The Power of Referential Advice

Steven Callander[1], Nicolas Lambert[2], Niko Matouschek[3]

[1]Stanford GSB

[2]MIT

[3]Northwestern Kellogg

Penn State Economic Theory Seminar

December 4, 2020

Motivation

- Experts play a role in all aspects of economic, social, and political life.
- Common form: expert simply offers a recommendation.
  - E.g., librarian recommends a book.
- But often, advice is more expansive: expert conveys contextual information beyond a recommendation.
  - E.g., doctor recommends a treatment and discuss alternative ones.
- Broad evidence that experts influence decisions to their personal advantage.
- Can experts use additional information to sway decisions in their favor?
  - Assume hard, verifiable information.
  - Done poorly, extra information hurts the expert.
  - Done well, expert is able to shape DM's beliefs in her favor.

Literature

- Rich information
  - Glazer and Rubinstein (2004), Shin (2003), Dziuda (2011), Hart et al. (2017), Ben-Porath et al. (2019), Rappaport (2020)
- Hard information
  - Grossman (1981), Milgrom (1981), Milgrom and Roberts (1986)
  - Dye (1985)
- Soft information & Bayesian Persuasion
  - Crawford and Sobel (1982)
  - Kamenica and Gentzkow (2011)
  - Lipnowski and Ravid (2017)
- Expertise as a Brownian motion
  - Callander (2008, 2011)

Model

Model

- One sender (the expert), one receiver (the DM).
- Set of options $\mathcal{D} = \{d_0, d_1, \ldots, d_n\}$, where $n \geq 2$ and $d_i = i \cdot \delta$.
- Option $d_0 = 0$ interpreted as the **default option**, with outcome $X(d_0) = 0$.
- Option $d > 0$ generates outcome $X(d)$, Gaussian random walk:

$$X(d_i + \delta) - X(d_i) = \mu\delta + \sigma\sqrt{\delta}\theta_i,$$

  where $\theta_i \sim \mathcal{N}(0, 1)$.

  - Outcome increment from d to $d + \Delta$ has mean $\mu\Delta$ and variance $\sigma^2\Delta$.
  - $\mu > 0$ measures the expected rate of change from one option to the next.
  - $\sigma > 0$ scales the variance of each option relative to its neighbors.
  - $\theta = (\theta_0, \ldots, \theta_{n-1}) \in \Theta \equiv \Re^n$ is the state of the world.

Our interest is in large option sets.

We normalize with $\delta = 1/\sqrt{n}$ so that as n grows, outcomes become realization of a Brownian motion with drift μ and scale σ.

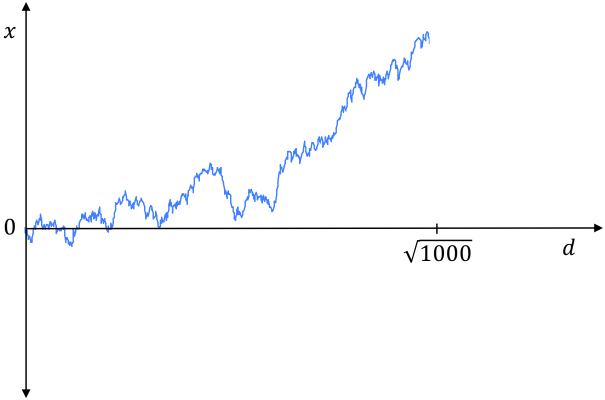

- Timing:
  1. Sender observes $X(d_i)$ for every $i = 1, \ldots, n$.
  2. Sender sends message m.
  3. Receiver updates beliefs and makes decision d.
  4. Sender & receiver realize their payoffs.
- Preferences:
  - Sender's utility: $u_S(d) = -d$.
  - Receiver's utility: $u_R(x) = -(x-b)^2 \implies \mathbb{E}[u_R] = -(\mathbb{E}[X] - b)^2 - \mathrm{Var}[X]$.
- Communication:
  - A message is a mapping $m : \mathcal{D} \to \mathbf{R} \cup \{\emptyset\}$.
  - $m(d) = \emptyset$ if sender hides the outcome of option d;
  - otherwise, she reveals it and $m(d) = X(d)$.

- Strategies:
  - Sender strategy $M : \Theta \rightarrow \mathcal{M}$.
  - Receiver strategy $D : \mathcal{M} \rightarrow \mathcal{D}$.
  - Receiver belief function $B : \mathcal{M} \rightarrow \Theta$.
- Solution concept:
  - Perfect Bayesian Equilibrium.
- Parameter restriction: $b > \dfrac{\sigma^2}{2\mu}$.

