Communication with Self-Interested Experts Part II: Models of Cheap Talk

Margaret Meyer

Nuffield College, Oxford

2020

Setting: Decision-maker (P) receives advice from an advisor (A) who possesses useful information but whose preferences over possible decisions don't match those of P. After receiving advice, P makes a decision.

Examples:

1. manager consulting supervisor about an employee
2. headquarters consulting a division manager about an investment project
3. investor evaluating information disclosed by a firm's mgt.
4. politician listening to a lobbyist representing an interest group
5. customer asking a salesman about a product's quality.

Maintained assumptions: Decision-maker *cannot commit* in advance as to how will use advisor's advice and hence cannot offer monetary transfers contingent on the advice. Contrast mechanism design.

NB: Signaling and signal-jamming models can also be used to study communication, but in these, emphasis is on *effort expended* as well as on information transmitted.

Cheap-Talk Models of Strategic Communication

In cheap-talk models, the privately-informed advisor (A) can *costlessly* make *any* report to the decision-maker (P) about the state of the world, and P has no way of verifying A's report.

- seminal paper: Crawford and Sobel, *Etrica*, 1982.
- widely applied to study "expert" advice (e.g. lobbying, legislative committee recommendations, financial advice to consumers) in settings where legal remedies for misleading advice are costly or unavailable

For cheap talk to be informative in equilibrium with fully rational and self-interested parties, some commonality of interests is necessary. In particular, A's ranking over P's decisions cannot be state-independent and monotonic.

Crawford and Sobel model: uniform-quadratic special case

$u^P(V, x) = -(V - x)^2$ $\quad$ P's ideal decision: $V^P(x) = x$
$u^A(V, x) = -(V - (x + b))^2$ $\quad$ A's ideal decision: $V^A(x) = x + b$

The parameter $b > 0$, which is common knowledge, measures divergence in preferences

- b is usually referred to as A's bias (relative to P's preferences)

State of the world $x \sim U[0, 1]$

Timing: i) A privately learns x; ii) A sends a costless but unverifiable report m to P; iii) P updates beliefs about x and then chooses a decision V.

Crawford and Sobel model: uniform-quadratic special case

Claim: $\exists$ a fully-pooling PBE, in which A's message reveals nothing about x.

- Such a "babbling eqm." always exists in **any** cheap talk game.

Claim: For any $b > 0$, a fully revealing PBE does not exist.

- If P expected A's message to be fully revealing, A would wish to deviate and mislead P. A could induce his ideal V by overstating x by b.

Proposition: *All PBE are payoff-equivalent to partially-pooling eqa in which, for some $n \geq 1$, n distinct messages $(m_1, m_2, \ldots, m_n)$ are sent, such that*

$$x \in [0, x_1) \Rightarrow m_1, \quad x \in [x_1, x_2) \Rightarrow m_2, \quad \ldots, \quad x \in [x_{n-1}, 1] \Rightarrow m_n.$$

Thus, in eqm., some communication may occur, but it is coarse.

- Coarseness allows informative communication in eqm. because a change in A's message induces a **large** change in P's response.

Example: PBE with two messages: $m = L$ if $x \in [0, x_1)$ and $m = H$ if $x \in [x_1, 1]$

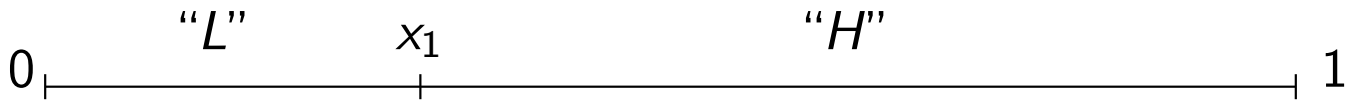

- The eqm. value of x_1 is the value of x at which A is indifferent between reporting "L" and "H", given that P expects A to use the strategy above.
- Suppose P conjectures that A reports "L" if and only if $x < \hat{x}_1$. Then P's optimal response to "L" is $V(L) = \frac{\hat{x}_1}{2}$, and to "$H$" is $V(H) = \frac{(1+\hat{x}_1)}{2}$.
- Then, given the symmetry of A's utility function around his ideal decision $x + b$, the value of x, $X^*(\hat{x}_1)$, at which A is indifferent between reporting "L" and "H" solves $X^*(\hat{x}_1) + b = \frac{1}{2}\left[\frac{\hat{x}_1}{2} + \frac{(1+\hat{x}_1)}{2}\right]$.
- The eqm value of x_1, x_1^*, solves the fixed-point condition $x_1^* = X^*(x_1^*)$.
- Hence $x_1^* = \frac{1}{2} - 2b$, so a two-message eqm. exists iff $b \in (0, \frac{1}{4})$. If $b \geq \frac{1}{4}$, then any PBE is payoff equivalent to the "babbling " one.

Example: PBE with two messages: $m = L$ if $x \in [0, x_1)$ and $m = H$ if $x \in [x_1, 1]$

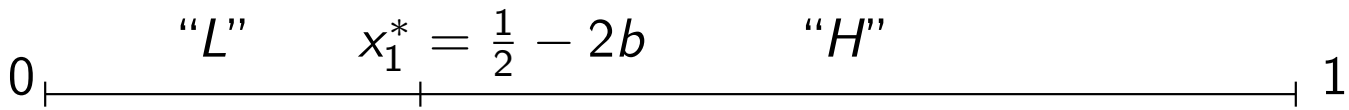

The interval of values for which A reports "H" is larger, by an amount $4b$, than the interval for which he reports "L".

