

Trembling mechanisms

João Correia-da-Silva^{*†}

Abstract. A principal can relax the participation constraints of the agents by designing the (otherwise off the equilibrium path) beliefs following disagreement using a mediation device that, with a small probability, manipulates the messages received from the agents to originate spurious rejections that are correlated with their announcements. The participation constraints can be further relaxed if the mediation device is also used as a joint randomization device that is designed to punish a genuine rejector. Applications to collusion under private information are provided.

Keywords: Learning from disagreement, Beliefs off-path, Mediation, Mechanism design.

JEL Classification Numbers: D82.

1 Introduction

Suppose that you want to establish a binding agreement with a privately informed rival, in an environment in which no commitment is possible regarding subsequent behavior in case of disagreement. Suppose also that you can use a mediator who is able to mimic rejection of the agreement by your rival, so that you never distinguish a genuine rejection from a spurious one. The purpose of this paper is to show how, and to what extent, such a mediator generates additional incentives for your rival to accept the agreement by credibly threatening to induce in you the disagreement beliefs and subsequent behavior that are the most adverse for him.

In the kind of economic relationship that is investigated, one of the parties, the principal, has the power to design the structure of the interaction, while the other party, the agent, can only accept to participate under the rules defined by the principal or reject to participate and take an outside option. In many environments where the agent has private information about his type, the form of interaction that is optimal for the principal consists in making a “take-it-or-leave-it” proposal of a menu of contracts that is incentive compatible and individually rational: each contract in the

I am grateful to Bruno Jullien, Daniel Garrett, Renato Gomes and, especially, Takuro Yamashita for very useful conversations and suggestions. I also thank Françoise Forges and audiences at Toulouse, Porto, and PEJ 2016. The usual disclaimer applies. This research has been funded by the European Commission through the Marie Skłodowska Curie Fellowship H2020-MSCA-IF-2014-657283; and by FEDER (through COMPETE) and FCT in the framework of projects PTDC/IIM-ECO/5294/2012 and PEst-OE/EGE/UI4105/2014.

^{*}Toulouse School of Economics. 21 Allée de Brienne, 31000 Toulouse, France. E-mail: joao.correia@tse-fr.eu.

[†]CEF.UP and Faculdade de Economia, Universidade do Porto.

menu is designed for one possible type of agent, in a way that makes it optimal for the agent to pick the contract designed for his actual type than choose another contract or reject the proposal.

We will focus on situations in which the outside option of the agent involves subsequent strategic interaction with the principal. For example: when a firm proposes a collusive agreement to a rival that has private information about its own efficiency or about demand; when a bidder proposes a pre-auction arrangement to a rival that has private information about the value of the good or procurement contract that will be auctioned; when a contender offers a bribe to an opponent that has private information about his skills; when a plaintiff proposes a settlement to a defendant that has private information about his degree of negligence, or when a defendant proposes a settlement to a plaintiff that has private information about the magnitude of the damages.

In such settings, the value of the outside option for the agent depends on the information that is inferred by the principal from the agent's rejection of the menu of contracts. This information may harm or benefit the agent, depending on whether it encourages or discourages aggressive behavior by the principal. For example, if rejection of a collusive agreement conveys the information that the agent is very efficient, then, in a subsequent interaction where firms simultaneously choose quantities to produce, the principal will expect a high output from the agent and will, therefore, choose a low output. This increases the value of the outside option for the agent, which means that the information content of the rejection tightens the participation constraints.

Unable to precommit to behave very aggressively in the outside game, the principal would like to at least be able to precommit to interpret a rejection in the worst possible way for the agent, that is, in the way that induces the principal's subsequent behavior to be as aggressive as possible. This would decrease the value of the outside option for the agent, and thus relax the participation constraints. However, it is hard to conceive how can the principal precommit to some future beliefs, or, to put it in another way, how can the principal control the information that is released from the rejection of a proposal. Perhaps surprisingly, this kind of commitment or control is shown to be viable if the principal can use a mediator who is able to mimic a rejection by the agent.

Before bringing the mediator to the picture, let us look more closely at the game that results from a given proposal: knowing his type, the agent picks a contract from the menu or rejects to participate; in case of rejection, principal and agent simultaneously choose an action.¹ In a sequential equilibrium in which the proposal is always accepted, we can interpret the strategy of the principal as a punishment threat whose credibility is sustained by the beliefs following rejection that are prescribed by the equilibrium. The fact that rejection never occurs on the equilibrium path implies that, if we adopt sequential equilibrium as the solution concept, any beliefs following rejection are admissible as long as they are common knowledge. This is good for the principal in the sense that there always exists a sequential equilibrium in which beliefs following disagreement are those that minimize the agent's payoff in the outside game. However, an appropriate refinement

¹The consideration of an outside game with multiple stages would obscure the analysis.

may rule out particularly incredible beliefs and the corresponding punishment strategies.²

For example, suppose that the outside game is a quantity-setting duopoly in which the agent has private information about his efficiency, and consider a sequential equilibrium in which the principal believes, following an off-the-equilibrium-path rejection of a menu of collusive agreements, that the agent has the lowest possible efficiency. Unfortunately for the principal, those beliefs will typically be untenable in the sense of not being neologism-proof. If there is any type of agent, other than the least efficient type, that is indifferent between accepting or rejecting the proposal, this type of agent will belong to a credible set of rejectors (set of types that are better off rejecting if the principal believes that this is the set of types that reject). Therefore, it would be unreasonable to believe that the agent has the lowest possible efficiency level. Whenever, as in this example, appropriate refinements rule out the beliefs following rejection that are the worst for the agent, the corresponding punishment loses its credibility, rendering the participation constraints tighter.

Now assume that the principal can offer a menu of contracts to the agent through an incentiveless mediator (a consultant, lawyer, machine or computer software) that is able to mimic a rejection by the agent. After the agent chooses a contract, with a probability that depends on the chosen contract and that is common knowledge, the mediator sends to the principal a message of rejection and the deal is off. Exactly the same message is sent if the agent rejects the menu of contracts, thus the principal will not be able to distinguish a genuine rejection from a spurious rejection. As we will see, the agent would always like the principal to believe that the rejection was genuine. It is crucial for the principal that the agent cannot prove that a rejection was genuine.

With recourse to such a mediator, the principal can design the (otherwise off-path) beliefs that result if the agent rejects the menu of contracts. The only restriction is that beliefs must be common knowledge. Turning disagreement into an event that occurs on the equilibrium path, the mediator allows the principal to bypass all refinements that are based on restrictions of beliefs formed off-path, which could constrain her by ruling out implausible beliefs following disagreement.

If the outside game is a quantity-setting oligopoly, the principal would like to commit to believe that an agent that rejects her proposal has the lowest possible efficiency. Therefore, she should instruct the mediator to mimic rejection (with a small probability) if and only if the agent picks the contract designed for the agent with the lowest possible efficiency. As a result, rejection – by the mediator, through misrepresentation of the agent’s decision – occurs in equilibrium, and thus beliefs following disagreement are no longer formed off-path. A genuine rejection is misinterpreted by the principal as a spurious rejection produced by the mediator through the manipulation of the choice of an agent with the lowest possible efficiency, as this is the only kind of rejection that occurs on-path. Hence, the principal will believe that the agent has the lowest possible efficiency.

The introduction of a mediator perturbs the incentive compatibility constraints whenever the

²For example, the intuitive criterion (Cho and Kreps, 1987), some version of divinity (Banks and Sobel, 1987; Cho and Sobel, 1990), neologism-proofness (Grossman and Perry, 1986; Farrell, 1993), stability (Kohlberg and Mertens, 1986), forward induction (van Damme, 1989), or undefeatedness (Mailath et al., 1993).

choice of contract influences the probability of spurious rejection.³ Fortunately, this perturbation can be made arbitrarily small because any strictly positive probability of spurious rejection, no matter how small, is infinitely greater than the probability of a genuine rejection (which is zero in equilibrium). Hence, any menu of contracts that is strictly incentive compatible without manipulation remains strictly incentive compatible after the introduction of a mediator that generates the desired beliefs. Strictness of incentive compatibility is not a severe requirement in the sense that a menu of contracts that is incentive compatible can be approximated by one that is strictly incentive compatible (as long as there exists a strictly incentive compatible menu of contracts).

The participation constraints of the agent can be further relaxed if the mediator is able to send private signals to the principal and the agent accompanying the (spurious or genuine) announcement of a rejection. The ability to use the mediator as a randomization device, when playing a mixed strategy in the outside game, may benefit the principal even if the agent has no private information. The principal should instruct the mediator to act as a trustworthy randomizer when the rejection is spurious but not when the rejection is genuine. After a genuine rejection, the mediator should recommend to the principal the worst action for the agent among those in the support of the mixed strategy that is played after a spurious rejection. The principal obeys the recommendation because she believes that the rejection has been spurious and that the mediator has drawn the recommendation according to the mixed strategy equilibrium distribution.⁴

For example, suppose that the outside game is a complete information all-pay auction with a cap on bids in which only principal and agent participate.⁵ Under the assumption that the valuation for winning of the agent is greater than that of the principal, and that the cap on bids is greater than half of the principal's valuation, the unique equilibrium of this game is in mixed strategies with mass points at the cap, such that the agent wins the auction with a probability that is greater than 50% and makes an expected payment that is lower than the cap. In such an environment, the principal can commit to the harshest possible punishment by instructing the mediator to draw a recommendation in accordance with the mixed strategy equilibrium if the outside game is reached after a spurious rejection, but recommend a bid equal to the cap with 100% probability whenever the rejection is genuine (the principal obeys the recommendation because she presumes that the rejection has been spurious, and that the mixed strategy equilibrium of the outside game is being played). Facing such a punishment, the agent may either bid the cap and win with 50% probability or bid zero. Such an extreme punishment would not be credible without the mediator.

More generally, the ability of the mediator to send private signals to principal and agent together with the announcement of a rejection allow the principal to use the mediator as an informed joint

³The exception is when the probability of spurious rejection is independent of the agent's choice of contract. In that case, disagreement reveals no information about the agent. This may be a foundation for *passive beliefs*.

⁴A similar mechanism has been used by Kandori (1991) and Mailath et al. (2002) in the context of repeated games with private monitoring and imperfect public monitoring, respectively.

⁵This game was analyzed by Che and Gale (1998b) in their study on political lobbying. See also Che and Gale (1998a), Gaviious et al. (2002), Szech (2015) and Olszewski and Siegel (2016).

randomization device that is designed to punish a genuine rejector (off the equilibrium path). In the presence of such a mediator, the outside game that follows a spurious rejection becomes an extended outside game in which the principal and the agent first receive private signals from the mediator, who has become informed about the type of the agent through the agent's choice of contract, and then choose a possibly mixed action.⁶ The private signals sent by the mediator after a spurious rejection can be assumed, w.l.o.g., to constitute a stochastic profile of recommendations of actions to be chosen in the outside game that is incentive compatible, i.e., that principal and agent have interest in obeying. Such a stochastic action profile can be designated as an Informed Mediator Bayesian Correlated Equilibrium (BCE^I) of the outside game, which is, by definition, a Bayesian Nash Equilibrium (BNE) of the extended outside game.⁷

Keep in mind that, after a spurious rejection, the set of incentive compatible distributions over profiles of recommendations is the set of BCE^I of the outside game for given beliefs (which can be designed by the principal). After a genuine rejection (which does not occur in equilibrium), the principal presumes that the rejection has been spurious and obeys any recommendation from the mediator that she would obey after a spurious rejection. This means that any action that belongs to the support of a BCE^I of the outside game for some beliefs is a credible punishment threat. Furthermore, since the recommendations and the induced beliefs can be concealed from the agent, any mixture over such actions is also a credible punishment threat.

When the principal interacts with multiple agents, she can use the mediator to design the disagreement beliefs of the principal and the acceptors about the types of all agents. For example, suppose that the disagreement beliefs that are worst for a rejector consist in each of the other players believing that everyone else has the lowest possible efficiency. To generate such beliefs (approximately), the principal should instruct the mediator to generate a spurious rejection with a small probability, $\epsilon_0 > 0$, if all agents choose the contract designed for the agent with the lowest possible efficiency (in short, announce the lowest efficiency); and with an even smaller probability, say $\epsilon_1 > 0$, if all agents except one announce the lowest efficiency. Following a rejection, an acceptor that did not announce the lowest efficiency will believe that all the other agents have the lowest efficiency. An acceptor that has announced the lowest efficiency will believe that, with a probability that converges to one as $\epsilon_1 \rightarrow 0$, all other agents have the lowest efficiency.

This means that, when there is more than one privately informed agent, the ability of the mediator to mimic a rejection allows the principal to generate disagreement beliefs that violate the “*no signaling what you don't know*” condition (Fudenberg and Tirole, 1991). Such beliefs would not be consistent in the absence of a mediator.⁸

⁶Such an extended game was the subject of recent research by Bergemann and Morris (2013, 2016).

⁷Bergemann and Morris (2013, 2016) designated this notion of correlated equilibrium in which an omniscient mediator is employed as *Bayes Correlated Equilibrium*. In the taxonomy provided by Forges (1993), where the mediator knows everything that agents know but not more than that, this notion is designated as *Bayesian solution*.

⁸In a sequential equilibrium, beliefs formed off-path must satisfy “*no signaling what you don't know*”. If types are independently distributed, this implies that a deviation by one agent does not influence beliefs about the other agents. See Fudenberg and Tirole (1991).

With multiple agents, as in the single-agent case, the ability of the mediator to send a profile of private signals with a distribution that depends on the input messages received from the agents is very powerful. It allows the mediator to induce principal and acceptors to obey any distribution over action profiles such that any action that is recommended to a player is also recommended to the same player in some BCE^I (each action can belong to the support of a different BCE^I and even BCE^I with different underlying disagreement beliefs).

Trembling mechanisms do not allow the principal to design the beliefs of the rejector. Moreover, since the anticipated behavior of a genuine rejector does not influence the behavior of acceptors (because a rejector is genuine with zero probability), any information that the mediator transmits to a genuine rejector can only increase his payoff, by allowing him to condition his best-response on that information. Hence, a genuine rejector should not receive any relevant signal from the mediator, and thus retain his prior belief and best-respond to the punishment.

Two caveats are worthwhile remarking. The first is that limited commitment is crucial: if the acceptors and the principal remained constrained by the mechanism in case of a rejection by some other agent, the harshest possible penalty could be imposed on the deviator. The design of disagreement beliefs explained above is only relevant in scenarios in which the principal cannot constrain neither himself nor the acceptors to follow the rules of the mechanism unless it is unanimously accepted. The second caveat is that we focus exclusively on equilibria in which agents always accept the principal's proposal. There may exist equilibria without full participation in which agents have higher expected payoffs. In such equilibria, spurious rejections that occur with a small probability would be insufficient to determine disagreement beliefs.

The remainder of the paper is structured as follows: the relation with the literature is described (Section 2); the model is introduced (Section 3); the single-agent case is analyzed (Section 4); the analysis is extended to the multiple-agent case (Section 5); illustrative examples are provided, and possible applications are described (Section 6); and concluding remarks are made (Section 7).

2 Literature review

2.1 Theory: mediated mechanism design

Myerson (1982) established the revelation principle for generalized principal-agent problems, where agents (besides having private information about their own characteristics) choose actions that cannot be contracted upon. This means that a principal can restrict her choice of a coordination mechanism to those that are direct and incentive compatible: each agent is asked to report his information, and receives a private recommendation of an action to choose; the profile of private recommendations is drawn according to a distribution that depends on the reported type profile, which is such that each agent is better off reporting his type truthfully and obeying the recommendation if all other agents report truthfully and obey the recommendations made to them.

To reconcile our setting with the environment of Myerson (1982), we can treat our mediator as the principal and our principal as an additional agent.⁹ The mediator asks each agent to report his type, and recommends to each agent a pure strategy in the subsequent game. Each agent is recommended to choose the contract designed for his type, and, in case another agent rejects the proposal, choose the action in the outside game that minimizes the payoff of the rejector. This relaxes participation constraints as much as possible. The problem is that obedience may be optimal when expecting that the proposal will be accepted by all agents but become suboptimal after a zero-probability rejection is observed (Selten, 1975; Kreps and Wilson, 1982).

Since we are concerned with situations in which the threat of minimaxing a rejector is indeed not credible, our setting should be framed as a multistage game with communication (Myerson, 1986; Forges, 1986).¹⁰ In this framework, the revelation principle still holds, but the fact that the concept of Nash equilibrium is insufficiently restrictive becomes more transparent. Myerson (1986) imposed an additional requirement of sequential rationality in the spirit of Kreps and Wilson (1982).¹¹ Addressing the critiques by Cho and Kreps (1987) and others, who proposed refinements of sequential equilibrium based on restrictions on the formation of beliefs off-path¹², we impose a stronger refinement by requiring that all information sets are reached on-path.

For all information sets to be reached on-path, we need a mediator who is able to generate any signal that agents can receive as a result of a deviation by the other agents. In our setting, this means that even if we allow an agent to transmit a public signal rejecting the proposal, we assume that the mediator can mimic this public signal. It is this signal jamming by the mediator that allows all information sets that result from a genuine rejection to be reached on-path.¹³

Mediation has been shown to be beneficial in mechanism design with imperfect commitment by the principal.¹⁴ Assuming that the principal can commit to a first-stage decision (as a function of the agent's report) but not to a second-stage decision, Bester and Strausz (2007) showed that the principal benefits from using a communication device that, with some probability, manipulates

⁹Limited commitment by the principal is transformed, through mediation, in moral hazard on the part of agent.

¹⁰In our setting, a non-trembling mechanism defines the following multistage game with communication. In the first stage: each agent reports his type and is recommended to choose the contract designed for his type. If all agents choose a contract, the resulting grand contract is enforced; if some agent rejects the proposal, his rejection is announced. In the second stage (which is only relevant in case of a rejection), each agent is recommended an action to carry out in the outside game. A trembling mechanism can be seen as a mediated stochastic mechanism in which rejection by some agent deterministically implies reversion to the outside game, while acceptance by all agents may either lead to the enforcement of the chosen grand contract or to reversion to the outside game.

¹¹Myerson (1986) showed that sequential rationality is satisfied if and only if *codominated* actions are never recommended to any player in any event. Loosely speaking, an action is *codominated* if and only if its recommendation implies that some player will gain from deviating, independently of the communication mechanism.

¹²See, for example, Grossman and Perry (1986), Kohlberg and Mertens (1986), Banks and Sobel (1987), van Damme (1989), Mertens (1989, 1991), Cho and Sobel (1990), Blume et al. (1991), Farrell (1993), and Mailath et al. (1993). See also Govindan and Wilson (2005) and the references therein.

¹³The term *signal jamming* was first used in game theory by Fudenberg and Tirole (1986).

¹⁴In our setting, limited commitment by the principal can be bypassed in the sense that the principal can be treated like any other agent that receives a recommendation and chooses an action that cannot be contracted upon.

the report of the agent.¹⁵ Mediation also helps to provide incentives for effort provision in partnerships (Rahman and Obara, 2010).¹⁶ By secretly appointing an agent to act as a budget-breaker by paying a large sum to the others if the outcome is good, even if the appointment only occurs with a small probability, effort can be incentivized while preserving budget-balance.¹⁷ In strategic information transmission (Crawford and Sobel, 1982), it is also beneficial to introduce a mediator that manipulates the communication between the informed party and the decision-maker (Goltsman et al., 2009). In fact, even *noise* can be beneficial (Blume et al., 2007).¹⁸ Finally, the benefits of mediation have also been highlighted in the literature on repeated games.¹⁹ Kandori (1991) and Mailath et al. (2002) pointed out that an imperfect monitoring technology can contribute to sustaining an equilibrium by working as a trustworthy randomization device on the equilibrium path, but an unreliable one off the equilibrium path (in a way that punishes a deviator).²⁰

2.2 Applications: mechanism design with an outside game

Based on the intuition that rejecting to participate in a mechanism signals characteristics that influence the outcome of the outside game, Cramton and Palfrey (1995) proposed a two-stage mechanism in which the decision to participate is prior to the actual play of the mechanism. In the first stage, firms either accept or reject the mechanism. If the mechanism is unanimously accepted, firms move to the second stage, in which they report their types and an outcome is enforced as a function of the type profile. If some firm rejects the mechanism, the set of rejectors is publicly announced, firms update their beliefs and play the outside game.²¹ They proposed the concept of *ratifiability*: an outcome is ratifiable if it is a sequential equilibrium of this two-stage game with beliefs that satisfy neologism-proofness (Grossman and Perry, 1986; Farrell, 1993).

