Manipulative Disclosure

Claudio Mezzetti

University of Warwick

Growth and Development Conference, ISI Delhi, December 16-18, 2009

Decision Makers and Experts

- An investor consults a financial adviser.
- The market reacts to news from a company.
- A patient goes to a doctor.
- A reader reads a journalist's article.
- A politician meets a lobbyist.
- Electors reacts to a politician's statement.
- A juror listens to an expert witness.
- A consumer talks to a salesperson.
- A firm asks a consultant about the optimal size of an investment project.

Decision Makers and Experts

- An investor consults a financial adviser.
- The market reacts to news from a company.
- A patient goes to a doctor.
- A reader reads a journalist's article.
- A politician meets a lobbyist.
- Electors reacts to a politician's statement.
- A juror listens to an expert witness.
- A consumer talks to a salesperson.
- A firm asks a consultant about the optimal size of an investment project.

Decision Makers and Experts

- An investor consults a financial adviser.
- The market reacts to news from a company.
- A patient goes to a doctor.
- A reader reads a journalist's article.
- A politician meets a lobbyist.
- Electors reacts to a politician's statement.
- A juror listens to an expert witness.
- A consumer talks to a salesperson.
- A firm asks a consultant about the optimal size of an investment project.

Decision Makers and Experts

- An investor consults a financial adviser.
- The market reacts to news from a company.
- A patient goes to a doctor.
- A reader reads a journalist's article.
- A politician meets a lobbyist.
- Electors reacts to a politician's statement.
- A juror listens to an expert witness.
- A consumer talks to a salesperson.
- A firm asks a consultant about the optimal size of an investment project.

Decision Makers and Experts

- An investor consults a financial adviser.
- The market reacts to news from a company.
- A patient goes to a doctor.
- A reader reads a journalist's article.
- A politician meets a lobbyist.
- Electors reacts to a politician's statement.
- A juror listens to an expert witness.
- A consumer talks to a salesperson.
- A firm asks a consultant about the optimal size of an investment project.

Decision Makers and Experts

- An investor consults a financial adviser.
- The market reacts to news from a company.
- A patient goes to a doctor.
- A reader reads a journalist's article.
- A politician meets a lobbyist.
- Electors reacts to a politician's statement.
- A juror listens to an expert witness.
- A consumer talks to a salesperson.
- A firm asks a consultant about the optimal size of an investment project.

Decision Makers and Experts

- An investor consults a financial adviser.
- The market reacts to news from a company.
- A patient goes to a doctor.
- A reader reads a journalist's article.
- A politician meets a lobbyist.
- Electors reacts to a politician's statement.
- A juror listens to an expert witness.
- A consumer talks to a salesperson.
- A firm asks a consultant about the optimal size of an investment project.

Decision Makers and Experts

- An investor consults a financial adviser.
- The market reacts to news from a company.
- A patient goes to a doctor.
- A reader reads a journalist's article.
- A politician meets a lobbyist.
- Electors reacts to a politician's statement.
- A juror listens to an expert witness.
- A consumer talks to a salesperson.
- A firm asks a consultant about the optimal size of an investment project.

Decision Makers and Experts

- An investor consults a financial adviser.
- The market reacts to news from a company.
- A patient goes to a doctor.
- A reader reads a journalist's article.
- A politician meets a lobbyist.
- Electors reacts to a politician's statement.
- A juror listens to an expert witness.
- A consumer talks to a salesperson.
- A firm asks a consultant about the optimal size of an investment project.

Main Questions

An expert that engages in manipulative disclosure of information is self-serving. Her goal is to induce the decision maker to choose what she wants.

- Can the expert succeed in inducing the DM to choose what the expert wants?
- When is she more likely to succeed?
- Is it important that the expert keeps a poker face, not letting the DM know her preferences?
- Should the DM consult an expert that is known to be biased? Or should she prefer an a priori unbiased expert?
- What information will the expert reveal? Will she reveal unfavourable information?
- Will the DM always benefit from an increase in his familiarity with the problem?

Main Questions

An expert that engages in manipulative disclosure of information is self-serving. Her goal is to induce the decision maker to choose what she wants.

- Can the expert succeed in inducing the DM to choose what the expert wants?
- When is she more likely to succeed?
- Is it important that the expert keeps a poker face, not letting the DM know her preferences?
- Should the DM consult an expert that is known to be biased? Or should she prefer an a priori unbiased expert?
- What information will the expert reveal? Will she reveal unfavourable information?
- Will the DM always benefit from an increase in his familiarity with the problem?

Main Questions

An expert that engages in manipulative disclosure of information is self-serving. Her goal is to induce the decision maker to choose what she wants.

- Can the expert succeed in inducing the DM to choose what the expert wants?
- When is she more likely to succeed?
- Is it important that the expert keeps a poker face, not letting the DM know her preferences?
- Should the DM consult an expert that is known to be biased? Or should she prefer an a priori unbiased expert?
- What information will the expert reveal? Will she reveal unfavourable information?
- Will the DM always benefit from an increase in his familiarity with the problem?

Main Questions

An expert that engages in manipulative disclosure of information is self-serving. Her goal is to induce the decision maker to choose what she wants.

- Can the expert succeed in inducing the DM to choose what the expert wants?
- When is she more likely to succeed?
- Is it important that the expert keeps a poker face, not letting the DM know her preferences?
- Should the DM consult an expert that is known to be biased? Or should she prefer an a priori unbiased expert?
- What information will the expert reveal? Will she reveal unfavourable information?
- Will the DM always benefit from an increase in his familiarity with the problem?

Main Questions

An expert that engages in manipulative disclosure of information is self-serving. Her goal is to induce the decision maker to choose what she wants.

- Can the expert succeed in inducing the DM to choose what the expert wants?
- When is she more likely to succeed?
- Is it important that the expert keeps a poker face, not letting the DM know her preferences?
- Should the DM consult an expert that is known to be biased? Or should she prefer an a priori unbiased expert?
- What information will the expert reveal? Will she reveal unfavourable information?
- Will the DM always benefit from an increase in his familiarity with the problem?

Main Questions

An expert that engages in manipulative disclosure of information is self-serving. Her goal is to induce the decision maker to choose what she wants.

- Can the expert succeed in inducing the DM to choose what the expert wants?
- When is she more likely to succeed?
- Is it important that the expert keeps a poker face, not letting the DM know her preferences?
- Should the DM consult an expert that is known to be biased? Or should she prefer an a priori unbiased expert?
- What information will the expert reveal? Will she reveal unfavourable information?
- Will the DM always benefit from an increase in his familiarity with the problem?

Main Questions

An expert that engages in manipulative disclosure of information is self-serving. Her goal is to induce the decision maker to choose what she wants.

- Can the expert succeed in inducing the DM to choose what the expert wants?
- When is she more likely to succeed?
- Is it important that the expert keeps a poker face, not letting the DM know her preferences?
- Should the DM consult an expert that is known to be biased? Or should she prefer an a priori unbiased expert?
- What information will the expert reveal? Will she reveal unfavourable information?
- Will the DM always benefit from an increase in his familiarity with the problem?

Received Wisdom

1: First Models (Persuasion Games): Grossman (1981), Milgrom (1981), Milgrom and Roberts (1986)

- Insight: Decision maker is highly skeptical. When evidence is incomplete, DM assumes that any missing information is likely to be unfavorable to the expert. As a consequence, expert will want to reveal all favorable information. This unravel any attempt to hide information and leads to full disclosure.

- It's important that the goal (bias) of the expert be known to DM

- Seidmann Winter (1997) - allow some uncertainty over the expert's preferences; they focus on the condition under which there is an equilibrium with full disclosure

Received Wisdom

1: First Models (Persuasion Games): Grossman (1981), Milgrom (1981), Milgrom and Roberts (1986)

- Insight: Decision maker is highly skeptical. When evidence is incomplete, DM assumes that any missing information is likely to be unfavorable to the expert. As a consequence, expert will want to reveal all favorable information. This unravel any attempt to hide information and leads to full disclosure.

- It's important that the goal (bias) of the expert be known to DM

- Seidmann Winter (1997) - allow some uncertainty over the expert's preferences; they focus on the condition under which there is an equilibrium with full disclosure

Received Wisdom

1: First Models (Persuasion Games): Grossman (1981), Milgrom (1981), Milgrom and Roberts (1986)

- Insight: Decision maker is highly skeptical. When evidence is incomplete, DM assumes that any missing information is likely to be unfavorable to the expert. As a consequence, expert will want to reveal all favorable information. This unravel any attempt to hide information and leads to full disclosure.

- It's important that the goal (bias) of the expert be known to DM

- Seidmann Winter (1997) - allow some uncertainty over the expert's preferences; they focus on the condition under which there is an equilibrium with full disclosure

Received Wisdom: Refinements

2: Uncertainty about whether the Expert has Information

- Shin (1994) - full disclosure my fail it there is uncertainty over the precision of the expert's information, i.e., whether the expert knows the true state; the expert conceals some information

- When there is uncertainty about whether the expert has information, the DM's skepticism is tempered. Evidence may be incomplete because expert does not have all the information. The expert will always withhold very bad news and report only relatively good news.