Two Sorts of Advice

Conative advice:

- The sender reveals one option only.
- The receiver chooses the only revealed option.

Referential advice:

- The sender reveals two or more options.
- The receiver chooses one of the revealed options.

Terminology borrowed from Jakobson's six functions of language.

We are primarily interested in **prescriptive advice**:

$\Longrightarrow$ *The receiver chooses an option revealed by the sender, interpreted as a recommendation.*

Outline

1. Can no advice be an equilibrium?
2. Full disclosure and other receiver optimal equilibria.
3. (How) Can the sender leverage his expertise?
    - ...with conative advice?
    - ...with referential advice?

On the Necessity of Equilibrium Advice

On the Necessity of Equilibrium Advice

Lemma

No-advice is not an equilibrium if the option set is large enough.

In all equilibria, there is a sender type who discloses the outcome of at least one non-default option.

Information Spillovers

When the sender reveals information on some options, this information impacts the receiver beliefs of unrevealed options.

Example 1: The sender reveals outcome of some given option d_a.

- Impacts beliefs of outcomes of options around d_a.
- This is a *direct* spillover: comes exclusively from the knowledge of $X(d_a)$.

In general, spillovers can be both direct and indirect.

Example 2: For some fixed d_a, d_b, the sender reveals $X(d_a)$ if $X(d_a) < b/2$, otherwise reveals $X(d_b)$.

Then, if the message is $X(d_b)$:

- Direct spillover from knowledge of $X(d_b)$.
- Indirect spillover from what can be inferred from the strategy itself.

Neutral Beliefs

Direct spillovers generate **neutral** beliefs.

$\implies$ Means the belief is equal to the initial state distribution conditional on the hard information in the message.

- Neutral beliefs are particularly tractable.

Strategies that only generate direct spillovers are the ones for which every possible deviation by the sender can be detected by the receiver.

- So that the receiver can take the message at face value.
- Formally, each state is compatible with only one message.
- Helps with message credibility.

Back to the Question: Can Sender Say Nothing?

The sender strategy of "no-advice" (empty message) generates neutral beliefs.

Receiver chooses d^{na} that optimality trades off $\mathbb{E}[X]$ against risk.

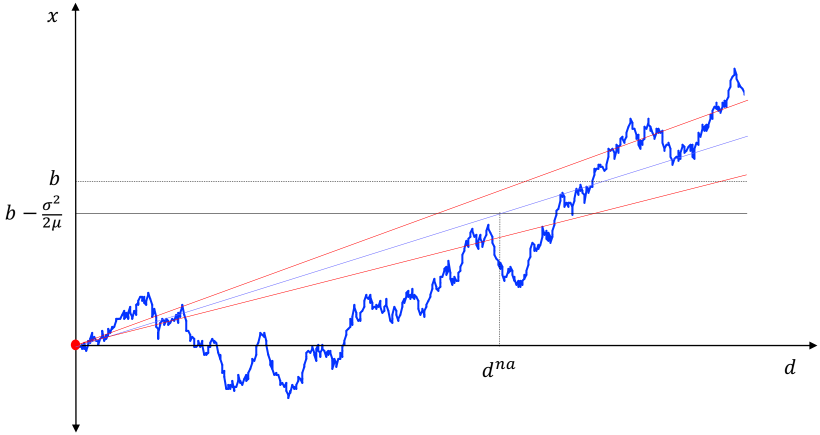

Receiver Optimal Equilibria

Fully Revealing Equilibria

Proposition

An equilibrium that is fully revealing exists.

- As long as off-path messages are sufficiently suspicious, deviations to off-path messages are unprofitable.
- But full revelation can only be sustained under extreme off-path beliefs.
- The equilibrium is also Pareto dominated.

$\Longrightarrow$ Fully revealing equilibria are not natural $\neq$ standard models (Milgrom).

Other Receiver Optimal Equilibria

Proposition

In all receiver-optimal equilibria, with probability one, the sender reveals all the options to the right of the equilibrium decision.

Proof sketch: *Referential advice helps provide credibility.* If a sender type who sends m hides an option d to the right of the decision d^*, all sender types compatible with m whose receiver-optimal option is d are best off sending m.

Bird-in-the-Hand refinement: for an off-path message, if the receiver chooses an option that is not revealed, it must be that the receiver believes the outcome of this option will deliver at least ε more utility than does the best option that is revealed.

