Rethinking Incentive Compatibility: The Hidden Cost of BDM Belief Elicitation

Christina Che Sun

ESA North American Conference

November 11, 2022

Overview

- Question: Does revealing BDM incentives increase truthful reporting?
- Main Result: No. False reports increases when incentives are revealed due to subject misunderstanding of incentives
- Effort (measured by time spent) on tasks is the same across treatments.

Motivation

Beliefs play an important role in economics research

- Game theory (Bellemare et al., 2008; Costa-Gomes and Weizsacker, 2008)
- Macroeconomics (Guiso and Parigi, 1999; Branch, 2004; Engelberg et al., 2011)
- Behavioral economics (Coffman et al., 2019; Zimmermann, 2020; Barron, 2020; Ashraf et al., 2006; Costa-Gomes, et al., 2015; Zimmermann and Schwerter, 2022)

Motivation

Key Problem: How do we encourage people to report their beliefs truthfully?

- Incentive Compatible Mechanisms: given the incentives, the dominant strategy for subjects is to report their beliefs truthfully
  - Quadratic Scoring Rule (QSR. Brier, 1950)
  - Binarized Scoring Rule (BSR. Hossain and Okui, 2013)
  - Becker-DeGroot-Marschack (BDM, aka Probability Matching, Stochastic BDM. Becker et a., 1964; DuCharme and Donnell, 1973; Karni, 2009)

Motivation

BDM Belief Elicitation is very popular

- Drobner (2022): motivated beliefs
- Kim and Kim (2022): game theory
- Castagnetti and Schmacker (2022): motivated beliefs
- Dustan et al. (2022): gender
- Casoria et al. (2022): social identity & discrimination
- Yaouanq and Schwardmann (2022): present bias
- Guo and Recalder (2022): gender
- Mobius et al. (2021): motivated beliefs
- Kessel et al. (2021): gender
- Valero (2021): motivated beliefs
- Barron (2020): motivated beliefs
- Aguirregabiria and Xie (2020): game theory
- Coutts (2019): belief updating
- Buser et al. (2018): belief updating

Motivation

Recent evidence shows that TIC mechanisms may not perform as advertised in practice

- Cason and Plott (2014) shows that subjects misunderstand BDM value elicitation
- Danz et al. (2022) shows that BSR incentives lead to more false reports

We need to think about another type of IC: Behavioral Incentive Compatibility

Behavioral Incentive Compatibility

A BIC mechanism satisfies these 2 weak conditions (Danz et al., 2022):

1. Information on deployed incentives increases truthful revelations
   - Focus of this experiment
2. When given a choice over pure incentives, most participants select the outcome that is uniquely maximizing under the mechanism
   - Not universally testable. BSR can be reduced to 2 compound lotteries given belief, but BDM cannot due to the randomization device

Violation of BIC

- Misreporting from confusion (Cason and Plott, 2014)
- Misreporting from violation of assumptions for TIC
  - BSR requires EU for TIC. Subjects could have non-EU preferences (Dustan et al., 2022)
  - BDM does not need EU, so ex-ante we expect less misreporting than BSR

BDM in Practice

How do we elicit beliefs in the real world?

Example Task: What is the probability of randomly selecting a red urn out of these 10 urns?

We could directly ask people to tell the truth.

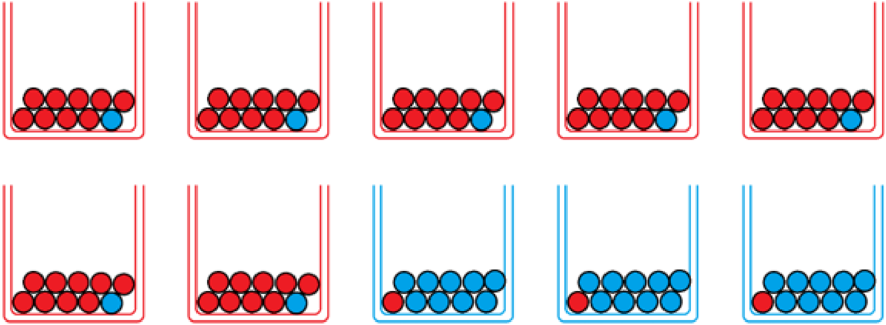

BDM in Practice

Now what if we use BDM?

Here is how the payment rule works:

After you report your guess in a question, the computer will use your guess to enter you into one of two lotteries: the **Event Lottery** and the **Number Lottery**. Your bonus earning for this question will be determined by playing the lottery you are entered into. The details are explained below.