Received Wisdom: Refinements

2: Uncertainty about whether the Expert has Information

- Shin (1994) - full disclosure my fail it there is uncertainty over the precision of the expert's information, i.e., whether the expert knows the true state; the expert conceals some information

- When there is uncertainty about whether the expert has information, the DM's skepticism is tempered. Evidence may be incomplete because expert does not have all the information. The expert will always withhold very bad news and report only relatively good news.

Recent Papers

- Wolinsky (2003) - the expert's bias is unknown, but it can only take two values; the expert can fully report, or underreport favorable information; the biased expert with favorable information above a threshold fully discloses; all other expert types play a mixed strategy

- Dziuda (2007) - (related to Glazer and Rubinstein (2001), (2004)) - uncertain whether expert is honest or biased in favor of one of two alternatives; the DM has a threshold of evidence needed to prefer one alternative over the other, which is not known to the experts. Result: the biased expert biases his report in favor of his preferences, but also presents some unfavourable information.

Recent Papers

- Wolinsky (2003) - the expert's bias is unknown, but it can only take two values; the expert can fully report, or underreport favorable information; the biased expert with favorable information above a threshold fully discloses; all other expert types play a mixed strategy

- Dziuda (2007) - (related to Glazer and Rubinstein (2001), (2004)) - uncertain whether expert is honest or biased in favor of one of two alternatives; the DM has a threshold of evidence needed to prefer one alternative over the other, which is not known to the experts. Result: the biased expert biases his report in favor of his preferences, but also presents some unfavourable information.

Cheap Talk Models

- Crawford and Sobel (1982) - known bias
- Sobel (1985) - analysis of credibility/reputation; dynamic model; it's uncertain if the expert is a friend or an enemy (one-sided uncertain bias)
- Benabou and Laroque (1992) - extend Sobel (1985) to allow for noisy information by the expert
- Dimitrakas and Sarafidis (2005), Morgan Stocken (2003), Morris (2001) - the bias is uncertain, but it is known to be positive; that is, the direction of the bias is known; obtain partition equilibria
- Li and Madarasz (2007) - the bias is uncertain; it takes two values; the direction of the bias can be unknown; focus on welfare effects of mandatory disclosure; the uncertain bias may make both expert and DM better off.

Cheap Talk Models

- Crawford and Sobel (1982) - known bias
- Sobel (1985) - analysis of credibility/reputation; dynamic model; it's uncertain if the expert is a friend or an enemy (one-sided uncertain bias)
- Benabou and Laroque (1992) - extend Sobel (1985) to allow for noisy information by the expert
- Dimitrakas and Sarafidis (2005), Morgan Stocken (2003), Morris (2001) - the bias is uncertain, but it is known to be positive; that is, the direction of the bias is known; obtain partition equilibria
- Li and Madarasz (2007) - the bias is uncertain; it takes two values; the direction of the bias can be unknown; focus on welfare effects of mandatory disclosure; the uncertain bias may make both expert and DM better off.

Cheap Talk Models

- Crawford and Sobel (1982) - known bias
- Sobel (1985) - analysis of credibility/reputation; dynamic model; it's uncertain if the expert is a friend or an enemy (one-sided uncertain bias)
- Benabou and Laroque (1992) - extend Sobel (1985) to allow for noisy information by the expert
- Dimitrakas and Sarafidis (2005), Morgan Stocken (2003), Morris (2001) - the bias is uncertain, but it is known to be positive; that is, the direction of the bias is known; obtain partition equilibria
- Li and Madarasz (2007) - the bias is uncertain; it takes two values; the direction of the bias can be unknown; focus on welfare effects of mandatory disclosure; the uncertain bias may make both expert and DM better off.

Cheap Talk Models

- Crawford and Sobel (1982) - known bias
- Sobel (1985) - analysis of credibility/reputation; dynamic model; it's uncertain if the expert is a friend or an enemy (one-sided uncertain bias)
- Benabou and Laroque (1992) - extend Sobel (1985) to allow for noisy information by the expert
- Dimitrakas and Sarafidis (2005), Morgan Stocken (2003), Morris (2001) - the bias is uncertain, but it is known to be positive; that is, the direction of the bias is known; obtain partition equilibria
- Li and Madarasz (2007) - the bias is uncertain; it takes two values; the direction of the bias can be unknown; focus on welfare effects of mandatory disclosure; the uncertain bias may make both expert and DM better off.

Cheap Talk Models

- Crawford and Sobel (1982) - known bias
- Sobel (1985) - analysis of credibility/reputation; dynamic model; it's uncertain if the expert is a friend or an enemy (one-sided uncertain bias)
- Benabou and Laroque (1992) - extend Sobel (1985) to allow for noisy information by the expert
- Dimitrakas and Sarafidis (2005), Morgan Stocken (2003), Morris (2001) - the bias is uncertain, but it is known to be positive; that is, the direction of the bias is known; obtain partition equilibria
- Li and Madarasz (2007) - the bias is uncertain; it takes two values; the direction of the bias can be unknown; focus on welfare effects of mandatory disclosure; the uncertain bias may make both expert and DM better off.

Some Evidence: Selective Reporting of Information

- "Merck's arthritis drug, Vioxx, was reportedly found to double the risk of heart attacks for its users, but although this dangerous side effect was suspected by scientists for years before the drug was banned, there were no full studies confirming that danger and no reports alerting users to the risks. The became apparent only after the drug had received approval from the Food and Drug Administration." (Milgrom, 2008)

- "Eli Lilly allegedly instructed its salespeople to downplay the risks associated with its schizophrenia drug, Zyprexa, although some evidence suggested that risks might be severe." (Milgrom, 2008)

Some Evidence: Selective Reporting of Information

- "Merck's arthritis drug, Vioxx, was reportedly found to double the risk of heart attacks for its users, but although this dangerous side effect was suspected by scientists for years before the drug was banned, there were no full studies confirming that danger and no reports alerting users to the risks. The became apparent only after the drug had received approval from the Food and Drug Administration." (Milgrom, 2008)

- "Eli Lilly allegedly instructed its salespeople to downplay the risks associated with its schizophrenia drug, Zyprexa, although some evidence suggested that risks might be severe." (Milgrom, 2008)

What's the Expert Goal?

- If, for example, the problem is multidimensional, how can be sure we know the expert's goal?

- Is the seller trying to sell me the the fancier air conditioning unit because she has a higher profit margin on it, or the more basic product that will require more service (from her) in the future?

- Is the politician interested in under- or over-estimate the danger of a serious epidemic, or economic crisis, or terrorism threat?

- Is the surgeon interested in pushing you towards an open or a keyhole (laparoscopic) inguinal hernia operation? ("Your surgeon will recommend which is suitable for you." Are we so sure?)

What's the Expert Goal?

- If, for example, the problem is multidimensional, how can be sure we know the expert's goal?
- Is the seller trying to sell me the the fancier air conditioning unit because she has a higher profit margin on it, or the more basic product that will require more service (from her) in the future?
- Is the politician interested in under- or over-estimate the danger of a serious epidemic, or economic crisis, or terrorism threat?
- Is the surgeon interested in pushing you towards an open or a keyhole (laparoscopic) inguinal hernia operation? ("Your surgeon will recommend which is suitable for you." Are we so sure?)

What's the Expert Goal?

- If, for example, the problem is multidimensional, how can be sure we know the expert's goal?
- Is the seller trying to sell me the the fancier air conditioning unit because she has a higher profit margin on it, or the more basic product that will require more service (from her) in the future?
- Is the politician interested in under- or over-estimate the danger of a serious epidemic, or economic crisis, or terrorism threat?
- Is the surgeon interested in pushing you towards an open or a keyhole (laparoscopic) inguinal hernia operation? ("Your surgeon will recommend which is suitable for you." Are we so sure?)

What's the Expert Goal?

- If, for example, the problem is multidimensional, how can be sure we know the expert's goal?
- Is the seller trying to sell me the the fancier air conditioning unit because she has a higher profit margin on it, or the more basic product that will require more service (from her) in the future?
- Is the politician interested in under- or over-estimate the danger of a serious epidemic, or economic crisis, or terrorism threat?
- Is the surgeon interested in pushing you towards an open or a keyhole (laparoscopic) inguinal hernia operation? ("Your surgeon will recommend which is suitable for you." Are we so sure?)

My Approach

- The Decision Maker does not know what his best choice is. It depends on the state of the world.
- The DM is uncertain about the bias of the expert. In particular, he does not know the direction of the bias.
- The expert reports some information.
- Lying is not allowed: the information reported by the expert must include the truth.
- The expert is free to decide how close to the truth her report is.
- After observing the expert's report, the DM updates his beliefs about the state of the world and the expert's bias, and chooses a course of action.
- I look at Perfect Bayesian Equilibrium.

My Approach

- The Decision Maker does not know what his best choice is. It depends on the state of the world.
- The DM is uncertain about the bias of the expert. In particular, he does not know the direction of the bias.
- The expert reports some information.
- Lying is not allowed: the information reported by the expert must include the truth.
- The expert is free to decide how close to the truth her report is.
- After observing the expert's report, the DM updates his beliefs about the state of the world and the expert's bias, and chooses a course of action.
- I look at Perfect Bayesian Equilibrium.