- The greater "coarseness" of the "H" report is the endogenous "signaling cost" borne by A that disciplines A to report "H" only for those states of the world where P expects him to do so.
- Same reasoning shows that in all eqa, the partition element corres. to message m_k is larger, by $4b$, than that corres. to message m_{k-1}, $\forall k > 1 \Longrightarrow$ communication gets coarser as the state moves in direction of A's bias.
- With A restricted to two reports, P would be better off if A used cutoff $x_1 = \frac{1}{2}$. A would also be better off (ex ante) if could commit to this cutoff.
  - But A using cutoff $x_1 = \frac{1}{2}$ would not constitute a PBE, because if P expected this, A would strictly prefer to report "H" for a larger set of states than $(\frac{1}{2}, 1]$.

Properties of PBE's in Crawford-Sobel model (uniform-quadratic case)
In a PBE with n distinct messages $(m_1, m_2, \ldots, m_n)$,

$$x \in [0, x_1) \Rightarrow m_1, \quad x \in [x_1, x_2) \Rightarrow m_2, \quad \ldots, \quad x \in [x_{n-1}, 1] \Rightarrow m_n,$$

and $\forall k > 1$, $(x_k - x_{k-1}) = (x_{k-1} - x_{k-2}) + 4b$. Hence such a PBE must satisfy $nx_1 + 4b\frac{n(n-1)}{2} = 1$, so $x_1 = \frac{1}{n} - 2b(n-1)$, which must be positive.

Proposition: *Given b, define $n^*(b)$ as max. number of messages in a PBE.*

1. $\forall\, n \in \{1, \ldots, n^*(b)\}$, $\exists$ *a unique partially pooling PBE of the above form.*
2. $n^*(b) \downarrow$ *in b: more communication can occur when preferences more closely aligned*
3. $\lim_{b\to 0}\ n^*(b) = \infty$ *and* $\lim_{b\to 0}$ [*P's exp. loss relative to full infor.*] $= 0$
4. $\forall b > 0, n^*(b) < \infty$: *for all $b > 0$, full revelation is impossible in a PBE*
5. $n^*(b)$ *is the largest integer n s.t. $2bn(n-1) < 1$, i.e. the largest integer less than* $\frac{1}{2}[1 + \sqrt{1 + \frac{2}{b}}]$.

Welfare analysis in Crawford and Sobel model (uniform-quadratic case)

Given a PBE, let σ_x^2 denote the expected posterior variance of x from P's perspective. Then ex ante expected payoffs are

$$EU^P = -\sigma_x^2 \qquad \text{and} \qquad EU^A = -(\sigma_x^2 + b^2)$$

(From A's pt. of view, P's decision is biased downward by b.)

Proposition:

1. *Given b, P's and A's ex ante expected payoffs are both $\uparrow$ in n.*
2. *Given n, P's payoff is $\downarrow$ in b.*
3. *For $b_0 > b_1$, $EU^P(b_0, n^*(b_0)) < EU^P(b_1, n^*(b_0)) \leq EU^P(b_1, n^*(b_1))$.*

Thus, in this analysis of cheap talk with a one-dimensional state and action space, both the maximal "amount" of information communicated, as measured by $n^*(b)$, and its value to the principal increase as the divergence in preferences, as measured by b, decreases.

With players fully rational and self-interested, as we have been assuming, some commonality of interests is necessary in order for cheap talk to be informative. In particular, A's ranking over P's decisions cannot be state-independent and monotonic.

1. Watson (*GEB*, 96) shows that even with no commonality of interests, if A is "confused" about what his private information means (i.e. uncertain how P would respond to it), then informative cheap-talk equilibria may exist.
   - Formally, Watson's model is one with fully rational players but two-sided private information.
2. Lipnowski and Ravid (*Ecta*, 20) show that there may exist informative cheap-talk equilibria, even with fully rational players, when A's payoff from P's decisions is state-independent but violates quasi-concavity.
3. Crawford (*AER*, 03), allowing for **bounded** strategic rationality, shows that even in a zero-sum game, cheap talk can be used for strategic advantage.
4. Some experiments also find a greater role for cheap talk than predicted by equilibrium analysis with fully rational, self-interested players.

Communication vs. delegation in the uniform-quadratic case

Dessein (*REStud*, 2002) asks whether a principal in the Crawford-Sobel environment might prefer (assuming it were possible) to *delegate* the decision to the agent, rather than *communicate* with him.

- For delegation to be credible, P must be able to *commit* not to overturn or modify A's decision.
- Under delegation, $\forall x \in [0,1]$, A chooses his own ideal point, that is, $V = x + b$. Thus $\forall x$, P's loss relative to his full-information optimum is b^2.
- If, rather than delegate, P chooses on the basis of only his prior information, his loss relative to full information is $\frac{1}{12} = var(x)$ for $x \sim U[0,1]$.
- Hence, for $b < \frac{1}{\sqrt{12}}$, P prefers to delegate rather than choose based only on his prior.
- But for $b \geq \frac{1}{4}$, communication produces a babbling eqm, in which P chooses based only on his prior.
- *Therefore, for* $b \in [\frac{1}{4}, \frac{1}{\sqrt{12}})$, P *prefers delegation to communication.*

Communication vs. delegation in the uniform-quadratic case

Proposition: *For all* $b \in (0, \frac{1}{\sqrt{12}})$, *delegation yields higher expected payoff for* P *than does the most informative PBE under communication*.

- Eqm. with communication yields a partition of $[0, 1]$ in which partition elements increase in size with x. This increase in size is harmful to P because it limits the rate of increase of $n^*(b)$ as b decreases.
- Define $\bar{A}(b) \equiv \frac{1}{n^*(b)}$ as the *average* interval length in the most informative PBE. P's expected loss from communication is (roughly) proportional to $(\bar{A}(b))^2$, while expected loss from delegation is b^2.
- Since $n^*(b)$ is the largest integer less than $\frac{1}{2}[1 + \sqrt{1 + \frac{2}{b}}]$, $(\bar{A}(b))^2 \to 0$ at rate b as $b \to 0$. Hence

$$\lim_{b \to 0} \frac{(\bar{A}(b))^2}{b^2} = \infty.$$

- Implication: As $b \to 0$, communication performs infinitely badly relative to delegation.