First, the computer randomly draws a number N between 1 and 100. All whole numbers between 1 and 100 have equal chances of being drawn. After you report your guess, the computer compares this random number N and your reported guess, and

- **If your guess is greater than or equal to the random number N**, the computer will enter you into the **Event Lottery**. When you play this lottery, you will earn the bonus of $ for this question if the selected urn in this question is indeed red. If the selected urn in this question is blue, you will earn $0 for this question. (For example, this means that *if you report 100, you will always be paid by the Event Lottery*.)
- **If your guess is less than the random number N**, the computer will enter you into the **Number Lottery**. When you play the Number Lottery, your bonus earnings for this question depends on random chance. This lottery gives you N chances out of 100 to earn the bonus of $, and with (100-N) chances out of 100, you will earn $0. (For example, if the random number N is 70 and your reported guess is 50, then you will play the lottery which gives you a 70% chance to win the bonus, and a 30% chance to win $0. And *if you report 0, you will always be paid by the Number Lottery*)

BDM in Practice

Explanation of BDM Incentives: Diagram for Event Lottery

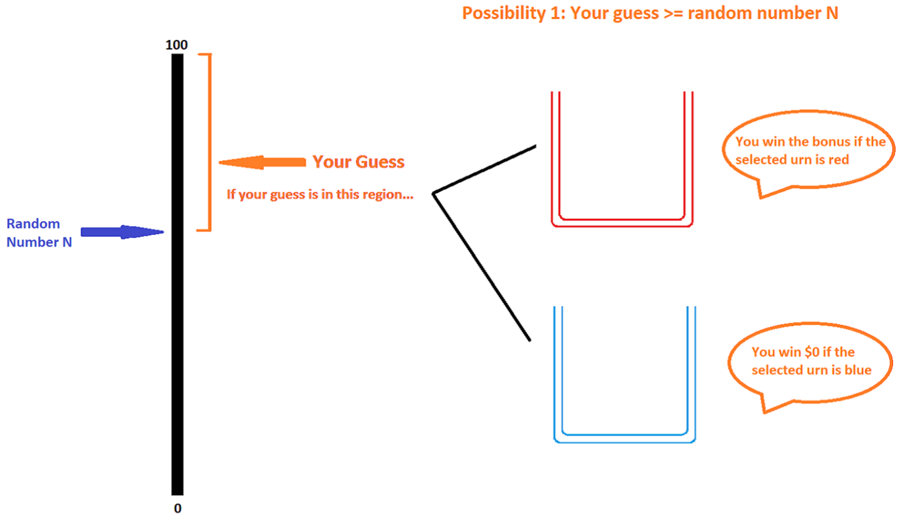

BDM in Practice

Explanation of BDM Incentives: Diagram for Number Lottery

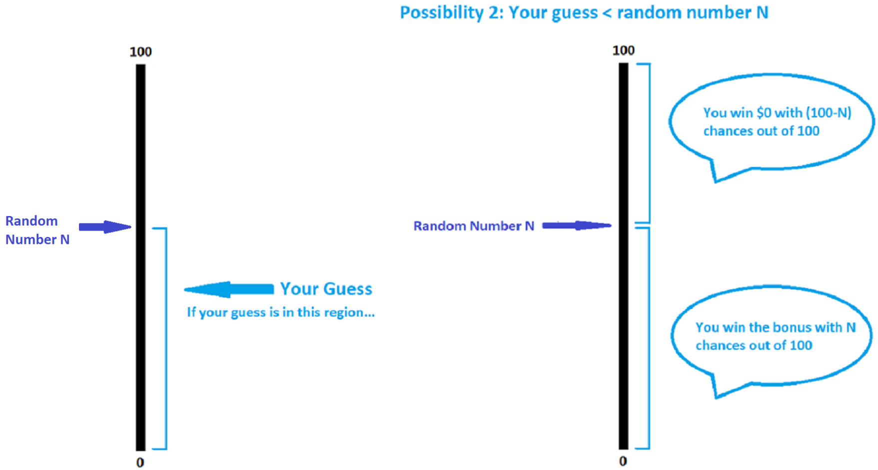

Design

Main treatments

- Full Information (real world)
  - Subjects learn the full incentive scheme of BDM
- No Information
  - Subjects do not receive incentive information
  - Simply told that truth-telling is in their best interest
  - Bonus is still determined by BDM

Design

- Endow subjects with objective priors
- 9 events (scenarios) with objective priors 10% to 90%
- Report priors + Bayesian updating task

Design

Incentives:

- Both treatments are incentivized, each question carries a bonus of \$0.5
- Bonus is calculated using BDM in both Full Info and No Info treatments
- All subjects are told truthful reporting maximizes chance of winning bonus
- *Only difference between treatments* is whether quantitative incentives are shown

Data

- Study conducted on Prolific
- First round of data collection in October 2022
- 99 subjects over 2 treatments