My Approach

- The Decision Maker does not know what his best choice is. It depends on the state of the world.
- The DM is uncertain about the bias of the expert. In particular, he does not know the direction of the bias.
- The expert reports some information.
- Lying is not allowed: the information reported by the expert must include the truth.
- The expert is free to decide how close to the truth her report is.
- After observing the expert's report, the DM updates his beliefs about the state of the world and the expert's bias, and chooses a course of action.
- I look at Perfect Bayesian Equilibrium.

My Approach

- The Decision Maker does not know what his best choice is. It depends on the state of the world.
- The DM is uncertain about the bias of the expert. In particular, he does not know the direction of the bias.
- The expert reports some information.
- Lying is not allowed: the information reported by the expert must include the truth.
- The expert is free to decide how close to the truth her report is.
- After observing the expert's report, the DM updates his beliefs about the state of the world and the expert's bias, and chooses a course of action.
- I look at Perfect Bayesian Equilibrium.

My Approach

- The Decision Maker does not know what his best choice is. It depends on the state of the world.
- The DM is uncertain about the bias of the expert. In particular, he does not know the direction of the bias.
- The expert reports some information.
- Lying is not allowed: the information reported by the expert must include the truth.
- The expert is free to decide how close to the truth her report is.
- After observing the expert's report, the DM updates his beliefs about the state of the world and the expert's bias, and chooses a course of action.
- I look at Perfect Bayesian Equilibrium.

My Approach

- The Decision Maker does not know what his best choice is. It depends on the state of the world.
- The DM is uncertain about the bias of the expert. In particular, he does not know the direction of the bias.
- The expert reports some information.
- Lying is not allowed: the information reported by the expert must include the truth.
- The expert is free to decide how close to the truth her report is.
- After observing the expert's report, the DM updates his beliefs about the state of the world and the expert's bias, and chooses a course of action.
- I look at Perfect Bayesian Equilibrium.

My Approach

- The Decision Maker does not know what his best choice is. It depends on the state of the world.
- The DM is uncertain about the bias of the expert. In particular, he does not know the direction of the bias.
- The expert reports some information.
- Lying is not allowed: the information reported by the expert must include the truth.
- The expert is free to decide how close to the truth her report is.
- After observing the expert's report, the DM updates his beliefs about the state of the world and the expert's bias, and chooses a course of action.
- I look at Perfect Bayesian Equilibrium.

Main Results of the Paper

- A positive measure of expert types obtain what they want: full manipulation.
- If the DM is totally unfamiliar with the problem and the expert has zero mean bias, then all expert types are able to fully manipulate the DM
- The DM is better off the more biased is the expert, i.e., the higher is the mean bias
- The expert is better off the lower is her mean bias
- The DM is made worse off by an increase in the bias variance
- The expert is better off if the bias variance increases
- Some types of experts send unfavorable unformation
- An increase in the DM's familiarity with the problem makes the DM better off

Main Results of the Paper

- A positive measure of expert types obtain what they want: full manipulation.
- If the DM is totally unfamiliar with the problem and the expert has zero mean bias, then all expert types are able to fully manipulate the DM
- The DM is better off the more biased is the expert, i.e., the higher is the mean bias
- The expert is better off the lower is her mean bias
- The DM is made worse off by an increase in the bias variance
- The expert is better off if the bias variance increases
- Some types of experts send unfavorable unformation
- An increase in the DM's familiarity with the problem makes the DM better off

Main Results of the Paper

- A positive measure of expert types obtain what they want: full manipulation.
- If the DM is totally unfamiliar with the problem and the expert has zero mean bias, then all expert types are able to fully manipulate the DM
- The DM is better off the more biased is the expert, i.e., the higher is the mean bias
- The expert is better off the lower is her mean bias
- The DM is made worse off by an increase in the bias variance
- The expert is better off if the bias variance increases
- Some types of experts send unfavorable unformation
- An increase in the DM's familiarity with the problem makes the DM better off

Main Results of the Paper

- A positive measure of expert types obtain what they want: full manipulation.
- If the DM is totally unfamiliar with the problem and the expert has zero mean bias, then all expert types are able to fully manipulate the DM
- The DM is better off the more biased is the expert, i.e., the higher is the mean bias
- The expert is better off the lower is her mean bias
- The DM is made worse off by an increase in the bias variance
- The expert is better off if the bias variance increases
- Some types of experts send unfavorable unformation
- An increase in the DM's familiarity with the problem makes the DM better off

Main Results of the Paper

- A positive measure of expert types obtain what they want: full manipulation.
- If the DM is totally unfamiliar with the problem and the expert has zero mean bias, then all expert types are able to fully manipulate the DM
- The DM is better off the more biased is the expert, i.e., the higher is the mean bias
- The expert is better off the lower is her mean bias
- The DM is made worse off by an increase in the bias variance
- The expert is better off if the bias variance increases
- Some types of experts send unfavorable unformation
- An increase in the DM's familiarity with the problem makes the DM better off

Main Results of the Paper

- A positive measure of expert types obtain what they want: full manipulation.
- If the DM is totally unfamiliar with the problem and the expert has zero mean bias, then all expert types are able to fully manipulate the DM
- The DM is better off the more biased is the expert, i.e., the higher is the mean bias
- The expert is better off the lower is her mean bias
- The DM is made worse off by an increase in the bias variance
- The expert is better off if the bias variance increases
- Some types of experts send unfavorable unformation
- An increase in the DM's familiarity with the problem makes the DM better off

Main Results of the Paper

- A positive measure of expert types obtain what they want: full manipulation.
- If the DM is totally unfamiliar with the problem and the expert has zero mean bias, then all expert types are able to fully manipulate the DM
- The DM is better off the more biased is the expert, i.e., the higher is the mean bias
- The expert is better off the lower is her mean bias
- The DM is made worse off by an increase in the bias variance
- The expert is better off if the bias variance increases
- Some types of experts send unfavorable unformation
- An increase in the DM's familiarity with the problem makes the DM better off

Main Results of the Paper

- A positive measure of expert types obtain what they want: full manipulation.
- If the DM is totally unfamiliar with the problem and the expert has zero mean bias, then all expert types are able to fully manipulate the DM
- The DM is better off the more biased is the expert, i.e., the higher is the mean bias
- The expert is better off the lower is her mean bias
- The DM is made worse off by an increase in the bias variance
- The expert is better off if the bias variance increases
- Some types of experts send unfavorable unformation
- An increase in the DM's familiarity with the problem makes the DM better off

Outline of the Talk

Outline of the Talk

Outline of the Talk

Outline of the Talk

Outline of the Talk

1. Introduction
2. The Model
3. The Totally Unfamiliar DM with a Zero-Mean-Bias Expert
4. The Totally Unfamiliar DM with a Positive-Mean-Bias Expert
5. The Partially Familiar DM
6. Conclusions

Outline of the Talk

Outline of the Talk

1. Introduction
2. The Model
3. The Totally Unfamiliar DM with a Zero-Mean-Bias Expert
4. The Totally Unfamiliar DM with a Positive-Mean-Bias Expert
5. The Partially Familiar DM
6. Conclusions

The Model...

- A decision maker must choose an alternative; the set of possible decisions is the real line.
- An expert observes the realization of the state of the world, $\omega \in [-M, M]$;the state of the world coincides with the first-best optimal decision of the DM.
- The decision maker does not observe the state of the world; she has prior beliefs that are uniformly distributed over $[-M, M]$.
- The parameter $1/M$ measures the prior familiarity of the DM with the decision problem.
- When $M = 0$, the DM is fully familiar with the problem; she knows that the state of the world is $\omega = 0$.
- When $M = \infty$, the DM is totally unfamiliar with the problem; her prior over ω is diffuse over the real line.

The Model...

- A decision maker must choose an alternative; the set of possible decisions is the real line.
- An expert observes the realization of the state of the world, $\omega \in [-M, M]$;the state of the world coincides with the first-best optimal decision of the DM.
- The decision maker does not observe the state of the world; she has prior beliefs that are uniformly distributed over $[-M, M]$.
- The parameter $1/M$ measures the prior familiarity of the DM with the decision problem.
- When $M = 0$, the DM is fully familiar with the problem; she knows that the state of the world is $\omega = 0$.
- When $M = \infty$, the DM is totally unfamiliar with the problem; her prior over ω is diffuse over the real line.

The Model...

- A decision maker must choose an alternative; the set of possible decisions is the real line.
- An expert observes the realization of the state of the world, $\omega \in [-M, M]$;the state of the world coincides with the first-best optimal decision of the DM.
- The decision maker does not observe the state of the world; she has prior beliefs that are uniformly distributed over $[-M, M]$.
- The parameter $1/M$ measures the prior familiarity of the DM with the decision problem.
- When $M = 0$, the DM is fully familiar with the problem; she knows that the state of the world is $\omega = 0$.
- When $M = \infty$, the DM is totally unfamiliar with the problem; her prior over ω is diffuse over the real line.

The Model...

- A decision maker must choose an alternative; the set of possible decisions is the real line.
- An expert observes the realization of the state of the world, $\omega \in [-M, M]$;the state of the world coincides with the first-best optimal decision of the DM.
- The decision maker does not observe the state of the world; she has prior beliefs that are uniformly distributed over $[-M, M]$.
- The parameter $1/M$ measures the prior familiarity of the DM with the decision problem.
- When $M = 0$, the DM is fully familiar with the problem; she knows that the state of the world is $\omega = 0$.
- When $M = \infty$, the DM is totally unfamiliar with the problem; her prior over ω is diffuse over the real line.