Communication with two experts

If P communicates with two experts, can he extract all of their infor., even without being able to commit, by using each one as a check on the other's truthfulness?

- Suppose the agents choose their messages simultaneously.
- Suppose the two agents both perfectly observe the state, so, crucially, their private information is perfectly correlated.
- *When does there exist a fully revealing PBE (henceforth, a "FRE")? In a FRE, for each state x, P learns his best feasible decision (henceforth, "decision" = "policy").*
- If a FRE exists, so does a truthful FRE, i.e. a FRE in which each agent reports P's best feasible policy truthfully.
- In a truthful FRE, if agents' reports agree, P chooses the policy specified by their reports.
- In a truthful FRE, if the agents' reports disagree, then P knows that one of the agents has deviated. But he does not know which one.

Communication with two identically-informed experts

In a truthful FRE, if the agents' reports disagree, then P knows that one of the agents has deviated. But he does not know which one.

Key question is: For each possible pair of non-matching reports, (x', x'') with $x' \neq x''$, is there a belief that P could hold, about his best feasible policy, s.t. his optimal response would be a "punishment" $V(x', x'')$ that would deter each A from deviating so as to generate this pair of reports?

- That is, does there exist a function $V(x', x'')$ such that,
  1. in state x', A_2 weakly prefers to report x' (thus inducing P to choose $V = x'$) to deviating to reporting x'' (thus inducing P to choose $V(x', x'')$); **and**
  2. in state x'', A_1 weakly prefers to report x'' (thus inducing P to choose $V = x''$) to deviating to reporting x' (thus inducing P to choose $V(x', x'')$)?

- *Since, given non-matching reports, P does not know which A deviated, the* **same** *"punishment" V must work whichever A deviated.*

Communication with two experts, one-dimensional policy space

Suppose the set of possible decisions for P ("policy space") and the set of states of the world ("state space) are **one-dimensional** (as in Crawford-Sobel model)

- One-dimensional models: Gilligan and Krehbiel (*AJPS*, 89), Krishna and Morgan (*QJE*, 01 and *APSR*, 01), Battaglini (*Etrica*, 02)

$u^P(V,x) = -(V-x)^2$ $\quad$ P's ideal policy: $V^P(x) = x$

$u^i(V,x) = -(V-(x+b^i))^2$ $\quad$ A_i's ideal policy: $V^i(x) = x + b^i$

State of the world $x \sim F$ on $[0,1]$; $b^i \neq 0$ is A_i's bias; b^1, b^2 are common knowledge

Crucial distinction: Do agents' biases have the same or opposite signs?

Communication with two experts, one-dimensional policy space

Crucial distinction: Do agents' biases have the same or opposite signs?

Proposition: *With a one-dimensional policy space and simultaneous reports, if agents' biases have the* **same sign**, *then for any magnitudes* $|b^1|$, $|b^2|$, *there exists a FRE.*

- Suppose $b^1, b^2 > 0$. If $x' \neq x''$, let P choose $V(x', x'') = \min\{x', x''\}$.
  - $\exists$ an off-eqm path belief that makes this "punishment" rule optimal for P.
  - Neither A gains by deviating, if he expects the other A to report truthfully. Hence there exists a PBE with truthful reporting.
  - For all x', x'', P sets V equal to whichever of the two reports **both** agents would like less, if the magnitudes of their biases were arbitrarily large.
  - But both agents slightly overstating the value of x would also be self-enforcing for them. So this FRE fails to be "collusion-proof".

Communication with two experts, one-dimensional policy space

Proposition: *With a one-dimensional policy space and simultaneous reports, if agents' biases have* **opposite signs**, *then a FRE exists if and only if* $|b^1| + |b^2|$ *is not too large relative to the size of the policy space.*

- With opposing biases, deviations from truthful reporting can be deterred only if P can respond to a discrepancy in the agents' reports by choosing a policy that is far enough away from both reports.
  - The larger are $|b^1|$ and $|b^2|$, the farther away from the reports P's response must be.
  - But the size of the policy space limits how far away this "punishment" imposed by P following a deviation can be.
  - If $\exists$ a FRE with opposing biases, it is collusion-proof.

  With opposing biases, just as in the Crawford-Sobel model, the *magnitudes* of preference divergences (btw. P and As) crucially determine how much information is communicated in eqm.

Communication with two experts, one-dimensional policy space

With opposing biases, if a FRE exists, deviations from truthful reporting are deterred by P responding to a discrepancy in the agents' reports by choosing a policy that is *far enough away* from both reports.

- But the out-of-eqm beliefs that support such a "punishment" rule are not very plausible: Small discrepancies in reports might be due to unanticipated mistakes by As, so P might want to respond by choosing an action *close to both reports*.
  - In contrast, with both biases positive, the beliefs supporting the "punishment" rule $V(x', x'') = \min\{x', x''\}$ are *not* vulnerable to this criticism.

- Battaglini (2002) proposed that one should focus on FRE that satisfy a **"robustness" criterion**, capturing the requirement that *small* deviations from eqm. (or small mistakes) cause only *small* punishments.

Battaglini (2002) shows that existence of a robust FRE is dramatically more likely **when the policy space and state space have two (or more) dimensions**.

Communication with two experts, multidimensional policy/state space

Battaglini (2002) shows that existence of a robust FRE is dramatically more likely **when the policy space and state space have two (or more) dimensions**.