As an application, Cramton and Palfrey (1995) showed that efficient collusion is not ratifiable

¹⁵A conclusion in the same spirit was obtained by Mitusch and Strausz (2005). Bester and Strausz (2001) had extended the revelation principle to the case of imperfect commitment, under the assumption that the principal perfectly observes the report of the single agent.

¹⁶Strausz (2012) showed that the kind of contracts considered by Rahman and Obara (2010) are permitted in the general framework of Myerson (1982).

¹⁷In related contributions: Rahman (2012) explained how an owner can incentivize a supervisor to exert effort by, with a small probability, secretly asking a worker to shirk and making a large payment to the supervisor if he detects the deviation; while Rahman (2014) showed that a mediator can induce firms to abide by a collusive agreement in an environment with imperfect monitoring by, with a small probability, secretly asking a firm to produce at capacity, and making a large payment to the other firm if the resulting market price is relatively high.

¹⁸In the work of Blume et al. (2007), *noise* refers to a limited form of manipulation: with some probability, the output message is drawn from a distribution that is independent of the input message. In the present model, a probability of spurious rejection that is independent of the input message profile would result in the outside game being played under *passive beliefs*: players would not learn anything from disagreement.

¹⁹See Forges (1988), Lehrer (1992), Renault and Tomala (2004), Tomala (2009), Mertens et al. (2015), and Sugaya and Wolitsky (2016), among others.

²⁰See Kandori (2002) and the references therein.

²¹The separation between a ratifying stage and a communication stage limits the payoff that the principal can attain. In our setting, it would imply a probability of spurious rejection that is independent of the type profile, and, hence, no learning from disagreement (*passive beliefs*).

when the outside game is Cournot competition with private information about costs.²² By rejecting to participate in the cartel, a firm credibly signals that it is efficient, and this increases its profit in the outside game (due to strategic substitutability). It is this anticipated increase in profit that makes it optimal for the firm to reject to collude.²³

Assuming a greater ability of the principal to structure the interaction, we obtain the opposite result. The principal can design the beliefs following disagreement if she is able to merge the ratification and participation stages and employ a mediator who is able to produce spurious rejections that mimic genuine rejections. It is not only the result of Cramton and Palfrey (1995) that is inverted, but all the non-ratifiability results in the literature that rest on the inability of the principal to control what is learned from disagreement. For example, non-ratifiability of efficient collusive agreements in second-price auctions with participation costs (Tan and Yilankaya, 2007).

More generally, trembling mechanisms enlarge the set of outcomes that a mechanism designer is able to induce when the outside option of the agents is a game.²⁴ This is relevant to the literature on collusion-proof mechanism design in the presence of informational frictions among the members of the coalition, where the collusive side-contract is itself a mechanism whose outside option is the status quo mechanism. Applications of collusion-proof mechanism design include the design of organizations (Tirole, 1986, 1992), supplier networks (Laffont and Martimort, 1997; Faure-Grimaud et al., 2003; Mookherjee and Tsumagari, 2004; Che and Kim, 2006), mechanisms for public good provision (Laffont and Martimort, 2000), or optimal auctions (Dequiedt, 2007; Pavlov, 2008; Che and Kim, 2009), under the threat of collusion.²⁵

Of course, collusion does not exhaust the scope for applications. Analogous issues arise in the design of any kind of agreement when disagreement leads to subsequent interaction under adverse selection. Applications may thus include the design of dispute resolution schemes (Cooter and Rubinfeld, 1989; Hörner et al., 2015), climate-change agreements (Martimort and Sand-Zantman, 2016), and other dynamic contracting settings (Philippon and Skreta, 2012; Tirole, 2012; Board and Pycia, 2014; Jullien et al., 2014).

²²Contrarily to what had been concluded by Cramton and Palfrey (1990) and Kihlstrom and Vives (1992) under the assumption that firms do not learn from disagreement.

²³Examining non-mediated two-stage mechanisms in which the only thing firms observe before playing the outside game is the set of rejectors, Celik and Peters (2011) showed that some outcomes are only implementable through mechanisms that are rejected on-path. Considering an uninformed firm and an informed firm whose cost may be low or high, they explained that: in an equilibrium with full participation, if the uninformed firm deviates and rejects to participate, the outside game is played with no belief updating; while, in an equilibrium in which the informed firm rejects to participate if it has a high cost, if the uninformed firm deviates and rejects to participate, the outside game is played with full information. In their example, disclosure hurts the uninformed firm in expectation, meaning that the equilibrium rejection relaxes its participation constraint. In mechanisms without a ratification stage (one-stage mechanisms), the benefits of equilibrium rejections can be achieved by making the type announcements public even if some agent rejects the mechanism. This would avoid the potential loss from inducing equilibrium rejections.

²⁴Jullien (2000) studied mechanism design with type-dependent outside options. In our setting, the outside option is not only type-dependent but also belief-dependent.

²⁵Particular forms of collusion that consist in bribing a rival have been considered, for example, by Schummer (2000), Eső and Schummer (2004), Chen and Tauman (2006), and Rachmilevitch (2013).

3 The model

A principal (P), without private information, designs and proposes a *trembling mechanism* to a finite set of agents, $I \equiv \{1, \dots, n\}$. Let $I_P \equiv I \cup \{P\}$. Each agent $i \in I$ is privately informed about his type, $\theta_i \in \Theta_i$. The set of possible type profiles is assumed to be finite and denoted by $\Theta \equiv \prod_{i \in I} \Theta_i$. The actual type profile of the agents is drawn according to a commonly known probability distribution, $\mu^0 \in \Delta(\Theta)$, assumed to have full support.²⁶

A trembling mechanism incorporates a *mediation device*, which is defined by: a set of private input message profiles from the agents, $M^I \equiv \prod_{i \in I} M_i^I$; a set of private output message profiles, $M^O \equiv \prod_{i \in I_P} M_i^O$; and transition probabilities from inputs to outputs, $\tau : M^I \rightarrow \Delta(M^O)$. We denote by $\tau(m^O | m^I)$ the probability of the output being $m^O \in M^O$ conditionally on the input message profile having been $m^I \in M^I$.

We consider mediation devices with the following structure. Each agent $i \in I$ can privately report a type or reject the mechanism: $M_i^I \equiv \Theta_i \cup \{R_i\}$, where R_i is the rejection message. The set of agents who send the rejection message is denoted by $I_R^I \equiv \{i \in I : m_i^I = R_i\}$. If there is at least one rejector ($I_R^I \neq \emptyset$), the output of the mediation device to each player $i \in I_P$ is the true set of rejectors plus a private signal: $m_i^O = (I_R^O, s_i)$ such that $I_R^O = I_R^I$ and $s_i \in S_i$, where S_i is finite. If all agents accept the mechanism ($I_R^I = \emptyset$), the mediation device either truthfully transmits the input message profile, $m_i^O = m_i^I, \forall i \in I_P$, or mimics rejection by one of the agents, transmitting $m_i^O = (I_R^O, s_i)$ such that $I_R^O \in I$ and $s_i \in S_i, \forall i \in I_P$.²⁷ The set of private output messages to player $i \in I_P$ is, therefore, $M_i^O \equiv \Theta \cup (\mathcal{I} \times S_i)$, where $\mathcal{I}$ is the set of nonempty subsets of I .

A mediation device of this kind can be decomposed into a trembling device and a correlating device. A *trembling device*, $\varepsilon : \Theta \rightarrow \Delta(I \cup \{0\})$, is defined by the probabilities of generating a spurious rejection by each agent $i \in I$ when the input message profile is $\theta \in \Theta$, denoted $\varepsilon^i(\theta)$.²⁸ With probability $\varepsilon^0(\theta) = 1 - \sum_{i \in I} \varepsilon^i(\theta)$, there is no spurious rejection. A *correlating device*, $\psi \equiv (\psi^i, \psi^{(i)})_{i \in I}$, is defined by the distributions over profiles of private signals that follow a spurious rejection by each agent $i \in I$ when the input message profile is $\theta \in \Theta$, denoted $\psi^i : \Theta \rightarrow \Delta(S)$, where $S \equiv \prod_{i \in I_P} S_i$; and by those that follow a genuine rejection by each agent $i \in I$ when the input message profile from acceptors is $\theta_{-i} \in \Theta_{-i}$, denoted $\psi^{(i)} : \Theta_{-i} \rightarrow \Delta(S)$.²⁹

A mechanism is said to be *strictly trembling* if and only if: any message that may be received by player $j \in I_P \setminus \{i\}$ of type $\theta_j \in \Theta_j$ when agent $i \in I$ genuinely rejects the mechanism (and others report truthfully), may also be received by player j of type $\theta_j \in \Theta_j$ when (all agents report

²⁶We denote by $\Delta(Z)$ the set of probability distributions over Z .

²⁷It is not useful for our purposes to generate spurious rejections by more than one agent, nor to generate a spurious rejection when some agent genuinely rejects the mechanism.

²⁸With probability $\varepsilon^i(\theta)$, the message “agent i has rejected” is transmitted to all players.

²⁹We will investigate several scenarios regarding the characteristics of the correlating device: the case in which there is no correlating device; the case in which players receive a public signal; the case in which signals are uncorrelated with the input message profile; and, finally, the general case in which the distribution over profiles of private signals depends on the input message profile.

truthfully and) there is a spurious rejection by agent i . Formally: $\sum_{\theta_{-ij} \in \Theta_{-ij}} \psi_j^{(i)}(s_j | (\theta_j, \theta_{-ij})) > 0 \Rightarrow \sum_{\theta_{-j} \in \Theta_{-j}} \varepsilon^i(\theta_j, \theta_{-j}) \psi_j^i(s_j | (\theta_j, \theta_{-j})) > 0, \forall s_j \in S_j, \forall \theta_j \in \Theta_j, \forall j \in I_P \setminus \{i\}, \forall i \in I$.³⁰ A mechanism is said to be *non-trembling* if and only if $\varepsilon^0(\theta) = 1, \forall \theta \in \Theta$.

A *trembling mechanism*, (ε, ψ, x) , is thus defined by a mediation device, (ε, ψ) , and an allocation, $x : \Theta \rightarrow \Delta(X)$, which yields a lottery over a finite set of consequences, X , as a function of the type announcements of the agents (unless the output of the trembling device is a rejection message). The resulting payoff of each player $i \in I_P$, denoted $\pi_i(x(m^O), \theta) \in \mathbb{R}$, is linear in the first variable.

If the output of the trembling device is a rejection message, principal and agents play a single-stage outside game by simultaneously choosing an action, $a_i \in A_i$, for $i \in I_P$.³¹ The resulting payoff of each player $i \in I_P$, denoted $\pi_i^R(y(\theta), \theta) \in \mathbb{R}$, is a linear function of the joint distribution over actions, $y(\theta) \in \Delta(A)$, where $A \equiv \prod_{i \in I_P} A_i$.

The timing of the interaction is the following:

1. Nature draws the type profile of the agents, $\theta \in \Theta$; each agent $i \in I$ observes θ_i .
2. The principal proposes a trembling mechanism, (ε, ψ, x) , to the agents.
3. Each agent $i \in I$ chooses a private input message, $m_i^I \in \Theta_i \cup \{R_i\}$.
4. The output message profile, $m^O \in M^O$, is drawn from $\tau(m^I)$; player $i \in I_P$ observes m_i^O .
5. If $I_R^O = \emptyset$, the outcome $x(m^I)$ is enforced. If $I_R^O \neq \emptyset$, principal and agents play the outside game by simultaneously choosing an action, $a_i \in A_i$, for $i \in I_P$.

We will focus on a simplified game in which the principal proposes a given trembling mechanism, (ε, ψ, x) , which is said to be *incentive compatible* and *individually rational* if and only if there exists a sequential equilibrium of this simplified game in which agents always accept to participate and report their types truthfully.

This will allow us to characterize the set of allocations, x , that are *virtually t -feasible* in the sense that there exists a sequence of incentive compatible and individually rational strictly trembling mechanisms, $(\varepsilon^m, \psi^m, x^m)_{m \in \mathbb{N}}$, converging to a non-trembling mechanism that proposes the desired allocation, $(0, \psi, x)$. Virtual t -feasibility of x means that an allocation that is arbitrarily close to x can be brought into effect with a probability that is arbitrarily close to one, in an equilibrium such that no unilateral deviation can lead to an information set off the equilibrium path.³²

³⁰We denote by ψ_j^i and $\psi_j^{(i)}$ the marginal distributions over S_j derived from the joint distributions ψ^i and $\psi^{(i)}$.

³¹A player $i \in I_P$ may not be active in the outside game (e.g., his action set, A_i , may be a singleton).

³²In the usual sense: an allocation x is *virtually feasible* if and only if there exists a feasible allocation that is arbitrarily close to x (Abreu and Matsushima, 1992).

4 The single-agent case

We start by analyzing the case in which the principal interacts with a single agent. This case is much simpler because we only need to worry about the belief of the principal about the agent, and because there are only two players and one-sided private information in the outside game.

4.1 Non-mediated mechanism

As a benchmark, consider a mechanism that is non-trembling (i.e., does not produce spurious rejections) and non-correlating (i.e., does not send any additional signals to principal and agent accompanying the announcement of a rejection). Such a mechanism, denoted $(0, 0, x)$, is completely defined by the allocation, $x : \Theta \rightarrow X$, that is proposed.

Expecting the distribution over action profiles in the outside game, as a function of his type, to be given by $y : \Theta \rightarrow \Delta(A)$, the agent finds it optimal to participate and report his type truthfully if and only if the following incentive compatibility and individual rationality constraints are satisfied:

$$\pi_1(x(\theta), \theta) \geq \pi_1(x(\theta'), \theta), \quad \forall \theta' \neq \theta, \quad \forall \theta \in \Theta \quad (\text{IC})$$

$$\pi_1(x(\theta), \theta) \geq \pi_1^R(y(\theta), \theta), \quad \forall \theta \in \Theta. \quad (\text{IR})$$

For the allocation x to be the outcome of a sequential equilibrium, in addition to (IC) and (IR) being satisfied for some outside option, y , there must exist consistent beliefs following disagreement such that there exists a Bayesian Nash Equilibrium (BNE) of the outside game that induces y .

Let $\mu \in \Delta(\Theta)$ denote the commonly known belief of the principal about the type of the agent when the outside game is played (we can restrict μ to be common knowledge because our solution concept is sequential equilibrium), and let $BNE(\mu)$ denote the set of distributions over action profiles that constitute a BNE of the outside game when the disagreement belief is μ .³³

In a sequential equilibrium in which the mechanism is never rejected, beliefs following disagreement are formed off the equilibrium path. Therefore, any μ that is common knowledge is consistent. Notice that the “*no signaling what you don’t know*” condition is trivially satisfied (there is nothing that the agent who deviates does not know).³⁴

An allocation $x : \Theta \rightarrow X$ is said to be *θ -feasible* if and only if the non-mediated mechanism, $(0, 0, x)$, is incentive compatible and individually rational.

³³Formally: $y \in BNE(\mu)$ if and only if $y(\theta, a_P, a_1) = \sigma_P(a_P) \sigma_1(\theta, a_1)$, $\forall (\theta, a_P, a_1) \in \Theta \times A_P \times A_1$, where the strategy of the agent, $\sigma_1 : \Theta \rightarrow \Delta(A_1)$, is such that $\sigma_1(\theta)$ is a best-response to σ_P if the agent has type θ , for all $\theta \in \Theta$; and the strategy of the principal, $\sigma_P \in \Delta(A_P)$ is the best-response to σ_1 if she believes that the agent is of type $\theta \in \Theta$ with probability $\mu(\theta)$.

³⁴In this benchmark case, the interaction can be described by a two-stage game with observable actions, therefore, a sequential equilibrium is equivalent to a perfect Bayesian equilibrium in which beliefs formed off the equilibrium path are common knowledge and satisfy “*no signaling what you don’t know*” (Fudenberg and Tirole, 1991).

Remark 1. An allocation x is 0-feasible if and only if x satisfies (IC) and (IR) for some $y \in \cup_{\mu \in \Delta(\Theta)} BNE(\mu)$.

With non-mediated mechanisms, the set of credible punishments is the union of the sets of Bayesian Nash Equilibria of outside game, $BNE(\mu)$, over all possible disagreement beliefs, $\mu \in \Delta(\Theta)$.

4.2 Non-trembling mechanism

Another benchmark is a non-trembling mechanism incorporating a correlating device. The difference with respect to a non-mediated mechanism is that, following a rejection, principal and agent receive private signals on which they may condition their strategies in the outside game.³⁵

This means that the outside game becomes an extended outside game, whose Bayesian Nash Equilibria are, by definition, *Uninformed Mediator Bayesian Correlated Equilibria* of the non-extended outside game. Denote by $BCE^U(\mu)$ the set of distributions over action profiles induced by the set of such equilibria, when the disagreement belief is μ .³⁶

An allocation $x : \Theta \rightarrow X$ is *nt-feasible* if and only if there is an incentive compatible and individually rational non-trembling mechanism proposing the allocation, $(0, \psi, x)$.

Remark 2. An allocation x is *nt-feasible* if and only if it satisfies (IC) and (IR) for some $y \in \cup_{\mu \in \Delta(\Theta)} BCE^U(\mu)$.

Hence, the set of credible punishments is the union of the sets of Uninformed Mediator Bayesian Correlated Equilibria of outside game, $BCE^U(\mu)$, over all possible disagreement beliefs, $\mu \in \Delta(\Theta)$.

It is trivial that $BNE(\mu) \subseteq BCE^U(\mu)$, but the inclusion may not be strict. A correlating device relaxes the participation constraints whenever there exists $y \in BCE^U(\mu)$, for some $\mu \in \Delta(\Theta)$, that yields a lower payoff to an agent of some type when compared with any $y' \in BNE^U(\mu')$, for any $\mu' \in \Delta(\Theta)$. An example is given in Appendix A.2.

A caveat regarding the credibility of the punishments that sustain 0-feasibility and *nt*-feasibility is that, as the solution concept is sequential equilibrium, any interpretation by the principal about the information content of the agent's off-path choice to reject the proposal is allowed, since consistency of beliefs only implies that μ is common knowledge. A stronger refinement, such as the intuitive criterion (Cho and Kreps, 1987), some version of divinity (Cho and Sobel, 1990), or neologism-proofness (Grossman and Perry, 1986; Farrell, 1993), could further restrict

³⁵Private signals may be recommendations of actions to make as a function of own type.

³⁶Forges (1993) describes several coherent definitions of correlated equilibrium in games with incomplete information. This one is designated as *strategic form correlated equilibrium* and is defined as follows. Considering the outside game in its strategic form, a pure strategy by player $i \in I_P$ is a mapping $\tilde{a}_i : \Theta_i \rightarrow A_i$, and a pure strategy profile is a vector $\tilde{a} \in \tilde{A} \equiv \prod_{i \in I_P} \tilde{A}_i$. A distribution over profiles of private recommendations, $y \in \Delta(\tilde{A})$, is a strategic form correlated equilibrium, $y \in BCE^U(\mu)$, if and only if it is optimal for players to obey the recommendations, i.e., $y(\tilde{a}_{-i} | \tilde{a}_i) \pi_i^R(\tilde{a}_i, \tilde{a}_{-i}) \geq y(\tilde{a}'_i | \tilde{a}_i) \pi_i^R(\tilde{a}'_i, \tilde{a}_{-i})$, $\forall \tilde{a}'_i \in \tilde{A}_i, \forall i \in I_P$.

beliefs formed off-path (and, as a result, the set of feasible allocations) by ruling out unreasonable conjectures of the principal about the possible types of the rejector and their relative probabilities.