No receiver-optimal equilibrium satisfies the BITH refinement.

Expert Power:
Can the Sender do Better?

First-Point Conative Strategy

The sender follows a **first-point conative strategy** when, for some $\Delta > 0$, the sender reveals the smallest option whose outcome falls in the range $[b - \Delta, b + \Delta]$ and if no such option exists, the sender reveals everything.

- Hybrid between conative and referential advice, but becomes fully conative for large option sets.
- How does it work? Transforms the relative standard into an absolute standard.

A first-point conative equilibrium exists if and only if $\Delta \in (0, \Delta^{\max}]$:

$$\Delta^{\max} = \frac{\sigma^2}{2\mu} + \frac{\mu}{2\sqrt{n}}.$$

- $\Delta \to 0 \implies$ optimal for receiver.
- $\Delta \to \Delta^{\max} \implies$ optimal for sender.

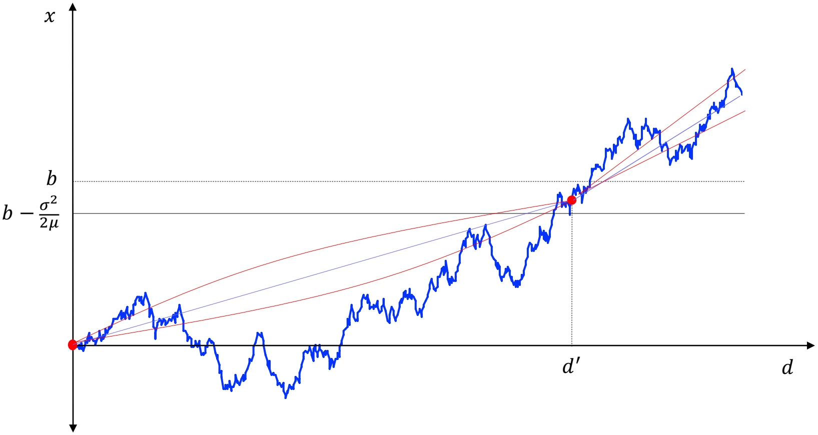

x
b
$b - \frac{\sigma^2}{2\mu}$
d'
d

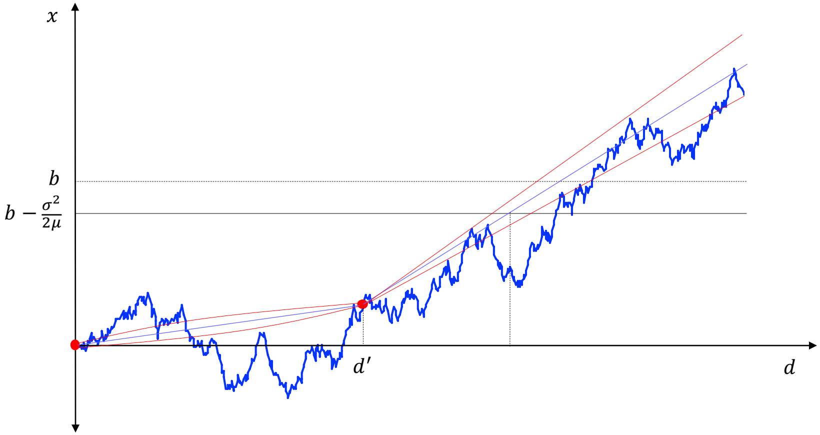

x
b
$b - \frac{\sigma^2}{2\mu}$
d'
d

About General Conative Advice

Theorem

1. For all sequences of equilibria that are conative in the limit, the equilibrium outcomes are in the range $[b - \sigma^2/(2\mu), b]$ in the limit.
2. Moreover, for any $x \in [b - \sigma^2/(2\mu), b]$ there exists a sequence of equilibria that are conative in the limit in which the outcomes converge to x.

Intuition for part (1):

- Let Ω_n be the set of states for which equilibrium advice is conative.
- Let $\mathcal{O}_n(d)$ be the set of outcomes $X(d)$ when receiver decides d.
- If $\theta \in \Omega_n$, sender reveals the left-most option d whose outcome is in $\mathcal{O}_n(d)$ $\implies$ conditionally on Ω_n, beliefs are neutral to the right of the decision.
- As $n \to \infty$, Ω_n grows to be the entire state space, so beliefs become neutral (unconditionally) in the limit.

Conative advice is good for the receiver, but not so good for the sender!