- Let both the policy space and the state space be $\mathbb{R}^2$
- State is $x \equiv (x_1, x_2)$, and P's policy choice is $V \equiv (V_1, V_2)$
- $u^P(V, x) = -(V_1 - x_1)^2 - (V_2 - x_2)^2$, so P's ideal policy is $V^P(x) = x$
- $u^i(V, x) = -(V_1 - (x_1 + b_1^i))^2 - (V_2 - (x_2 + b_2^i))^2$, for $i = 1, 2$, so A_i's ideal policy is $V^i(x) = x + b^i$
  - The **vectors** $b^1 = (b_1^1, b_2^1)$ and $b^2 = (b_1^2, b_2^2)$ represent how much, and in which direction, the agents' preferences diverge from P's.
  - A_i's indifference curves are circles centered at $x + b^i$.

Proposition: *Whenever the vectors b^1 and b^2 are linearly independent, and regardless of their magnitudes, there exists a fully revealing, robust PBE.*

Communication with two experts, multidimensional policy/state space

Proposition: *Whenever the vectors b^1 and b^2 are linearly independent, and regardless of their magnitudes, there exists a fully revealing, robust PBE.*

- Consider first the special case where b^1 and b^2 are *orthogonal* to each other: Let x_i represent the efficient level of investment in division i and V_i represent the actual level of investment in division i. Assume there is no upper limit on total investment.
  - Let $b^1 = (k, 0)$ and $b^2 = (0, m)$, where k, m are known, positive constants.
  - A_i's prefs. differ from P's only w.r.t. the ith dimension, i.e. the choice of V_i: each A is biased towards more investment in his own division.
- There is a fully revealing, robust PBE in which P asks each agent how much he should invest in the **other** agent's division and follows the agents' advice.
  - Since all pairs of reports arise on the eqm path, out-of-eqm beliefs are irrelevant, so robustness holds.
  - In this special case with b^1 and b^2 *orthogonal*, it is a *dominant strategy* for each agent to tell the truth.

Communication with two experts, multidimensional policy/state space

Battaglini's eqm construction when b^1 and b^2 are *orthogonal*: P asks agent each agent how much he should invest in other division and follows the advice. It is a *dominant strategy* for each to tell the truth.

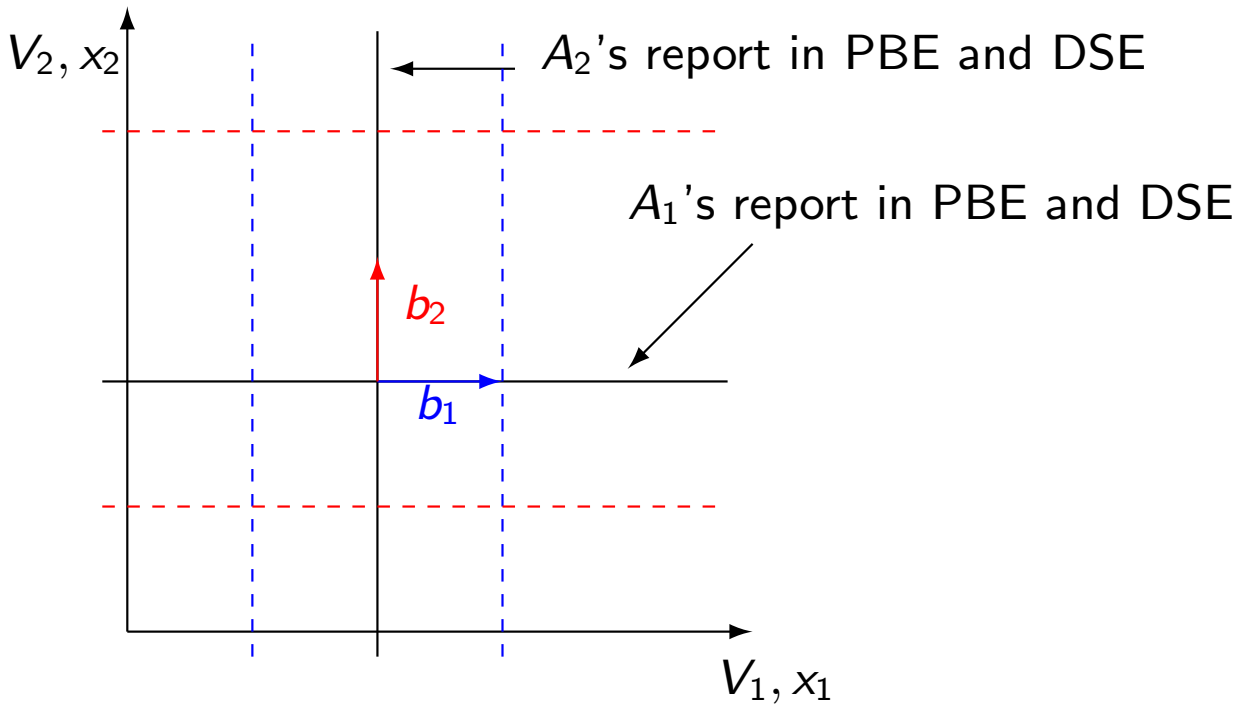

Battaglini's eqm construction for b^1, b^2 not orthogonal: P can still achieve full revelation as a PBE (though not a DSE). Essentially, P limits each agent's influence to a one-dimensional subspace over which the preferences of P and that agent coincide.