4.3 Trembling devices and strict incentive compatibility

Now consider a trembling device that, when the agent reports type $\theta \in \Theta$, produces a spurious rejection with probability $\varepsilon(\theta) = \frac{\mu(\theta)}{\mu^0(\theta)}\epsilon$, where $0 < \epsilon < \min_{\theta \in \Theta} \mu^0(\theta)$ and $\mu \in \Delta(\Theta)$.³⁷ With such a device, we do not need to worry about beliefs formed off-path following a rejection, because, in an equilibrium in which the agent always accepts the mechanism and reports truthfully, rejection occurs with positive probability. The equilibrium belief of the principal following disagreement is determined by Bayesian updating, which yields the probability distribution $\mu \in \Delta(\Theta)$.

Let $v_1(\theta', \theta)$ be the agent's expected payoff after a spurious rejection, where θ' is his report and θ is his true type. The agent finds it optimal to participate and report truthfully if and only if:

$$\pi_1(x(\theta), \theta) + \frac{\varepsilon(\theta)}{1-\varepsilon(\theta)} v_1(\theta, \theta) \geq \frac{1-\varepsilon(\theta')}{1-\varepsilon(\theta)} \pi_1(x(\theta'), \theta) + \frac{\varepsilon(\theta')}{1-\varepsilon(\theta)} v_1(\theta', \theta), \quad \forall \theta' \neq \theta, \quad \forall \theta \in \Theta \quad (\text{IC}')$$

$$\pi_1(x(\theta), \theta) \geq \pi_1^R(y(\theta), \theta), \quad \forall \theta \in \Theta. \quad (\text{IR})$$

By generating spurious rejections that are correlated with the agent's report, trembling devices distort the incentives for truth-telling. Fortunately, the distortion can be made arbitrarily small, since disagreement beliefs, μ , do not depend on the *ex ante* probability of spurious rejection, $\epsilon > 0$. Given any allocation, x , that strictly satisfies all the incentive compatibility (IC) conditions in a non-trembling mechanism, there is a trembling probability, $\epsilon > 0$, that is small enough for all the incentive compatibility (IC') conditions to remain strictly satisfied in the trembling mechanism.

The following assumption ensures that any allocation that satisfies (IR) and (IC) can be approximated by a sequence of allocations that satisfy (IR) and strictly satisfy (IC).³⁸

Assumption 1. *There exists an allocation, x^f , that strictly satisfies (IC) and satisfies (IR) for any distribution y .*

The approximation can be made through a sequence of (IR) and strictly (IC) allocations that are weighted averages between the allocation that we are approximating (whose weight increases along the sequence) and x^f (whose weight vanishes in the limit).

A trembling device allows the principal to design the disagreement belief, μ , which becomes the result of Bayesian updating in equilibrium (instead of being a postulated off-the-equilibrium-path belief that could violate appropriate refinements). The costs of this design are the strictly positive probability of reversion to the outside game and the possible necessity of strict incentive

³⁷The upper bound on ϵ guarantees that the conditional probability of a spurious rejection is smaller than 100%.

³⁸I am grateful to Takuro Yamashita for suggesting this assumption and the subsequent argument, which he used in Yamashita (2014).

compatibility.³⁹ As a result, an allocation that is feasible with a non-trembling mechanism becomes only *virtually* feasible if the principal uses a strictly trembling mechanism.

4.4 Trembling mechanism with no correlating device

Let us start by considering the simplest kind of trembling mechanism, which is a trembling mechanism with no correlating device: the mediator produces a spurious rejection with strictly positive probability, $\epsilon > 0$, but does not transmit any signals besides the announcement of rejection.

An allocation $x : \Theta \rightarrow X$ is said to be virtually *tnc-feasible* if and only if there exists a sequence of incentive compatible and individually rational strictly trembling mechanisms with no correlating device, $(\epsilon^m, 0, x^m)_{m \in \mathbf{N}}$, converging to a non-mediated mechanism proposing the allocation, $(0, 0, x)$.

Proposition 1. *An allocation is virtually tnc-feasible if and only if it is 0-feasible.*

Despite the equivalence expressed in Proposition 1, strictly trembling mechanisms have an advantage: beliefs are formed on-path, according to Bayes' rule. Thus, refinements based on restrictions of beliefs formed off-path are trivially satisfied because there are no information sets off-path.

4.5 Trembling mechanism with public correlating device

Allowing the mediator to send a public signal after announcing a rejection expands the set of credible threats to the convex hull of the set of Bayesian Nash Equilibria of the outside game.

The principal becomes able to induce any mixture of BNE distributions over action profiles, $y \in \text{Conv}(\cup_{\mu \in \Delta(\Theta)} \text{BNE}(\mu))$, because the disagreement belief, $\mu^k \in \Delta(\Theta)$, and the continuation equilibrium, $y^k \in \text{BNE}(\mu^k)$, can be made dependent on the public signal, $k \in K$.⁴⁰

An allocation $x : \Theta \rightarrow X$ is said to be virtually *tpc-feasible* if and only if there exists a sequence of incentive compatible and individually rational strictly trembling mechanisms with public correlation, $(\epsilon^m, \psi^{pc,m}, x^m)_{m \in \mathbf{N}}$, that converges to a non-trembling mechanism that proposes the allocation, $(0, \psi^{pc}, x)$.

Proposition 2. *An allocation x is virtually tpc-feasible if and only if it satisfies (IC) and (IR) for some $y \in \text{Conv}(\cup_{\mu \in \Delta(\Theta)} \text{BNE}(\mu))$.*

A public correlating device allows the principal to choose the disagreement belief and the associated continuation equilibrium as a function of the agent's report when the rejection is spurious.

³⁹Strictness of all incentive compatibility conditions is sufficient, but not necessary, to accommodate the distortions caused by the trembling device.

⁴⁰From Caratheodory's theorem, any given $y \in \text{Conv}(\cup_{\mu \in \Delta(\Theta)} \text{BNE}(\mu))$ can be written as a convex combination of a finite number, $|A|$, of distributions.

However, when the rejection is genuine, since the agent does not report his type, the principal cannot condition the punishment on the type reported by the agent.

Allowing the agent to announce his type when he rejects the mechanism would not enable the principal to harshen the punishment, because the agent would not be willing to provide information that would only be used to punish him. On the other hand, as shown by Lehrer and Sorin (1997), if principal and agent could send input messages that are rich enough to define an encryption code, the common signal could become in practice a profile of private signals because each player would not be able to decode the parcel of the common signal intended to be decoded by the rival.⁴¹

4.6 Trembling mechanism with extraneous correlating device

If the trembling device, when announcing a rejection, also sends a profile of private signals that is uncorrelated with the announcement made by the agent (i.e., is extraneous), the outside game becomes an extended game in which principal and agent observe correlated private signals whose distribution is independent of the agent's type, and may condition their strategies on these signals. As already mentioned, a Bayesian Nash Equilibrium of this extended outside game is, by definition, an Uninformed Mediator Bayesian Correlated Equilibrium of the non-extended outside game.

An allocation $x : \Theta \rightarrow X$ is said to be virtually *tec-feasible* if and only if there exists a sequence of incentive compatible and individually rational strictly trembling mechanisms with extraneous correlation, $(\varepsilon^m, \psi^{ec,m}, x^m)_{m \in \mathbf{N}}$, that converges to a non-trembling mechanism that proposes the allocation, $(0, \psi^{ec}, x)$.

Proposition 3. *An allocation is virtually tec-feasible if and only if it is nt-feasible.*

The set of credible punishments is the same as in the case of non-trembling mechanisms: the union of the sets of Uninformed Mediator Bayesian Correlated Equilibria of the outside game, $BCE^U(\mu)$, across the set of possible (common knowledge) disagreement beliefs, $\mu \in \Delta(\Theta)$.

This means that an extraneous profile of private signals can be used to punish the agent more harshly, whenever the outside game has some $y \in BCE^U(\mu)$, for some $\mu \in \Delta(\Theta)$, that yields a lower payoff to an agent of some type when compared to any $y' \in BNE(\mu')$, for any $\mu' \in \Delta(\Theta)$.⁴²

4.7 Trembling mechanism with general correlating device

We now consider the general case in which the mediator is able to send a profile of private signals to principal and agent according to a probability distribution that depends on the input message.

⁴¹The scenario in which the mediator sends a profile of private signals is studied in Section 4.7.

⁴²See Appendix A.2 for an example.

In an equilibrium with participation and truth-telling: after a spurious rejection, since the agent reported truthfully (on-path), the mediator is informed and thus can send private recommendations to principal and agent that depend on the type of the agent; if the rejection is genuine (off-path), the mediator is uninformed and thus cannot condition the recommendations on the agent's type. Unable to distinguish genuine from spurious rejections, the principal presumes that a rejection is spurious and obeys any recommendation she would obey after a spurious rejection.⁴³

Formally, after announcing a (spurious or genuine) rejection, the mediator sends a profile of private signals distributed according to $\psi : M^I \rightarrow \Delta(T \times A)$, where $T \equiv \prod_{i \in I_P} T_i$ is a set of private signal profiles, and $A \equiv \prod_{i \in I_P} A_i$ is the set of action profiles in the outside game. It is equivalent to think of private signals as being sent in two stages: first, through $\psi^T : M^I \rightarrow \Delta(T)$; second, through $\psi^A : M^I \times T \rightarrow \Delta(A)$. If the input message is $\theta \in \Theta$, the trembling device generates a spurious rejection with probability $\varepsilon(\theta) \equiv \frac{\mu(\theta)}{\mu^0(\theta)}\epsilon$, with $\epsilon > 0$. A spurious rejection is thus generated with *ex ante* probability ϵ , and the common knowledge disagreement belief is $\mu \in \Delta(\Theta)$. As will be shown, the common disagreement belief is not relevant as long as it has full support.⁴⁴

The continuation game after a spurious rejection is an extended outside game in which principal and agent receive private signals that may be correlated with the type of the agent. This game was studied by Bergemann and Morris (2016), with their common prior corresponding to the common disagreement belief, μ . By definition, a Bayesian Nash Equilibrium of this extended outside game is an *Informed Mediator Bayesian Correlated Equilibrium*, $BCE^I(\mu)$, of the non-extended game.⁴⁵

Bergemann and Morris (2016) showed that any relevant private signal besides a recommendation of an action to be chosen shrinks the set of $BCE^I(\mu)$, i.e., shrinks the set of recommendation rules that are obeyed. Providing additional information to the players generates additional incentive compatibility conditions, which, taken together, are typically stronger and are never weaker than

⁴³After a genuine rejection, the mediator should send to the principal the recommendation that is the most harmful to the agent, among those that the principal obeys. Whether the principal has incentives to obey it conditionally on a genuine rejection is not relevant. The principal obeys a recommendation if and only if she has incentives to obey it conditionally on a spurious rejection (which is infinitely more likely than a genuine rejection).

⁴⁴The mediator could generate a mixture over disagreement beliefs, but it will be clear that this would not benefit the principal. As long as it has full support, the common disagreement belief, μ , is irrelevant because the support of the set of recommendations that are obeyed in a BCE^I does not depend on μ (see Lemma 1 in Appendix A.4).

⁴⁵The work of Bergemann and Morris (2016) is instrumental for our purposes. They considered an environment in which an informed mediator commits *ex ante* to an information structure, $\psi^T : \Theta \rightarrow \Delta(T)$, according to which a profile of private signals, $t \in T$, with $t = (t_P, t_1, \dots, t_n)$ and $T \equiv \prod_{i \in I_P} T_i$, is drawn according to a distribution that depends on the state of nature, $\theta \in \Theta$. The state of nature is observed by the mediator but not by the players, who share a common prior, $\mu \in \Delta(\Theta)$, with full support. The mediator also commits *ex ante* to a decision rule, $\psi^A : T \times \Theta \rightarrow \Delta(A)$, according to which it sends a profile of private recommendations of actions to be made by each player, $a = (a_P, a_1, \dots, a_n) \in A \equiv \prod_{i \in I_P} A_i$. The distribution over profiles of private recommendations may depend on the state of nature, $\theta \in \Theta$, and on the realization of the information signal, $t \in T$. Players obey the recommendations if this is in their interest. A decision rule that is always obeyed induces an outcome $\psi : \Theta \rightarrow \Delta(A)$ such that $\psi(a|\theta) = \sum_{t \in T} \psi^A(a|t, \theta) \psi^T(t, \theta)$, which they designate as a *Bayes Correlated Equilibrium*.

In our setting, the game that follows a spurious rejection is equivalent to the game considered by Bergemann and Morris (2016) under the restriction that ψ^T must, at least, transmit to each agent the information about his own type (which, in our setting, agents already possess). An *Informed Mediator Bayesian Correlated Equilibrium* is, therefore, a *Bayes Correlated Equilibrium* in which each agent is at least informed about his own type. Equivalently, it is a *Bayesian Solution* (Forges, 1993, 2006).

the pooled incentive compatibility condition that arises from a given recommendation when no further information is provided. Therefore, we can restrict private output messages to consist of a private recommendation of an action to be chosen, $a_i \in A_i$.

Off-path, after a genuine rejection, the messages received by principal and agent should be uncorrelated because it is advantageous to conceal the behavior of the principal from the agent. Any relevant signal sent to the agent after a genuine rejection can only increase his payoff by allowing him to condition his best-response (notice that such signals do not influence the decision of the principal, who attributes zero probability to the rejection being genuine). Therefore, the trembling device should completely conceal from the agent the recommendation made to the principal (for example, by always sending the same signal to the agent after a genuine rejection).

In sum: on-path (after a spurious rejection), the trembling device privately recommends actions according to some $y \in BCE^I(\mu)$; off-path (after a genuine rejection), the trembling device privately sends to the principal a recommendation, $a_P \in A_P$, that is also sent on-path, drawn according to $\sigma_P \in \Delta(\text{supp}(y_P|\mu))$.⁴⁶ After a genuine rejection, the principal wrongly presumes that she is on-path with a commonly known disagreement belief, μ , and obeys any recommendation that is also made on-path; while the agent knows that he is off-path and best-responds to the distribution over recommendations sent to the principal in case of a genuine rejection. The distribution over recommendations, $\sigma_P \in \Delta(\text{supp}(BCE_P^I))$, which is necessarily independent of the agent's type, should be constructed in the way that minimizes the agent's best-response payoff.⁴⁷

An allocation $x : \Theta \rightarrow X$ is said to be virtually t -feasible if and only if there exists a sequence of incentive compatible and individually rational strictly trembling mechanisms, $(\varepsilon^m, \psi^m, x^m)_{m \in \mathbf{N}}$, converging to a non-trembling mechanism that proposes it, $(0, \psi, x)$.

Proposition 4. *An allocation x is virtually t -feasible if and only if it satisfies (IC) and (IR) for some y induced by (σ_P, σ_1) such that $\sigma_P \in \Delta(\text{supp}(BCE_P^I))$ and $\sigma_1(\theta)$ is a best-response to σ_P , $\forall \theta \in \Theta$.*

With general correlating devices, the set of credible punishments is the set of mixed strategies, $\sigma_P \in \Delta(A_P)$, whose support is contained in the union across disagreement beliefs, $\mu \in \Delta(\Theta)$, of the union of the supports of marginal distributions $y_P \in BCE_P^I(\mu)$. Equivalently, whose support is contained in the support of the marginal distribution of a BCE^I with maximal support.

Comparing the set of credible threats with general correlating devices with the one with extraneous correlating devices, we can distinguish three advantages of sending private signals according to a distribution that depends on the input message. First, the potential support of the principal's

⁴⁶Given a distribution over profiles of actions, $y(\theta) \in \Delta(A_P \times A_1)$: we denote the marginal distribution over A_P by $y_P(\theta) \in \Delta(A_P)$. We define $\text{supp}(y_P|\mu) \equiv \cup_{\theta \in \text{supp}(\mu)} \text{supp}(y_P(\theta))$, and, writing $y_P \in BCE_P^I(\mu)$ if and only if $y \in BCE^I(\mu)$, we also define $\text{supp}(BCE_P^I(\mu)) \equiv \cup_{y \in BCE^I(\mu)} \text{supp}(y_P|\mu)$ and $\text{supp}(BCE_P^I) \equiv \cup_{\mu \in \Delta(\Theta)} \text{supp}(BCE_P^I(\mu))$.

⁴⁷The restriction that recommended actions belong to the support of a single BCE^I is without loss of generality. Given a set of actions such that each action belongs to the support of a different BCE^I under possibly different beliefs, we can construct a single BCE^I whose support contains all those actions (see Lemma 1 in Appendix A.4).

mixed action becomes larger (because any BCE^U is also a BCE^I).⁴⁸ Second, any distribution over that support is allowed (not only the equilibrium distribution). For example, if there is an action that punishes the agent more than any other, the principal can play it with 100% probability. Third, if the principal plays a mixed action, the correlating signal to the agent can be shut off (preventing the agent from conditioning his best-response).⁴⁹

The second advantage is illustrated in the example of an all-pay auction with bid caps (Section 6.2.1). The first and the third advantages do not appear in our single-agent examples, where there is only one-sided private information. They do materialize in the Cournot triopoly example, where there are multiple-agents, and thus private information is multi-sided (Section 6.1.3).

4.8 Summary

The benefit for the principal of using trembling mechanisms is the expansion of the set of credible punishments. Below, we compare the sets of credible punishments for different mechanism formats:

- Non-trembling, non-correlating: $y \in \cup_{\mu \in \Delta(\Theta)} BNE(\mu)$
- Non-trembling, correlating signals: $y \in \cup_{\mu \in \Delta(\Theta)} BCE^U(\mu)$
- Trembling, non-correlating: $y \in \cup_{\mu \in \Delta(\Theta)} BNE(\mu)$
- Trembling, public correlating signals: $y \in Conv(\cup_{\mu \in \Delta(\Theta)} BNE(\mu))$
- Trembling, extraneous correlating signals: $y \in \cup_{\mu \in \Delta(\Theta)} BCE^U(\mu)$
- Trembling, correlating signals: $y_P \in \Delta(\cup_{\mu \in \Delta(\Theta)} supp(BCE_P^I(\mu)))$.

It may seem that, in the absence of general correlating signals, trembling mechanisms do not improve on non-trembling mechanisms. There is actually an important improvement. Since a strictly trembling mechanism generates disagreement on-path, feasibility is robust to the introduction of refinements based on restrictions of beliefs formed off-path.

Regarding the usefulness of correlating signals, we verify that public signals only convexify the set of credible threats, while extraneous private signals are only advantageous in relatively specific outside games (by enlarging the space of punishments from the set of BNE to the possibly larger set of BCE^U). The usefulness of general private signals is more evident: credibility of a punishment only requires that any action recommended to the principal is also recommended in some BCE^I .⁵⁰

⁴⁸A $BCE^U(\mu)$ is a $BCE^I(\mu)$ in which recommendations are independent of the type of the agent.

⁴⁹Shutting down the correlating signal to the agent is advantageous because it prevents the agent from conditioning his best-response. Since this only occurs off-path, it is not taken into account by the principal. If the principal knew that the agent was not receiving his correlating signal, she could prefer disobey the recommendation.

⁵⁰It is important to keep in mind that the punishment strategy cannot depend on the agent's type.

