. Section 2 presents the model, and discusses the interpretation of its key features. Section 3 analyses the model.

[7]We do not normally think of hate speech or incitement of violence as conveying information. However, we can imagine that some types of e.g. racist incitement may be convincing to the public, while others serve only to discredit the speaker and the racist message. In this sense, an effective racist demagogue is "bad news", while an ineffective speaker is "good news".

[8]A study of press censorship in 19th century Europe by Goldstein (2000) mentions a large number of censorship laws and bureaucratic circulars issued to newspapers by various governments. Kris (1941) describes a twenty-page set of instructions, given to Czechoslovak newspaper editors by the Nazi German occupation authorities in 1939, explaining which kinds of news stories would be allowed to be published in future. In either case, the authorities had to commit to a specific set of instructions, rather than examining every article that the newspapers wanted to publish - probably because the latter approach would be too time-consuming.

It first progressively narrows down the set of possible equilibrium strategies of the Sender, showing that the optimal news management policy comes from a relatively small set of simple strategies. It then derives optimal news management strategies for some kinds of the Receiver's preference distributions. Finally, Section 4 concludes.

2 Model

There are two players: Sender (he) and Receiver (she). State of the world is a pair $(\tau, \omega) \in [0,1]^2$; ω is the news what the Sender chooses whether to disclose, and τ is the Receiver's preference parameter, which is her private information. Higher values of ω indicate "better" news, and higher values of τ suggest that the Receiver is predisposed against the Sender. Nature draws τ from a distribution F and ω from a distribution G; the associated densities are f and g. Assume that τ and ω are independent. Further, assume that f is continuously differentiable, and that g is strictly positive everywhere on $[0,1]$.

The Sender has a set of messages M available to him. The Sender's strategy is a function $h : [0,1] \to \Delta(M)$ that associates every realisation of the state with a probability distribution over the set of messages. I will refer to a probability distribution $p = h(\omega) \in \Delta(M)$ as a *signal*, and I will say that a subset K of the news space $[0,1]$ *induces a signal* p if $h(\omega) = p \in \Delta(M)$ for every $\omega \in K$. A signal $h(\omega) = p \in \Delta(M)$ *includes* a message $m \in M$ if m occurs with a strictly positive probability when the news are ω.

The Receiver chooses an action $c \in \{0,1\}$, where 1 is the action that the Sender prefers[9].

The timing of the game is as follows. First, the Sender commits to a disclosure policy by choosing h, which is announced to the Receiver. Then, Nature draws τ and ω from F and G, respectively; and the Receiver learnsτ. Next, the Receiver learns the element of $\mathcal{P}$ to which ω belongs. The Receiver then chooses $c \in C$. Finally, payoffs are realised in the following way: the Sender's payoff equals c, while the Receiver's payoff is $c(\omega - \tau)$.

Let us pause for a moment to examine the intuition behind the model. In the context of this model, the Sender can be a government deciding how to communicate politically relevant news, a central bank picking its level of transparency, or a firm choosing whether to disclose the quality of its product. The Sender's choice of a signal represents his chosen revelation strategy.

We can think of the Sender's revelation strategy as partitioning the news space into subsets, with each subset inducing a different signal. For example, the Sender can opt for full disclosure, in which the Receiver always learns the exact value of news. This corresponds to partitioning the news space into singletons, so that each element of the partition induces a signal that includes a unique message. On the other extreme, a partition that contains only one set (the entire $[0,1]$ interval) reveals no information at all to the Receiver. A more complicated

[9] The analysis can be easiliy extended to a case when the Receiver's action set is an arbitrary compact subset of $\mathbb{R}$.

partition might consist of sets $S_1 = [0, 0.1] \bigcup [0.9, 1]$, $S_2 = [0.4, 0.6]$, and all singletons that do not belong to these sets, and then choosing a degenerate signal (i.e. a signal that only includes one message) for each element of the partition. Under such a partition, the Receiver can learn that the news is "exceptional" (very far from 0.5); that it is "average" (very close to 0.5); and if it is somewhere else, she learns the news exactly. Additionally, since signals need not be degenerate, the Sender can choose an even more complex revelation strategy.

Thus, the game gives the Sender a very large strategy space. The partition that he chooses can pool the news into a large number of sets; these sets need not be connected - each can consist of a large number of disjoint intervals; and each set can induce a signal that includes a large number of messages.

The Sender always wants to choose h that would encourage the Receiver to pick higher c. The Receiver, on the other hand, is better off with larger c if and only if the news is good enough - namely, when it exceeds a threshold given by her preferences (i.e. when $\omega \geq \tau$). Higher τ thus means that the Receiver is more reluctant to choose an action that the Sender prefers. The distribution F can thus be seen as a distribution of public opinion. For example, when F has a larger mass on its right tail, this implies there is a large number of individuals that would vote against the government unless the news is very good.

Several assumptions are implied when the model is applied to specific situations. First, it is assumed that the Receiver's preferences are independent of the news. This may not hold in the long term - for example, a citizen may become less inclined to support the government if bad news keep coming - but in a one-shot interaction this should hold. Alternatively, we may think of τ as indicating the Receiver's preferences affected by factors other than those captured in ω - for instance, if the news ω relate to how well the government is conducting its foreign policy, then τ can show a citizen's level of support for its economic programme.

Furthermore, we can think of ω as representing the valence of the incumbent candidate relative to that of the challenger. The incumbent can manipulate information to convince voters that his valence is high. On the other hand, τ can represent a voter's political position, which is not affected by news about the candidate's valence[10].

Additionally, the model assumes that the Sender can hide the news or reveal it up to a subset of the news space, but he cannot lie. In some situations, this is straightforward - for instance, a government can either ban hate speech or permit it. In other instances, reputational losses to governments, firms or central banks that are caught lying may be so large that lying is never an optimal strategy. Finally, we may think of revealing information as providing hard evidence of the news (photos, videos, testimony by independent media) - evidence which cannot be falsified. In this setting, lying (or telling the truth with no evidence to support it) leads the Receiver to ignore the Sender's message, and therefore corresponds to a no disclosure case (a partition consisting of one set).

[10] I thank Miguel Ballester for suggesting this interpretation.

3 Analysis and Results

3.1 General Results

Suppose the Receiver knows that the Sender has chosen a disclosure policy h. Then if the Receiver gets a message m, her expected payoff from taking an action c, given her preference parameter τ, equals $\mathrm{E}\left[c(\omega - \tau) \mid m\right] = c\mathrm{E}\left[\omega \mid m\right] - c\tau$. If $\tau < \mathrm{E}\left[\omega \mid m\right]$, this expression is maximised at $c = 1$, while if $\tau > \mathrm{E}\left[\omega \mid m\right]$, it is maximised at $c = 0$. This describes the Receiver's best response.

Thus, given any realisation of the signal, the Sender's expected payoff equals the probability that the Receiver's preference parameter τ is below the expeced value of the news given that realisation - in other words, it equals $F\left(\mathrm{E}\left[\omega \mid m\right]\right)$.