The Model...

- A decision maker must choose an alternative; the set of possible decisions is the real line.
- An expert observes the realization of the state of the world, $\omega \in [-M, M]$;the state of the world coincides with the first-best optimal decision of the DM.
- The decision maker does not observe the state of the world; she has prior beliefs that are uniformly distributed over $[-M, M]$.
- The parameter $1/M$ measures the prior familiarity of the DM with the decision problem.
- When $M = 0$, the DM is fully familiar with the problem; she knows that the state of the world is $\omega = 0$.
- When $M = \infty$, the DM is totally unfamiliar with the problem; her prior over ω is diffuse over the real line.

The Model...

- A decision maker must choose an alternative; the set of possible decisions is the real line.
- An expert observes the realization of the state of the world, $\omega \in [-M, M]$;the state of the world coincides with the first-best optimal decision of the DM.
- The decision maker does not observe the state of the world; she has prior beliefs that are uniformly distributed over $[-M, M]$.
- The parameter $1/M$ measures the prior familiarity of the DM with the decision problem.
- When $M = 0$, the DM is fully familiar with the problem; she knows that the state of the world is $\omega = 0$.
- When $M = \infty$, the DM is totally unfamiliar with the problem; her prior over ω is diffuse over the real line.

...The Model...

- The expert chooses a verifiable message m to send to the DM.
- The message consists of a closed interval $[a, b]$, which must include the true state of the world ω, $\omega \in [a, b]$.
- After receiving a message, the DM forms posterior beliefs about the state of nature and takes the decision d that corresponds to the expected value of the state of nature, given his posterior beliefs: $d = E[\omega | m = [a, b]]$.

...The Model...

- The expert chooses a verifiable message m to send to the DM.

- The message consists of a closed interval $[a, b]$, which must include the true state of the world ω, $\omega \in [a, b]$.

- After receiving a message, the DM forms posterior beliefs about the state of nature and takes the decision d that corresponds to the expected value of the state of nature, given his posterior beliefs: $d = E[\omega | m = [a, b]]$.

...The Model...

- The expert chooses a verifiable message m to send to the DM.

- The message consists of a closed interval $[a, b]$, which must include the true state of the world ω, $\omega \in [a, b]$.

- After receiving a message, the DM forms posterior beliefs about the state of nature and takes the decision d that corresponds to the expected value of the state of nature, given his posterior beliefs: $d = E[\omega | m = [a, b]]$.

...The Model...

- When the state of the world is ω, the ideal choice for the expert is $\theta = \omega + \beta$.
- The expert's goal is to minimize the distance between the decision d and θ.
- The parameter β, which is the expert's private information, represent the expert's bias.
- The decision maker has prior beliefs over β, which are uniformly distributed over the interval $[-\Delta + b, \Delta + b]$.
- The parameter $b \geq 0$ measures the mean bias of the expert.
- The parameter $\Delta \geq 0$ measures the expert's bias variance.

...The Model...

- When the state of the world is ω, the ideal choice for the expert is $\theta = \omega + \beta$.
- The expert's goal is to minimize the distance between the decision d and θ.
- The parameter β, which is the expert's private information, represent the expert's bias.
- The decision maker has prior beliefs over β, which are uniformly distributed over the interval $[-\Delta + b, \Delta + b]$.
- The parameter $b \geq 0$ measures the mean bias of the expert.
- The parameter $\Delta \geq 0$ measures the expert's bias variance.

...The Model...

- When the state of the world is ω, the ideal choice for the expert is $\theta = \omega + \beta$.
- The expert's goal is to minimize the distance between the decision d and θ.
- The parameter β, which is the expert's private information, represent the expert's bias.
- The decision maker has prior beliefs over β, which are uniformly distributed over the interval $[-\Delta + b, \Delta + b]$.
- The parameter $b \geq 0$ measures the mean bias of the expert.
- The parameter $\Delta \geq 0$ measures the expert's bias variance.

...The Model...

- When the state of the world is ω, the ideal choice for the expert is $\theta = \omega + \beta$.
- The expert's goal is to minimize the distance between the decision d and θ.
- The parameter β, which is the expert's private information, represent the expert's bias.
- The decision maker has prior beliefs over β, which are uniformly distributed over the interval $[-\Delta + b, \Delta + b]$.
- The parameter $b \geq 0$ measures the mean bias of the expert.
- The parameter $\Delta \geq 0$ measures the expert's bias variance.

...The Model...

- When the state of the world is ω, the ideal choice for the expert is $\theta = \omega + \beta$.
- The expert's goal is to minimize the distance between the decision d and θ.
- The parameter β, which is the expert's private information, represent the expert's bias.
- The decision maker has prior beliefs over β, which are uniformly distributed over the interval $[-\Delta + b, \Delta + b]$.
- The parameter $b \geq 0$ measures the mean bias of the expert.
- The parameter $\Delta \geq 0$ measures the expert's bias variance.

...The Model...

- When the state of the world is ω, the ideal choice for the expert is $\theta = \omega + \beta$.
- The expert's goal is to minimize the distance between the decision d and θ.
- The parameter β, which is the expert's private information, represent the expert's bias.
- The decision maker has prior beliefs over β, which are uniformly distributed over the interval $[-\Delta + b, \Delta + b]$.
- The parameter $b \geq 0$ measures the mean bias of the expert.
- The parameter $\Delta \geq 0$ measures the expert's bias variance.

...The Model

- Let $T = \{(\omega, \beta) : \omega \in [-M, M], \beta \in [-\Delta + b, -\Delta + b]\}$; T is the set of possible types of the expert.
- If an expert type (ω, β) with $\beta > 0$ sends a message $[\omega, y]$, we say that she only reports favorable information; if she sends a message $[x, \omega]$ she only reports unfavorable information.
- Similarly, an expert type with $\beta < 0$ only reports favorable information if she sends a message $[x, \omega]$, and only reports unfavorable information if she send $[\omega, y]$.
- An expert type with $\beta = 0$ only reports favorable information if she sends the message $[\omega, \omega]$.
- In all other cases, we say that the expert reports both favorable and unfavorable information.

...The Model

- Let $T = \{(\omega, \beta) : \omega \in [-M, M], \beta \in [-\Delta + b, -\Delta + b]\}$; T is the set of possible types of the expert.
- If an expert type (ω, β) with $\beta > 0$ sends a message $[\omega, y]$, we say that she only reports favorable information; if she sends a message $[x, \omega]$ she only reports unfavorable information.
- Similarly, an expert type with $\beta < 0$ only reports favorable information if she sends a message $[x, \omega]$, and only reports unfavorable information if she send $[\omega, y]$.
- An expert type with $\beta = 0$ only reports favorable information if she sends the message $[\omega, \omega]$.
- In all other cases, we say that the expert reports both favorable and unfavorable information.

...The Model

- Let $T = \{(\omega, \beta) : \omega \in [-M, M], \beta \in [-\Delta + b, -\Delta + b]\}$; T is the set of possible types of the expert.
- If an expert type (ω, β) with $\beta > 0$ sends a message $[\omega, y]$, we say that she only reports favorable information; if she sends a message $[x, \omega]$ she only reports unfavorable information.
- Similarly, an expert type with $\beta < 0$ only reports favorable information if she sends a message $[x, \omega]$, and only reports unfavorable information if she send $[\omega, y]$.
- An expert type with $\beta = 0$ only reports favorable information if she sends the message $[\omega, \omega]$.
- In all other cases, we say that the expert reports both favorable and unfavorable information.

...The Model

- Let $T = \{(\omega, \beta) : \omega \in [-M, M], \beta \in [-\Delta + b, -\Delta + b]\}$; T is the set of possible types of the expert.
- If an expert type (ω, β) with $\beta > 0$ sends a message $[\omega, y]$, we say that she only reports favorable information; if she sends a message $[x, \omega]$ she only reports unfavorable information.
- Similarly, an expert type with $\beta < 0$ only reports favorable information if she sends a message $[x, \omega]$, and only reports unfavorable information if she send $[\omega, y]$.
- An expert type with $\beta = 0$ only reports favorable information if she sends the message $[\omega, \omega]$.
- In all other cases, we say that the expert reports both favorable and unfavorable information.

...The Model

- Let $T = \{(\omega, \beta) : \omega \in [-M, M], \beta \in [-\Delta + b, -\Delta + b]\}$; T is the set of possible types of the expert.
- If an expert type (ω, β) with $\beta > 0$ sends a message $[\omega, y]$, we say that she only reports favorable information; if she sends a message $[x, \omega]$ she only reports unfavorable information.
- Similarly, an expert type with $\beta < 0$ only reports favorable information if she sends a message $[x, \omega]$, and only reports unfavorable information if she send $[\omega, y]$.
- An expert type with $\beta = 0$ only reports favorable information if she sends the message $[\omega, \omega]$.
- In all other cases, we say that the expert reports both favorable and unfavorable information.