5 The multiple-agent case

With multiple agents, further complexity arises from multi-sided private information. A genuine rejector faces uncertainty about the characteristics of the other agents, and his beliefs are not amenable to manipulation by the trembling device.⁵¹ As a result, after a genuine rejection, the outside game is played under beliefs that exhibit a particular violation of the common prior assumption: acceptors (presuming that the rejection was spurious) update their beliefs according to Bayes' rule after observing the rejection, and believe that all players do so; while the genuine rejector retains his prior belief, and knows that all other players have updated their beliefs.⁵²

We will continue to consider candidate sequential equilibria in which agents always participate and report truthfully, and focus on whether any agent of any type can gain by unilaterally rejecting the mechanism. To extend the incentive compatibility (IC') and individual rationality (IR) conditions to the multiple-agent case, we must consider expected payoffs according to the *interim* expectation of agent $i \in I$ of type $\theta_i \in \Theta_i$ about the type profile, $\theta \in \Theta$, denoted $\mathbb{E}_{\theta|\theta_i} [\pi(\theta)] \equiv \sum_{\theta \in \Theta} \mu^0(\theta|\theta_i) \pi(\theta)$. We also introduce further notation: let $y^i(\theta)$ be the distribution over action profiles that agent $i \in I$ expects if he genuinely rejects the mechanism, as a function of the type profile, $\theta \in \Theta$; and let $v_i^j(\theta'_i, \theta)$ be the payoff that agent $i \in I$ expects if there is a spurious rejection by agent $j \in I$, as a function of his report, $\theta'_i \in \Theta_i$, and the type profile, $\theta \in \Theta$.

If all other agents participate and report truthfully, it is optimal for agent $i \in I$ to participate and report truthfully if and only if:

$$\mathbb{E}_{\theta|\theta_i} \left[\pi_i(x(\theta), \theta) + \sum_{j \in I} \frac{\varepsilon^j(\theta)}{\varepsilon^0(\theta)} v_i^j(\theta_i, \theta) \right] \geq \mathbb{E}_{\theta|\theta_i} \left[\frac{\varepsilon^0(\theta'_i, \theta_{-i})}{\varepsilon^0(\theta)} \pi_i(x(\theta'_i, \theta_{-i}), \theta) + \sum_{j \in I} \frac{\varepsilon^j(\theta'_i, \theta_{-i})}{\varepsilon^0(\theta)} v_i^j(\theta'_i, \theta) \right], \quad \forall \theta'_i \neq \theta_i, \quad \forall \theta_i \in \Theta_i \quad (\text{IC}')$$

$$\mathbb{E}_{\theta|\theta_i} [\pi_i(x(\theta), \theta)] \geq \mathbb{E}_{\theta|\theta_i} [\pi_i^R(y^i(\theta), \theta)], \quad \forall \theta_i \in \Theta_i. \quad (\text{IR})$$

We say that (IC) is satisfied if and only if (IC') is satisfied with $\varepsilon^0(\theta) = 1, \forall \theta \in \Theta$.

As in the single-agent case, we start by considering non-mediated and non-trembling mechanisms. Then, we focus on trembling mechanisms without correlating devices and with general correlating devices. The cases in which correlating devices are public or extraneous are only briefly discussed because they are complex and do not help to clarify the general case.⁵³

We will work under the natural extension of Assumption 1 to the multiple-agent case.

⁵¹In the single-agent case, the single rival of the rejector does not have any private information.

⁵²Here, and henceforth, the principal is referred to as an acceptor, although she proposes the mechanism and does not explicitly accept it.

⁵³In the case of public correlation, it is necessary to trade off tailoring the punishment to the characteristics of acceptors with the fact that such tailoring may convey information to the rejector that allows him to better respond to the punishment. In the case of extraneous correlation, it is necessary to deal with the mentioned violation of the common prior assumption: after a genuine rejection, acceptors update their beliefs and believe that all other players do, while the rejector maintains his prior belief knowing that other players have updated their beliefs.

Assumption 1’. *There exists an allocation, x^f , that, $\forall i \in I$, strictly satisfies (IC) and satisfies (IR) for any distribution y^i .*

5.1 Non-mediated mechanism

In the absence of mediation, as we consider sequential equilibria with participation and truthful reporting, a rejection by agent $i \in I$ can only directly influence beliefs about himself. The set of possible disagreement beliefs is composed by those that satisfy “no signaling what you don’t know”:⁵⁴ $\mathcal{B}_0^i \equiv \{\mu^i \in \Delta^*(\Theta) : \mu^i(\theta|Z) = \mu^i(\theta_i|Z) \mu^0(\theta|\theta_i \cap Z), \forall \theta \in \Theta\}$.

After a rejection by agent $i \in I$, the common belief becomes $\mu^i \in \mathcal{B}_0^i$: it is common knowledge that the probability that each player $j \in I_P$ attributes to the type profile being $\theta \in \Theta$ is $\mu^i(\theta|\theta_j)$.⁵⁵

Proposition 5. *An allocation x is 0-feasible if and only if, for each $i \in I$, it satisfies (IC) and (IR) for some $y^i \in \cup_{\mu \in \mathcal{B}_0^i} BNE(\mu)$.*

5.2 Non-trembling mechanism

Even without producing spurious rejections, a mediator can enlarge the set of credible punishments. Being informed about the type profile of acceptors, $\theta_{-i} \in \Theta_{-i}$ (but not about the type of the rejector), the mediator can make private recommendations that, besides correlating players’ actions (as in the single-agent case), may convey information about the type profile of acceptors.

Designate a BNE of the extended outside game in which a mediator makes recommendations that depend on θ_{-i} but not on θ_i as a *Partially Informed Mediator Bayesian Correlated Equilibrium*, and denote by $BCE^{-i}(\mu)$ the set of such equilibria when the disagreement belief is $\mu \in \Delta^*(\Theta)$.

Proposition 6. *An allocation x is nt-feasible if and only if, for each $i \in I$, it satisfies (IC) and (IR) for some $y^i \in \cup_{\mu \in \mathcal{B}_0^i} BCE^{-i}(\mu)$.*

5.3 Trembling mechanism with no correlating device

When a strictly trembling mechanism is used, although we focus on equilibria with participation and truth-telling, beliefs of acceptors in the face of disagreement are determined on-path. Acceptors always presume that the rejection is spurious, and update their beliefs accordingly.

⁵⁴With multiple agents and non-trembling mechanisms, it is convenient to define a common belief as a conditional probability system (Myerson, 1986), because the disagreement belief may attribute zero marginal probability to some type of rejector and it may be necessary to consider the *interim* belief of a genuine rejector of this type. A conditional probability system is a function $\mu \in \Delta^*(\Theta)$ such that, for every Z that is a nonempty subset of Θ , the conditional probability function $\mu(\cdot|Z) \in \Delta(\theta)$ is such that $\mu(Z|Z) = 1$ and $\mu(Z''|Z) = \mu(Z''|Z') \mu(Z'|Z)$, for all $Z'' \subseteq Z' \subseteq Z$. We will continue to denote $\mu(Z') \equiv \mu(Z'|\Theta)$, $\forall Z' \subseteq \Theta$.

⁵⁵Notice that $\mu^i \in \mathcal{B}_0^i$ implies that the posterior of player i remains unchanged: $\mu^i(\theta|\theta_i) = \mu^i(\theta_i|\theta_i) \mu^0(\theta|\theta_i \cap \theta_i) = \mu^0(\theta|\theta_i)$, $\forall \theta \in \Theta$. Working with conditional probability systems allows $\mu^i(\theta|\theta_i)$ to be well defined even if $\mu^i(\theta_i) = 0$.

If the mediator does not send additional signals to the players besides the announcement of rejection, the beliefs of an acceptor when playing the outside game result from two pieces of information: the rejection message, which is public; and the information about own type, which is private. Acceptors can be seen as having a common disagreement belief, which results from Bayesian updating of the common prior, μ^0 , after the public observation of the rejection message.⁵⁶

The principal can induce any disagreement belief that is compatible with each acceptor's private information. For an example of incompatibility, suppose that μ^i attributed zero probability to agent $j \in I \setminus \{i\}$ being of type $\hat{\theta}_j$, which means that a spurious rejection by agent i is never produced if agent j reports type $\hat{\theta}_j$. Then, if agent j is of type $\hat{\theta}_j$, he is able to infer that a rejection by agent i has been genuine, and forms beliefs off-path. To rule out this possibility, the disagreement belief must attribute strictly positive marginal probabilities to all types of all acceptors (not necessarily to all type profiles of acceptors): $\mu^i \in \mathcal{B}_\varepsilon^i \equiv \{\mu \in \Delta(\Theta) : \mu_j(\theta_j) > 0, \forall \theta_j \in \Theta_j, \forall j \neq i\}$.

The principal designs the disagreement belief through the trembling device. For the common belief after a (spurious or genuine) rejection by agent $i \in I$ to be $\mu^i \in \Delta(\Theta)$, the probability of spurious rejection by agent i should be given by $\varepsilon^i(\theta) = \frac{\mu^i(\theta)}{\mu^0(\theta)}\epsilon^i$, with $0 < \epsilon^i < \frac{1}{|I|} \min_{\theta \in \Theta} \mu^0(\theta)$, $\forall i \in I$. Following a spurious rejection, players share a common disagreement belief, $\mu^i \in \mathcal{B}_\varepsilon^i$, which means that the posterior of player $j \in I_P$ of type $\theta_j \in \Theta_j$ is $\mu^i(\cdot|\theta_j) \in \Delta(\Theta)$. This is common knowledge. Following a genuine rejection by agent $i \in I$, the beliefs of acceptors are exactly the same as after a spurious rejection, and this is also common knowledge. In contrast, the genuine rejector maintains his prior belief, $\mu^0(\cdot|\theta_i) \in \Delta(\Theta)$, and this is only known by the rejector himself.

Let $R_{BNE(\mu)}^i$ be the set of mappings $y^{(i)} : \Theta \rightarrow \Delta(A)$ that can be induced by a mixed strategy profile $(\sigma_{-i}^i, \sigma_i^{(i)})$, such that $\sigma_{-i}^i \in BNE_{-i}(\mu)$ and $\sigma_i^{(i)}(\theta_i)$ is a best-response to σ_{-i}^i under the prior belief $\mu^0(\cdot|\theta_i)$, $\forall \theta_i \in \Theta_i$.⁵⁷

Proposition 7. *If an allocation x , for each $i \in I$, satisfies (IC) and (IR) for some $y^{(i)} \in \cup_{\mu \in \mathcal{B}_\varepsilon^i} R_{BNE(\mu)}^i$, then x is virtually tnc-feasible.*

The exact converse is not true because lower hemi-continuity of $BNE(\mu)$ payoffs with respect to μ may fail at the boundary (Einy et al., 2012). Since $\mathcal{B}_\varepsilon^i$ is not closed, a sequence of punishments that are enacted along a sequence of strictly trembling mechanisms may converge to a $BNE_{-i}(\mu)$ with $\mu \notin \mathcal{B}_\varepsilon^i$, which may be a significantly harsher punishment than any $BNE_{-i}(\mu)$ with $\mu \in \mathcal{B}_\varepsilon^i$.

Proposition 8. *If an allocation x is virtually tnc-feasible, then, for each $i \in I$, it satisfies (IC) and (IR) for some $y^{(i)} \in \cup_{\mu \in \Delta(\Theta)} R_{BNE(\mu)}^i$.*

⁵⁶Agents receive their private information before publicly observing the rejection. Still, since posterior beliefs do not depend on the order in which public and private information is processed, they can be thought of as resulting from a common component (the disagreement belief) that each agent then combines with his private information.

⁵⁷We say that $z_{-i} \in BNE_{-i}(\mu)$ if and only if there exists $z \in BNE(\mu)$ and $z_i : \Theta_i \rightarrow \Delta(A_i)$ such that $z(\theta) = z_{-i}(\theta_{-i}) z_i(\theta_i)$, $\forall \theta \in \Theta$.

The set of credible punishments is the union of the sets of BNE strategy profiles (restricted to acceptors), $BNE_{-i}(\mu)$, at least across beliefs that satisfy the interiority condition, $\mu \in \mathcal{B}_\varepsilon^i$.

Perhaps the main takeaway is that since the probability of a spurious rejection by one agent can depend on the types announced by all agents, the resulting disagreement beliefs do not have to satisfy “*no signaling what you don’t know*”.⁵⁸ With multiple agents, therefore, trembling mechanisms enlarge the set of credible disagreement beliefs (relatively to non-trembling mechanisms).

5.4 Trembling mechanism with public correlating device

With multiple agents, public signals allow the mediator to condition the disagreement belief and the associated punishment on the information provided by acceptors. However, such conditioning transmits information to the rejector about the type profile of his rivals. There is thus a trade-off between: on the one hand, tailoring the punishment to the characteristics of acceptors; and, on the other hand, concealing from the rejector the punishment and the characteristics of acceptors.

Even if we restrict the public signals to be independent of type profile of acceptors, a public correlating device convexifies the set of credible threats (as in the single-agent case).⁵⁹ Without that restriction, an even larger set of punishments can be made credible.⁶⁰

5.5 Trembling mechanism with extraneous correlating device

If the mediator can send an extraneous profile of private signals, the set of credible punishments becomes the set of Uninformed Mediator Bayesian Correlated Equilibria under beliefs that violate the common prior assumption in the way already described. After a genuine rejection: while the common belief of acceptors can be chosen by design, and acceptors believe that it is common to all players, the rejector retains his prior belief and knows the belief of acceptors.⁶¹

5.6 Trembling mechanism with general correlating device

The set of credible punishments can become larger (and further relax the participation constraints) if the mediator is able to transmit a profile of private signals whose distribution depends on the

⁵⁸If the probability of spurious rejection by agent $i \in I$ only depended on his report, the mediator would only be able to influence beliefs about the rejector. Disagreement beliefs would satisfy “*no signaling what you don’t know*”.

⁵⁹Although independent of the type profile, such a public signal is not extraneous because its distribution depends on whether the rejection has been spurious or genuine.

⁶⁰A complete characterization would require the analysis of an extended outside game in which an informed mediator publicly recommends actions, taking into account the fact that a genuine rejector retains his prior belief and makes inferences about the type profile of acceptors from the observation of these public recommendations. Again, it is worth mentioning the result by Lehrer and Sorin (1997), according to which public recommendations become *de facto* private if players can send input messages that are rich enough to define an encryption code.

⁶¹The common disagreement belief of acceptors must satisfy the same interiority condition as in the scenario in which the mediator does not send any additional signals: $\mu^i \in \mathcal{B}_\varepsilon^i$.

input message profile. As in the single-agent case, these private signals can be restricted w.l.o.g. to consist of recommendations of actions to choose in the outside game. A correlating device can thus be defined by the distributions over profiles of private signals that follow a spurious and a genuine rejection by agent $i \in I$, denoted $\psi^i : \Theta \rightarrow \Delta(A)$ and $\psi^{(i)} : \Theta_{-i} \rightarrow \Delta(A)$, respectively.

Suppose that, when the input is $\theta \in \Theta$, the mediator: produces a spurious rejection by agent $i \in I$ with probability $\varepsilon^i(\theta) = \frac{\mu^i(\theta)}{\mu^0(\theta)}\epsilon^i$, where $0 < \epsilon^i < \frac{1}{|I|} \min_{\theta \in \Theta} \mu^0(\theta)$, and μ^i has full support; conditionally on a spurious rejection by agent i , sends a profile of private recommendations, $a \in A$, according to a distribution $\psi^i(\theta)$ such that $\psi^i \in BCE^I(\mu^i)$.⁶² Since, after rejection by agent $i \in I$ is observed, the common belief is μ^i (derived using Bayes' rule), by definition, players always find it optimal to obey the recommendations if and only if, given the common belief μ^i , the mapping $\psi^i : \Theta \rightarrow \Delta(A)$ is an Informed Mediator Bayesian Correlated Equilibrium, $BCE^I(\mu^i)$.

After a genuine rejection by agent $i \in I$ (off-path), the set of recommendations that acceptor $j \in I_P \setminus \{i\}$ of type $\theta_j \in \Theta_j$ obeys is the set of recommendations that acceptor j obeys after a spurious rejection (on-path) under the same circumstances: the union across $\theta_{-j} \in \Theta_{-j}$ of the supports of marginal distributions $\psi_j^i(\theta_j, \theta_{-j})$. Off-path, acceptor $j \in I_P \setminus \{i\}$ of type $\theta_j \in \Theta_j$ can thus be induced to obey a recommendation if and only if it is a recommendation that he receives with positive probability in ψ^i . From Lemma 1, there exists $z^i \in BCE^I(\mu^i)$ whose support contains the supports of all BCE^I . Let us, therefore: set $\psi^i = z^i$, let $\mathcal{A}_j(\theta_j) \equiv \cup_{\theta_{-j} \in \Theta_{-j}} \text{supp} [\psi_j^i(\theta_j, \theta_{-j})]$, and let $\mathcal{A}_{-i}(\theta_{-i}) \equiv \Pi_{j \in I_P \setminus \{i\}} \mathcal{A}_j(\theta_j)$, $\forall \theta_{-i} \in \Theta_{-i}$.⁶³ Acceptors can be induced to obey any punishment $y_{-i}^{(i)} : \Theta_{-i} \rightarrow \Delta(A_{-i})$ such that $y_{-i}^{(i)}(\theta_{-i}) \in \Delta(\mathcal{A}_{-i}(\theta_{-i}))$.⁶⁴ The recommendation made to each acceptor $j \in I_P \setminus \{i\}$ will typically not accord with ψ^i , but each acceptor will believe that it does.

Finally, it is important to understand what should the mediator recommend to a genuine rejector. It is straightforward that, since the private signal sent to a genuine rejector does not affect the incentives of acceptors (because genuine rejections occur with zero probability), an informative signal may benefit – but never harm – the rejector by allowing him to condition his best-response. Hence, a genuine rejector should be sent an irrelevant signal (e.g., a constant signal or pure noise).

Proposition 9. *An allocation x is virtually t -feasible if and only if, for each $i \in I$, it satisfies (IC) and (IR) for some $y^{(i)}$ induced by $(y_{-i}^{(i)}, \sigma_i^{(i)})$ such that $y_{-i}^{(i)}(\theta_{-i}) \in \Delta(\mathcal{A}_{-i}(\theta_{-i}))$, $\forall \theta_{-i} \in \Theta_{-i}$, and $\sigma_i^{(i)}(\theta_i)$ is a best-response to $y_{-i}^{(i)}$, $\forall \theta_i \in \Theta_i$.⁶⁵*

⁶²From Lemma 1 (in Appendix A.4), the belief μ^i is irrelevant (as long as it has full support) because there exists $z^i \in BCE^I(\mu^i)$ whose support contains all the actions in the support of a BCE^I . Only the support of the BCE^I played after a spurious rejection is relevant because the recommendations that acceptors obey after a genuine rejection are exactly those in the support of that BCE^I . Hence, after a spurious rejection, the mediator should recommend a BCE^I with maximal support, as z^i . It also follows from Lemma 1 that it is not useful to induce different acceptors to have different beliefs off-path (given that they have a common disagreement belief on-path).

⁶³From Lemma 1: $\mathcal{A}_j(\theta_j) = \cup_{\mu^i \in \Delta(\Theta)} \cup_{\phi^i \in BCE^I(\mu^i)} \cup_{(\theta_j, \theta_{-j}) \in \text{supp}(\mu^i)} \text{supp} [\phi_j^i(\theta_j, \theta_{-j})]$.

⁶⁴The principal should choose the one that maximizes her payoff by relaxing the participation constraints as much as possible. The punishment that is the most harmful to the rejector may depend on the type of the rejector. However, the punishment cannot be contingent on the rejector's type (which is only known by the rejector himself).