Now we can narrow down the set of possible optimal strategies of the Sender. To proceed further, I will make an assumption about the possible signals that are available to the Sender.

Assumption 1. The Sender's strategy h includes a finite number of non-degenerate signals, and each signal includes a finite number of messages.

With this assumption in mind, we can obtain the following result about the Sender's strategy.

Proposition 1. Suppose that the Sender has chosen a disclosure strategy h. Then there exists a a strategy $\hat{h}$ consisting only of degenerate signals that gives the Sender the same equilibrium payoff.

Proof. See Appendix

Proposition 1 says that the Sender's payoffs under any disclosure strategy of the Sender can be obtained under some strategy that only contains degenerate signals. Thus, we can restrict the Sender's strategy space to strategies in which the Sender assigns a message to every element of the news space with probability one.

This result makes it possible to represent any strategy of the Sender by a partition $\mathcal{P}$, such that all news that are associated with the same message belong to the same element of $\mathcal{P}$, and news that are associated with different messages belong to different elements of $\mathcal{P}$.

Assumption 1. The number of non-singleton sets in $\mathcal{P}$ is bounded by some finite number N.

This can be seen as a restriction on the language that is available to the Sender. Denote by $\mathfrak{P}$ the set of all partitions of the $[0,1]$ interval that meet this assumption.

Given that $\omega \in S \in \mathcal{P}$, the Sender's expected payoff equals the probability that $\tau < \mathrm{E}\left[\omega \mid \omega \in S\right]$, which is $F\left(\mathrm{E}\left[\omega \mid \omega \in S\right]\right)$. Denote by μ_S the probability that ω falls in S (i.e. the measure of S associated with g). Then,

$\mu_S \equiv \int\limits_{\omega \in S} g(\omega)\, d\omega$. Then the Sender's overall expected payoff from choosing $\mathcal{P}$ equals

$$v(\mathcal{P}) = \int\limits_{S \in \mathcal{P}} F(\mathrm{E}[\omega \mid \omega \in S])\, d\mu_S$$

which the Sender maximises by choosing $\mathcal{P} \in \mathfrak{P}$. To proceed with the analysis, we must first check whether the maximum of $v(\mathcal{P})$ is well-defined on $\mathfrak{P}$. This is not a trivial problem, since $\mathfrak{P}$, the set of partitions with a bounded number of non-singleton elements, is an infinite set. Fortunately, it is possible to show that $\mathfrak{P}$ can be represented by a compact set, hence ensuring the existence of a maximum of $v(\mathcal{P})$. This is captured in the following proposition, the proof of which is given in the Appendix:

Proposition 2. $\max\limits_{\mathcal{P} \in \mathfrak{P}} \{v(\mathcal{P})\}$ *exists.*

This proposition ensures the existence of a pure-strategy equilibrium.

Recall that by assumption, the number of non-singleton sets is finite. From now on, we can restrict our attention only to those partitions for which all non-singleton sets have positive measure. This is without loss of generality - if a partition includes a finite number of zero-measure sets, their overall measure is zero as well, and hence they can be split into singletons without a change in the Sender's payoff.

By similar reasoning, we can make a further assumption that all positive-measure elements of $\mathcal{P}$ are collections of intervals. This is because points that are not attached to intervals[11] can be converted into singletons without a change in payoff.

Note that if a set $S \in \mathcal{P}$ is a singleton $\{\omega\}$, then $d\mu_S = g(\omega)$, and $\mathrm{E}[\omega \mid \omega \in S] = \omega$. If S is not a singleton, i.e. if $\mu_S > 0$, let $t_S \equiv \mathrm{E}[\omega \mid \omega \in S] = \frac{1}{\mu_S} \int\limits_{\omega \in S} \omega g(\omega)\, d\omega$. The expression for the Sender's expected payoff then becomes:

$$v(\mathcal{P}) = \sum_{S \in \mathcal{P}\,:\,\mu_s > 0} F(t_S)\, \mu_S + \int\limits_{\omega \in S \in \mathcal{P}\,:\,\mu_S = 0} F(\omega) g(\omega) d\omega$$

Clearly, the strategy that maximises it depends on the shapes of f and g. In order to characterise it, define for every positive-measure set S a function $z_S(\omega) \equiv \int_{t_S}^{\omega} [f(t_S) - f(x)]\, dx$. Now consider a partition $\mathcal{P}$ that is a candidate for an optimal partition. We can check for several kinds of deviations from $\mathcal{P}$. First, we can take some news ω that are pooled into a positive-measure set A and disclose them, i.e. turn them into singleton elements of the partition. We can also remove them from A and pool them with some other set B instead. Finally, we can also take some other news that under $\mathcal{P}$ are not pooled into any positive-measure set (i.e. that forms a singleton element of $\mathcal{P}$), and merge them

[11] I.e. all $\omega \in S$ for which there exists a neighborhood T that contains no other elements of S besides ω.

with some $A \in \mathcal{P}$. If $\mathcal{P}$ is optimal, none of these deviations can be beneficial. This is captured in the following necessary condition for an optimum:

Proposition 3. *Suppose that $\mathcal{P}$ maximises $v(\cdot)$. Then the following must hold for every positive-measure set $A \in \mathcal{P}$:*

1. $z_A(\omega) \geq 0$ *for any* $\omega \in A$
2. $z_A(\omega) \geq z_B(\omega)$ *for any* $\omega \in A$ *and any positive-measure set* $B \in \mathcal{P}$
3. $z_A(\omega) \leq 0$ *for any* ω *such that* $\{\omega\}$ *forms a singleton element of* $\mathcal{P}$

The proof of this proposition is in the Appendix, but the reasoning behind it is a follows. Suppose that for some positive-measure $A \in \mathcal{P}$ we take a set $M \subseteq A$ (e.g. an interval) and remove it from A, converting it into singletons instead. At the margin, when the measure of M is close to zero - i.e. when M is close to being a single point ω - the change in the Sender's payoff comes in two ways. First, ω is no longer pooled with A, so the contribution of the news ω to the Sender's payoff is now $F(\omega)$ instead of $F(t_A)$. This change in payoff equals $F(\omega) - F(t_A) = \int_{t_A}^{\omega} f(x)\,dx$. Second, t_A - the expected news over A - change because ω is removed from A, and this change is represented by $(t_A - \omega) f(t_A) = -\int_{t_A}^{\omega} f(t_A)\,dx$. The sum of these two effects equals $\int_{t_A}^{\omega} [f(x) - f(t_S)]\,dx = -z_A(\omega)$, which must be negative if the deviation is not profitable. Since we can select M to be anywhere within A, this must hold for any $\omega \in A$. Similarly, any $\omega \notin A$ can be attached (together with its neighbourhood M) to A; this would give rise to the opposite effects on payoff, and by similar reasoning, $z_A(\omega)$ must be negative. Finally, removing ω from A and attaching it to some other set $B \in \mathcal{P}$ creates a combination of these effects for A and B; hence, $z_A(\omega) \geq z_B(\omega)$ must hold for $\mathcal{P}$ to be optimal.