- In the worst case for the receiver, the equilibrium outcome still converges to $b - \sigma^2/(2\mu)$ = the average outcome the receiver would get without advice.

  $\implies$ receiver is always strictly better off than without advice.

- On average, the equilibrium option is equal to the option chosen by the receiver without advice.

  $\implies$ sender is always worse off than without advice, except in the sender-optimal equilibrium where she is as well off.

Interval Strategy

The sender follows the **interval strategy** when the sender reveals the outcomes of options $\{d_0, d_1, \ldots, d^r\}$, where d^r is the smallest option that satisfies

$$\max_{d \leq d^r} u_R(X(d)) \geq \max_{d > d^r} \mathbb{E}\left[u_R(X(d)) \mid X(d^r)\right],$$

with $d^r = d_n$ if no such option exists.

Interpretation: sender follows a stopping rule, revealing options starting from the default option until the receiver would not want to experiment to the right of the right-most revealed option, assuming neutral beliefs.

The interval strategy is an equilibrium strategy.

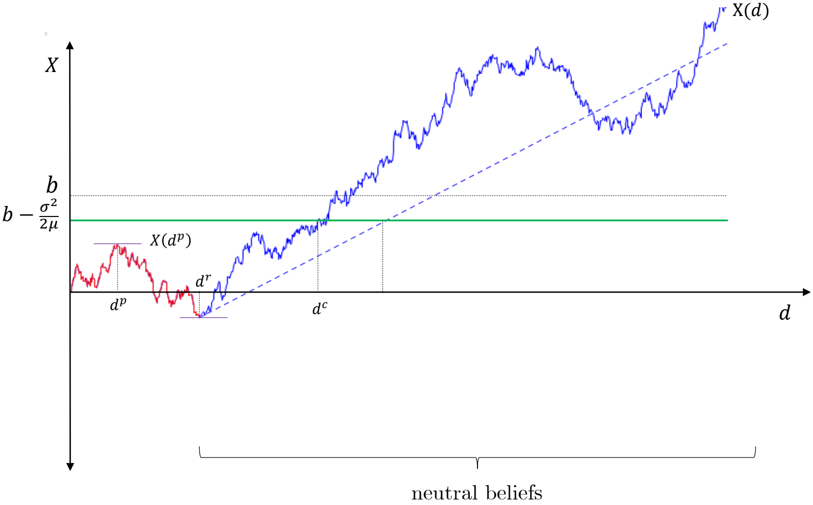

X(d)
X
b
$b - \frac{\sigma^2}{2\mu}$
$X(d^p)$
d^r
d^p
d^c
d
neutral beliefs

Domination of the Interval Strategy

An equilibrium (M, D, B) *weakly (resp. strictly) dominates* an equilibrium (M', D', B') in state θ if $D(M(\theta))$ is no greater than (resp. less than) $D'(M'(\theta))$.

Theorem

1. This interval equilibrium strictly dominates the sender-optimal conative equilibrium with positive probability, and weakly dominates for all states.
2. For every sequence of equilibria $\Sigma_1, \Sigma_2, \ldots$ that are conative in the limit, the interval equilibrium strictly dominates Σ_n in the limit as $n \to \infty$.

Proof Intuition (1)

By contradiction:

- Focus on the limiting case of a Brownian outcome path with drift μ, scale σ.
- For the quadratic receiver utility, d^r is defined as the smallest option d whose outcome reaches either the upper barrier

$$\overline{X} \equiv b - \frac{\sigma^2}{2\mu}$$

  or the (moving) lower barrier

$$\underline{X}(d) \equiv b - \frac{\sigma^2}{4\mu} - \frac{\mu}{\sigma^2} \min\{(b - X(d'))^2 : d' \in [0, d^r]\}.$$

- Look at a case where the interval decision is same as sender-optimal conative decision.
- Then the outcome path crosses $\overline{X}$ which then sets d^r.

Proof Intuition (2)

- As $X(d)$ approaches $b - \sigma^2/(2\mu)$ from below before crossing it, we have $X(d) = b - \sigma^2/(2\mu) - \delta$ for small $\delta > 0$, and

$$\begin{aligned}\underline{X}(d) &\geq b - \frac{\sigma^2}{4\mu} - \frac{\mu}{\sigma^2}(b - X(d))^2 \\ &= b - \frac{\sigma^2}{2\mu} - \delta - \frac{\mu}{\sigma^2}\delta^2 \\ &= X(d) - \frac{\mu}{\sigma^2}\delta^2\end{aligned}$$

- So, $X(d)$ is δ-far from the upper barrier, and approx δ^2-far from the lower barrier.