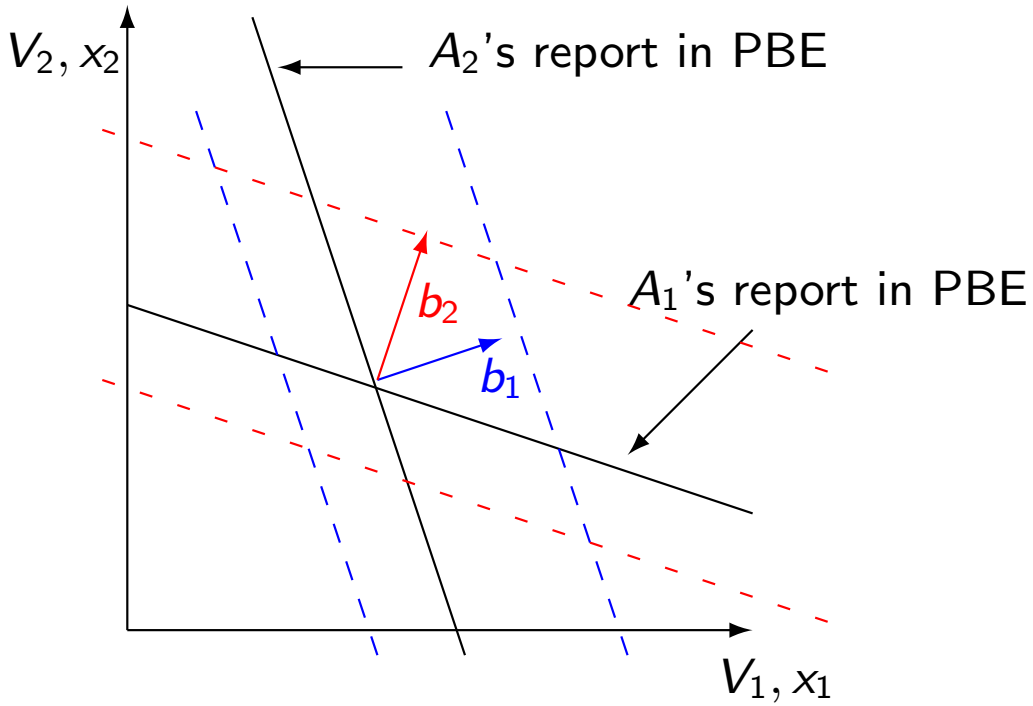

Message from Battaglini's construction: With **multidimensional** policy and state spaces, a fully revealing, robust PBE exists as long as agents' bias vectors are neither perfectly opposed nor perfectly aligned; **magnitudes** of the biases are irrelevant.

Assumptions underlying Battaglini's (2002) analysis are very strong:

1. Both agents are *perfectly* informed about *both* components of the state.
   - In investment ex., A_i might be better informed about own division.
2. There are *no constraints* on the policy space that force a tradeoff between the two dimensions.
   - In investment ex., a budget constraint might limit total investment.
3. The principal is *perfectly* informed about the *directions* of the agents' bias vectors.
   - Battaglini's construction fails if P is even slightly wrong about the directions of these biases.

Battaglini (2004) relaxes Assn. 1, allowing agents to observe only noisy versions of the state:

- He shows that full revelation of agents' information never arises in eqm. But eqm. outcomes approach the full-information benchmark as i) the noise in the agents' signals disappears or as ii) the number of agents increases.

Bounds on policy space: When does a FRE exist?

Ambrus + Takahashi (2008) and Meyer, Moreno, + Nafziger (2016; 2019) relax Assn. 2, analyzing when bounds on the policy space prevent existence of a FRE.

- Suppose support of distrib. of state x is $\mathbb{R}^2$, but the policy space, Y, is a **strict subset** of $\mathbb{R}^2$.
- Maintain same assumptions as above on preferences, information structure, and timing.
- In a FRE, for each state x, P learns from the agents his best *feasible* policy.
- With Y a strict subset of $\mathbb{R}^2$, a FRE can no longer be constructed by P limiting each agent's influence to a one-dimensional subspace: some pairs of reports are *incompatible*. So now P has each agent make a report about *all* dimensions of P's best feasible policy.

Bounds on policy space: When does a robust FRE exist?

Meyer, Moreno, and Nafziger (2016; 2019) focus on **robust FRE**, that is, FRE in which out-of-eqm. beliefs are such that *small* deviations from eqm. (or small mistakes) cause only *small* punishments.