Data

Sample Summary Statistics

	Full Info		No Info		Overall	
	mean	sd	mean	sd	mean	sd
Age from Prolific	36.73	11.62	34.69	11.00	35.71	11.31
Gender: Woman	0.46	0.50	0.65	0.48	0.56	0.50
Gender: Man	0.52	0.50	0.31	0.47	0.41	0.50
Bachelor's degree or above	0.44	0.50	0.47	0.50	0.45	0.50
Taken probability/statistics course	0.44	0.50	0.39	0.49	0.41	0.50
Total approvals on Prolific	1278.78	962.20	1107.47	836.11	1193.12	900.83
Observations	50		49		99	

Results

Data patterns: fraction of false reports in Full Info and No Info treatments by scenario (left) and by prior (right)

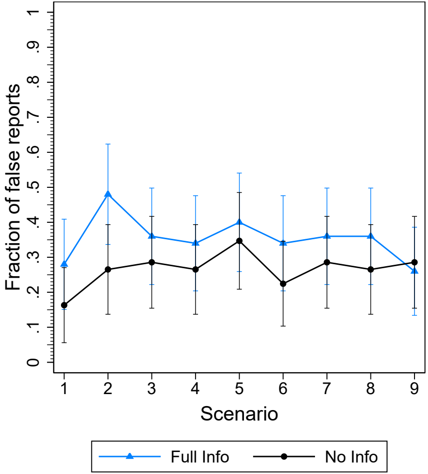

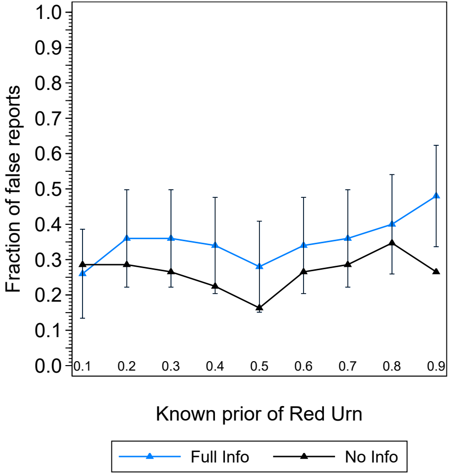

Results

What is driving the higher percentage of false reports in the Full Info treatment?

- Classify subjects by whether they understood the incentives
- Use the 4 questions that check incentive understanding at the end of survey

Results

Sample understanding check question 1:

Suppose in hypothetical question of a scenario, the randomly drawn number is 20, and you reported a guess of 15. Please answer the following using this information. You can see the payment rule here for more reference.

How will your bonus payment be determined?

- Event Lottery
- Number Lottery

Results

Sample understanding check question 2:

What are your chances of winning the bonus for this question?

- I win the bonus with 20% chance
- I win the bonus with 80% chance
- I win the bonus with 85% chance
- I win the bonus with 15% chance
- I win the bonus if the selected urn is red
- I win the bonus if the selected urn is blue

Results

Classification of understanding:

- If passed all 4 incentive understanding check questions, classify as understanding incentives
- If answered 1 or more incorrectly, classify as did not understand incentives

Results

Main Result: False reports in priors in the Full Info treatment are driven by subjects who did not understand incentives.

	False Report for Prior
Understanding	0.0127
	(0.0826)
Confused	0.226*
	(0.100)
Gender: Woman	0.106
	(0.0763)
Bachelor's degree or above	-0.142
	(0.0817)
Taken probability/statistics course	-0.0193
	(0.0806)
Observations	846

Standard errors in parentheses

OLS model, SE clustered at individual level

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Baseline: No Info treatment. Probit model gives the same results.

Conclusion

"Totally convoluted directions and way to much to digest and memorize for the pay. Honest feedback, Id skip the lottery part unless there is some purpose to it Im unaware of." - Subject Feedback

Conclusion

The Bottom line:

- Instructing subjects on incentives is more work for both the experimenter and the subject
  - In the Full Info treatment, mean time spent on the incentives page ~ 2.5 minutes (sd 3 min, max 20 min!)
  - In the No Info treatment, mean time spent is 25 seconds (max 1.3 minutes)
- The hidden cost of BDM is misreporting driven by confusion from complicated incentives

The End

Thank You!

ucsun@ucdavis.edu

Appendix

What is QSR

- Continuous payoff function
- Quadratic loss function
- Payoff is $[a - b(1-\mu)^2]$ if event obtains $(a > b > 0)$
- Payoff is $(a - b\mu^2)$ if event does not obtain

What is BSR

- Binary payoff
- Same loss function as QSR
- Receives payoff with probability $[1-(1-\mu)^2]$ if event obtains
- Receives payoff with probability $(1-\mu^2)$ otherwise

What is BDM?