⁶⁵When agent $i \in I$ genuinely rejects, a punishment can be induced if and only if: for each acceptor $j \in I_P \setminus \{i\}$,

6 Examples

6.1 Collusion in oligopoly

6.1.1 Cournot duopoly

Consider a duopoly in which firms simultaneously choose the quantities to supply under one-sided private information. Firm A (agent) has private information about its unit cost, $\theta_A \in \{0, \frac{1}{3}\}$, while the unit cost of firm P (principal) is zero ($\theta_P = 0$). Demand is such that the profit function of firm $i \in \{A, P\}$ is given by $\pi_i = (1 - q_i - \frac{1}{2}q_j - \theta_i)q_i$, where $\{i, j\} = \{A, P\}$.⁶⁶

Suppose that the principal can propose an enforceable “take-it-or-leave-it” collusive agreement involving side-payments.⁶⁷ As we will see, there is a menu of contracts that is incentive compatible, maximizes joint-profit and leaves each type of agent indifferent between accepting or rejecting.

Joint-profit maximization under complete information yields: $(q_A, q_P) = (\frac{1}{3}, \frac{1}{3})$ if $\theta_A = 0$, and $(q_A, q_P) = (\frac{1}{9}, \frac{4}{9})$ if $\theta_A = \frac{1}{3}$. The resulting profits of the agent (before side-payments) are $\pi_A(0) = \frac{1}{6}$ and $\pi_A(\frac{1}{3}) = \frac{1}{27}$. In case of disagreement, denoting by $\mu_A \in [0, 1]$ the commonly known probability that the principal attributes to $\theta_A = \frac{1}{3}$, the principal produces $q_P^R = \frac{18+2\mu_A}{45}$, and the profits of the agent are $\pi_A^R(0) = (\frac{36-\mu_A}{90})^2$ and $\pi_A^R(\frac{1}{3}) = (\frac{21-\mu_A}{90})^2$. The side-payments to the agent that leave him indifferent between accepting or rejecting are $t(0) = \frac{-54-72\mu_A+\mu_A^2}{90^2}$ and $t(\frac{1}{3}) = \frac{141-42\mu_A+\mu_A^2}{90^2}$.

The best disagreement belief for the principal consists in believing that the agent has the highest cost with 100% probability if he rejects the agreement ($\mu_A = 1$). This is due to strategic substitutability. A high-cost agent produces less, leaving a greater residual demand. Expecting to face a high-cost agent, the principal produces more, which reduces the profit of the agent.

In the sequential equilibrium in which the payoff of the principal is maximal: the disagreement belief is $\mu_A = 1$, transfers are $t(0) = -\frac{5}{324}$ and $t(\frac{1}{3}) = \frac{1}{81}$, and the agent’s payoffs are $\pi_A(0) = \frac{49}{324}$ and $\pi_A(\frac{1}{3}) = \frac{4}{81}$.⁶⁸ However, this equilibrium does not resist the neologism-proof refinement. The full set of agents is a credible veto set: if the principal believes that the rejector may be of either type, both types have, in fact, strict incentives to reject. This means that a rejector would likely be able to convince the principal to change his belief from $\mu_A = 1$ to the prior probability.

A mediator with the ability to mimic rejection by the agent allows the principal to commit to

the support of the recommendations made to j is contained in $\mathcal{A}_j(\theta_j)$. Contrarily to the scenario with no correlating device, lower hemi-continuity of the support of BCE^I with respect to μ holds at the boundary.

⁶⁶This is a version of the model of Singh and Vives (1984) with one-sided private information. We consider a very fine grid of possible output choices, including all the values that are obtained below.

⁶⁷The enforceability assumption is more tenable in legal cartels, such as export cartels (for example, those operating in the U.S. under the Webb-Pomerene Export Trade Act of 1918 or the Export Trading Company Act of 1982) or agriculture cartels (for example, those operating in the U.S. under a Federal Marketing Order).

⁶⁸Strict incentive compatibility can easily be checked: if the agent announces $\theta_A = \frac{1}{3}$ when $\theta_A = 0$, his payoff becomes $\pi_A = (1 - \frac{4}{9} - \frac{1}{9})\frac{1}{9} + \frac{1}{81} = \frac{5}{81} < \frac{49}{324}$; if he announces $\theta_A = 0$ when $\theta_A = \frac{1}{3}$, his payoff becomes $\pi_A = (1 - \frac{1}{3} - \frac{1}{3} - \frac{1}{3})\frac{1}{3} - \frac{5}{324} = -\frac{5}{324} < \frac{4}{81}$.

her preferred off-path beliefs, $\mu_A = 1$, in a way that resists refinements based on restrictions of beliefs formed off-path. The principal should use a mediator to propose two contracts to the agent: $(q_A, q_P, t) = (\frac{1}{3}, \frac{1}{3}, -\frac{5}{324})$ and $(q_A, q_P, t) = (\frac{1}{9}, \frac{4}{9}, \frac{1}{81})$. If the agent chooses the second contract, it is enforced; if the agent chooses the first contract, it is enforced with probability $1 - \epsilon$, with $\epsilon > 0$. Upon choice of the second contract, the mediator generates a spurious disagreement with probability ϵ , and Bayesian updating implies that disagreement beliefs are $\mu_A = 1$.

6.1.2 Cournot triopoly (without correlating device)

Consider a similar model, but with three firms. The unit cost of firm P (principal) is zero ($\theta_P = 0$), while each of the other firms (agents), $i \in \{A, B\}$, has private information about its unit cost, $\theta_i \in \{0, \frac{1}{3}\}$. Suppose that costs are independently and uniformly distributed, and that the profit function of firm $i \in \{A, B, P\}$ is given by $\pi_i = (1 - q_i - \frac{1}{2}q_j - \frac{1}{2}q_k - \theta_i)q_i$, where $\{i, j, k\} = \{A, B, P\}$.

Joint-profit maximization yields: $(q_A, q_B, q_P) = (\frac{1}{4}, \frac{1}{4}, \frac{1}{4})$ if $(\theta_A, \theta_B) = (0, 0)$; $(q_i, q_j, q_P) = (\frac{1}{3}, 0, \frac{1}{3})$ if $(\theta_i, \theta_j) = (0, \frac{1}{3})$; $(q_i, q_j, q_P) = (0, \frac{1}{3}, \frac{1}{3})$ if $(\theta_i, \theta_j) = (\frac{1}{3}, 0)$; and $(q_A, q_B, q_P) = (\frac{1}{12}, \frac{1}{12}, \frac{5}{12})$ if $(\theta_A, \theta_B) = (\frac{1}{3}, \frac{1}{3})$. The resulting profits (before side-payments) of firm A , $\pi_A(\theta_A, \theta_B)$, are $\pi_A(0, 0) = \frac{1}{8}$, $\pi_A(0, \frac{1}{3}) = \frac{1}{6}$, $\pi_A(\frac{1}{3}, 0) = 0$, and $\pi_A(\frac{1}{3}, \frac{1}{3}) = \frac{1}{36}$.

In case of disagreement, firms choose quantities non-cooperatively. Suppose that the common posterior belief is that $\theta_i = \frac{1}{3}$ with probability μ_i , for $i \in \{A, B\}$, and that θ_A and θ_B are independent. With these beliefs: the output of the principal is $q_P = \frac{9 + \mu_A + \mu_B}{27}$; the output of firm A is $q_A(\theta_A) = \frac{18 - \mu_A + 2\mu_B}{54} - \frac{\theta_A}{2}$; and, analogously, the output of firm B is $q_B(\theta_B) = \frac{18 - \mu_B + 2\mu_A}{54} - \frac{\theta_B}{2}$. The aggregate output of firms P and B is: $q_P + q_B(\theta_B) = \frac{36 + 4\mu_A + \mu_B}{54} - \frac{\theta_B}{2}$.

The “*no signaling what you don’t know*” property of sequential equilibrium restricts beliefs about firm B to remain unchanged ($\mu_B = \frac{1}{2}$) as a result of a deviation by firm A . The worst possible beliefs for firm A when it rejects the agreement are $(\mu_A, \mu_B) = (1, \frac{1}{2})$, i.e., 100% probability that firm A has high cost and unchanged belief about firm B , yielding an expected output from rivals $q_P + \frac{1}{2}[q_B(0) + q_B(\frac{1}{3})] = \frac{2}{3}$, own output $q_A(\theta_A) = \frac{1}{3} - \frac{\theta_A}{2}$, and own profit $\pi_A^R(\theta_A) = (\frac{1}{3} - \frac{\theta_A}{2})^2$.

A trembling mechanism (without a correlating device) allows the principal to avoid the “*no signaling what you don’t know*” restriction and induce, approximately, her preferred disagreement beliefs: $(\mu_A, \mu_B) = (1, 1)$.⁶⁹ These are the beliefs that maximize the sum of the outputs of the principal and the agent that accepts the agreement, yielding $q_P + \frac{1}{2}[q_B(0) + q_B(\frac{1}{3})] = \frac{73}{108}$. The best-response of the rejector is $q_A(\theta_A) = \frac{143}{216} - \frac{\theta_A}{2}$, and the disagreement profit is $\pi_A^R(\theta_A) = (\frac{143}{216} - \frac{\theta_A}{2})^2$.

To induce these beliefs, the mediator should offer a contract whose terms are contingent on the profile of reports, enforcing the first-best output profile and the side-payments that leave each agent indifferent between accepting or rejecting.⁷⁰ If both agents report high cost, the mediator

⁶⁹Beliefs could depend on the identity of the rejector, but this is not useful in this example.

⁷⁰Let us verify that such a menu of contracts is strictly incentive compatible when disagreement beliefs are $(\mu_A, \mu_B) = (1, 1)$, which implies that this is also the case with slightly different beliefs. Truth-telling yields the

should generate a spurious rejection by agent $i \in \{A, B\}$ with probability $\frac{\epsilon}{2(1+\epsilon)}$, with $\epsilon > 0$; if only one agent reports a high cost, the mediator should generate a spurious rejection by this agent with probability $\frac{\epsilon^2}{2(1+\epsilon)}$. The total probability of spurious rejection is ϵ , and the common disagreement belief following a rejection by firm A is $(\mu_A, \mu_B) = (1, \frac{1}{1+\epsilon})$, which converges to $(1, 1)$ as $\epsilon \rightarrow 0$.

When there is more than one informed agent, trembling mechanisms allow the principal to generate a disagreement belief that not only is robust to refinements that restrict beliefs formed off-path, but can also be more favorable to the principal than any belief that is consistent in the absence of a mediator (because it does not have to satisfy “*no signaling what you don’t know*”).

6.1.3 Cournot triopoly (with correlating device)

Now suppose that, after a spurious or genuine rejection, the mediator can send private recommendations to the principal and to the agents concerning how much to produce. The ability of the mediator to work as an informed correlating device allows the principal to relax the participation constraints even further.⁷¹ In our setting, the mediator can be assumed to be perfectly informed when the rejection is spurious (on-path), because agents have truthfully reported their costs.

In an equilibrium in which recommendations are obeyed, denote by $\underline{q}_i$ and $\bar{q}_i$ the minimum and maximum output of firms $i \in \{A, B, P\}$. The maximum output of firm $i \in \{A, B, P\}$ cannot be higher than the best-response to the minimum outputs of its rivals when $\theta_i = 0$. That is: $\bar{q}_i \leq \frac{1}{2} - \frac{1}{4}(\underline{q}_j + \underline{q}_k)$, where $\{i, j, k\} = \{A, B, P\}$. Similarly, the minimum output of firm $i \in \{A, B\}$ cannot be lower than the best-response to the maximum outputs of its rivals when firm i has the highest possible cost: $\underline{q}_i \geq \frac{1}{3} - \frac{1}{4}(\bar{q}_j + \bar{q}_P)$, where $\{i, j\} = \{A, B\}$; and $\underline{q}_P \geq \frac{1}{2} - \frac{1}{4}(\bar{q}_A + \bar{q}_B)$.

As a result of these restrictions, the following system of equations yields upper and lower bounds on the minimum and maximum outputs that firms produce in equilibrium:

$$\left\{ \begin{array}{l} \bar{q}_A = \frac{1}{2} - \frac{1}{4}(\underline{q}_B + \underline{q}_P) \\ \bar{q}_B = \frac{1}{2} - \frac{1}{4}(\underline{q}_A + \underline{q}_P) \\ \bar{q}_P = \frac{1}{2} - \frac{1}{4}(\underline{q}_A + \underline{q}_B) \\ \underline{q}_A = \frac{1}{3} - \frac{1}{4}(\bar{q}_B + \bar{q}_P) \\ \underline{q}_B = \frac{1}{3} - \frac{1}{4}(\bar{q}_A + \bar{q}_P) \\ \underline{q}_P = \frac{1}{2} - \frac{1}{4}(\bar{q}_A + \bar{q}_B) \end{array} \right. \Leftrightarrow \left\{ \begin{array}{l} \bar{q}_A = \bar{q}_B = \frac{53}{135} \\ \bar{q}_P = \frac{59}{135} \\ \underline{q}_A = \underline{q}_B = \frac{17}{135} \\ \underline{q}_P = \frac{41}{135} \end{array} \right.$$

The punishment that can be inflicted on firm A if it genuinely rejects the agreement cannot be harsher than having, with 100% probability, firm P producing $q_P = \frac{59}{135}$ and firm B producing

same profit as rejection: $\pi_A(0) = \frac{143^2}{432^2}$ and $\pi_A(\frac{1}{3}) = \frac{71^2}{432^2}$. Deviating by mimicking the other type yields: $\pi_A^D(0) = \frac{71^2}{432^2} + \frac{1}{2} [q_A(\frac{1}{3}, 0) + q_A(\frac{1}{3}, \frac{1}{3})] \frac{1}{3} = \frac{71^2}{432^2} + \frac{1}{72} < \frac{143^2}{432^2}$, and $\pi_A^D(\frac{1}{3}) = \frac{143^2}{432^2} - \frac{1}{2} [q_A(0, 0) + q_A(0, \frac{1}{3})] \frac{1}{3} = \frac{143^2}{432^2} - \frac{7}{72} < \frac{71^2}{432^2}$.

⁷¹Such an extended game, in which a perfectly informed mediator commits to a distribution over profiles of private recommendations that is conditional on the type profile has been studied by Bergemann and Morris (2016).

$q_B = \frac{53}{135} - \frac{\theta_B}{2}$. These bounds seem impossible to attain because they are best-responses under beliefs that are extremely asymmetric. However, surprisingly, the mediator can induce obedience of recommendations that approximate these bounds. As a result, if firm A genuinely rejects the agreement, the output from its rivals will be approximately equal to $\bar{q}_B - \frac{\theta_B}{2} + \bar{q}_P = \frac{112}{135} - \frac{\theta_B}{2}$. With the expected output from the rivals being equal to $\frac{112}{135} - \frac{1}{12} = \frac{403}{540}$, the best-response of firm A is to produce $q_A^R(\theta_A) = \frac{677}{2160} - \frac{\theta_A}{2}$ and the resulting disagreement profit is $\pi_A^R(\theta_A) = \left(\frac{677}{2160} - \frac{\theta_A}{2}\right)^2$.

To understand how the existence of a mediator allows these bounds to be approximated, keep in mind that the mediator is informed (we are considering a spurious rejection after both firms have truthfully reported their unit costs). This allows the distribution over recommendations to be contingent on the type profile. Observe also that the recommendations that firms B and P obey after a spurious rejection by firm A (on-path) are also obeyed after a genuine rejection by firm A (off-path), because the two kinds of rejection are undistinguishable.

We will construct a stochastic recommendation that is always obeyed after a spurious rejection by firm A , such that the maximum recommendations made to firms B and P , denoted $\bar{q}_B$ and $\bar{q}_P$, are arbitrarily close to the upper and lower bounds, $\bar{q}_B$ and $\bar{q}_P$. The mediator will be instructed to recommend $\bar{q}_B$ and $\bar{q}_P$ with 100% probability after a spurious rejection. The stochastic recommendation to be made after a spurious rejection by firm A is constructed as follows.

1. With probability $1 - \epsilon_0$, recommend the full-information Nash equilibrium outputs: $(q_A, q_B, q_P) = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ if $(\theta_A, \theta_B) = (0, 0)$; $(q_A, q_B, q_P) = (\frac{13}{27}, \frac{13}{27}, \frac{5}{54})$ if $(\theta_A, \theta_B) = (0, \frac{1}{3})$; $(q_A, q_B, q_P) = (\frac{5}{54}, \frac{13}{27}, \frac{13}{27})$ if $(\theta_A, \theta_B) = (\frac{1}{3}, 0)$; $(q_A, q_B, q_P) = (\frac{5}{27}, \frac{5}{27}, \frac{11}{27})$ if $(\theta_A, \theta_B) = (\frac{1}{3}, \frac{1}{3})$. With $\epsilon_0 = 0$, these recommendations would be obeyed, and the (incentive compatibility) obedience constraints would be strictly satisfied.⁷² Choose $\epsilon_0 > 0$ that is small enough for the obedience constraints associated with these recommendations to be remain strictly satisfied for any possible recommendations made with probability ϵ_0 . This guarantees that these output recommendations are obeyed independently of the remainder of the construction. Let $q^0 \equiv (\underline{q}_A^0, \bar{q}_A^0, \underline{q}_B^0, \bar{q}_B^0, \underline{q}_P^0, \bar{q}_P^0) = (\frac{5}{27}, \frac{13}{27}, \frac{5}{27}, \frac{13}{27}, \frac{1}{3}, \frac{11}{27})$.

2. Construct the sequence $\{q^n\}_{n \in \mathbb{N}}$ as follows. Let: $\underline{q}_i^n$ be the best-response by firm i , when it has the highest possible cost, to $(\bar{q}_j^{n-1}, \bar{q}_k^{n-1})$, where $\{i, j, k\} = \{A, B, P\}$. Explicitly: $\underline{q}_A^n = \frac{1}{3} - \frac{1}{4}(\bar{q}_B^{n-1} + \bar{q}_P^{n-1})$, $\underline{q}_B^n = \frac{1}{3} - \frac{1}{4}(\bar{q}_A^{n-1} + \bar{q}_P^{n-1})$ and $\underline{q}_P^n = \frac{1}{2} - \frac{1}{4}(\bar{q}_A^{n-1} + \bar{q}_B^{n-1})$. Similarly, let $\bar{q}_i^n$ be the best-response by firm i , when it has the lowest possible cost, to $(\underline{q}_j^{n-1}, \underline{q}_k^{n-1})$, where $\{i, j, k\} = \{A, B, P\}$. Explicitly: $\bar{q}_A^n = \frac{1}{2} - \frac{1}{4}(\underline{q}_B^{n-1} + \underline{q}_P^{n-1})$, $\bar{q}_B^n = \frac{1}{2} - \frac{1}{4}(\underline{q}_A^{n-1} + \underline{q}_P^{n-1})$ and $\bar{q}_P^n = \frac{1}{2} - \frac{1}{4}(\underline{q}_A^{n-1} + \underline{q}_B^{n-1})$. This transformation is a contraction with modulus $\frac{1}{2}$. Therefore, it has a single fixed point (the bounds that we wish to approximate), which is the limit of the sequence, and convergence is very fast.⁷³

3. With probability given by $(1 - \epsilon_n)\prod_{m=0}^{n-1}\epsilon_m$, make the following recommendations: $(\underline{q}_A^n, \bar{q}_B^{n-1}, \bar{q}_P^{n-1})$

⁷²We are considering a very fine, but finite, grid of possible output choices.