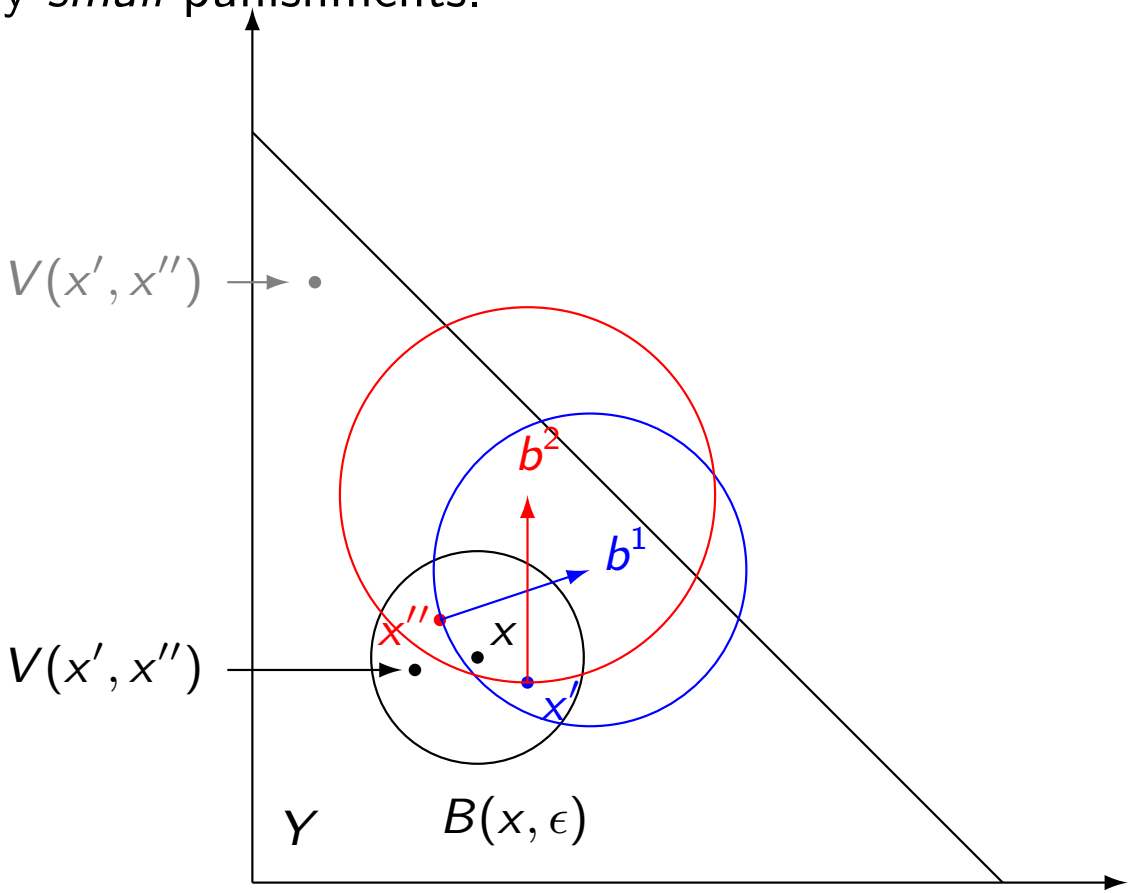

Bounds on policy space: When does a robust FRE exist?

Meyer et al seek a robust FRE that is also *independent of the magnitudes of the biases*.

- For given, finite magnitude of bias vector, an agent's indifference curves are circles.
- As magnitude of bias vector becomes arbitrarily large, an agent's indifference curve through a given point approaches a straight line.

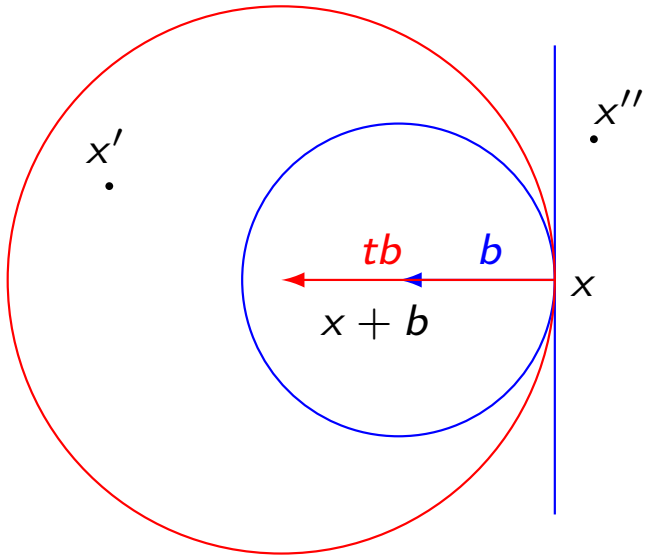

Bounds on policy space: When does a robust FRE exist?

Proposition (Meyer et al): *Assume that the set Y of feasible actions is convex. Whenever there exists a FRE that is independent of the magnitudes of the biases, there exists such an FRE that is robust.*

- When punishments must "work" no matter how large the biases, then for any small deviation, a small feasible punishment exists whenever a large one does.