- The outcome path goes up on average, but for small increments, the path movement due to the white noise component dominates the drift component.

$\Longrightarrow$ the likelihood of reaching the lower barrier first converges to one as δ vanishes.

Can Sender Do (Even) Better?

Advice is **strongly prescriptive** when the sender reveals, at least, the options $d_0, d_1, \ldots, d^*$, where d^* is the receiver decision.

Theorem

If Σ is an equilibrium whose advice is strongly prescriptive with probability one and that weakly dominates the sender-optimal conative equilibrium, then the interval equilibrium weakly dominates Σ.

As a special case, this result applies to the equilibria whose communication is with direct information spillovers only.

Proof for Dominance Theorem for $n = 2$

Consider a strongly prescriptive equilibrium (M, D, B) that

- does weakly better for sender than the best conative eqm in every state,
- does strictly better than interval eqm for a positive mass of states.

Consider the set of messages $\mathcal{M}_+ \subseteq M(\Theta)$ for which the sender does strictly better than interval eqm.

1. Can the message that reveals only d_0 belong to $\mathcal{M}_+$?
2. Can a message that only reveals d_0 and d_1 belong to $\mathcal{M}_+$?
3. Can a message that only reveals d_0 and d_2 belong to $\mathcal{M}_+$?

$\Longrightarrow$ If the answer is always no then we have a contradiction.

We prove each in turn and by contradiction.

Step 1

Can the message that reveals only d_0 belong to $\mathcal{M}_+$?

- Note that $\Delta^{\max} < b$.
- So $u_R(0) < \mathbb{E}[u_R(X(d_1))]$.
- If answer is true, then eqm decision is d_0 always.
- And $\mathbb{E}[u_R(X(d_1))]$ is a convex combination of terms of the form $\mathbb{E}[u_R(X(d_1)) \mid M(\theta) = m]$ where m yields receiver decision d_0.
- So for some on-path message m that yields decision d_0, $u_R(0) < \mathbb{E}[u_R(X(d_1)) \mid M(\theta) = m]$.
- And upon observing such message, receiver is better off choosing d_1.

$\Longrightarrow$ Contradiction.

Step 2

Can a message that only reveals d_0 and d_1 belong to $\mathcal{M}_+$?

Suppose $m_1 \in \mathcal{M}_+$ reveals exactly d_0 and d_1.

$\implies m_1(d_1) < b - \Delta^{\text{max}}$

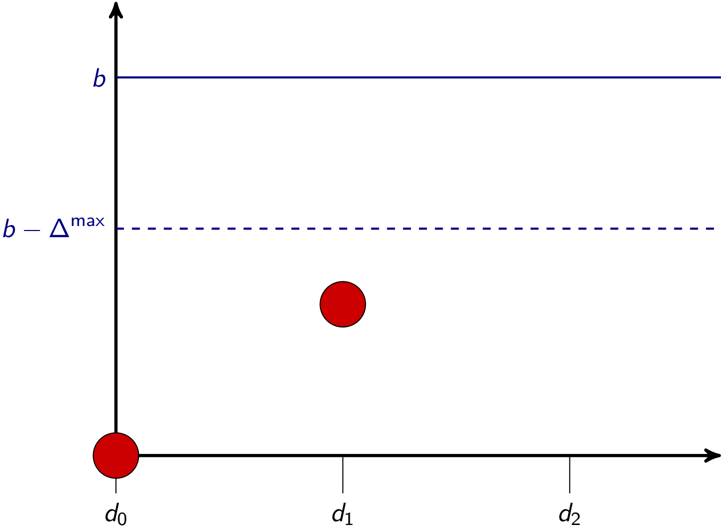

b
$b - \Delta^{\max}$
d_0
d_1
d_2

Step 2

Can a message that only reveals d_0 and d_1 belong to $\mathcal{M}_+$?

Let $x_1 = m_1(d_1)$.