On a technical note, observe that any positive-measure set $S \in \mathcal{P}$ can be split into sets S_1 and S_2 such that $t_S = t_{S_1} = t_{S_2}$; the resulting partition will yield the same payoff to the Sender as $\mathcal{P}$.[12] This means that the Sender has, strictly speaking, infinitely many equilibrium strategies that all yield the same expected payoff. To simplify the analysis, we can assume that, when choosing between such strategies, the Sender will always choose a partition with the smaller number of positive-measure sets - perhaps because, all other things being equal, he has a preference for "simpler" strategies. Thus, all sets $S \in \mathcal{P}$ for which the expected news t_S are the same will be combined into one set.

Using the condition in Proposition 3, and the simplification described above, we can substantially narrow down the set of possible equilibrium strategies, in the following way:

[12] To see that this is the case, denote by $\tilde{P}$ the partition that is similar to $\mathcal{P}$ except that S is split into S_1 and S_2 such that $t_S = t_{S_1} = t_{S_2}$. Note that $\mu_S = \mu_{S_1} + \mu_{S_2}$. Then,

$$\begin{aligned} v\left(\tilde{P}\right) - v(\mathcal{P}) &= F(t_S)\mu_S - F(t_{S_1})\mu_{S_1} - F(t_{S_2})\mu_{S_2} = \\ &= F(t_S)(\mu_{S_1} + \mu_{S_2}) - F(t_{S_1})\mu_{S_1} - F(t_{S_2})\mu_{S_2} = 0 \end{aligned}$$

Thus, $\tilde{P}$ and $\mathcal{P}$ give the Sender the same expected payoff.

Proposition 4. *For any g, and for any f that has no horizontal sections, every equilibrium partition $\mathcal{P}$ contains at most one positive-measure set.*

Proof: see Appendix.

Proposition 4 ensures that at the optimum, almost all shapes of f induce a disclosure strategy under which all the news that do not belong to singleton elements of $\mathcal{P}$ are pooled into at most one set. In other words, some news are revealed precisely, while others are not revealed at all (i.e. when they occur, the only thing the Receiver learns is that they are in a set of news that the Sender prefers to pool together). Thus, the Sender's equilibrium strategy is the strategy of *censorship* - some news are disclosed, while others are hidden. The Sender will never prefer a strategy under which the news are revealed imprecisely (i.e. pooled into several positive-measure subsets of the news space), even though such a strategy is possible in this setting.

Let us denote by S the positive-measure set that is a part of $\mathcal{P}$ at the equilibrium. By eliminating all partitions with more than one positive-measure set, Proposition 5 reduces the problem of determining the optimal partition to finding the optimal set S. This simplifies the problem substantially; but nevertheless, there are still many potential shapes of S. In particular, S can be considered more or less complex depending on the number of disjoint intervals it includes. The following proposition puts a restriction on the complexity of S:

Proposition 5. *If f has $m < \infty$ local weak maxima, then at the equilibrium, S includes no more than m disjoint intervals.*

Proof: see Appendix.

This proposition underscores the importance of f, the distribution of the Receiver's preference parameter, in determining the optimal revelation strategy. It shows that in most cases, we should not expect a very complicated disclosure policy. Optimal disclosure strategy will only be "complex" - i.e. include a set of censored news consisting of many disjoint intervals - when the distribution f of the Receiver's preference parameter is "complex" as well (i.e. has many peaks). For distributions with a small number of peaks, this greatly reduces the space of possible optimal strategies. This result suggest that for the most part, we are unlikely to see very complex disclosure and censorship policies in real-life situations, as long as the Sender is optimising.

3.2 Optimal Disclosure Policies

With these results in mind, we can look at the actual equilibrium disclosure policies for some specific distributions of the Receiver's preferences. We can start by looking at two specific cases - the case of full disclosure, and the case when all news are hidden.

Proposition 6.1. *Full disclosure is an equilibrium strategy if and only if f is weakly increasing. Furthermore, full disclosure is the unique equilibrium strategy if and only if f is strictly increasing.*

Proposition 6.2. *Pooling all the news is the unique equilibrium strategy if f is strictly decreasing.*

Proofs: see Appendix.

Recall that τ is a threshold above which the news are good enough for the Receiver to take action $c = 1$, which the Sender prefers. Increasing f implies that the Receiver's τ is more likely to be high - i.e. she is predisposed against the Sender. Similarly, decreasing τ means that the population includes a large number of Receivers with a low threshold for supporting the Sender. Thus, full revelation of information is an equilibrium strategy for a government that faces a skeptical public, while total censorship is an optimal choice if the public is largely willing to support the government.

This result may seem somewhat paradoxical, but it has an intuitive explanation. If the news is disclosed, the Receiver chooses $c = 1$ when the news is above her threshold τ. If the news is censored, the Receiver selects $c = 1$ when t_S, the expected value of news over the set of censored news, is above τ. When the distribution of τ is increasing, τ is likely to be high, and hence there is a high probability that, if news are pooled over some set S, their expected value over S ends up below τ. On the other hand, if the news is not censored, there is some probability that it ends up above any given threshold - thus giving the Sender a higher chance of getting $c = 1$. Similarly, if f is decreasing, t_S is more likely to be above τ, and censorship becomes a safe way for the Sender to secure the Receiver's support.

We can also see that full revelation is a rather special case. More typically, some information will be censored.

We can now look at more general cases. Two types of preference distributions are particularly interesting - the one in which f is unimodal, and the one in which it is U-shaped. The former case corresponds, for example, to a society in which most individuals tend to be moderate with respect to their willingness to support the government. The latter case describes a polarised society with a large number of die-hard supporters and opponents of the government.

Proposition 7.1. *If f strictly increasing on $(0, k)$ and strictly decreasing on $(k, 1)$ for some $k \in (0, 1)$, then there is a unique equilibrium strategy $S = [a, 1]$, such that $0 \leq a < k$, and $t_S > k$.*

Proposition 7.2. *If f strictly decreasing on $(0, k)$ and strictly increasing on $(k, 1)$ for some $k \in (0, 1)$, then there is a unique equilibrium strategy $S = [0, b]$, such that $k < b \leq 1$, and $t_S < k$.*

Proofs: see Appendix

Hence, the unimodal and the U-shaped distributions induce disclosure policies under which either the best news, or the worst news are not disclosed. It is easy to see that full disclosure and full pooling emerge as special cases of either the unimodal f or the U-shaped f when k equals 1 and 0.