⁷³To prove that it is a contraction, let $d(x, y) \equiv |\underline{x}_A - \underline{y}_A| + |\bar{x}_A - \bar{y}_A| + |\underline{x}_B - \underline{y}_B| + |\bar{x}_B - \bar{y}_B| + |\underline{x}_P - \underline{y}_P| + |\bar{x}_P - \bar{y}_P|$ and $Tx \equiv (\frac{1}{3} - \frac{\bar{x}_B + \bar{x}_P}{4}, \frac{1}{3} - \frac{\bar{x}_A + \bar{x}_P}{4}, \frac{1}{3} - \frac{\bar{x}_A + \bar{x}_P}{4}, \frac{1}{2} - \frac{\bar{x}_A + \bar{x}_B}{4}, \frac{1}{2} - \frac{\bar{x}_A + \bar{x}_B}{4}, \frac{1}{2} - \frac{\bar{x}_A + \bar{x}_B}{4})$. Then: $d(Tx, Ty) = \frac{1}{4}|\bar{x}_B - \bar{y}_B| + \frac{1}{4}|\bar{x}_P - \bar{y}_P| + \frac{1}{4}|\bar{x}_A - \bar{y}_A| + \frac{1}{4}|\bar{x}_A - \bar{y}_A| + \frac{1}{4}|\bar{x}_A - \bar{y}_A| + \frac{1}{4}|\bar{x}_A - \bar{y}_A|$. By the triangular inequality: $d(Tx, Ty) \leq \frac{1}{2}(|\bar{x}_B - \bar{y}_B| + |\bar{x}_P - \bar{y}_P| + |\bar{x}_A - \bar{y}_A| + |\bar{x}_A - \bar{y}_A| + |\bar{x}_A - \bar{y}_A| + |\bar{x}_A - \bar{y}_A|) = \frac{1}{2}d(x, y)$.

or $(\underline{q}_A^{n-1}, \bar{q}_B^n, \underline{q}_P^{n-1})$, equiprobably, if $(\theta_A, \theta_B) = (\frac{1}{3}, 0)$; $(\bar{q}_A^n, \underline{q}_B^{n-1}, \underline{x}_P^{n-1})$ or $(\bar{q}_A^{n-1}, \underline{q}_B^n, \bar{q}_P^{n-1})$, equiprobably, if $(\theta_A, \theta_B) = (0, \frac{1}{3})$; $(\bar{q}_A^{n-1}, \bar{q}_B^{n-1}, \underline{q}_P^n)$, if $(\theta_A, \theta_B) = (0, 0)$; $(\underline{q}_A^{n-1}, \underline{q}_B^{n-1}, \bar{q}_P^n)$, if $(\theta_A, \theta_B) = (\frac{1}{3}, \frac{1}{3})$. Choose $\epsilon_n > 0$ sufficiently small for all the obedience constraints to remain strictly satisfied independently of the remainder of the construction (i.e., independently of what is recommended with the remaining probability, $\Pi_{m=0}^n \epsilon_m$). Recommendations $\underline{q}_i^{n-1}$ and $\bar{q}_i^{n-1}$, not being best-responses in this case, are obeyed because the obedience constraints are pooled with the much more likely cases (previous iterations) in which they are best-responses. Recommendations $\underline{q}_i^n$ and $\bar{q}_i^n$ are obeyed because they are best-responses, and the obedience constraints are strictly satisfied.

A sufficient number of iterations allows the support of the recommendations made after a spurious rejection (on-path) to approximate the bounds as closely as desired. Recommendations made after a genuine rejection by firm A are deterministic, being given by the maximal recommendations in that support: $(\bar{q}_B^N - \frac{\theta_B}{2}, \bar{q}_P^N)$, where N is the number of performed iterations.

6.1.4 Bertrand oligopoly

In a Bertrand oligopoly with differentiated goods (imperfect substitutes) and private information about costs, since prices are strategic complements, the price set by an acceptor is the lowest when the prices it expects the rivals to set are the lowest. Of course, the lower are prices set by the acceptors, the harsher is the punishment of the rejector.

In a duopoly, the principal chooses the lowest price when she believes that the agent has the lowest cost and this belief is common knowledge. To induce these beliefs, the mediator generates a spurious rejection (with a small probability) if and only if the agent announces the lowest cost.

The case of a triopoly is slightly more involved. To induce the harshest possible punishment, the common belief should be that both agents have the lowest cost. The mediator cannot induce this belief by generating a spurious rejection if and only if both agents announce the lowest cost, because an acceptor that does not have the lowest cost would discover that the rejection was genuine. The best the principal can do is induce approximately optimal beliefs. The mediator should generate a spurious rejection by agent $i \in \{A, B\}$: with a small probability, $\epsilon > 0$, if both agents announce the lowest cost; and with an even smaller probability, ϵ^2 , if agent i announces the lowest cost but agent $j \neq i$ does not. Then, after observing a rejection by firm i : the principal and firm j believe that the rejector has the lowest cost with 100% probability; and the principal believes that firm j has the lowest cost with probability $\frac{1}{1+\epsilon}$.

6.2 Collusion in auctions

6.2.1 All-pay auction with bid cap

Two lobbyists, $i \in \{A, P\}$, compete for a government contract by simultaneously choosing how much to spend, $b_i \in [0, m]$, where m is the cap on spending.⁷⁴ The one who spends the most gets the contract, worth $v_P > 0$ to lobbyist P (principal) and $v_A > v_P$ to lobbyist A (agent). In case of tie, each lobbyist wins with 50% probability. Supposing that the cap on bids has an intermediate value, $m \in (\frac{v_P}{2}, v_P)$, non-cooperative spending levels are as follows: lobbyist A randomizes over $(0, 2m - v_P] \cup \{m\}$, with a mass point at m ; while lobbyist P randomizes over $[0, 2m - v_P] \cup \{m\}$, with mass points at 0 and m . Their expected payoffs are $\pi_A = v_A - v_P$ and $\pi_P = 0$, respectively.⁷⁵

The principal proposes a collusive agreement to the agent, offering to withdraw from the contest in exchange for a payment of $v^P - \delta$. If $\delta > 0$, the agent strictly gains from the agreement, as his payoff increases to $v_A - v_P + \delta$. If the agent rejects to collude, non-cooperative bidding ensues.

If the principal has access to a mediator that can mimic the rejection of a proposal, she can obtain a higher payment. The mediator should generate a spurious rejection with probability $\epsilon > 0$, and, after a spurious rejection, recommend a bid to the principal according to the non-cooperative equilibrium distribution over $[0, 2m - v_P] \cup \{m\}$. If the agent genuinely rejects the agreement, the mediator should recommend a bid equal to m with 100% probability.

In the absence of a trembling mechanism, the principal would not obey such recommendation. After observing a rejection, the principal would know that the randomization is not proper and, therefore, would perform the randomization herself. With a trembling mechanism, the principal, after observing a rejection, believes with 100% probability that it is a spurious rejection and, thus, that the randomization is proper. It is the strictly positive probability of spurious rejection that makes the principal obey the faulty randomization, by concealing that it is faulty.

We conclude that, using the mediator as a randomization device, the principal can credibly threaten to bid m in case of a genuine rejection. Then, if $m > \frac{v_A}{2}$, the agent withdraws from the contest and gets a payoff of zero; if $m < \frac{v_A}{2}$, the agent bids at the cap and gets a payoff of $\frac{v_A}{2} - m$. The disagreement payoff of the agent is reduced to $\max\{0, \frac{v_A}{2} - m\}$, which means that the agent becomes willing to pay the principal $v_A - \max\{0, \frac{v_A}{2} - m\} - \delta$, which is clearly more than $v^P - \delta$.

6.2.2 Second-price auction with participation costs

In a second-price auction with participation costs, trembling mechanisms may be crucial for a cartel to be established. Without mediation, as shown by Tan and Yilankaya (2007), a cartel agreement is not ratifiable. Rejecting the agreement credibly signals a high valuation, which makes

⁷⁴Since our theory of trembling mechanisms is developed for discrete action sets, the conclusions of this example hold only asymptotically. It is also worthwhile remarking that, in this example, there is no private information.

⁷⁵This model was proposed and analyzed by Che and Gale (1998b).

it unprofitable for the rivals to support the participation cost. As a result, a high-value bidder is better off rejecting the collusive agreement than ratifying it. However, the same agreement, proposed using a trembling mechanism, would be accepted by all agents.⁷⁶

6.3 Other possible applications

6.3.1 Dispute resolution

Rejection of a settlement proposal can release information that is relevant for a future decision about whether to litigate and about the level of effort to make in the litigation process.⁷⁷ Therefore, it may be important to generate the disagreement beliefs that are the most adverse for a rejector.⁷⁸

6.3.2 Sequential contracting

When two principals contract sequentially with the same agent, the first principal may extract additional surplus from the agent by manipulating the otherwise off-path beliefs of the second principal.⁷⁹ For example, consider an agent with private information about his valuation for a good that is supplied by two principals, in an environment in which the first principal proposes a “take-it-or-leave-it” menu of contracts to the agent, and, if the agent rejects, the second principal has the opportunity to contract with the agent. Employing a mediator that, with a small probability, generates disagreement if the agent picks the contract designed for an agent with high valuation, the first principal reduces the outside payoff of the agent by inducing in the second principal the belief that an agent that rejects the proposal of the first principal must have a high valuation.

6.3.3 Dynamic contracting

In a two-period dynamic screening model with limited commitment, Deb and Said (2015) showed that a seller can gain from inducing rejections in the first period to incentivize herself to charge higher prices in the second period. They constrained rejections to be genuine (the proposal must be designed in a way that makes it optimal for some types to reject), while trembling mechanisms dispense with restrictions on the utility offered to the types that the principal wants to reject.⁸⁰

⁷⁶Without participation costs, beliefs would be irrelevant and thus trembling mechanisms would not be useful.

⁷⁷See Cooter and Rubinfeld (1989).

⁷⁸In Hörner et al. (2015), a mediator can recommend war, with some of the flavor of a spurious rejection.

⁷⁹See Calzolari and Pavan (2006a,b, 2008), and Pavan and Calzolari (2009).

⁸⁰In the model of Deb and Said (2015), some buyers arrive early while others arrive at the last-minute. Those who arrive early can wait for the last-minute to contract. When contracting early, the seller is not able to commit to the contract that she will offer at the last-minute. Therefore, the last-minute proposal is an endogenous outside option for the early proposal. Rejections of the early proposal change the composition of demand at the last-minute and, therefore, may mitigate the seller’s loss due to limited commitment. Deb and Said (2015) conclude that the seller should distort the early contract in a way that induces rejection by an intermediate set of types.

6.3.4 Informed principal

Although outside our scope, the use of a mediator can also allow an informed principal to manipulate the off-path beliefs about her own characteristics. The informed principal must be able to transmit her proposal through a communication device that, with a small probability, modifies the proposal. As a result, an agent that receives a proposal which is different from the one that he should have received according to the candidate equilibrium, will update his beliefs about the principal in a Bayesian way (amenable of design through the communication device).⁸¹

For example, in models of multilateral vertical contracting where a supplier is constrained to make private bilateral offers to multiple downstream firms, belief updating off-path by a downstream firm occurs if it receives an offer that is different from the one prescribed by the candidate equilibrium. It is common to assume *passive beliefs*: continue to expect that the proposals made to the other downstream firms are in accordance with the candidate equilibrium (Hart and Tirole, 1990; Segal, 1999). However, an alternative known as *wary beliefs* has been put forward: expect that the proposals made to the other downstream firms are optimal for the supplier, given the proposal just received (McAfee and Schwartz, 1994; Rey and Vergé, 2004).

Suppose that the supplier can commit to use a mediation device that transforms her proposal into another proposal with a small probability. Then, if a downstream firm receives an unexpected proposal, it may infer that it was due to a tremble by the mediation device and will not update its beliefs about the proposals that the supplier has sent to the other downstream firms. In this limited sense, trembling mechanisms provide a foundation for assuming passive beliefs.

7 Conclusion

In principal-agent environments where the outside payoff of an agent depends on what the other players infer from his rejection of the mechanism, the principal can relax the participation constraints of the agents using trembling mechanisms. These are a specific kind of mediated stochastic mechanisms. Stochastic because even if all agents accept to participate, (with an arbitrarily small probability) the agreement may not be enforced. Mediated because it is crucial that if an agent rejects to participate, this decision is not observed – the other players cannot distinguish a genuine rejection from a spurious rejection. Besides signal-jamming genuine rejections, the mediator also collects information and makes non-binding recommendations. Being always trustworthy on-path allows the mediator to be trusted off-path when it induces acceptors to punish a genuine rejector.

⁸¹In an example presented by Rabin and Sobel (1996), the informed agent, who moves first, has a higher payoff in a pooling equilibrium which does not satisfy reasonable refinements than in a partially separating equilibrium that satisfies the refinements. The informed agent would gain from using a trembling mechanism to choose the otherwise off-path beliefs of the uninformed agent.

A Some remarks on correlated equilibrium

A.1 BCE^U vs BNE

An outside game in which a BCE^U can entail a harsher punishment than any BNE is described in Figure 1.⁸² Despite being a complete information game, the example illustrates the point.

	C1	C2	C3	C4
R1	0 0	3 6	6 3	4 4
R2	6 3	0 0	3 6	4 4
R3	3 6	6 3	0 0	4 4
R4	4 4	4 4	4 4	5 5

Figure 1: A correlated equilibrium may be a harsher punishment than any Nash equilibrium.

In this degenerate outside game, a BCE^U is simply a correlated equilibrium, while a BNE is simply a Nash equilibrium. While the unique Nash equilibrium payoffs are (5,5), there is a correlated equilibrium in which expected payoffs are (5,4).⁸³

A.2 BCE^I vs BCE^U

It is easy to find examples in which a $BCE^I(\mu)$ strategy can punish the agent more than any $BCE^U(\mu)$ strategy. Figure 2 presents a degenerate one: the agent (column player) does not play.

In the game depicted in Figure 2, the principal plays $R3$ in the single BCE^U (which is the single BNE), yielding a payoff of 4 to the agent (column player). The BCE^I in which the principal plays $R1$ if the agent is of type 1 and $R2$ if the agent is of type 2 yields a payoff of 0 to the agent.

Notice that the principal cannot implement such a harsh punishment using trembling mechanisms, because, after a genuine rejection, the mediator cannot make a recommendation that is contingent on the type of the agent. Still, the mediator can recommend $R1$ with probability $\alpha \in [0, 1]$ and $R2$ with probability $1 - \alpha$, yielding an outside payoff of $3(1 - \alpha)$ to the agent of type 1 and an outside payoff of 3α to the agent of type 2.

⁸²This example is attributed to Robert J. Aumann. Versions of it appeared in the works of Nau and McCardle (1990) and Evangelista and Raghavan (1996), from where this version was adapted.

⁸³Observe that a profile of correlated equilibrium strategies that places probability $\frac{2}{9}$ on each outcome whose payoffs are (6, 3), i.e., on $(R1, C2)$, $(R2, C3)$ and $(R3, C1)$, and probability $\frac{1}{9}$ on each outcome whose payoffs are (3, 6), i.e., on $(R1, C3)$, $(R2, C1)$ and $(R3, C2)$, yields expected payoffs of (5, 4).

	(type 1, 50%)		(type 2, 50%)								
R1	<table><tr><td></td><td>0</td></tr><tr><td>3</td><td></td></tr></table>		0	3		R1	<table><tr><td></td><td>3</td></tr><tr><td>0</td><td></td></tr></table>		3	0	
	0										
3											
	3										
0											
R2	<table><tr><td></td><td>3</td></tr><tr><td>0</td><td></td></tr></table>		3	0		R2	<table><tr><td></td><td>0</td></tr><tr><td>3</td><td></td></tr></table>		0	3	
	3										
0											
	0										
3											
R3	<table><tr><td></td><td>4</td></tr><tr><td>2</td><td></td></tr></table>		4	2		R3	<table><tr><td></td><td>4</td></tr><tr><td>2</td><td></td></tr></table>		4	2	
	4										
2											
	4										
2											

Figure 2: A BCE^I may be a harsher punishment than any BCE^U .

A.3 Choosing an action in the support

While, after a spurious rejection, the recommendations of the trembling device must constitute a $BCE^I(\mu)$ strategy profile; after a genuine rejection, the principal will obey any action in the support of a $BCE_P^I(\mu)$ strategy.

There are many examples in which a strategy with support contained in that of an equilibrium strategy can punish the agent more than any equilibrium strategy. In the (complete information) example of Figure 3, the pure strategy $R1$ reduces the payoff of the column player to -7 . The minimum payoff of the column player in a BCE^I is -4 .

	C1	C2				
R1	<table><tr><td>-9</td></tr><tr><td>1</td></tr></table>	-9	1	<table><tr><td>-7</td></tr><tr><td>-1</td></tr></table>	-7	-1
-9						
1						
-7						
-1						
R2	<table><tr><td>1</td></tr><tr><td>-1</td></tr></table>	1	-1	<table><tr><td>-1</td></tr><tr><td>1</td></tr></table>	-1	1
1						
-1						
-1						
1						

Figure 3: A strategy in $\Delta(\text{supp}(BCE^I))$ may punish more than any BCE^I strategy.

A.4 On the structure of BCE^I

There exists a single BCE^I whose support is equal to the union of the supports of all BCE^I for all common beliefs. On-path, the mediator could randomize over beliefs and over different BCE^I , but this would not enlarge the set of recommendations obeyed off-path (which is the support of the distribution over recommendations made on-path) relatively to that single BCE^I .

Lemma 1. *Let $\mu \in \Delta(\Theta)$ have full support. There exists $y \in BCE^I(\mu)$ such that, for all $i \in I_P$: $\text{supp}(y_i|\mu) = \text{supp}(BCE_i^I)$.*⁸⁴

⁸⁴Recall that $\text{supp}(BCE_i^I) \equiv \cup_{\mu \in \Delta(\Theta)} \cup_{y \in BCE^I(\mu)} \text{supp}(y_P|\mu)$.

Proof. Given a common belief with full support, $\mu \in \Delta(\Theta)$, let us construct $y \in BCE^I(\mu)$ whose support contains all actions $a_j \in A_j$, for all $j \in I_P$, that are played with strictly positive probability in some BCE^I . Collect all such a_j , for all $j \in I_P$, and denote this set by $\mathcal{A}$. For each action $a^k \in \mathcal{A}$, there exists a common belief, μ^k , and a BCE^I under that belief, $y^k \in BCE^I(\mu^k)$, in which a^k is played with strictly positive probability. There exists a set of signals that transforms μ into a distribution over posteriors whose support contains $\{\mu^1, \dots, \mu^K\}$, where K is the cardinality of $\mathcal{A}$. Concretely, consider a random signal with possible realizations $\{s^1, \dots, s^K\}$ where the probability of s^k if the type profile is θ is given by $\frac{\mu^k(\theta)}{\mu(\theta)}\epsilon^k$, with $\epsilon^k > 0$, for $k \in K$. Imposing $\sum_k \epsilon^k < \min_\theta \mu(\theta)$ guarantees that the conditional probability of some signal being released is smaller than unity. The posterior probability that results from the observation of s^k is μ^k , for $k \in K$. Conditionally on signal s^k being released, the mediator recommends $y^k \in BCE^I(\mu^k)$. As a result, all actions $a^k \in \mathcal{A}$ are played with strictly positive probability. $\square$

B Proofs

B.1 Single-agent case

Proof of Proposition 1. ($\Rightarrow$) If x is virtually *tnc*-feasible, there exists a sequence of strictly trembling mechanisms, $(\varepsilon^m, 0, x^m)_{m \in \mathbb{N}}$, converging to $(0, 0, x)$, each having a sequential equilibrium in which the agent participates and reports truthfully. Consider a sequence of such equilibria, which are characterized by: an *ex ante* probability of spurious rejection, ϵ^m ; a disagreement belief, μ^m ; and a type-dependent distribution over action profiles, $y^m \in BNE(\mu^m)$, induced by a strategy profile, (σ_P^m, σ_1^m) , where $\sigma_1^m : \Theta \rightarrow \Delta(A_1)$ is such that $\sigma_1^m(\theta)$ is a best-response to σ_P^m if the agent has type θ , $\forall \theta \in \Theta$, while $\sigma_P^m \in \Delta(A_P)$ is a best-response to σ_1^m if the principal believes that the agent is of type $\theta \in \Theta$ with probability $\mu(\theta)$. Consider a subsequence, $(\epsilon^m, \mu^m, y^m, \sigma_P^m, \sigma_1^m)_{m \in M}$, that converges to $(0, \mu, y, \sigma_P, \sigma_1)$. The limit strategy profile, (σ_P, σ_1) , induces the limit distribution over action profiles, y . Since $\pi_P^R(y(\theta), \theta)$ and $\pi_1^R(y(\theta), \theta)$ are continuous with respect to y , the limit strategies are mutual best-responses, thus, $y \in BNE(\mu)$. Since (IC') and (IR) are satisfied along the sequence, x satisfies (IC) and (IR) when the outside option is y .