- $x_1 < b - \Delta^{\max}$ so $u_R(x_1) < \mathbb{E}[u_R(X(d_2)) \mid X(d_1) = x_1]$.
- $\mathbb{E}[u_R(X(d_2)) \mid X(d_1) = x_1]$ is convex combination of $\mathbb{E}[u_R(X(d_2)) \mid X(d_1) = x_1, M(\theta) = m_1]$, and terms of the form $\mathbb{E}[u_R(X(d_2)) \mid X(d_1) = x_1, M(\theta) = m]$ for $m \neq m_1$.
- By Step 1, any on-path $m \neq m_1$ sent when $X(d_1) = x_1$ must reveal d_2.
- Receiver decision must be d_0 or d_1, so for such m, $u_R(x_1) \geq u_R(X(d_2))$.
- Hence, $u_R(x_1) \geq \mathbb{E}[u_R(X(d_2)) \mid X(d_1) = x_1, M(\theta) = m]$.
- So $u_R(x_1) < \mathbb{E}[u_R(X(d_2)) \mid M(\theta) = m_1]$.
- Upon observing m_1, receiver is better off choosing d_2.

$\implies$ Contradiction.

Step 3

Can a message that only reveals d_0 and d_2 belong to $\mathcal{M}_+$?

First, we show that the message m_0 that reveals

1. $X(d_0) = 0$ and
2. $X(d_1) = b - \Delta^{\max}$

is on path: $m_0 \in M(\Theta)$.

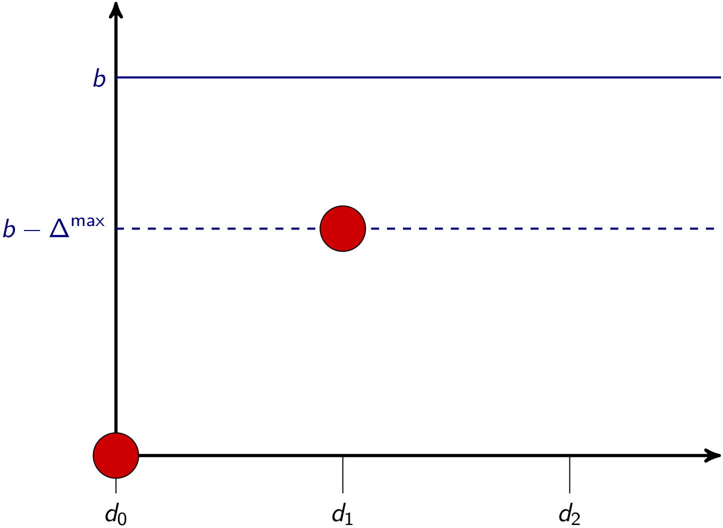

b
$b - \Delta^{\max}$
d_0
d_1
d_2

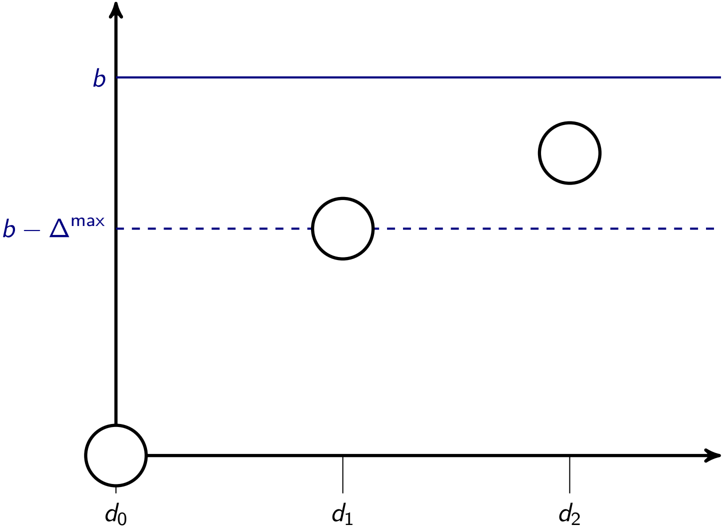

b
$b - \Delta^{\max}$
d_0
d_1
d_2

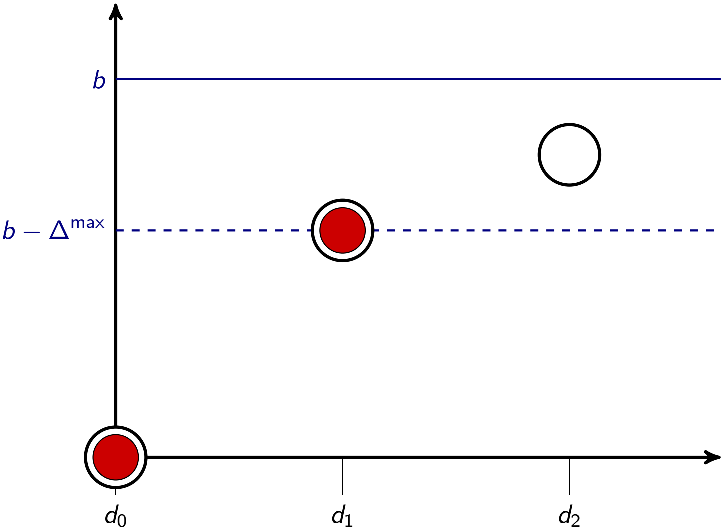
b
$b - \Delta^{\max}$
d_0
d_1
d_2

Step 3

Can a message that only reveals d_0 and d_2 belong to $\mathcal{M}_+$?