Now take the case when f is unimodal with a peak at k and thus $S = [a, 1]$. Suppose that the density of Receiver's preferences changes to another $\hat{f}$, which is also unimodal and induces a censorship strategy $\hat{S} = [\hat{a}, 1]$, with an expected value $t_{\hat{S}}$. But suppose that $\hat{f}$ has a peak $\hat{k}$ further to the right - far enough that $\hat{k} > t_S$. From Proposition 6.1 it follows that $t_{\hat{S}} > \hat{k}$. Thus, $t_{\hat{S}} > t_S$, i.e. the expected value of the news over $\hat{S}$ is greater than over S - which means that fewer news are being censored under $\hat{f}$ than under f. Similarly, if f is replaced by a unimodal $\tilde{f}$ with a peak $\tilde{k} < a$, and a new set of censored news $\tilde{S} = [\tilde{a}, 1]$, then we will have $\tilde{a} < \tilde{k} < a$ - meaning that more news will be censored under $\tilde{f}$.

Thus, a change that preserves the general shape of the density of preferences but moves its peak far enough will change the extent to which news are censored at equilibrium. A move to a density with a peak further to the right - i.e. to a distribution corresponding to a more skeptical audience - will reduce the amount of news that are being hidden, and vice versa. It can be easily shown that similar effects will happen when the original density is U-shaped. In this case, a change to a density with k further on the left (and thus a greater mass of τ on the right, implying a more skeptical audience) will result in a smaller set of censored news.

The idea that more extensive revelation is optimal if the Receiver is predisposed against the Sender has several implications. On the normative side, consider the often raised question[13] of whether hate speech should be restricted. Suppose that the government is trying to minimise the number of citizens who choose to adopt a racist mentality, and that the citizens' choice depends on how eloquent and persuasive the message is. The analysis above suggests that a restriction on hate speech will be effective if citizens are ex ante inclined to reject the hate message. If they are predisposed to believe it, restrictions will be counterproductive - seeing that the racist message is suppressed, citizens are likely to assume that it is sufficiently convincing for them.

On a similar note, a central bank that tries to prevent investors with selling the currency should opt for less transparency if the investors have a higher ex ante trust in the currency stability. This is because if investors are enthusiastic, the posterior belief that they form upon receiving no news is still sufficiently good for them to keep the currency. On the other hand, a central bank facing reluctant investors is better off with revealing more information, since, seeing that the news are not revealed, they will conclude that they must be bad enough for them to start selling.

On the positive side, these results indicate a crucial link between information disclosure and the distribution of public opinion. Much of the previous research[14] has looked at censorship as a determinant of public opinion. In contrast, the results above suggest that a reverse link may also be in place - the optimal (from a government's point of view) degree of censorship depends on the distribution of views within the population.

In particular, we can expect more political censorship in countries where the public is less inclined to oppose the regime. On the other hand, societies in which citizens have less trust for the government are likely to enjoy greater

[13] For example, by Herz and Molnar, eds (2012).

[14] See references in the Introduction.

media freedom, even when the regime is authoritarian. When the population is skeptical but the government nevertheless insists (for some exogenous reasons) on censoring the news, such censorship may be counterproductive for the regime, and may undermine its credibility - which is probably what happened following the Moscow Metro accident described in the beginning of the paper.

As noted above, an increase in the share of citizens who are opposed to the government shifts the optimal disclosure strategy towards greater openness, and vice versa. Thus, a change in the public's attitude towards the government may be a reason for changes in media freedom. For example, suppose that a regime experiences a suddent boost in its popularity - due to, for instance, better than expected economic performance, or due to an outbreak of war causing the population to rally behind the government. How would the government respond? Traditionaly[15], it has been thought that in the former case, economic development would lead to political liberalisation. On the other hand, some recent observers - e.g. Bueno de Mesquita and Downs (2005) discussing China - have questioned this conclusion. This paper suggests one possible channel through which strong economic performance may lead to *less* political freedom - by making citizens less inclined to oppose the government when no information is released, economic growth might lead an optimising ruler to tighten censorship.

On the other hand, when facing a decline in popular support, a rational authoritarian ruler may choose to increase media freedom. Hence, a rise in popular opposition to a regime may coincide in time with political liberalisation. This fact has been noted long ago[16], and was usually interpreted as suggesting that liberalisation makes an authoritarian government more vulnerable. The analysis above suggests that an alternative mechanism may be in place: political liberalisation, at least in the area of media freedom, may be a rational reaction to a loss of public trust in the government.

4 Conclusions

This paper has examined optimal information disclosure by a Sender, such as a government, who commits to a particular disclosure strategy to induce a Receiver to take a favourable action.

In general, the Sender's strategy space can be quite large - subject to a few technical restrictions, the Sender can send different messages for different subsets of the news space, thus pooling news in many possible ways. These subsets need not be connected. Furthermore, the messages need not be sent in a deterministic way - Sender can randomise over messages. Nevertheless, the set of optimal strategies was found to be quite small - and, in general, optimal strategies are simple. First, the Sender does not need to mix over messages -

[15] See a highly influential work by Lipset (1959).

[16] Writing about the French Revolution, Alexis de Tocqueville has famously noted that "the regime which is destroyed by a revolution is almost always an improvement on its immediate predecessor, and experience teaches that the most critical moment for bad governments is the one which witnesses their first steps toward reform" (Goldhammer and Elster, 2011).

at the optimum he can send a specific message with probability one for every value of the news. Second, the Sender will not want to pool the news into many different sets. Instead, he will disclose some news precisely, and hide other news completely by pooling them into one set. It was thus established that censoring news is the equilibrium strategy. Other possible strategies - such as revealing news imperfectly (up to some subsets of the news space) - are not optimal for the Sender. Third, the set of news that are pooled together is likely to be simple as well. In fact, unless the Receiver's preference distribution is very complex, the set of news that are suppressed will consist of only a small number of disjoint intervals.

In practice, this means that, for example, a government that has the ability to restrict the flow of news will choose the news that will be suppressed, and will allow the media to disclose the rest. It will not try to force the media to send random messages. Nor will it require that the media disclose the news approximately (for example, by vaguely distinguishing between good and bad news without giving further detail). Even though these strategies are theoretically possible, they will not make the government better off than a simple censorship strategy.

Optimal revelation strategy will largely depend on the distribution of the audience's preferences. Much of the previous research on censorship has analysed its effect on the views of the public. This paper suggests an alternative link between political communication and public opinion: the distribution of political views in the society affects the optimal disclosure policy.

In general, more information is likely to be revealed if the Receiver is predisposed against the Sender. For instance, we can expect less political censorship when the distribution of views in the population is skewed towards opposing the government. Alternatively, when censorship is present under such circumstances, we can expect it to hurt the Sender. Thus, policies such as banning racist hate speech are effective when the public opinion is skewed towards rejecting racism; when the public is largely sympathetic to the racist message, a ban can do more harm than good. Similarly, a central bank should opt for more transparency when facing skeptical public.