($\Leftarrow$) From Remark 1, if an allocation x is 0-feasible, there exists a belief, $\mu \in \Delta(\Theta)$, and a type-dependent distribution over action profiles, $y \in BNE(\mu)$, such that x satisfies (IC) and (IR). Given Assumption 1 and linearity of $\pi_1(x(\theta), \theta)$ in the first variable, we can construct a sequence of allocations, $(x^m)_{m \in \mathbb{N}}$, converging to x , that strictly satisfy (IC'), and satisfy (IR) given the same outside option, y . For example, letting $x^m \equiv \frac{m}{m+1}x + \frac{1}{m+1}x^f$, $\forall m \in \mathbb{N}$. Each allocation in the sequence, x^m , satisfies (IC') and (IR) for sufficiently small $\epsilon^m > 0$ (since π_1^R is bounded, v_1 is also bounded, thus ϵ^m can be chosen irrespectively of the behavior in the outside game after a spurious rejection). Therefore, we can construct a sequence of pairs $(x^m, \epsilon^m)_{m \in \mathbb{N}}$, converging to $(x, 0)$, such that (IC') and (IR) are satisfied along the sequence. The sequence of strictly trembling mechanisms, $(\varepsilon^m, 0, x^m)_{m \in \mathbb{N}}$, becomes completely defined by letting $\varepsilon^m(\theta) = \frac{\mu(\theta)}{\mu^0(\theta)}\epsilon^m$, which implies that the disagreement belief is μ along the sequence. As a result, following a genuine rejection, the independent strategies that induce $y \in BNE(\mu)$ are sequentially rational. Hence, each mechanism in the sequence has a sequential equilibrium in which the

agent participates and reports truthfully. As $\lim_{m \rightarrow \infty} \epsilon^m = 0$, the sequence of mechanisms converges to the non-mediated mechanism, $(0, x)$. $\square$

Proof of Proposition 2. ($\Rightarrow$) Let x be virtually *tpc*-feasible. There exists a sequence of strictly trembling mechanisms with public correlation, $(\epsilon^m, \psi^{pc,m}, x^m)_{m \in \mathbf{N}}$, converging to $(0, \psi^{pc}, x)$, each having a sequential equilibrium in which the agent participates and reports truthfully. In mechanism $(\epsilon^m, \psi^{pc,m}, x^m)_{m \in \mathbf{N}}$, if the agent reports type $\theta \in \Theta$, a spurious rejection is produced and a public signal, $k \in K$, is sent with probability $\epsilon_k^m(\theta) = \frac{\mu_k^m(\theta)}{\mu_0^m(\theta)} \epsilon_k^m$. Observation of $k \in K$ induces the disagreement belief $\mu_k^m \in \Delta(\Theta)$, and the distribution over action profiles, $y_k^m \in BNE(\mu_k^m)$. If there is a genuine rejection (off-path), signal $k \in K$, such that $\epsilon_k^m > 0$ (so that the principal cannot detect that the rejection has been genuine), is sent with probability p_k^m . This implies that the same belief, μ_k^m , and the same behavior, $y_k^m \in BNE(\mu_k^m)$, is induced. The principal behaves in the same way because she does not distinguish a genuine rejection from a spurious rejection. The agent could behave differently, but is not interested in deviating because he is best-responding to the behavior of the principal. Let $y^m \in \text{Conv}(\cup_{\mu \in \Delta(\Theta)} BNE(\mu))$ be the mixture over $\{y_k^m\}_{k \in K}$ induced by probabilities p_k^m . Each (x^m, ϵ^m) satisfies (IC') and (IR) with y^m as the outside option. Taking limits, as $(x^m, \epsilon^m, y^m) \rightarrow (x, 0, y)$, we find that x satisfies (IC) and (IR), with y as the outside option. Continuity of $\pi_P^R(\cdot, \theta)$ and $\pi_1^R(\cdot, \theta)$ implies that $y \in \text{Conv}(\cup_{\mu \in \Delta(\Theta)} BNE(\mu))$.

($\Leftarrow$) Let y be a mixture over a set of distributions over action profiles, $\{y_k\}_{k \in K}$, where $y_k \in BNE(\mu_k)$, with probability weights p_k (from Caratheodory's theorem, K can be assumed to be finite). As explained in the proof of Proposition 1, construct a sequence of allocations, $(x^m)_{m \in \mathbf{N}}$, that satisfy (IR) and strictly satisfy (IC), whose limit is x . Let the associated sequence of strictly trembling mechanisms, $(\epsilon^m, \psi^{pc,m}, x^m)_{m \in \mathbf{N}}$, be such that, following report $\theta \in \Theta$, a spurious rejection is generated and signal $k \in K$ is released with probability $\epsilon_k(\theta) = \frac{\mu_k(\theta)}{\mu_0(\theta)} p_k \epsilon^m$, where $\epsilon^m > 0$ is sufficiently small for (IC) to remain strictly satisfied. After a genuine rejection (off-path), signal $k \in K$ is released with probability p_k . The resulting belief is μ_k , and the resulting behavior is $y_k \in BNE(\mu_k)$. Each of the mechanisms in the sequence has, therefore, a sequential equilibrium in which: the agent participates and reports truthfully; a spurious rejection is generated with probability ϵ^m ; and the outside option is y . As $\epsilon^m \rightarrow 0$, the sequence of strictly trembling mechanisms converges to a non-trembling mechanism, $(0, \psi^{pc}, x)$. $\square$

Proof of Proposition 3. ($\Rightarrow$) If x is virtually *tec*-feasible, there exists a sequence of strictly trembling mechanisms, $(\epsilon^m, \psi^{ec,m}, x^m)_{m \in \mathbf{N}}$, converging to $(0, \psi^{ec}, x)$, each having a sequential equilibrium in which the agent participates and reports truthfully. Such an equilibrium is characterized by: a probability of spurious rejection, ϵ^m ; a disagreement belief, μ^m ; and a joint distribution over actions for each type of the agent, $y^m \in BCE^U(\mu^m)$. Consider a subsequence, (ϵ^m, μ^m, y^m) , that converges to $(0, \mu, y)$. Since $\pi_P^R(y(\theta), \theta)$ and $\pi_1^R(y(\theta), \theta)$ are continuous with respect to y , optimality of obeying the extraneous private recommendations along the sequence implies optimality of obeying the extraneous private recommendations in the limit: $y \in BCE^U(\mu)$. Since (IC') and (IR) are satisfied along the sequence, they are also satisfied in the limit. Hence, x satisfies (IC) and (IR) with y as the outside option.

($\Leftarrow$) From Remark 2, if an allocation x is *nt*-feasible, there exists a belief, $\mu \in \Delta(\Theta)$, and a distribution over action profiles in the outside game, $y \in BCE^U(\mu)$, such that (IC) and (IR) are satisfied. The fact that $y \in BCE^U(\mu)$ means that there exists a distribution over profiles of private signals, $\psi^{ec} \in \Delta(A_P \times A_1^{|\Theta|})$,

recommending an action for the principal to execute and a plan of actions for the agent to execute as a function of his type, which is such that $y(\theta, a_P, a_1) = \psi^{ec}(a_P, a_1(\theta))$, $\forall(\theta, a_P, a_1) \in \Theta \times A_P \times A_1$, and such that neither the principal nor the agent gain from disobeying. Notice that ψ^{ec} must be the same after a spurious and a genuine rejection (because it is extraneous). Assumption 1 and linearity of $\pi_1(\cdot, \theta)$ allow us to construct a sequence of allocations, $(x^m)_{m \in \mathbb{N}}$, converging to x , that strictly satisfy (IC), and satisfy (IR) when the outside option is y . Let $\epsilon^m(\theta) = \frac{\mu(\theta)}{\mu^0(\theta)} \epsilon^m$ be the probability of a spurious rejection following report $\theta \in \Theta$, which implies that the disagreement belief is μ . Since each x^m satisfies (IC') and (IR) as long as $\epsilon^m > 0$ is sufficiently small, we can construct a sequence $(x^m, \epsilon^m)_{m \in \mathbb{N}}$, converging to $(x, 0)$, such that (IC') and (IR) are satisfied along the sequence. All mechanisms in the sequence send extraneous profiles of private signals according to ψ^{ec} . Thus, the outside option is y along the sequence. We conclude that each mechanism in the sequence has a sequential equilibrium in which the agent participates and reports truthfully. As $\epsilon^m \rightarrow 0$, the sequence converges to a non-trembling mechanism, $(0, \psi^{ec}, x)$. $\square$

Proof of Proposition 4. ($\Rightarrow$) Let $(\epsilon^m, \psi^m, x^m)_{m \in \mathbb{N}}$ be a sequence of strictly trembling mechanisms converging to a non-trembling mechanism, $(0, \psi, x)$, such that each mechanism in the sequence has a sequential equilibrium in which the agent participates and reports truthfully. Consider mechanism $m \in \mathbb{N}$ in the sequence. After a spurious rejection, common disagreement beliefs are μ^m and principal and agent play some $z^m \in BCE^I(\mu^m)$. After a genuine rejection (off-path): the principal is recommended an action in $\text{supp}(z_P^m | \mu^m)$ (she obeys, as after a spurious rejection); the agent best-responds to the distribution of recommendations made to the principal. Let $\sigma_P^m \in \Delta(\text{supp}(z_P^m | \mu^m))$ be the mixed action that the principal plays off-path, following a genuine rejection; and let $\sigma_1^m : \Theta \rightarrow \Delta(A_1)$ be a best-response to σ_P^m . Denote by $y^m : \Theta \rightarrow \Delta(A)$ the resulting joint distribution over action profiles. Considering a subsequence that converges, denote the limit of σ_P^m by $\sigma_P \in \Delta(\text{supp}(BCE_P^I))$, and let $\sigma_1(\theta)$ be a best-response to σ_P when the agent is of type θ . Since there is participation and truth-telling along the sequence, x satisfies (IC) and (IR) when rejection leads to (σ_P, σ_1) .

($\Leftarrow$) Let σ_P consist in playing $a^k \in \text{supp}(BCE_P^I)$ with probability p_k . Consider a common disagreement belief with full support, $\mu \in \Delta(\Theta)$, and a BCE^I distribution over action profiles, $y \in BCE^I(\mu)$, whose support contains the support of σ_P . Existence of y follows from Lemma 1. Construct the trembling device of each mechanism $m \in \mathbb{N}$ so that if the input message is θ , a spurious rejection is generated with probability $\epsilon(\theta) \equiv \frac{\mu(\theta)}{\mu^0(\theta)} \epsilon^m$, where $\epsilon^m > 0$. Along the sequence, the common disagreement belief is μ . Instruct the mediator to recommend $y \in BCE^I(\mu)$ after a spurious rejection (in all mechanisms in the sequence). Following a genuine rejection, the mediator recommends a^k with probability p_k to the principal, and sends no meaningful signal to the agent. Construct $(x^m, \epsilon^m)_{m \in \mathbb{N}}$ as in the previous proofs, so that (IC') and (IR) are satisfied with y induced by $(\sigma_P, \sigma_1(\theta))$, where σ_1 is a best-response to σ_P . Consider a candidate sequential equilibrium in which the agent participates and reports truthfully. It is optimal for principal and agent to obey the recommendations made after a spurious rejection. Therefore, the principal also obeys the recommendation following a genuine rejection (believing that the rejection has been spurious, which is infinitely more likely). After a genuine rejection, the agent anticipates that the principal will play σ_P , and plays a best-response, $\sigma_1(\theta)$. This means that the distribution over action profiles is induced by $(y_P, y_1(\theta))$, $\forall \theta \in \Theta$. Since (x^m, ϵ^m) satisfies (IC') and (IR), no deviation from the candidate sequential equilibrium is beneficial. Letting, $\epsilon^m \rightarrow 0$, the mechanism converges to a

non-trembling mechanism, $(0, \psi, x)$. □

B.2 Multiple-agent case

Proof of Proposition 5. ($\Rightarrow$) If x is 0-feasible, there exists a sequential equilibrium of the non-mediated mechanism, $(0, 0, x)$, in which agents participate and report their types truthfully. In this sequential equilibrium, the common belief following a rejection by agent $i \in I$ (off-path), denoted μ^i , belongs to $\mathcal{B}_0^i$; and the resulting distribution over action profiles, y^i , belongs to $BNE(\mu^i)$. Since agents participate and report their types truthfully, (IC) and (IR) are satisfied for each $i \in I$, given each y^i .

($\Leftarrow$) Since, for each $i \in I$, $\mu^i \in \mathcal{B}_0^i$ and $y^i \in BNE(\mu^i)$, we can construct a sequential equilibrium of the non-mediated mechanism, $(0, 0, x)$, in which agents participate and report their types truthfully. For each $i \in I$, let μ^i be the common belief following a rejection by agent $i \in I$ (off-path), and let y^i be the resulting distribution over action profiles. □

Proof of Proposition 6. ($\Rightarrow$) Since x is *nt*-feasible, there exists a sequential equilibrium of a non-trembling mechanism, $(0, \psi, x)$, in which agents participate and report their types truthfully. The common belief following a rejection by agent $i \in I$ (off-path), denoted μ^i , belongs to $\mathcal{B}_0^i$ (remember that under $\mu^i \in \mathcal{B}_0^i$, the belief of agent i of type $\theta_i \in \Theta_i$ is the same as under μ^0); and the continuation distribution over action profiles, y^i , constitutes a $BCE^{-i}(\mu^i)$, by definition. Since agents participate and report truthfully, (IC) and (IR) are satisfied for each $i \in I$, given y^i .

($\Leftarrow$) Let $\mu^i \in \mathcal{B}_0^i$ and $y^i \in BCE^{-i}(\mu^i)$ be such that x satisfies (IC) and (IR), $\forall i \in I$. We will construct a sequential equilibrium of a non-trembling mechanism, $(0, \psi, x)$, in which agents participate and report their types truthfully. Let a pure strategy by player $i \in I$ be a mapping $\tilde{a}_i : \Theta_i \rightarrow A_i$, and denote the finite set of such pure strategies by $\tilde{A}_i$. For each $i \in I$, let μ^i be the common belief following a genuine rejection by agent i , and let $y^i : \Theta_{-i} \rightarrow \Delta(A_{-i} \times \tilde{A}_i)$ be the subsequent distribution over profiles of private recommendations (agent i receives a recommendation for each of his possible types). The recommendations are obeyed because $y^i \in BCE^{-i}(\mu^i)$, by definition. Since, $\forall i \in I$, (IC) and (IR) are satisfied with y^i being the outside option, there is a sequential equilibrium of $(0, \psi, x)$ in which agents participate and report their types truthfully. □

Proof of Proposition 7. Using Assumption 1' and linearity of $\pi_i(x(\theta), \theta)$ in the first variable, we can construct a sequence of allocations, $(x^m)_{m \in \mathbb{N}}$, converging to x , that strictly satisfy (IC) and satisfy (IR), for each $i \in I$, given the outside option, $y^{(i)}$. We know that $y^{(i)}$ is induced by $(\sigma_{-i}^i, \sigma_i^{(i)})$ such that $\sigma_{-i}^i \in BNE_{-i}(\mu^i)$, with $\mu^i \in \mathcal{B}_\varepsilon^i$. Construct a sequence of trembling devices by setting $\varepsilon^{i,m}(\theta) = \frac{\mu^i(\theta)}{\mu^0(\theta)} \varepsilon^m$, for each $i \in I$, with $\lim_{m \rightarrow \infty} \varepsilon^m = 0$. The only difference across trembling devices along the sequence is the value of $\varepsilon^m > 0$. For each $m \in \mathbb{N}$, $\varepsilon^m > 0$ must be sufficiently small for (IC') and (IR) to be satisfied, $\forall i \in I$, given the outside option $y^{(i)}$. This completes the construction of a sequence of strictly trembling mechanisms, $(\varepsilon^m, 0, x^m)_{m \in \mathbb{N}}$, that converges to $(0, 0, x)$. Consider a candidate sequential equilibrium in which: agents participate truthfully; a spurious rejection by agent $i \in I$ leads to a mixed action profile

$(\sigma_{-i}^i, \sigma_i^i) \in BNE(\mu^i)$; and a genuine rejection by agent i leads to the same mixed action profile by the acceptors, $\sigma_{-i}^i \in BNE_{-i}(\mu^i)$, and a best-response by the rejector, $\sigma_i^{(i)}$, which induces $y^{(i)}$. Since (IC') and (IR) are satisfied along the sequence, this is a sequential equilibrium. $\square$

Proof of Proposition 8. Since x is virtually *tnc*-feasible, there exists a sequence of strictly trembling mechanisms, $(\varepsilon^m, 0, x^m)_{m \in \mathbf{N}}$, converging to $(0, 0, x)$, each having a sequential equilibrium in which agents participate and report truthfully. Each equilibrium in the sequence is characterized by: the probability of spurious rejection by each agent $i \in I$, denoted $\varepsilon^{i,m} : \Theta \rightarrow [0, 1]$; the resulting disagreement belief, $\mu^{i,m}$; the strategy of each acceptor $j \in I_P \setminus \{i\}$, following a rejection by agent $i \in I$, denoted $\sigma_j^{i,m} : \Theta_j \rightarrow \Delta(A_j)$; the strategy of a spurious rejector $i \in I$, denoted $\sigma_i^{i,m} : \Theta_i \rightarrow \Delta(A_i)$; and the strategy of a genuine rejector $i \in I$, denoted $\sigma_i^{(i),m} : \Theta_i \rightarrow \Delta(A_i)$. Since the mechanisms in the sequence are strictly trembling, all types of all players play the outside game on-path. Notice also that $\sigma^{i,m} \equiv (\sigma_j^{i,m})_{j \in I_P}$ is a $BNE(\mu^{i,m})$ strategy profile: each $\sigma_j^{i,m} : \Theta_j \rightarrow \Delta(A_j)$ is such that $\sigma_j^{i,m}(\theta_j)$ is a best-response to the other players' strategies, $(\sigma_l^{i,m})_{l \in I_P \setminus \{j\}}$, if agent j has type θ_j , $\forall \theta_j \in \Theta_j$. The best-response of agent j of type θ_j is calculated given the belief that the type profile of the other agents is $\theta_{-j} \in \Theta_{-j}$ with probability $\mu^{i,m}(\theta|\theta_j)$. On the other hand, the strategy profile that results from a genuine rejection, denoted $\sigma^{(i),m} \equiv (\sigma_i^{(i),m}, (\sigma_j^{i,m})_{j \in I_P})$, is not typically a $BNE(\mu^{i,m})$. The genuine rejector best-responds to $(\sigma_j^{i,m})_{j \in I_P}$ given his belief that the type profile of the other players is $\theta_{-i} \in \Theta_{-i}$ with probability $\mu^0(\theta|\theta_i)$. The genuine rejector does not update his belief from μ^0 to $\mu^{i,m}$ as a result of his own deviation (although acceptors believe, and the rejector knows that acceptors believe, that $\mu^{i,m}$ is the common disagreement belief).