Take $\mathcal{M}_g$ the set of all such "messages with a gap."

If $m \in \mathcal{M}_g$ then $D(m) = d_0$ (by strongly prescriptive) so $m(d_2) \not\in (0, 2b)$.

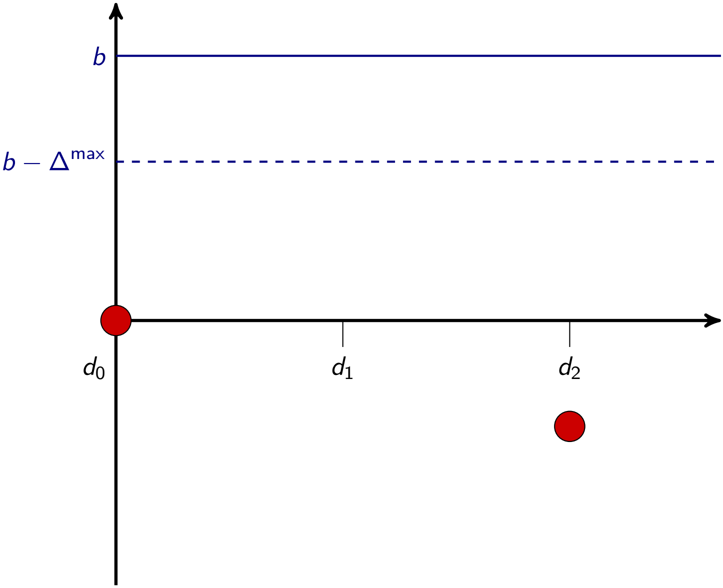

b
$b - \Delta^{\max}$
d_0
d_1
d_2

Step 3

Can a message that only reveals d_0 and d_2 belong to $\mathcal{M}_+$?

Take $\mathcal{M}_g$ the set of all such "messages with a gap."

Let Ω = set of all states where

- $X(d_1) = b - \Delta^{\max}$, and
- $X(d_2) = m(d_2)$ for some $m \in \mathcal{M}_g$.