By affecting optimal censorship, a shift in the distribution of public opinion can become a cause of political change. Thus, we can expect an optimising government to restrict freedom of information as a reaction to an increase in public support. In contrast, a loss of popularity can lead the government to relax censorship.

References

Boas, Taylor C., "Television and Neopopulism in Latin America: Media Effects in Brazil and Peru," *Latin American Research Review*, 2005, *40* (2), 27–49.

Cho, Charles H. and Dennis M. Patten, "The role of environmental disclosures as tools of legitimacy: A research note," *Accounting, Organizations and Society*, 2007, *32* (7-8), 639–647.

de Mesquita, B. Bueno and G.W. Downs, "Development and democracy," *Foreign Aff.*, 2005, *84*, 77.

Dyczok, Marta, "Was Kuchma's censorship effective? mass media in Ukraine before 2004," *Europe-Asia Studies*, 2006, *58* (2), 215–238.

Eijffinger, Sylvester C.W. and Petra M. Geraats, "How transparent are central banks?," *European Journal of Political Economy*, March 2006, *22* (1), 1–21.

Enikolopov, Ruben, Maria Petrova, and Ekaterina Zhuravskaya, "Media and Political Persuasion: Evidence from Russia," *American Economic Review*, December 2011, *101* (7), 3253–85.

Geraats, Petra M., "Central Bank Transparency," *The Economic Journal*, 2002, *112* (483), F532–F565.

Giovannoni, Francesco and Daniel J. Seidmann, "Secrecy, two-sided bias and the value of evidence," *Games and Economic Behavior*, May 2007, *59* (2), 296–315.

Goldhammer, A. and J. Elster, *Tocqueville: The Ancien Régime and the French Revolution*, Cambridge University Press, 2011.

Goldstein, Robert Justin, "The Debate over Censorship of Caricature in Nineteenth-Century France," *Art Journal*, 1989, *48* (1), pp. 9–15.

_ , *The war for the public mind: political censorship in nineteenth-century Europe*, Praeger Publishers, 2000.

Grossman, Sanford J, "The Informational Role of Warranties and Private Disclosure about Product Quality," *Journal of Law & Economics*, December 1981, *24* (3), 461–83.

Healy, Paul M. and Krishna G. Palepu, "Information asymmetry, corporate disclosure, and the capital markets: A review of the empirical disclosure literature," *Journal of Accounting and Economics*, September 2001, *31* (1-3), 405–440.

Herz, Michael and Peter Molnar, eds, *The Content and Context of Hate Speech: Rethinking Regulation and Responses*, Cambridge University Press, 2012.

Jin, Ginger Zhe and Phillip Leslie, "The Effect Of Information On Product Quality: Evidence From Restaurant Hygiene Grade Cards," *The Quarterly Journal of Economics*, May 2003, *118* (2), 409–451.

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

Kern, Holger Lutz and Jens Hainmueller, "Opium for the Masses: How Foreign Media Can Stabilize Authoritarian Regimes," *Political Analysis*, 2009, *17* (4), 377–399.

Koessler, Frederic, "Persuasion games with higher-order uncertainty," *Journal of Economic Theory*, June 2003, *110* (2), 393–399.

Kris, Ernst, "German Censorship Instructions for the Czech Press," *Social Research*, 1941, *8* (2), pp. 238–246.

Lawson, Chappell and James A. McCann, "Television News, Mexico's 2000 Elections and Media Effects in Emerging Democracies," *British Journal of Political Science*, 2005, *35* (01), 1–30.

Lipset, Seymour Martin, "Some Social Requisites of Democracy: Economic Development and Political Legitimacy," *The American Political Science Review*, 1959, *53* (1), pp. 69–105.

Mathis, Jérôme, "Full revelation of information in Sender-Receiver games of persuasion," *Journal of Economic Theory*, November 2008, *143* (1), 571–584.

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

Patten, Dennis M., "The relation between environmental performance and environmental disclosure: a research note," *Accounting, Organizations and Society*, November 2002, *27* (8), 763–773.

Pérignon, Christophe and Daniel R. Smith, "The level and quality of Value-at-Risk disclosure by commercial banks," *Journal of Banking & Finance*, February 2010, *34* (2), 362–377.

Seidmann, Daniel J. and Eyal Winter, "Strategic Information Transmission with Verifiable Messages," *Econometrica*, January 1997, *65* (1), 163–170.

Shin, Hyun Song, "The Burden of Proof in a Game of Persuasion," *Journal of Economic Theory*, October 1994, *64* (1), 253–264.

Wolinsky, Asher, "Information transmission when the sender's preferences are uncertain," *Games and Economic Behavior*, February 2003, *42* (2), 319–326.

5 Appendix

5.1 Proof of Proposition 1

Under Assumption 1, the number of non-degenerate signals induced by h is finite. Take all non-degenerate signals induced by zero-measure subsets of the news space $[0,1]$. We can replace each of these signals by some degenerate signal - since their number is finite, their total measure is zero, and hence this operation does not change the Sender's expected payoff.

Now take a non-degenerate signal p that is induced by a positive-measure set of news $K \subseteq [0,1]$. Suppose p includes n messages $m_1, ..., m_n$ that occur with probabilities $p_1, ..., p_n$. Let $\mathrm{E}[\omega \mid m_i]$ be the expected value of ω given that a message m_i is received. Then if message m_i is sent, the Sender receives an expected payoff of $F(\mathrm{E}[\omega \mid m_i])$. Hence, if news happen to be in K, the Sender's expected payoff is $\sum\limits_{i=1}^{n} p_i F(\mathrm{E}[\omega \mid m_i])$. Now let us modify h in the following way. Divide K into n subsets $K_1, .., K_n$ such that $\mu(K_i) = \frac{p_i}{\mu(K)}$ and $\mathrm{E}[\omega \mid \omega \in K_i] = \mathrm{E}[\omega \mid \omega \in K]$ for all $i \in \{1, ..., n\}$ - this can always be done, since K is a positive-measure set. Let each K_i induce a degenerate signal that includes message m_i only. Then the Sender's expected payoff if the news happen to be in K is the same as before. We can do this for every positive-measure subset that induces a non-degenerate signal. Once this is done, call the new strategy $\hat{h}$. It will give the Sender the same payoff as h.

5.2 Proof of Proposition 2

Consider a partition $\mathcal{P} \in \mathfrak{P}$. We know that the number of non-singleton sets included in $\mathcal{P}$ is at most some finite number N. We can assign a number $\{1, ..., N\}$ to every non-singleton set that is part of $\mathcal{P}$, and also assign zero to the union of all singletons in $\mathcal{P}$. Now every point in the $[0,1]$ interval can be assigned a number $\{0, 1, ..., N\}$, depending on the set $S \in \mathcal{P}$ to which it belongs under our chosen partition (with zero indicating that this point is a singleton element of the partition). Listing these numbers for all points in the unit interval (i.e. a mapping from $[0,1]$to $\{0, 1, ..., N\}$) can describe any partition $\mathcal{P} \in \mathfrak{P}$.