...The Model

- Let $D(\omega, \beta)$ be the (perfect Bayesian) equilibrium decision when the true state and bias are ω, β.
- Disclosure games typically have multiple equilibria associated with a given decision map $D(\cdot)$. Since equilibria that always induce the same decision are payoff equivalent, I will not distinguish among them and focus on the equilibrium decision map D.
- For example, when $\Delta = 0$ and $b > 0$, as in the classic model of Grossman (1981) and Milgrom (1981), for all $y \in [\omega, M]$ there is an equilibrium in which $m(\omega, \beta) = [\omega, y]$ and $d([x, y]) = x$. In all these equilibria the decision map is $D(\omega, \beta) = \omega$ for all (ω, θ) (there is full disclosure).

...The Model

- Let $D(\omega, \beta)$ be the (perfect Bayesian) equilibrium decision when the true state and bias are ω, β.
- Disclosure games typically have multiple equilibria associated with a given decision map $D(\cdot)$. Since equilibria that always induce the same decision are payoff equivalent, I will not distinguish among them and focus on the equilibrium decision map D.
- For example, when $\Delta = 0$ and $b > 0$, as in the classic model of Grossman (1981) and Milgrom (1981), for all $y \in [\omega, M]$ there is an equilibrium in which $m(\omega, \beta) = [\omega, y]$ and $d([x, y]) = x$. In all these equilibria the decision map is $D(\omega, \beta) = \omega$ for all (ω, θ) (there is full disclosure).

…The Model

- Let $D(\omega, \beta)$ be the (perfect Bayesian) equilibrium decision when the true state and bias are ω, β.
- Disclosure games typically have multiple equilibria associated with a given decision map $D(\cdot)$. Since equilibria that always induce the same decision are payoff equivalent, I will not distinguish among them and focus on the equilibrium decision map D.
- For example, when $\Delta = 0$ and $b > 0$, as in the classic model of Grossman (1981) and Milgrom (1981), for all $y \in [\omega, M]$ there is an equilibrium in which $m(\omega, \beta) = [\omega, y]$ and $d([x, y]) = x$. In all these equilibria the decision map is $D(\omega, \beta) = \omega$ for all (ω, θ) (there is full disclosure).

Full Disclosure

- The classic model in which the (upward) bias of the expert is known to the DM, corresponds to the case in which $b > 0$ and $\Delta = 0$. In the classic model, the expert fully discloses the state of the world in equilibrium. The full disclosure result extends to the case in which the bias of the expert is not fully known by the DM, but it is known that the bias β is non-negative.

Theorem

If $\Delta \leq b$, then the equilibrium decision map is $D(\omega, \beta) = \omega$ for all $(\omega, \beta) \in T$; the expert fully reveals the true state of the world and the DM achieves her first-best outcome. In equilibrium the expert only reveals favorable information.

- In the remainder of the paper, we assume that $\Delta > b$.

Full Disclosure

- The classic model in which the (upward) bias of the expert is known to the DM, corresponds to the case in which $b > 0$ and $\Delta = 0$. In the classic model, the expert fully discloses the state of the world in equilibrium. The full disclosure result extends to the case in which the bias of the expert is not fully known by the DM, but it is known that the bias β is non-negative.

Theorem

If $\Delta \leq b$, then the equilibrium decision map is $D(\omega, \beta) = \omega$ for all $(\omega, \beta) \in T$; the expert fully reveals the true state of the world and the DM achieves her first-best outcome. In equilibrium the expert only reveals favorable information.

- In the remainder of the paper, we assume that $\Delta > b$.

Full Disclosure

- The classic model in which the (upward) bias of the expert is known to the DM, corresponds to the case in which $b > 0$ and $\Delta = 0$. In the classic model, the expert fully discloses the state of the world in equilibrium. The full disclosure result extends to the case in which the bias of the expert is not fully known by the DM, but it is known that the bias β is non-negative.

Theorem

If $\Delta \leq b$, then the equilibrium decision map is $D(\omega, \beta) = \omega$ for all $(\omega, \beta) \in T$; the expert fully reveals the true state of the world and the DM achieves her first-best outcome. In equilibrium the expert only reveals favorable information.

- In the remainder of the paper, we assume that $\Delta > b$.

Totally Unfamiliar DM, Zero-Mean-Bias Expert...

- When $M = \infty$, the DM is totally unfamiliar with the problem. When the expert's mean bias b is zero, the prior expected optimal decision of the expert coincides with the state of the world and the first best decision of the DM. The expert is ex-ante unbiased.

Theorem

If $M = \infty$ and $b = 0$, then $D(\omega, \beta) = \omega + \beta = \theta$ for all $(\omega, \beta) \in T$ is an equilibrium decision map; the decision taken always coincides with the expert's optimal (first-best) decision. A set of equilibrium strategies that implements this decision map is: $m(\omega, \beta) = [\omega - \Delta, \omega + \Delta]$; $d([x, y]) = \frac{x+y}{2}$.

Totally Unfamiliar DM, Zero-Mean-Bias Expert...

- When $M = \infty$, the DM is totally unfamiliar with the problem. When the expert's mean bias b is zero, the prior expected optimal decision of the expert coincides with the state of the world and the first best decision of the DM. The expert is ex-ante unbiased.

Theorem

If $M = \infty$ and $b = 0$, then $D(\omega, \beta) = \omega + \beta = \theta$ for all $(\omega, \beta) \in T$ is an equilibrium decision map; the decision taken always coincides with the expert's optimal (first-best) decision. A set of equilibrium strategies that implements this decision map is: $m(\omega, \beta) = [\omega - \Delta, \omega + \Delta]$; $d([x, y]) = \frac{x+y}{2}$.

Totally Unfamiliar DM, Zero-Mean-Bias Expert - Isodecision curves (in red)

All expert types on an isodecision curve have the same preferred decision.
All get what they want; all expert types fully manipulate the DM.

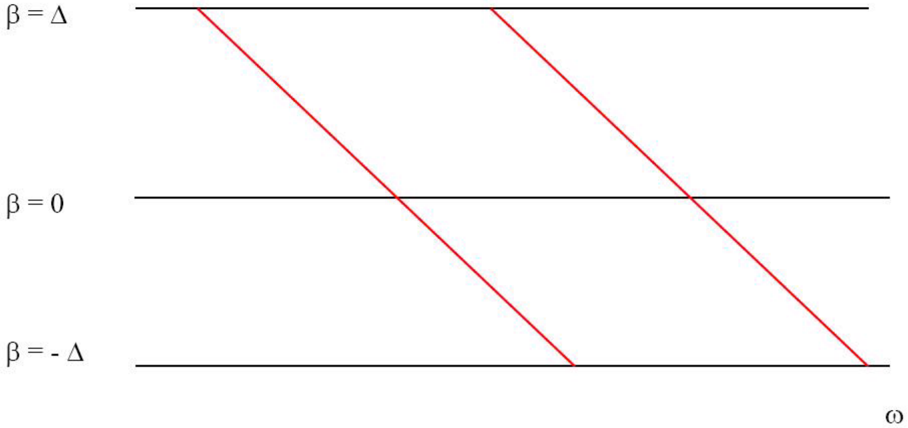

...Totally Unfamiliar DM, Zero-Mean-Bias Expert

- The DM's posterior beliefs attach equal probability to all states of nature belonging to the message interval sent by the expert, and hence the DM chooses the mid point of the message interval. By sending an interval centered around his ideal point, the expert induces the DM to choose the expert's ideal point. The expert always gets what she wants. Full manipulation.

- In the canonical equilibrium described in the theorem, except for maximally biased types ($\beta = \pm\Delta$), the expert's message contains both favorable and unfavorable information. However, there is also an equilibrium, with the same decision map, in which all types only send favorable information: types (ω, β), with $\beta > 0$ sends the message $[\omega, \omega + 2\beta]$, while types (ω, β), with $\beta < 0$ sends the message $[\omega + 2\beta, \omega]$.

...Totally Unfamiliar DM, Zero-Mean-Bias Expert

- The DM's posterior beliefs attach equal probability to all states of nature belonging to the message interval sent by the expert, and hence the DM chooses the mid point of the message interval. By sending an interval centered around his ideal point, the expert induces the DM to choose the expert's ideal point. The expert always gets what she wants. Full manipulation.

- In the canonical equilibrium described in the theorem, except for maximally biased types ($\beta = \pm\Delta$), the expert's message contains both favorable and unfavorable information. However, there is also an equilibrium, with the same decision map, in which all types only send favorable information: types (ω, β), with $\beta > 0$ sends the message $[\omega, \omega + 2\beta]$, while types (ω, β), with $\beta < 0$ sends the message $[\omega + 2\beta, \omega]$.

Totally Unfamiliar DM, Positive-Mean-Bias Expert...

- When the mean bias is positive, a positive measure of expert types get their first-best outcome, while all other types get less than their first-best choice.

- It is useful to define the parameter ℓ^* :

$$\ell^* = \frac{\Delta + b - \sqrt{(\Delta + b)^2 - \left(2\sqrt{2} - 2\right)(\Delta - b)^2}}{2 - \sqrt{2}}$$

Totally Unfamiliar DM, Positive-Mean-Bias Expert...

- When the mean bias is positive, a positive measure of expert types get their first-best outcome, while all other types get less than their first-best choice.