Given the sequence $((\varepsilon^{i,m}, \mu^{i,m}, \sigma^{i,m}, \sigma_i^{(i),m})_{i \in I})_{m \in \mathbf{N}}$, pick a subsequence that converges and denote the limit by $(0, \mu^i, \sigma^i, \sigma_i^{(i)})_{i \in I}$. The *ex ante* payoff function in the limit, denoted $\mathbb{E}_{\mu^i, m} [\pi_j^R(z(\theta), \theta)]$, is bilinear in z and μ , and is uniformly bounded. This implies that, if $\mu_j^i(\theta_j) > 0$, the limit strategy $\sigma_j^i(\theta_j)$ is a best-response to σ_{-j}^i : if an alternative strategy of agent j of type θ_j , denoted $\sigma_j^{i'}(\theta_j)$, was strictly preferred in the limit, the same strategy would be strictly preferred for sufficiently large m , which would be a contradiction. Although $\mu_j^{i,m}(\theta_j) > 0$, along the sequence, it is possible that $\mu_j^i(\theta_j) = 0$ in the limit. In that case, considering the limit conditional probability, the objective of agent j of type θ_j becomes well defined and the same kind of contradiction is reached. We conclude that $\sigma^i \in BNE(\mu^i)$. Acceptors $j \in I_P \setminus \{i\}$ play the same strategy profile, $\sigma_{-j}^i \in BNE_{-i}(\mu^i)$, after a genuine rejection. The limit strategy of a genuine rejector, $\sigma_i^{(i)}$, is also a best-response in the limit (for the same reason as before). Let $y^{(i),m}$ and $y^{(i)}$ denote the distributions over joint actions induced by the strategy profiles played after a genuine rejection along the sequence and in the limit, respectively. Since (IC') and (IR) are satisfied along the sequence (with the outside option being $y^{(i),m}$), (IC) and (IR) are also satisfied in the limit (with the outside option being $y^{(i)}$). $\square$

Proof of Proposition 9. ($\Rightarrow$) Let $(\varepsilon^m, \psi^m, x^m)_{m \in \mathbf{N}}$ be a sequence of strictly trembling mechanisms converging to a non-trembling mechanism, $(0, \psi, x)$, such that each mechanism in the sequence has a sequential equilibrium in which the agent participates and reports truthfully. Consider the equilibrium of mechanism $(\varepsilon^m, \psi^m, x^m)$. After a spurious rejection by agent $i \in I$, given the type profile, $\theta \in \Theta$: a common belief, $\mu^{i,m} \in \mathcal{B}^\theta \equiv \{\mu \in \Delta(\Theta) : \mu_j(\theta_j) > 0, \forall j \in I\}$ results from Bayesian updating; and players receive, and obey, recommendations according to $y^{i,m} \in BCE^I(\mu^{i,m})$. After a genuine rejection

by agent $i \in I$ (off-path): each acceptor $j \in I_P \setminus \{i\}$ is recommended, and obeys (presuming that the rejection has been spurious), an action, $a_j \in A_j$, in the support of $y_j^{i,m}$; the profile of recommendations, a_{-i} , is drawn from a distribution $y_{-i}^{(i),m}(\theta_{-i})$; the rejector, i , chooses a best-response, $a_i \in A_i$, or follows a recommendation that allows him to obtain a higher payoff (the worst that can happen to player i is to receive a meaningless private signal, which prevents him from conditioning his best-response). Since there is a sequential equilibrium with truthful participation, x^m satisfies (IC') and (IR) for a joint distribution over action profiles, $y^{(i),m} : \Theta \rightarrow \Delta(A)$, induced by $y_{-i}^{(i),m} \in \Delta(\mathcal{A}_{-i}(\theta_{-i}))$, for each $\theta_{-i} \in \Theta_{-i}$ (which is the distribution over acceptors' action profiles following a genuine rejection by agent i) and by a best-response from the rejector, $\sigma_i^{(i),m} : \Theta \rightarrow \Delta(A_i)$.

Consider a subsequence $(y_{-i}^{(i),m})_{m \in M}$ that converges and denote the limit by $y_{-i}^{(i)}$. It is clear that $y_{-i}^{(i)} \in \Delta(\mathcal{A}_{-i}(\theta_{-i}))$, because the support of a sequence of distributions "cannot increase in the limit". Denote the limit of $\sigma_i^{(i),m}$ by $\sigma_i^{(i)}$. Since the best-response mapping is upper hemi-continuous, $\sigma_i^{(i)}$ is a best-response to $y_{-i}^{(i)}$. Since there is participation and truth-telling along the sequence, there is also participation and truth-telling in the limit: x satisfies (IC) and (IR) when rejection leads to $(y_{-i}^{(i)}, \sigma_i^{(i)})$.

($\Leftarrow$) Since $y_{-i}^{(i)}(\theta_{-i}) \in \Delta(\mathcal{A}_{-i}(\theta_{-i}))$, $\forall \theta_{-i} \in \Theta_{-i}$, from Lemma 1, $\exists y^i \in BCE^I(\mu^i)$, where μ^i is any common belief with full support, s.t., for each $j \in I_P \setminus \{i\}$: $a_j \in \text{supp}[y_j^{(i)}(\theta_{-i})]$ implies $a_j \in \text{supp}(y_j^i)$.

Construct a trembling device such that if the input message is $\theta \in \Theta$, a spurious rejection by agent $i \in I$ is generated with probability $\varepsilon^i(\theta) \equiv \frac{\mu^i(\theta)\epsilon^m}{\mu^0(\theta)|I|}$, where $0 < \epsilon^m < \min_{\theta \in \Theta} \mu^0(\theta)$. Following a spurious rejection by agent i , the disagreement belief is μ^i , and the mediator recommends $y^i(\theta)$. Observe that all actions in the support of $y_j^{(i)}(\theta_{-i})$ are chosen with positive probability by player j when he has type θ_j , after a spurious rejection by agent i . After a genuine rejection by agent i , let the mediator send recommendations to the acceptors according to $y_{-i}^{(i)}(\theta_{-i})$. These recommendations are obeyed, because acceptors $j \in I_P \setminus \{i\}$ believe that a spurious rejection has occurred and the recommendations are made according to $y^i \in BCE^I(\mu^i)$.

Since x satisfies (IC) and (IR) for the given $y^{(i)}$, we can define (in the usual way) a sequence of pairs, $(x^m, \epsilon^m)_{m \in \mathbb{N}}$, such that (IC') and (IR) are satisfied along the sequence for the given $y^{(i)}$. Consider a candidate sequential equilibrium in which agents always participate and report truthfully. As we have seen, it is optimal for players to obey the recommendations made after a spurious rejection, and, therefore, acceptors also obey the recommendations made after a genuine rejection (believing that the rejection has been spurious, which is infinitely more likely). After a genuine rejection, the rejector anticipates that the acceptors will play $y_{-i}^{(i)} : \Theta_{-i} \rightarrow \Delta(A_{-i})$, and plays a best-response, $\sigma_i(\theta_i)$. This means that the distribution over action profiles is in fact induced by $(y_{-i}^{(i)}, \sigma_i^{(i)})$ such that $y_{-i}^{(i)} \in \Delta(\mathcal{A}_{-i}(\theta_{-i}))$, $\forall \theta_{-i} \in \Theta_{-i}$, and $\sigma_i^{(i)}(\theta_i)$ is a best-response to $y_{-i}^{(i)}$, $\forall \theta_i \in \Theta_i$. Since (x^m, ϵ^m) satisfies (IC') and (IR), no deviation from the candidate sequential equilibrium is beneficial. Letting, $\epsilon^m \rightarrow 0$, the mechanism converges to a non-trembling mechanism. $\square$

References

- Abreu, D. and Matsushima, H. (1992). Virtual implementation in iteratively undominated strategies: complete information. *Econometrica*, 60(5):993–1008.
- Banks, J. S. and Sobel, J. (1987). Equilibrium selection in signaling games. *Econometrica*, 55(3):647–661.
- Bergemann, D. and Morris, S. (2013). Robust predictions in games with incomplete information. *Econometrica*, 81(4):1251–1308.
- Bergemann, D. and Morris, S. (2016). Bayes correlated equilibrium and the comparison of information structures in games. *Theoretical Econ*, 11(2):487–522.
- Bester, H. and Strausz, R. (2001). Contracting with imperfect commitment and the revelation principle: the single agent case. *Econometrica*, 69(4):1077–1098.
- Bester, H. and Strausz, R. (2007). Contracting with imperfect commitment and noisy communication. *J Econ Theory*, 136(1):236–259.
- Blume, A., Board, O. J., and Kawamura, K. (2007). Noisy talk. *Theoretical Econ*, 2(4):395–440.
- Blume, L., Brandenburger, A., and Dekel, E. (1991). Lexicographic probabilities and choice under uncertainty. *Econometrica*, 59(1):61–79.
- Board, S. and Pycia, M. (2014). Outside options and the failure of the coase conjecture. *American Econ Rev*, 104(2):656–671.
- Calzolari, G. and Pavan, A. (2006a). Monopoly with resale. *RAND J Econ*, 37(2):362–375.
- Calzolari, G. and Pavan, A. (2006b). On the optimality of privacy in sequential contracting. *J Econ Theory*, 130(1):168–204.
- Calzolari, G. and Pavan, A. (2008). On the use of menus in sequential common agency. *Games Econ Behavior*, 64(1):329–334.
- Celik, G. and Peters, M. (2011). Equilibrium rejection of a mechanism. *Games Econ Behavior*, 73(2):375–387.
- Che, Y.-K. and Gale, I. (1998a). Standard auctions with financially constrained bidders. *Rev Econ Studies*, 65(1):1–21.
- Che, Y.-K. and Gale, I. L. (1998b). Caps on political lobbying. *American Econ Rev*, 88(3):643–651.
- Che, Y.-K. and Kim, J. (2006). Robustly collusion-proof implementation. *Econometrica*, 74(4):1063–1107.
- Che, Y.-K. and Kim, J. (2009). Optimal collusion-proof auctions. *J Econ Theory*, 144(2):565–603.
- Chen, C.-L. and Tauman, Y. (2006). Collusion in one-shot second-price auctions. *Econ Theory*, 28(1):145–172.
- Cho, I.-K. and Kreps, D. M. (1987). Signaling games and stable equilibria. *Quarterly J Econ*,

102(2):179–221.

- Cho, I.-K. and Sobel, J. (1990). Strategic stability and uniqueness in signaling games. *J Econ Theory*, 50(2):381–413.
- Cooter, R. D. and Rubinfeld, D. L. (1989). Economic analysis of legal disputes and their resolution. *J Econ Literature*, 27(3):1067–1097.
- Cramton, P. C. and Palfrey, T. R. (1990). Cartel enforcement with uncertainty about costs. *Int Econ Rev*, 31(1):17–47.
- Cramton, P. C. and Palfrey, T. R. (1995). Ratifiable mechanisms: learning from disagreement. *Games Econ Behavior*, 10(2):255–283.
- Crawford, V. P. and Sobel, J. (1982). Strategic information transmission. *Econometrica*, 50(6):1431–1451.
- Deb, R. and Said, M. (2015). Dynamic screening with limited commitment. *J Econ Theory*, 159:891–928.
- Dequiedt, V. (2007). Efficient collusion in optimal auctions. *J Econ Theory*, 136(1):302–323.
- Einy, E., Haimanko, O., and Tumendemberel, B. (2012). Continuity of the value and optimal strategies when common priors change. *Int J Game Theory*, 41(4):829–849.
- Eső, P. and Schummer, J. (2004). Bribing and signaling in second price auctions. *Games Econ Behavior*, 47(2):299–324.
- Evangelista, F. S. and Raghavan, T. (1996). A note on correlated equilibrium. *Int J Game Theory*, 25(1):35–41.
- Farrell, J. (1993). Meaning and credibility in cheap-talk games. *Games Econ Behavior*, 5(4):514–531.
- Faure-Grimaud, A., Laffont, J.-J., and Martimort, D. (2003). Collusion, delegation and supervision with soft information. *Rev Econ Studies*, 70(2):253–279.
- Forges, F. (1986). An approach to communication equilibria. *Econometrica*, 54(6):1375–1385.
- Forges, F. (1988). Communication equilibria in repeated games with incomplete information. *Mathematics of Operations Research*, 13(2):191–231.
- Forges, F. (1993). Five legitimate definitions of correlated equilibrium in games with incomplete information. *Theory and Decision*, 35(3):277–310.
- Forges, F. (2006). Correlated equilibrium in games with incomplete information revisited. *Theory and Decision*, 61(4):329–344.
- Fudenberg, D. and Tirole, J. (1986). A “signal-jamming” theory of predation. *RAND J Econ*, 17(3):366–376.
- Fudenberg, D. and Tirole, J. (1991). Perfect Bayesian equilibrium and sequential equilibrium. *J*

- Econ Theory*, 53(2):236–260.
- Gavious, A., Moldovanu, B., and Sela, A. (2002). Bid costs and endogenous bid caps. *RAND J Econ*, 33(4):709–722.
- Goltsman, M., Hörner, J., Pavlov, G., and Squintani, F. (2009). Mediation, arbitration and negotiation. *J Econ Theory*, 144(4):1397–1420.
- Govindan, S. and Wilson, R. (2005). Refinements of Nash equilibrium. In Durlauf, S. and Blume, L., editors, *The New Palgrave Dictionary of Economics*. Palgrave Macmillan, London, 2nd edition.
- Grossman, S. J. and Perry, M. (1986). Perfect sequential equilibrium. *J Econ Theory*, 39(1):97–119.
- Hart, O. and Tirole, J. (1990). Vertical integration and market foreclosure. *Brookings papers on economic activity: Microeconomics*, 1990:205–286.
- Hörner, J., Morelli, M., and Squintani, F. (2015). Mediation and peace. *Rev Econ Studies*, 82(4):1483–1501.
- Jullien, B. (2000). Participation constraints in adverse selection models. *J Econ Theory*, 93(1):1–47.
- Jullien, B., Pouyet, J., and Sand-Zantman, W. (2014). An offer you can’t refuse: Early contracting with endogenous threat. *mimeo*.
- Kandori, M. (1991). Cooperation in finitely repeated games with imperfect private information. *mimeo*.
- Kandori, M. (2002). Introduction to repeated games with private monitoring. *J Econ Theory*, 102(1):1–15.
- Kihlstrom, R. and Vives, X. (1992). Collusion by asymmetrically informed firms. *J Econ & Management Strategy*, 1(2):371–396.
- Kohlberg, E. and Mertens, J.-F. (1986). On the strategic stability of equilibria. *Econometrica*, 54(5):1003–1037.
- Kreps, D. M. and Wilson, R. (1982). Sequential equilibria. *Econometrica*, 50(4):863–894.
- Laffont, J.-J. and Martimort, D. (1997). Collusion under asymmetric information. *Econometrica*, 65(4):875–911.
- Laffont, J.-J. and Martimort, D. (2000). Mechanism design with collusion and correlation. *Econometrica*, 68(2):309–342.
- Lehrer, E. (1992). Correlated equilibria in two-player repeated games with nonobservable actions. *Mathematics of Operations Research*, 17(1):175–199.
- Lehrer, E. and Sorin, S. (1997). One-shot public mediated talk. *Games Econ Behavior*, 20(2):131–

- Mailath, G. J., Matthews, S. A., and Sekiguchi, T. (2002). Private strategies in finitely repeated games with imperfect public monitoring. *Contributions in Theoretical Econ*, 2(1).
- Mailath, G. J., Okuno-Fujiwara, M., and Postlewaite, A. (1993). Belief-based refinements in signalling games. *J Econ Theory*, 60(2):241–276.
- Martimort, D. and Sand-Zantman, W. (2016). A mechanism design approach to climate-change agreements. *J European Econ Association*, 14(3):669–718.
- McAfee, R. P. and Schwartz, M. (1994). Opportunism in multilateral vertical contracting: Nondiscrimination, exclusivity, and uniformity. *American Econ Rev*, 84(1):210–230.
- Mertens, J.-F. (1989). Stable equilibria – a reformulation: Part I. Definition and basic properties. *Mathematics of Operations Research*, 14(4):575–625.
- Mertens, J.-F. (1991). Stable equilibria – a reformulation: Part II. Discussion of the definition, and further results. *Mathematics of Operations Research*, 16(4):694–753.
- Mertens, J.-F., Sorin, S., and Zamir, S. (2015). *Repeated games*. Cambridge University Press.
- Mitusch, K. and Strausz, R. (2005). Mediation in situations of conflict and limited commitment. *J Law, Econ, and Organization*, 21(2):467–500.
- Mookherjee, D. and Tsumagari, M. (2004). The organization of supplier networks: effects of delegation and intermediation. *Econometrica*, 72(4):1179–1219.
- Myerson, R. B. (1982). Optimal coordination mechanisms in generalized principal–agent problems. *J Mathematical Econ*, 10(1):67–81.
- Myerson, R. B. (1986). Multistage games with communication. *Econometrica*, 54(2):323–358.
- Nau, R. F. and McCardle, K. F. (1990). Coherent behavior in noncooperative games. *J Econ Theory*, 50(2):424–444.
- Olszewski, W. and Siegel, R. (2016). Bid caps in contests. *mimeo*.
- Pavan, A. and Calzolari, G. (2009). Sequential contracting with multiple principals. *J Econ Theory*, 144(2):503–531.
- Pavlov, G. (2008). Auction design in the presence of collusion. *Theoretical Econ*, 3(3):383–429.
- Philippon, T. and Skreta, V. (2012). Optimal interventions in markets with adverse selection. *American Econ Rev*, 102(1):1–28.
- Rabin, M. and Sobel, J. (1996). Deviations, dynamics, and equilibrium refinements. *J Econ Theory*, 68(1):1–25.
- Rachmilevitch, S. (2013). Endogenous bid rotation in repeated auctions. *J Econ Theory*, 148(4):1714–1725.
- Rahman, D. (2012). But who will monitor the monitor? *American Econ Rev*, 102(6):2767–2797.

- Rahman, D. (2014). The power of communication. *American Econ Rev*, 104(11):3737–3751.
- Rahman, D. and Obara, I. (2010). Mediated partnerships. *Econometrica*, 78(1):285–308.
- Renault, J. and Tomala, T. (2004). Communication equilibrium payoffs in repeated games with imperfect monitoring. *Games Econ Behavior*, 49(2):313–344.
- Rey, P. and Vergé, T. (2004). Bilateral control with vertical contracts. *RAND J Econ*, 35(4):728–746.
- Schummer, J. (2000). Manipulation through bribes. *J Econ Theory*, 91(2):180–198.
- Segal, I. (1999). Contracting with externalities. *Quarterly J Econ*, 114(2):337–388.
- Selten, R. (1975). Reexamination of the perfectness concept for equilibrium points in extensive games. *Int J Game Theory*, 4(1):25–55.
- Singh, N. and Vives, X. (1984). Price and quantity competition in a differentiated duopoly. *RAND J Economics*, 15(4):546–554.
- Strausz, R. (2012). Mediated contracts and mechanism design. *J Econ Theory*, 147(3):1280–1290.
- Sugaya, T. and Wolitsky, A. (2016). Bounding equilibrium payoffs in repeated games with private monitoring. *Theoretical Econ*.
- Szech, N. (2015). Tie-breaks and bid-caps in all-pay auctions. *Games Econ Behavior*, 92:138–149.
- Tan, G. and Yilankaya, O. (2007). Ratifiability of efficient collusive mechanisms in second-price auctions with participation costs. *Games Econ Behavior*, 59(2):383–396.
- Tirole, J. (1986). Hierarchies and bureaucracies: On the role of collusion in organizations. *J Law, Econ, and Organization*, 2(2):181–214.
- Tirole, J. (1992). Collusion and the theory of organizations. In Laffont, J.-J., editor, *Advances in Economic Theory: Sixth World Congress*, volume II. Cambridge University Press.
- Tirole, J. (2012). Overcoming adverse selection: How public intervention can restore market functioning. *American Econ Rev*, 102(1):29–59.
- Tomala, T. (2009). Perfect communication equilibria in repeated games with imperfect monitoring. *Games Econ Behavior*, 67(2):682–694.
- van Damme, E. (1989). Stable equilibria and forward induction. *J Econ Theory*, 48(2):476–496.
- Yamashita, T. (2014). Strategic and structural uncertainties in robust implementation. *mimeo*.