The set $\mathfrak{P}$ can therefore be fully described by a collection of such listings. This collection is a set $\{0, 1, ..., N\}^{[0,1]}$. We know that the set $\{0, 1, ..., N\}$ is compact, and Tikhonov theorem states that a Cartesian product of compact sets is compact as well. Therefore, the set $\{0, 1, ..., N\}^{[0,1]}$, and by extension $\mathfrak{P}$, is compact. Hence, $v(\cdot)$ - being a continuous function with compact support $\mathfrak{P}$ - must, by Weierstrass theorem, have a maximum.

5.3 Proof of Proposition 3

To prove (1), take a partition $\mathcal{P}$ containing a positive-measure set A. Now take a w belonging to the interior of A and suppose that $z_A(w) < 0$.[17] We want to prove that $\mathcal{P}$ is not optimal.

Consider a deviation from $\mathcal{P}$ to a partition $\hat{\mathcal{P}}$ that differs from $\mathcal{P}$ in that an interval $[w, r]$ is removed from A and instead all the news in $[w, r]$ are disclosed (i.e. turned into singleton elements of the partition). If $r = w$, then $v\left(\hat{\mathcal{P}}\right) = v(\mathcal{P})$. Recall that $v(\mathcal{P}) = \sum\limits_{S \in \mathcal{P}\,:\,\mu_s > 0} F(t_S)\,\mu_S + \int\limits_{\omega \in S \in \mathcal{P}\,:\,\mu_S = 0} F(\omega)g(\omega)d\omega$. Then

$$v\left(\hat{\mathcal{P}}\right) - v(\mathcal{P}) = F\left(t_{A\setminus[w,r]}\right)\mu_{A\setminus[w,r]} + \int_w^r F(\omega)g(\omega)d\omega - F(t_A)\,\mu_A$$

Note that that

$$\mu_{A\setminus[w,r]} = \int\limits_{\omega \in A} g(\omega)\,d\omega - \int_w^r g(\omega)\,d\omega$$

and

$$t_{A\setminus[w,r]} = \frac{\int\limits_{\omega \in A} \omega g(\omega)\,d\omega - \int_w^r \omega g(\omega)\,d\omega}{\int\limits_{\omega \in A} g(\omega)\,d\omega - \int_w^r g(\omega)\,d\omega}$$

Taking the derivative of $v\left(\hat{\mathcal{P}}\right) - v(\mathcal{P})$ with respect to r yields

$$\frac{\partial\left[v\left(\hat{\mathcal{P}}\right) - v(\mathcal{P})\right]}{\partial r} = g(r)\left[f\left(t_{A\setminus[w,r]}\right)\left(t_{A\setminus[w,r]} - r\right) - F\left(t_{A\setminus[w,r]}\right) + F(r)\right] = -g(r)\,z_{A\setminus[w,r]}(r)$$

If $r = w$, then $A \setminus [w, r] = A$, and $v\left(\hat{\mathcal{P}}\right) - v(\mathcal{P}) = 0$. If $\mathcal{P}$ is an optimal strategy, that difference must be weakly decreasing in r at $r = w$. But if $z_A(w) < 0$, then

$$\left.\frac{\partial\left[v\left(\hat{\mathcal{P}}\right) - v(\mathcal{P})\right]}{\partial r}\right|_{r=w} = -g(w)\,z_A(w) > 0$$

as g is assumed to be strictly positive everywhere. Therefore, the Sender benefits from increasing r, which means that $\mathcal{P}$ is not optimal.

Parts (2) and (3) are proved analogously. To prove (2), suppose that for some positive-measure sets $A, B \in \mathcal{P}$, $z_A(w) < z_B(w)$ for some $w \in A$. Consider a

[17] The assumption that w is in the interior of A is without loss of generality, since for every w on the boundary of A such that $z_A(w) < 0$, there must - since $z_A(w)$ is continuous - be another w' in the neighborhood of w for which this property also holds.

deviation from $\mathcal{P}$ to $\hat{\mathcal{P}}$ in which an interval $[w, r]$ is removed from A and pooled with B. Again, $v\left(\hat{\mathcal{P}}\right) = v(\mathcal{P})$ for $r = w$. But then

$$\left.\frac{\partial\left[v\left(\hat{\mathcal{P}}\right) - v(\mathcal{P})\right]}{\partial r}\right|_{r=w} = -g(w)\, z_A(w) + g(w)\, z_B(w) > 0$$

Hence, the Sender benefits from a deviation in which r is increased, and thus $\mathcal{P}$ cannot be an equilibrium strategy.

Finally, to prove (3), assume that $z_A(w) > 0$ for some positive-measure set A and for some w that is not part of any positive-measure set. Now take some interval $[w, r]$ such that every $\omega \in [w, r]$ is a singleton element of the partition, and consider a change from $\mathcal{P}$ to $\hat{\mathcal{P}}$ in which this interval is pooled with A. Then

$$\left.\frac{\partial\left[v\left(\hat{\mathcal{P}}\right) - v(\mathcal{P})\right]}{\partial r}\right|_{r=w} = g(w)\, z_A(w) > 0$$

so again there is a profitable deviation.

5.4 Proof of Proposition 4

Take a distribution f that has no horizontal sections. Suppose that there exists a partition $\mathcal{P}$ which includes two positive-measure sets A and B, and assume without loss of generality that $t_A < t_B$. We can now prove in two steps that $\mathcal{P}$ is not an optimal partition for the Sender, i.e. that there exists a partition which gives him a higher expected payoff.

Step 1. *For $\mathcal{P}$ to be an optimal partition, f must be increasing on $[t_A, t_B]$.*

Suppose that $\mathcal{P}$ is an optimum partition, and consider the following deviation: take $C \subseteq A$ such that $t_C = t_A = t_{A\setminus C}$.[18] Now remove it from A and pool with B; call the resulting partition $\mathcal{P}'$. Then

$$\begin{aligned} v(\mathcal{P}) - v(\mathcal{P}') &= F(t_A)\,\mu_A + F(t_B)\,\mu_B - F\left(t_{A\setminus C}\right)\mu_{A\setminus C} - F\left(t_{B\bigcup C}\right)\mu_{B\bigcup C} = \\ &= F(t_A)\,\mu_A + F(t_B)\,\mu_B - F(t_A)(\mu_A - \mu_C) - F\left(t_{B\bigcup C}\right)(\mu_B + \mu_C) = \\ &= F(t_A)\,\mu_C + F(t_B)\,\mu_B - F\left(\tfrac{\mu_B}{\mu_B+\mu_C} t_B + \tfrac{\mu_C}{\mu_B+\mu_C} t_A\right)(\mu_B + \mu_C) \end{aligned}$$