- It is useful to define the parameter ℓ^* :

$$\ell^* = \frac{\Delta + b - \sqrt{(\Delta + b)^2 - \left(2\sqrt{2} - 2\right)(\Delta - b)^2}}{2 - \sqrt{2}}$$

...Totally Unfamiliar DM, Positive-Mean-Bias Expert...

Theorem

If $M = \infty$ and $\Delta > b \geq 0$, then the following is an equilibrium decision map:

$$D(\omega, \beta) = \begin{cases} \omega + \beta & \text{if} \quad \beta \leq \ell^* \\ \omega + \ell^* & \text{if} \quad \beta > \ell^* \end{cases}$$

An equilibrium that implements this decision map is:

$$m(\omega, \theta) = \begin{cases} [\omega + \beta - \ell^*, \omega + \beta + \Delta - b] & \text{if} \quad \beta \leq \ell^* \\ [\omega, \omega + \ell^* + \Delta - b] & \text{if} \quad \beta > \ell^* \end{cases}$$

$$d([x, y]) = \min\{x + \ell^*, y\}.$$

Totally Unfamiliar DM, Positive-Mean-Bias Expert

Expert types on the diagonal part of the isocurve get what they want (full manipulation). All types on the vertical part of an isocurve get a lower decision than their most preferred decision (partial manipulation).

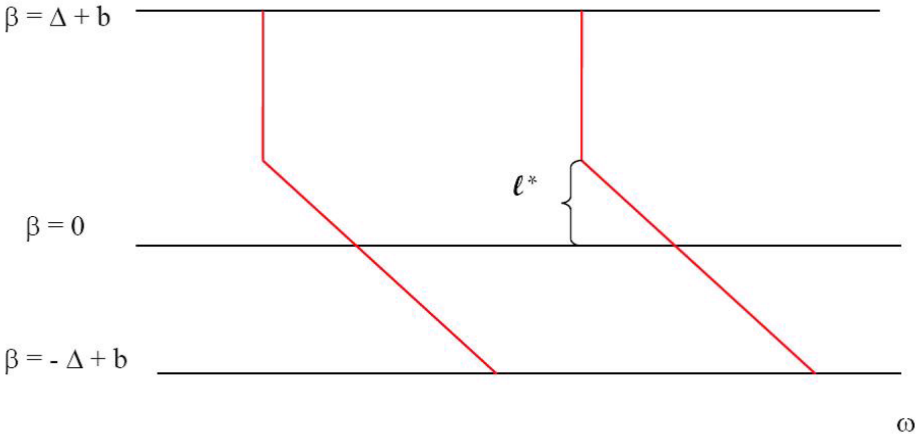

...Totally Unfamiliar DM, Positive-Mean-Bias Expert...

- All expert types (ω, β) whose bias is not greater than ℓ^*, or, equivalently, whose preferred choice θ is either below or above the true state ω by at most ℓ^*, induce the DM to take their favorite decision.
- Strongly upward biased expert types with $\beta > \ell^*$,or $\theta > \omega + \ell^*$, induce the decision maker to choose $\omega + \ell^*$.
- Since $\ell^* > 0$, all experts with a downward bias, and all experts with an upward bias less than ℓ^*, are able to induce the DM to choose their favorite decision (full manipulation).
- In all equilibria there must be at least a positive measure of downward biased types, who send both favorable and unfavorable information (in order to pool with types with an upward bias above ℓ^*).

...Totally Unfamiliar DM, Positive-Mean-Bias Expert...

- All expert types (ω, β) whose bias is not greater than ℓ^*, or, equivalently, whose preferred choice θ is either below or above the true state ω by at most ℓ^*, induce the DM to take their favorite decision.
- Strongly upward biased expert types with $\beta > \ell^*$,or $\theta > \omega + \ell^*$, induce the decision maker to choose $\omega + \ell^*$.
- Since $\ell^* > 0$, all experts with a downward bias, and all experts with an upward bias less than ℓ^*, are able to induce the DM to choose their favorite decision (full manipulation).
- In all equilibria there must be at least a positive measure of downward biased types, who send both favorable and unfavorable information (in order to pool with types with an upward bias above ℓ^*).

...Totally Unfamiliar DM, Positive-Mean-Bias Expert...

- All expert types (ω, β) whose bias is not greater than ℓ^*, or, equivalently, whose preferred choice θ is either below or above the true state ω by at most ℓ^*, induce the DM to take their favorite decision.
- Strongly upward biased expert types with $\beta > \ell^*$,or $\theta > \omega + \ell^*$, induce the decision maker to choose $\omega + \ell^*$.
- Since $\ell^* > 0$, all experts with a downward bias, and all experts with an upward bias less than ℓ^*, are able to induce the DM to choose their favorite decision (full manipulation).
- In all equilibria there must be at least a positive measure of downward biased types, who send both favorable and unfavorable information (in order to pool with types with an upward bias above ℓ^*).

...Totally Unfamiliar DM, Positive-Mean-Bias Expert...

- All expert types (ω, β) whose bias is not greater than ℓ^*, or, equivalently, whose preferred choice θ is either below or above the true state ω by at most ℓ^*, induce the DM to take their favorite decision.
- Strongly upward biased expert types with $\beta > \ell^*$,or $\theta > \omega + \ell^*$, induce the decision maker to choose $\omega + \ell^*$.
- Since $\ell^* > 0$, all experts with a downward bias, and all experts with an upward bias less than ℓ^*, are able to induce the DM to choose their favorite decision (full manipulation).
- In all equilibria there must be at least a positive measure of downward biased types, who send both favorable and unfavorable information (in order to pool with types with an upward bias above ℓ^*).

...Totally Unfamiliar DM, Positive-Mean-Bias Expert...

- We can think of ℓ^* as a measure of the set of upward biased experts that obtain what they want (fully manipulate)

Proposition

The measure ℓ^* of the upward biased experts that obtain their first best choice is a decreasing function of the mean bias b, and an increasing function of the bias variance Δ.

...Totally Unfamiliar DM, Positive-Mean-Bias Expert...

- We can think of ℓ^* as a measure of the set of upward biased experts that obtain what they want (fully manipulate)

Proposition

The measure ℓ^* of the upward biased experts that obtain their first best choice is a decreasing function of the mean bias b, and an increasing function of the bias variance Δ.

...Totally Unfamiliar DM, Positive-Mean-Bias Expert...

- Define $L^E(\omega)$ as the ex-ante expected distance (loss) of the decision from the expert's optimal choice, conditional on the true state being ω. Let $L^E = E\left[L^E(\omega)\right]$ be the ex-ante expected loss

Proposition

If $M = \infty$ and $\Delta > b \geq 0$, then the ex-ante expected loss of the expert is

$$L^E = \frac{(\Delta + b - \ell^*)^2}{4\Delta}.$$

The ex-ante expected loss of the expert's is a decreasing function of the expert's bias variance Δ; it is an increasing function of the expert's mean bias b.

...Totally Unfamiliar DM, Positive-Mean-Bias Expert...

- Define $L^E(\omega)$ as the ex-ante expected distance (loss) of the decision from the expert's optimal choice, conditional on the true state being ω. Let $L^E = E\left[L^E(\omega)\right]$ be the ex-ante expected loss

Proposition

If $M = \infty$ and $\Delta > b \geq 0$, then the ex-ante expected loss of the expert is

$$L^E = \frac{(\Delta + b - \ell^*)^2}{4\Delta}.$$

The ex-ante expected loss of the expert's is a decreasing function of the expert's bias variance Δ; it is an increasing function of the expert's mean bias b.

...Totally Unfamiliar DM, Positive-Mean-Bias Expert...

- Define $L^{DM}(\beta)$ as the expected distance (loss) of the decision from the DM's optimal choice, conditional on the expert's bias being β:

$$L^{DM}(\beta) = \begin{cases} \ell^* & \text{if} \quad \beta \geq \ell^* \\ \beta & \text{if} \quad 0 \leq \beta \leq \ell^* \\ -\beta & \text{if} \quad \beta \leq 0 \end{cases}$$

Proposition

If $M = \infty$ and $\Delta > b \geq 0$, the ex-ante expected loss of the DM is

$$L^{DM} = E[L^{DM}(\beta)] = \frac{(\Delta + b)^2 + (\Delta - b)^2 - (\Delta + b - \ell^*)^2}{4\Delta}.$$

The expected loss increases with the variance in the expert's bias Δ; it decreases with the expert's mean bias b.

...Totally Unfamiliar DM, Positive-Mean-Bias Expert...

- Define $L^{DM}(\beta)$ as the expected distance (loss) of the decision from the DM's optimal choice, conditional on the expert's bias being β:

$$L^{DM}(\beta) = \begin{cases} \ell^* & \text{if} \quad \beta \geq \ell^* \\ \beta & \text{if} \quad 0 \leq \beta \leq \ell^* \\ -\beta & \text{if} \quad \beta \leq 0 \end{cases}$$

Proposition

If $M = \infty$ and $\Delta > b \geq 0$, the ex-ante expected loss of the DM is

$$L^{DM} = E[L^{DM}(\beta)] = \frac{(\Delta + b)^2 + (\Delta - b)^2 - (\Delta + b - \ell^*)^2}{4\Delta}.$$

The expected loss increases with the variance in the expert's bias Δ; it decreases with the expert's mean bias b.