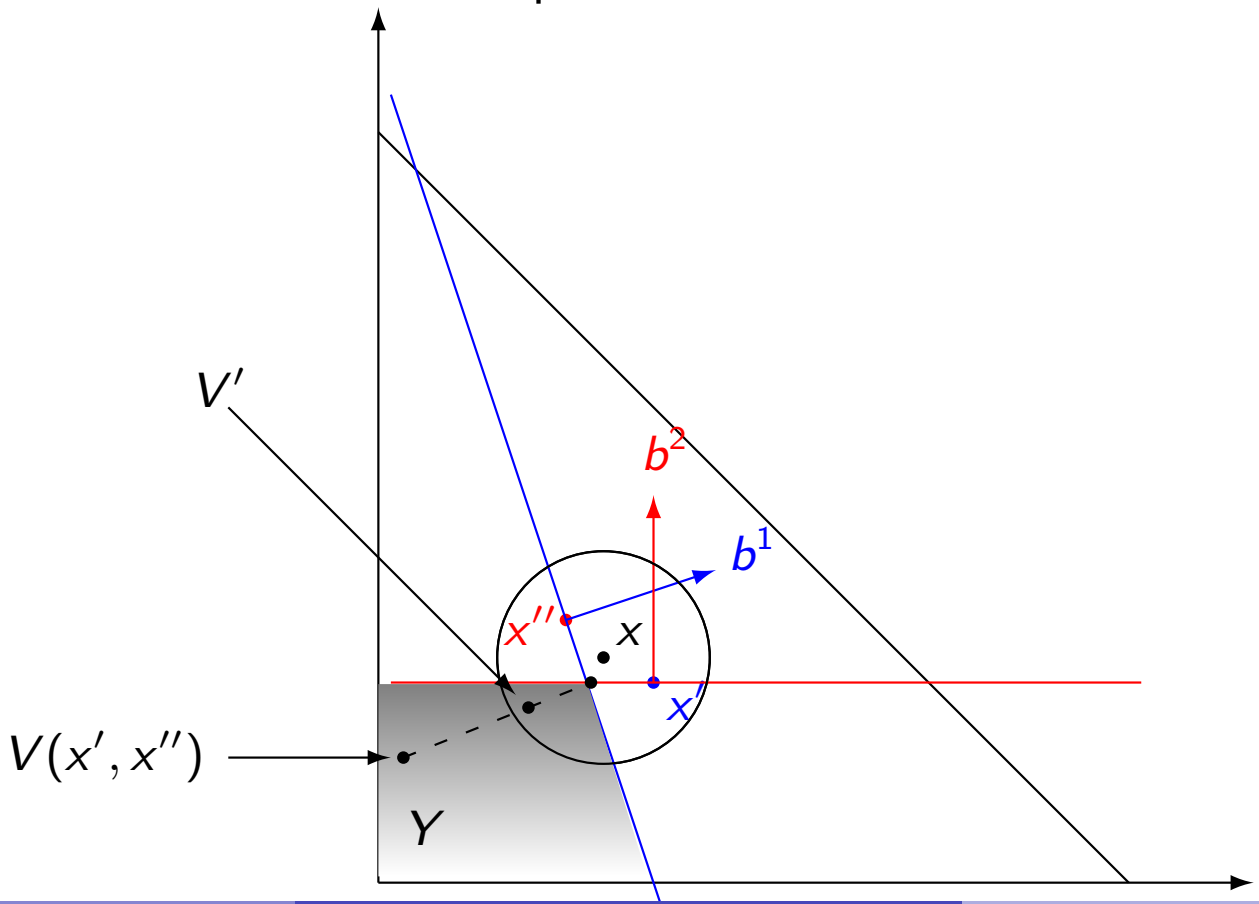

Bounds on policy space: When does a robust FRE exist?

Observation: If the bias vectors are *perfectly aligned*, then there exists a (robust) FRE that is independent of their magnitudes.

- For all x', x'', P sets V equal to whichever of the two reports **both** agents would like less, if the magnitudes of their biases were arbitrarily large.
- This "punishment rule" for P generalizes the rule $V(x', x'') = \min\{x', x''\}$ from the one-dimensional setting with both biases positive.
- This rule, like its one-dimensional version $V(x', x'') = \min\{x', x''\}$, is independent of the magnitudes of the biases.
- This rule satisfies robustness, since it selects one of the two reports.

Bounds on policy space: When does a robust FRE exist?

We need some definitions:

- Define $Fr(Y) \equiv$ frontier of Y, and $n_Y(x) \equiv$ *inward*-pointing normal vector to $Fr(Y)$ at a *smooth* point $x \in Fr(Y)$.
- Given b^1, b^2 linearly indep., $C(b^1, b^2) \equiv$ open convex cone spanned by b^1, b^2.
- Define (generalized) "Min Rule" for P to select the policy which is the **coordinate-wise minimum** of x' and x'', using the coordinate system formed by the normal vectors to the biases.
  - For orthogonal bias directions $b^1 = (1, 0)$, $b^2 = (0, 1)$, Min Rule selects $V(x', x'') = (\min\{x_1', x_1''\}, \min\{x_2', x_2''\})$.
  - Min Rule is independent of the magnitudes of the biases.
  - Min Rule punishes incompatible reports in the **least severe way** that deters deviations for all magnitudes of the biases and that does not depend on who sent which report.
  - Min Rule satisfies robustness, since for x', x'' close to each other, policy selected is close to both.

Bounds on policy space: When does a robust FRE exist?

Proposition (Meyer et al): *Given $Y \subset \mathbf{R}^2$ closed and convex and $b^1, b^2 \in \mathbf{R}^2$ linearly independent, the following statements are equivalent:*

(i) $\exists$ *a (robust) FRE that is independent of the magnitudes of the biases.*

(ii) $\forall$ *smooth points* $x \in Fr(Y)$, $n_Y(x) \notin C(b^1, b^2)$.

(iii) $\forall$ *pairs of reports* $x', x'' \in Y$, *the policy selected by the Min Rule is feasible.*

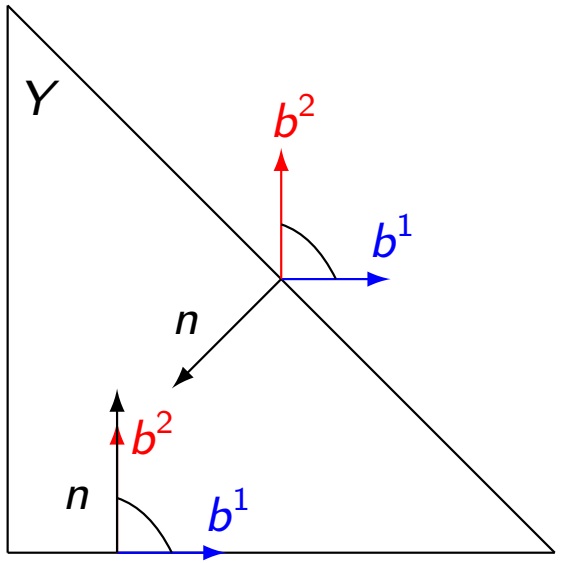

Example: Allocating goods (e.g. investment funds) subject to upper limit on total:

Here, for every smooth point $x \in Fr(Y)$, $n_Y(x) \notin C(b^1, b^2)$, so $\exists$ a robust FRE that is independent of the magnitudes of the biases.