The expression above must be non-negative for $\mathcal{P}$ to be the optimal partition. Denote $\gamma \equiv \frac{\mu_C}{\mu_B+\mu_C}$, and note that we can choose μ_C to be of any value between 0 and μ_A. Then

$$\gamma F(t_A) + (1-\gamma)\, F(t_B) \geq F(\gamma t_A + (1-\gamma)\, t_B)\ ,\ \forall \gamma \in \left[0, \frac{\mu_A}{\mu_A + \mu_B}\right]$$

Now consider a deviation from $\mathcal{P}$ to $\mathcal{P}''$ of the following form: take $D \subseteq B$ such that $t_D = t_B = t_{B\setminus D}$, remove D from B and pool it with A. Then

[18] For example, if A is an interval, we can choose C to be a middle part of A.

$$v(\mathcal{P}) - v(\mathcal{P}'') = F(t_A)\mu_A + F(t_B)\mu_B - F\left(t_{B\setminus D}\right)\mu_{B\setminus D} - F\left(t_{A\bigcup D}\right)\mu_{A\bigcup D} =$$
$$= F(t_A)\mu_A + F(t_B)\mu_B - F(t_B)(\mu_B - \mu_D) - F\left(t_{A\bigcup D}\right)(\mu_A + \mu_D) =$$
$$= F(t_A)\mu_A + F(t_B)\mu_D - F\left(\frac{\mu_A}{\mu_A+\mu_D}t_A + \frac{\mu_D}{\mu_A+\mu_D}t_D\right)(\mu_A + \mu_D)$$

Denote $\delta \equiv \frac{\mu_A}{\mu_A+\mu_D}$; note that μ_D can be chosen between 0 and μ_B. Then

$$\delta F(t_A) + (1-\delta)F(t_B) \geq F(\delta t_A + (1-\delta)t_B)\ ,\ \forall\delta \in \left[\frac{\mu_A}{\mu_A+\mu_B}, 1\right]$$

This inequality and the previous one, taken together, imply that F must be convex on $[t_A, t_B]$, and thus f must be increasing on that interval. Since we have assumed that f has no horizontal sections, it must be strictly increasing in the neighbourhood of t_A.

Step 2. *If f is strictly increasing in the neighbourhood of t_A, then $\mathcal{P}$ cannot be optimal.*

If f must be increasing on (t_A, t_B), then $z_A(\omega) = \int_{t_A}^{\omega} f(t_A) - f(x)\,dx < 0$ for every $\omega \in (t_A, t_B)$. If $\mathcal{P}$ is an optimal partition, then by Proposition 2, no news in the interval (t_A, t_B) belong to A.

On the other had, increasing f implies that $f(t_B) > f(t_A)$. Then for every $\omega \geq t_B$, $z_A(\omega) = \int_{t_A}^{\omega} f(t_A) - f(x)\,dx < \int_{t_B}^{\omega} f(t_A) - f(x)\,dx < \int_{t_B}^{\omega} f(t_B) - f(x)\,dx = z_B(\omega)$. Hence, if $\mathcal{P}$ is an optimal partition, then by Proposition 2, no $\omega \geq t_B$ can belong to A.

Therefore, optimal $\mathcal{P}$ implies that $t_A \geq \max\{A\}$. But this is not possible, since t_A is the expected value of ω over A.

To summarise, we have assumed that $\mathcal{P}$ contains two positive-measure sets, A and B. If f is not increasing between t_A and t_B, then $\mathcal{P}$ cannot be optimal. If f is increasing on that interval, $\mathcal{P}$ cannot be optimal either. Thus, a partition with two positive-measure sets cannot from part of an equilibrium. But Proposition 1 establishes that an equilibrium must exist, So we can conclude that an equilibrium partition has at most one positive-measure set.

5.5 Proof of Proposition 5

We have earlier established that without loss of generality S can be thought of as a countable union of disjoint intervals. Thus, we can write $S = \bigcup\limits_{i\in I}[a_i, b_i]$ such that $0 \leq a_i \leq b_i \leq a_{i+1} \leq 1$, $\forall i \in I$, where I is countable.

Let us start by taking a set S comprising $|I|$ disjoint intervals. What is the smallest number of local maxima that f needs to have for S to be the Sender's equilibrium strategy?

Observe that since we have assumed the number of weak local maxima to be finite, f cannot be constant at any interval. This means that $z_S(\cdot)$ cannot equal zero on any interval $[p, q] \subseteq [0, 1]$, since if it was zero, this would mean that $f(\omega) = f(t_S)$, $\forall \omega \in [p,q]$, i.e. that f is horizontal. The fact that $z_S(\cdot)$ cannot be zero on an interval, together with Proposition 2, implies that $z_S(\omega)$ is increasing at $\omega = a_i$ and decreasing at $\omega = b_i$, $\forall i \in I$. This means - since $z_S(\omega)$ is continuously differentiable - that $\frac{dz_S}{d\omega}(a_i) > 0$ and $\frac{dz_S}{d\omega}(b_i) < 0$.

Hence, for every $i \in I$, $z_S(\omega)$ must have at least one local maximum $c_i \in (a_i, b_i)$ and at least one local minimum $d_i \in (b_i, a_{i+1})$. At a local maximum, $\frac{dz_S^2}{d\omega^2}(c_i) = -f'(c_i) < 0$, while at a local minimum, $\frac{dz_S^2}{d\omega^2}(d_i) = -f'(d_i) > 0$. But f is assumed to be continuously differentiable, and thus for every $i \in I$ there must be news $w_i \in (c_i, d_i)$ such that f' is positive to the left of w_i and negative to the right of it. This w_i is therefore a local maximum of f, and there must be such a point in every interval (a_i, a_{i+1}). There are $|I|-1$ such intervals, which gives us $|I|-1$ local maxima.

To see that another maximum of f must exist, note that $z_S(t_S) = 0$, and also, $\frac{dz_S}{d\omega}(t_S) = 0$. This gives several possibilities. If $t_S \in (b_i, a_{i+1})$ for some $i \in I$, then it is a local maximum of z_S, and there are not one but at least two local minima in (b_i, a_{i+1}) - one to the left of t_S and one to the right. So there is one more pair of a maximum and a minimum of $z_S(\omega)$, and by the above logic, there must be a local maximum of f in addition to the ones found in the previous paragraph. If $t_S \in (a_i, b_i)$ for some $i \in I$, then it is a local minimum, and there are two local maxima in (a_i, b_i) - again, f must have an extra local maximum. If $t_S = a_i$ for some $i \in I$, then the shape of $z_S(\omega)$ implies that in some neighbourhood of a_i, $\frac{dz_S^2}{d\omega^2}(\omega) = -f'(\omega) < 0$ for $\omega < a_i$ and $\frac{dz_S^2}{d\omega^2}(\omega) = -f'(\omega) > 0$ for $\omega > a_i$. Consequently, f must have an additional local maximum at a_i. Finally, if $t_S = b_i$ for some $i \in I$, then the shape of $z_S(\omega)$ implies that f must have not one but at least two local maxima in (c_i, d_i).