...Totally Unfamiliar DM, Positive-Mean-Bias Expert...

- An increase in the variance of the expert's bias hurts the decision maker, but it helps the expert.
- An interesting result is that an increase in the mean bias of the expert benefits the DM and it hurts the expert.
- Experts should strive to look ex-ante unbiased, even if this implies an increase in the variance of their perceived bias.
- Decision makers should look for experts with low variance in their bias, and/or experts with high mean bias.
- Intuitively, messages from highly biased experts are easier to read.

...Totally Unfamiliar DM, Positive-Mean-Bias Expert...

- An increase in the variance of the expert's bias hurts the decision maker, but it helps the expert.
- An interesting result is that an increase in the mean bias of the expert benefits the DM and it hurts the expert.
- Experts should strive to look ex-ante unbiased, even if this implies an increase in the variance of their perceived bias.
- Decision makers should look for experts with low variance in their bias, and/or experts with high mean bias.
- Intuitively, messages from highly biased experts are easier to read.

...Totally Unfamiliar DM, Positive-Mean-Bias Expert...

- An increase in the variance of the expert's bias hurts the decision maker, but it helps the expert.
- An interesting result is that an increase in the mean bias of the expert benefits the DM and it hurts the expert.
- Experts should strive to look ex-ante unbiased, even if this implies an increase in the variance of their perceived bias.
- Decision makers should look for experts with low variance in their bias, and/or experts with high mean bias.
- Intuitively, messages from highly biased experts are easier to read.

...Totally Unfamiliar DM, Positive-Mean-Bias Expert...

- An increase in the variance of the expert's bias hurts the decision maker, but it helps the expert.
- An interesting result is that an increase in the mean bias of the expert benefits the DM and it hurts the expert.
- Experts should strive to look ex-ante unbiased, even if this implies an increase in the variance of their perceived bias.
- Decision makers should look for experts with low variance in their bias, and/or experts with high mean bias.
- Intuitively, messages from highly biased experts are easier to read.

...Totally Unfamiliar DM, Positive-Mean-Bias Expert...

- An increase in the variance of the expert's bias hurts the decision maker, but it helps the expert.
- An interesting result is that an increase in the mean bias of the expert benefits the DM and it hurts the expert.
- Experts should strive to look ex-ante unbiased, even if this implies an increase in the variance of their perceived bias.
- Decision makers should look for experts with low variance in their bias, and/or experts with high mean bias.
- Intuitively, messages from highly biased experts are easier to read.

Partially Familiar DM

On an isocurve (region): Types on the diagonal get what they want; Types on the vertical line (region) get a compromise decision between their and the DM's ideal.

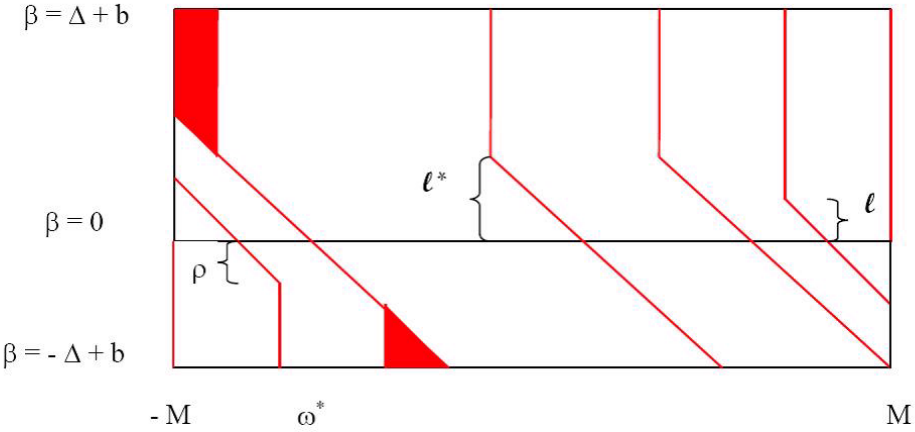

Partially Familiar DM...

$$\rho(\omega) = \frac{\Delta - b - \sqrt{(\Delta - b)^2 - \left(2\sqrt{2} - 2\right)\left(M + \min\{\omega, -M + \Delta - b\}\right)^2}}{2 - \sqrt{2}}$$

$$\ell(\omega) = \frac{\Delta + b - \sqrt{(\Delta + b)^2 - \left(2\sqrt{2} - 2\right)\left(M - \max\{\omega, M - \Delta + b\}\right)^2}}{2 - \sqrt{2}}$$

- $\ell(\omega)$ is strictly decreasing in ω for $\omega > M - \Delta + b$; $\ell(M) = 0$ and $\ell(\omega) = \ell^*$ for $\omega \leq M - \Delta + b$.
- $\rho(\omega)$ is strictly increasing in ω for $\omega < -M + \Delta - b$; $\rho(-M) = 0$ and $\rho(\omega) = \Delta - b$ for $\omega \geq -M + \Delta - b$.

$$\psi(\omega) = \int_{\rho(\omega)}^{\Delta - b} x\left(\Delta - b - x\right) dx + \int_{-M-\omega}^{-\ell(\omega)} x\left(\Delta + b + x\right) dx.$$

Lemma

There exists a unique state of the world ω^* *such that* $\psi(\omega^*) = 0$.

Partially Familiar DM...

$$\rho(\omega) = \frac{\Delta - b - \sqrt{(\Delta - b)^2 - \left(2\sqrt{2} - 2\right)(M + \min\{\omega, -M + \Delta - b\})^2}}{2 - \sqrt{2}}$$

$$\ell(\omega) = \frac{\Delta + b - \sqrt{(\Delta + b)^2 - \left(2\sqrt{2} - 2\right)(M - \max\{\omega, M - \Delta + b\})^2}}{2 - \sqrt{2}}$$

- $\ell(\omega)$ is strictly decreasing in ω for $\omega > M - \Delta + b$; $\ell(M) = 0$ and $\ell(\omega) = \ell^*$ for $\omega \leq M - \Delta + b$.
- $\rho(\omega)$ is strictly increasing in ω for $\omega < -M + \Delta - b$; $\rho(-M) = 0$ and $\rho(\omega) = \Delta - b$ for $\omega \geq -M + \Delta - b$.

$$\psi(\omega) = \int_{\rho(\omega)}^{\Delta - b} x(\Delta - b - x)\, dx + \int_{-M-\omega}^{-\ell(\omega)} x(\Delta + b + x)\, dx.$$

Lemma

There exists a unique state of the world ω^* such that $\psi(\omega^*) = 0$.

Partially Familiar DM...

$$\rho(\omega) = \frac{\Delta - b - \sqrt{(\Delta - b)^2 - \left(2\sqrt{2} - 2\right)(M + \min\{\omega, -M + \Delta - b\})^2}}{2 - \sqrt{2}}$$

$$\ell(\omega) = \frac{\Delta + b - \sqrt{(\Delta + b)^2 - \left(2\sqrt{2} - 2\right)(M - \max\{\omega, M - \Delta + b\})^2}}{2 - \sqrt{2}}$$

- $\ell(\omega)$ is strictly decreasing in ω for $\omega > M - \Delta + b$; $\ell(M) = 0$ and $\ell(\omega) = \ell^*$ for $\omega \leq M - \Delta + b$.
- $\rho(\omega)$ is strictly increasing in ω for $\omega < -M + \Delta - b$; $\rho(-M) = 0$ and $\rho(\omega) = \Delta - b$ for $\omega \geq -M + \Delta - b$.

$$\psi(\omega) = \int_{\rho(\omega)}^{\Delta - b} x(\Delta - b - x)\,dx + \int_{-M-\omega}^{-\ell(\omega)} x(\Delta + b + x)\,dx.$$

Lemma

There exists a unique state of the world ω^* such that $\psi(\omega^*) = 0$.

...Partially Familiar DM...

Theorem

If $M < \infty$ and $\Delta > b \geq 0$, then the equilibrium decision map is:

$$D(\omega, \beta) =$$

$$\begin{cases} \omega + \beta & \text{if } \omega \geq \omega^* - \beta;\ \beta \leq \ell(\omega + \beta) \\ \omega + \ell(\omega) & \text{if } \omega \geq \max\{\omega^* - \beta, \omega^* - \ell(\omega^*)\};\ \beta > \ell(\omega + \beta) \\ \omega^* & \text{if } \omega^* - \ell(\omega^*) \geq \omega \geq \omega^* - \beta;\ \beta > \ell(\omega + \beta) \\ \omega^* & \text{if } \omega^* - \beta \geq \omega \geq \omega^* + \rho(\omega^*);\ \beta < -\rho(\omega + \beta) \\ \omega + \beta & \text{if } \omega \leq \omega^* - \beta;\ \beta \geq -\rho(\omega + \beta) \\ \omega + \rho(\omega) & \text{if } \omega \leq \min\{\omega^* - \beta, \omega^* + \rho(\omega^*)\};\ \beta < -\rho(\omega + \beta) \end{cases}$$

...Partially Familiar DM...