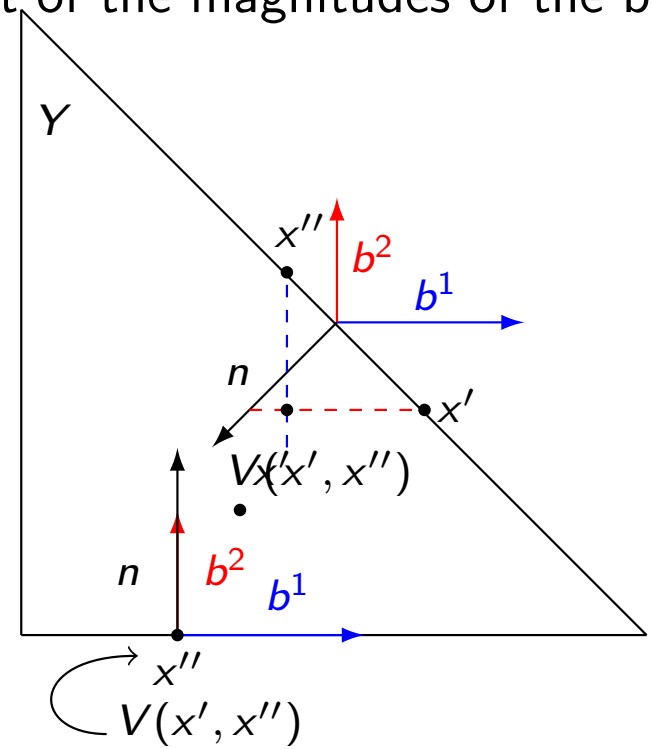

Here, the "Min Rule" $V(x', x'') \equiv (\min\{x_1', x_1''\}, \min\{x_2', x_2''\})$ is feasible $\forall\, x', x''$.

- Here, the Min Rule is also collusion-proof, in contrast to the one-dimensional setting.

Ex.: Allocating bads (e.g. budget cuts, pollution reductions) subj. to upper limit on total:

Here, $\exists x \in Fr(Y)$ s.t. $n_Y(x) \in C(b^1, b^2)$, so $\nexists$ a FRE.

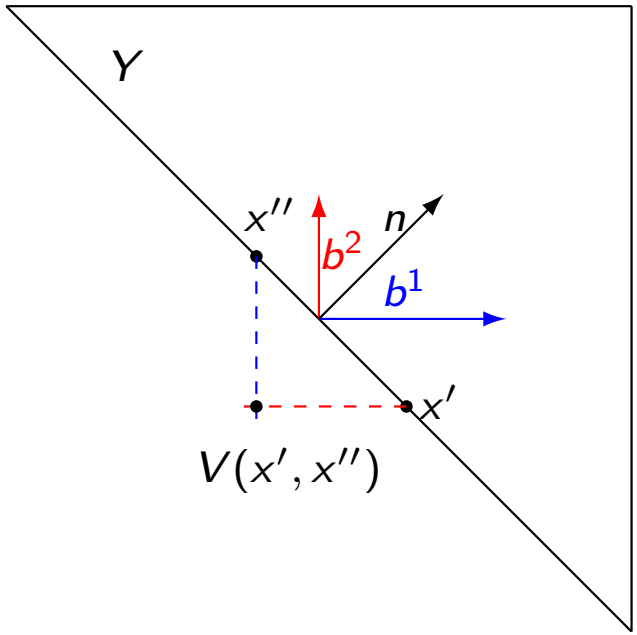

Here, for the deviation (x', x'') in the figure, the punishment selected by the Min Rule is not feasible.

Recall the strong assumptions underlying Battaglini's (2002) analysis:

1. Both agents are *perfectly* informed about *both* components of the state.
2. There are *no constraints* on the action space that force a tradeoff between the two dimensions.
3. **The principal is *perfectly* informed about the *directions* of the agents' bias vectors.**

Meyer et al's approach, in contrast to Battaglini's, can easily handle uncertainty on the principal's part about the directions of the agents' biases:

- The *least aligned* possible realizations of the biases, $(\underline{b}, \overline{b})$, are what matter for existence of a robust FRE that is independent of bias magnitudes.

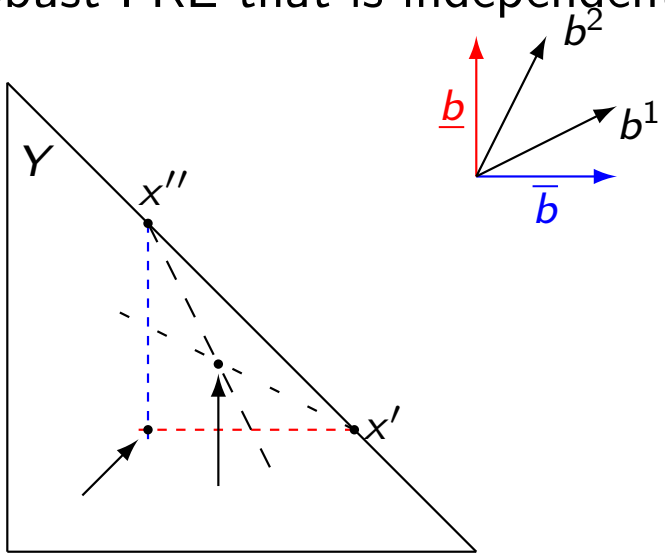

Dynamic Models of Communication with Self-Interested Experts

Suppose that self-interested, privately informed experts are consulted repeatedly. Does concern for their reputation help to discipline such experts?

- Applications: policy advisors; financial market insiders; managers; referees for job candidates

Usually studied with models in which decision-makers learn about experts' intrinsic characteristics.

What characteristics of experts might decision-makers want to learn about?

- How closely their preferences are aligned with those of decision-maker
  - Sobel, *REStud*, 85; Benabou + Laroque, *QJE*, 92; Morris, *JPE*, 01; Ely + Valimaki, *QJE*, 03; Avery + Meyer, W.P., 12
- Quality of their information
  - Scharfstein + Stein, *AER*, 90; Prendergast + Stole, *JPE*, 96; Ottaviani + Sorensen, *JET*, 06, and *Rand*, 06; Prat, *AER*, 05; Grubb, *JEMS*, 11

Dynamic Models of Communication with Self-Interested Experts

On what information are experts' reputations based?

Is more information necessarily better for disciplining the experts? **No**.

- Prat shows that an expert's incentives can be worse when both his action and its outcome are observed, compared to when only its outcome is observed.

- Morris, Ely+Valimaki, and Avery+Meyer show that an expert's incentives early in his career **may** be worse when the expert expects to be consulted repeatedly (so is concerned about his reputation) than when he expects to be consulted only once.