Hence, if S forms part of an equilibrium, f must have at least $|I|-1$ local maxima plus one more. Therefore, $m \geq |I|$.

5.6 Proof of Proposition 6.1

First statement. To prove necessity, suppose that f is not weakly increasing - this implies that is is strictly decreasing on some interval $[p,q]$. Let $\hat{\mathcal{P}} \equiv \left\{[p,q], \{\omega\}_{\omega \in [0,1] \setminus [p,q]}\right\}$ be a partition consisting of the interval $[p,q]$ and singletons. Then the difference in the Sender's payoff from $\hat{\mathcal{P}}$ and from the fully revealing partition $\left\{\{\omega\}_{\omega \in [0,1]}\right\}$ equals:

$$\begin{aligned}
&v\left(\hat{\mathcal{P}}\right) - v\left(\left\{\{\omega\}_{\omega\in[0,1]}\right\}\right) = F\left(t_{[p,q]}\right)\mu_{[p,q]} - \textstyle\int_p^q F(\omega)\, g(\omega)\, d\omega = \\
&= F\left[\mathrm{E}\left(\omega \mid \omega \in [p,q]\right)\right]\Pr\left(\omega \in [p,q]\right) - \mathrm{E}\left(F[\omega] \mid \omega \in [p,q]\right)\Pr\left(\omega \in [p,q]\right) = \\
&= \Pr\left(\omega \in [p,q]\right)\left[F\left[\mathrm{E}\left(\omega \mid \omega \in [p,q]\right)\right] - \mathrm{E}\left(F[\omega] \mid \omega \in [p,q]\right)\right] > 0
\end{aligned}$$

where the last inequality sign follows from Jensen's inequality and the fact that decreasing f implies concave F. Hence, if f is not weakly increasing, full disclosure cannot be optimal.

To prove sufficiency, consider a weakly increasing f. Pick an arbitrary partition $\tilde{\mathcal{P}}$ containing one[19] positive-measure set S. Now consider a partition $\mathcal{P}'$ that differs from $\tilde{\mathcal{P}}$ by having singletons instead of S - i.e. every $\omega \in S$ is a singleton element of $\tilde{\mathcal{P}}$. If $f(\omega) = f(t_S)$ for all $\omega \in S$, then

$$\begin{aligned}
&v\left(\tilde{\mathcal{P}}\right) - v(\mathcal{P}') = F(t_S)\mu_S - \int_{\omega\in S} F(\omega)g(\omega)d\omega = \\
&= \int_{\omega\in S} [F(t_S) - F(\omega)]\, g(\omega)d\omega = \int_{\omega\in S} f(t_S)[t_S - \omega]\, g(\omega)d\omega = \\
&= f(t_S) \int_{\omega\in S} [t_S - \omega]\, g(\omega)d\omega = f(t_S)[t_S\mu_S - t_S\mu_S] = 0
\end{aligned}$$

On the other hand, if $f(w) \neq f(t_S)$ for some $w \in S$, then either $w > t_S$ and $f(w) > f(t_S)$, or $w < t_S$ and $f(w) < f(t_S)$. In either case, $z_S(w) = \int_{t_S}^{w} f(t_S) - f(x)\,dx < 0$. From Proposition 3 it follows that $\tilde{\mathcal{P}}$ containing S cannot be an optimal.

Hence, a partition containing a positive-measure set S can only be optimal if $f(\omega) = f(t_S)$ for all $\omega \in S$. But every strategy that fits this criterion yields the same expected payoff to the Sender as the fully revealing strategy. Thus, full disclosure must be an equilibrium strategy.

Second statement. To prove necessity, suppose, that f is not strictly increasing. Then f is weakly decreasing on some interval $[p,q] \subseteq [0,1]$. Defining $\hat{\mathcal{P}} \equiv \left\{[p,q], \{\omega\}_{\omega\in[0,1]\setminus[p,q]}\right\}$ as above and using the same reasoning, we can prove that $v\left(\hat{\mathcal{P}}\right) - v\left(\left\{\{\omega\}_{\omega\in[0,1]}\right\}\right) \geq 0$, so full disclosure cannot be a unique equilibrium strategy.

To prove sufficiency, note that if f is strictly increasing, then for any partition $\mathcal{P}$ containing a positive-measure set S, we have $\omega > t_S \Leftrightarrow f(\omega) > f(t_S)$. Pick news $w \in S$ such that $w > t_S$ (such w exists as $t_S < \max(S)$), and observe that $z_S(w) = \int_{t_S}^{w} f(t_S) - f(x)\,dx < 0$. Proposition 2 then ensures that $\mathcal{P}$ is not an equilibrium strategy.

5.7 Proof of Proposition 6.2

Take a strictly decreasing f, and consider a partition $\mathcal{P}$. If $\mathcal{P}$ is fully revealing, it cannot be optimal by Proposition 5. Now suppose $\mathcal{P}$ is not fully revealing, i.e. it contains a positive-measure set S. If there exists $\omega \notin S$, then $z_S(\omega) = \int_{t_S}^{\omega} f(t_S) - f(x)\,dx > 0$ - so $\mathcal{P}$ cannot be optimal. But Proposition 1 states that an optimal partition must exist. Therefore, $\mathcal{P} = \{[0,1]\}$ is an equilibrium, and as such it is unique.

[19] Proposition 3 has already established that an optimal partition chosen will have at most one positive-measure set.

5.8 Proof of Proposition 7.1

From Propositions 4 and 5.1 it follows that under a unimodal density f with a peak on $(0, 1)$, the set S will consist of exactly one interval. Therefore, t_S must be in the interior of S. If $t_S \leq k$, this would mean that f is increasing on some neighbourhood of t_S, implying that $z_S(\omega) < 0$ for news in that neighborhood. Since this neighbourhood belongs to S, by Proposition 2 this cannot hold at an equilibrium. Thus, $t_S \in (k, 1]$. Then $z_S(\omega) > 0$ for all $\omega > t_S$. Similarly, $z_S(k) > 0$, and $z_S(\omega) > 0$ for some $\omega < k$. Hence, $S = [a, 1]$ for some $a \in [0, k)$.

5.9 Proof of Proposition 7.2

From Propositions 4 and 5.1 it follows that under a unimodal density f with a peak on $(0, 1)$, the set S will consist of exactly one interval. Therefore, t_S must be in the interior of S. If $t_S \geq k$, this would mean that f is increasing on some neighbourhood of t_S, implying that $z_S(\omega) < 0$ for news in that neighborhood. Since this neighbourhood belongs to S, by Proposition 2 this cannot hold at an equilibrium. Thus, $t_S \in [0, k)$. Then $z_S(\omega) > 0$ for all $\omega < t_S$. Similarly, $z_S(k) > 0$, and $z_S(\omega) > 0$ for some $\omega > k$. Hence, $S = [0, b]$ for some $b \in (k, 1]$.