Theorem

Equilibrium strategies that implements this decision map are:

1) $m(\omega,\beta) =$

$$
\begin{cases}
[\omega+\beta-\ell(\omega+\beta), \min\{\omega+\beta+\Delta-b, M\}] & \text{if} \quad \omega \geq \omega^*-\beta;\ \beta \leq \ell(\omega+\beta) \\
[\omega, \min\{\omega+\ell(\omega)+\Delta-b, M\}] & \text{if} \quad \omega \geq \max\{\omega^*-\beta, \omega^*-\ell(\omega^*)\};\ \beta > \ell(\omega+\beta) \\
[-M, \min\{\omega^*+\Delta-b, M\}] & \text{if} \quad \omega^*-\ell(\omega^*) \geq \omega \geq \omega^*-\beta;\ \beta > \ell(\omega+\beta) \\
[-M, \min\{\omega^*+\Delta-b, M\}] & \text{if} \quad \omega^*-\beta \geq \omega \geq \omega^*+\rho(\omega^*);\ \beta < -\rho(\omega+\beta) \\
[-M, \omega+\beta+\rho(\omega+\beta)] & \text{if} \quad \omega \leq \omega^*-\beta;\ \beta \geq -\rho(\omega+\beta) \\
[-M, \omega] & \text{if} \quad \omega \leq \min\{\omega^*-\beta, \omega^*+\rho(\omega^*)\};\ \beta < -\rho(\omega+\beta)
\end{cases}
$$

...Partially Familiar DM...

Theorem

2) $d([x,y]) =$

$$\begin{cases} \min\{\lambda^{-1}(x), y\} & \textit{if} \quad x \geq \omega^* - \ell(\omega^*) \\ w^* & \textit{if} \quad x < \omega^* - \ell(\omega^*);\ y \geq \min\{\omega^* + \Delta - b, M\} \\ \max\{r^{-1}(y), x\} & \textit{if} \quad x < \omega^* - \ell(\omega^*);\ y < \min\{\omega^* + \Delta - b, M\} \end{cases}$$

where λ^{-1} is the inverse of the increasing function $\lambda(\omega) = w - \ell(\omega)$, and r^{-1} is the inverse of the increasing function $r(\omega) = \omega + \rho(\omega)$.

- Note that

$$\lambda^{-1}(x) = x + \ell(\lambda^{-1}(x)),$$

$$r^{-1}(y) = y - \rho(r^{-1}(x)).$$

...Partially Familiar DM...

Theorem

2) $d([x, y]) =$

$$\begin{cases} \min\{\lambda^{-1}(x), y\} & \textit{if} \quad x \geq \omega^* - \ell(\omega^*) \\ w^* & \textit{if} \quad x < \omega^* - \ell(\omega^*);\ y \geq \min\{\omega^* + \Delta - b, M\} \\ \max\{r^{-1}(y), x\} & \textit{if} \quad x < \omega^* - \ell(\omega^*);\ y < \min\{\omega^* + \Delta - b, M\} \end{cases}$$

where λ^{-1} is the inverse of the increasing function $\lambda(\omega) = w - \ell(\omega)$, and r^{-1} is the inverse of the increasing function $r(\omega) = \omega + \rho(\omega)$.

- Note that

$$\lambda^{-1}(x) = x + \ell(\lambda^{-1}(x)),$$

$$r^{-1}(y) = y - \rho(r{-}^{1}(x)).$$

...Partially Familiar DM

Proposition

The ex-ante expected loss of the expert is an increasing function of the DM's familiarity with the problem $1/M$.

- An increase in the DM's familiarity with the problem shrinks the size of the equilibrium region in the middle of the state space, where the expert's loss is smaller and the DM's loss is greater.

Proposition

The ex-ante expected loss of the DM is a decreasing function of the DM's familiarity with the problem $1/M$.

...Partially Familiar DM

Proposition

The ex-ante expected loss of the expert is an increasing function of the DM's familiarity with the problem $1/M$*.*

- An increase in the DM's familiarity with the problem shrinks the size of the equilibrium region in the middle of the state space, where the expert's loss is smaller and the DM's loss is greater.

Proposition

The ex-ante expected loss of the DM is a decreasing function of the DM's familiarity with the problem $1/M$*.*

Conclusions 1

- A positive measure of expert types obtain what they want (fully manipulates).
- If the DM is totally unfamiliar with the problem and the expert has zero mean bias, then the expert always get what she wants
- The DM is better off the more biased the expert is, i.e., the higher is the mean bias
- The expert is better off the lower is her mean bias
- The DM is made worse off by an increase in the bias variance
- The expert is better off if the bias variance increases
- Some types of experts send unfavorable unformation
- An increase in the DM's familiarity with the problem makes the DM better off

Conclusions 1

- A positive measure of expert types obtain what they want (fully manipulates).
- If the DM is totally unfamiliar with the problem and the expert has zero mean bias, then the expert always get what she wants
- The DM is better off the more biased the expert is, i.e., the higher is the mean bias
- The expert is better off the lower is her mean bias
- The DM is made worse off by an increase in the bias variance
- The expert is better off if the bias variance increases
- Some types of experts send unfavorable unformation
- An increase in the DM's familiarity with the problem makes the DM better off

Conclusions 1

- A positive measure of expert types obtain what they want (fully manipulates).
- If the DM is totally unfamiliar with the problem and the expert has zero mean bias, then the expert always get what she wants
- The DM is better off the more biased the expert is, i.e., the higher is the mean bias
- The expert is better off the lower is her mean bias
- The DM is made worse off by an increase in the bias variance
- The expert is better off if the bias variance increases
- Some types of experts send unfavorable unformation
- An increase in the DM's familiarity with the problem makes the DM better off

Conclusions 1

- A positive measure of expert types obtain what they want (fully manipulates).
- If the DM is totally unfamiliar with the problem and the expert has zero mean bias, then the expert always get what she wants
- The DM is better off the more biased the expert is, i.e., the higher is the mean bias
- The expert is better off the lower is her mean bias
- The DM is made worse off by an increase in the bias variance
- The expert is better off if the bias variance increases
- Some types of experts send unfavorable unformation
- An increase in the DM's familiarity with the problem makes the DM better off

Conclusions 1

- A positive measure of expert types obtain what they want (fully manipulates).
- If the DM is totally unfamiliar with the problem and the expert has zero mean bias, then the expert always get what she wants
- The DM is better off the more biased the expert is, i.e., the higher is the mean bias
- The expert is better off the lower is her mean bias
- The DM is made worse off by an increase in the bias variance
- The expert is better off if the bias variance increases
- Some types of experts send unfavorable unformation
- An increase in the DM's familiarity with the problem makes the DM better off

Conclusions 1

- A positive measure of expert types obtain what they want (fully manipulates).
- If the DM is totally unfamiliar with the problem and the expert has zero mean bias, then the expert always get what she wants
- The DM is better off the more biased the expert is, i.e., the higher is the mean bias
- The expert is better off the lower is her mean bias
- The DM is made worse off by an increase in the bias variance
- The expert is better off if the bias variance increases
- Some types of experts send unfavorable unformation
- An increase in the DM's familiarity with the problem makes the DM better off

Conclusions 1

- A positive measure of expert types obtain what they want (fully manipulates).
- If the DM is totally unfamiliar with the problem and the expert has zero mean bias, then the expert always get what she wants
- The DM is better off the more biased the expert is, i.e., the higher is the mean bias
- The expert is better off the lower is her mean bias
- The DM is made worse off by an increase in the bias variance
- The expert is better off if the bias variance increases
- Some types of experts send unfavorable unformation
- An increase in the DM's familiarity with the problem makes the DM better off

Conclusions 1

- A positive measure of expert types obtain what they want (fully manipulates).
- If the DM is totally unfamiliar with the problem and the expert has zero mean bias, then the expert always get what she wants
- The DM is better off the more biased the expert is, i.e., the higher is the mean bias
- The expert is better off the lower is her mean bias
- The DM is made worse off by an increase in the bias variance
- The expert is better off if the bias variance increases
- Some types of experts send unfavorable unformation
- An increase in the DM's familiarity with the problem makes the DM better off

Some Open Issues

- Do experts have the proper incentives to acquire information? Or do they tend to acquire to little or too much information?
- What if the DM has private information? Could it be the case that the DM does not want to reveal his preferences to the expert?
- Is there a role for multiple experts?
- Should the DM visit experts sequentially? Should the DM reveal to the current expert the report of previously visited experts?

Some Open Issues

- Do experts have the proper incentives to acquire information? Or do they tend to acquire to little or too much information?
- What if the DM has private information? Could it be the case that the DM does not want to reveal his preferences to the expert?
- Is there a role for multiple experts?
- Should the DM visit experts sequentially? Should the DM reveal to the current expert the report of previously visited experts?

Some Open Issues

- Do experts have the proper incentives to acquire information? Or do they tend to acquire to little or too much information?
- What if the DM has private information? Could it be the case that the DM does not want to reveal his preferences to the expert?
- Is there a role for multiple experts?
- Should the DM visit experts sequentially? Should the DM reveal to the current expert the report of previously visited experts?

Some Open Issues

- Do experts have the proper incentives to acquire information? Or do they tend to acquire to little or too much information?

- What if the DM has private information? Could it be the case that the DM does not want to reveal his preferences to the expert?

- Is there a role for multiple experts?

- Should the DM visit experts sequentially? Should the DM reveal to the current expert the report of previously visited experts?