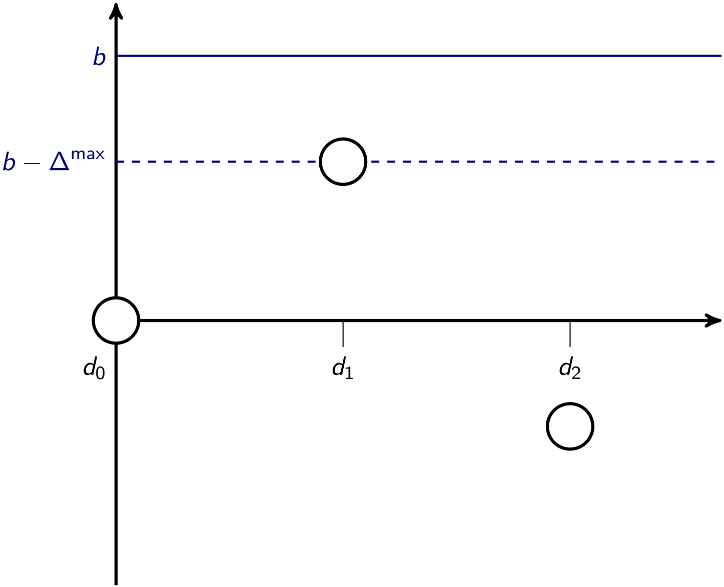

b
$b - \Delta^{\max}$
d_0
d_1
d_2

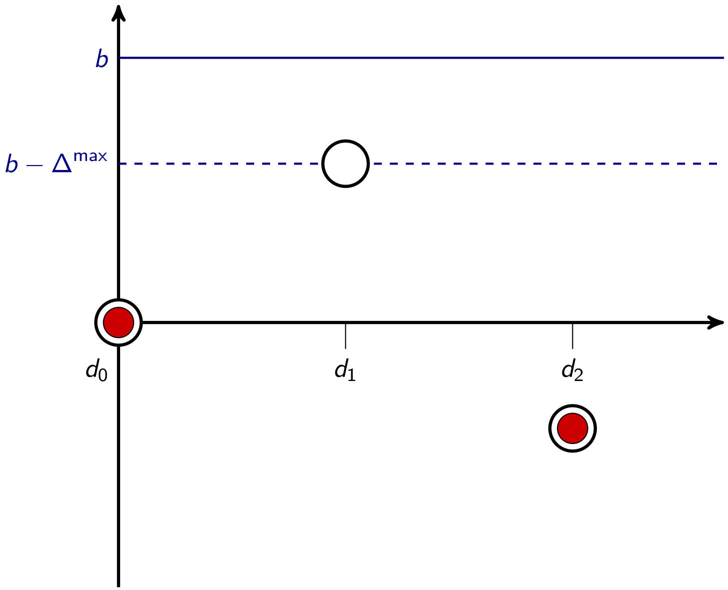
b
$b - \Delta^{\max}$
d_0
d_1
d_2

Step 3

Can a message that only reveals d_0 and d_2 belong to $\mathcal{M}_+$?

- If $\theta \in \Omega$ then sender sends a gap message of $\mathcal{M}_g$ to get decision d_0.
- $\Delta^{\max}$ is defined so that $u_R(b - \Delta^{\max}) = \mathbb{E}[u_R(X(d_2)) \mid X(d_1) = b - \Delta^{\max}]$.
- RHS is convex combination of $\mathbb{E}[u_R(X(d_2)) \mid \theta \in \Omega]$ and terms of the form $\mathbb{E}[u_R(X(d_2)) \mid X(d_1) = b - \Delta^{\max}, M(\theta) = m]$ where $m \notin \mathcal{M}_g$.
- Note $u_R(b - \Delta^{\max}) > \mathbb{E}[u_R(X(d_2)) \mid \theta \in \Omega]$.
- So there exists at least one on-path m revealing $X(d_1) = b - \Delta^{\max}$, and where $u_R(b - \Delta^{\max}) < \mathbb{E}[u_R(X(d_2)) \mid M(\theta) = m]$.
- Under m, the receiver chooses d_2.
  So sender is better off sending m_0 to get the receiver to decide d_1.

$\Longrightarrow$ Contradiction.

Conclusion

Conclusion

- We build a model of strategic communication that accounts for different types of advice.
  - Key ingredient: rich private information modeled as imperfectly correlated outcomes across options.
- Referential advice, by itself, does not help experts.
  - Revealing too much makes the expert worse off.
- When done right however, it allows experts to sway decisions in their favor, compared to conative advice.
  - Under some general conditions, the optimal expert communication is to reveal the recommendation along with related options.
- Rich advice can be so persuasive that more often than not the DM is worse off than if he got no advice at all. The reason he still seeks advice is that it prevents rare disasters.
- The DM benefits from rules & norms that require simple advice.

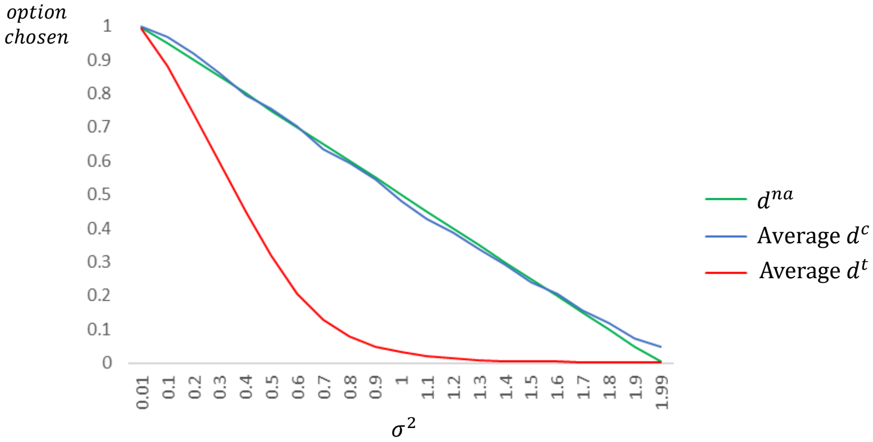

option chosen
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
0.01
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1
1.1
1.2
1.3
1.4
1.5
1.6
1.7
1.8
1.9
1.99
σ^2
d^{na}
Average d^c
Average d^t