We are interested in the subject's belief of event E, $\pi(E)$

- The subject is asked to report their belief $\mu \in [0, 1]$
- A random number r is generated from the uniform distribution $[0, 1]$
  - If $r \leq \mu$, the mechanism awards the subject the event bet $x_E y$ where $x > y$ (event lottery)
  - If $r > \mu$, the subject is awarded the number bet $l(r, x, y)$ (number lottery)

Overview of BDM

TIC of BDM requires 3 assumptions

Assumption (Probabilistic Sophistication)

A subject's preference relation $\succsim$ *over the set of finite acts and lotteries* D *is said to exhibit probabilistic sophistication if it ranks acts or lotteries solely on the basis of their probability distributions over outcomes.* *(Machina and Schmeidler, 1995)*

This means that for all acts f and lotteries $l(p, x, y) = [x, p; y, (1 - p)]$, $p \in [0, 1]$, $\pi(f^{-1}(x)) = p$ implies $x_{f^{-1}(x)} y \sim l(p, x, y)$

Overview of BDM

What does it mean for an agent to be probabilistically sophisticated?

- The agent **can** evaluate subjective probabilities (and does so)
- Preferences **need not** satisfy expected utility
- The agent is **not** ambiguity averse
- The agent **cannot** have preferences for certain sources of ambiguity

Overview of BDM

Assumption (Dominance)

A subject's preference relation over the set of finite acts and lotteries D exhibits dominance if for $x > y$, $p \geq p'$ implies $l(p, x, y) \succsim l(p', x, y)$, with strict preference if $p > p'$.

- This is a much simplified version of Axiom 5 (First-Order Stochastic Dominance Preference) in Machina and Schmeidler (1995).
- Subjects ranks acts and lotteries only based on their expected payoffs and never prefers a money lottery that's FOSD'd

Overview of BDM

Assumption (No Stakes)

The No Stakes assumption requires that subjects have no financial stakes in the realization of event E apart from the incentive payments on the belief elicitation.

TIC of BDM

Why is BDM attractive?

- It does not require risk neutrality
- Strictly speaking, EU is not required (Dominance is weaker than Independence Axiom)

Are there downsides?

- Difficult to explain
- Implementation can be cumbersome

Implementation of BDM

There are 4 main ways of implementing BDM:

- Direct elicitation (what I do)
- Random Binary Choice/MPL (Healy, 2020)
- Synchronized 2-Stage Lottery Choice Menu (Holt and Smith, 2016)
- Clock Auction (Karni, 2009; Hao and Hauser, 2012)

Implementation of BDM

Random Binary Choice (Healy, 2020):

Q#		Option A		Option B
1	Would you rather have:	\$20 if E occurs	or	1% chance of \$20
2	Would you rather have:	\$20 if E occurs	or	2% chance of \$20
3	Would you rather have:	\$20 if E occurs	or	3% chance of \$20
⋮	⋮	⋮	⋮	⋮
99	Would you rather have:	\$20 if E occurs	or	99% chance of \$20
100	Would you rather have:	\$20 if E occurs	or	100% chance of \$20

- A row is randomly selected and the choice is implemented
- Switching point is the belief of the event
- Downside: almost impossible to implement

Implementation of BDM

Two-Stage Synchronized Lottery Choice Menu (SC) (Holt and Smith, 2016):

- Two-stage version of RBC/MPL
- First present a coarse table, present a finer table depending on the previous answer
- Less burdensome than the full MPL but difficult to explain

Implementation of BDM

First stage of SC:

Option A: Earn $4 if the top scorer in the group has odd ID	A	B	**Option B**: Earn $4 if the number on the ball drawn from Basket B is in the range indicated on each line.
Option A	○	○	1-10
Option A	○	○	1-20
Option A	○	○	1-30
Option A	○	○	1-40
Option A	○	○	1-50
Option A	○	○	1-60
Option A	○	○	1-70
Option A	○	○	1-80
Option A	○	○	1-90

Implementation of BDM

Second stage of SC (if subject switched at 20 in the first stage):

Option A: Earn \$4 if the top scorer in the group has odd ID	A	B	**Option B**: Earn \$4 if the number on the ball drawn from Basket B is in the range indicated on each line.
Option A	○	○	1-11
Option A	○	○	1-12
Option A	○	○	1-13
Option A	○	○	1-14
Option A	○	○	1-15
Option A	○	○	1-16
Option A	○	○	1-17
Option A	○	○	1-18
Option A	○	○	1-19

Implementation of BDM

Clock Auction (Karni, 2009; Hao and Hauser, 2012):

- Subject plays in an English auction against a robot bidder who exits at a randomly drawn number $r \in [0, 100]$
- If the subject exits first, they are rewarded the lottery $l(r, x, 0)$
- If the robot exits first, subject is rewarded $x_E 0$
- The subject's exit point is $\pi(E)$
- More intuitive to explain, but subject beliefs are censored if robot exits first

Results: Priors

Result 3: False reports by subjects who did not understand BDM incentives pull to the center.

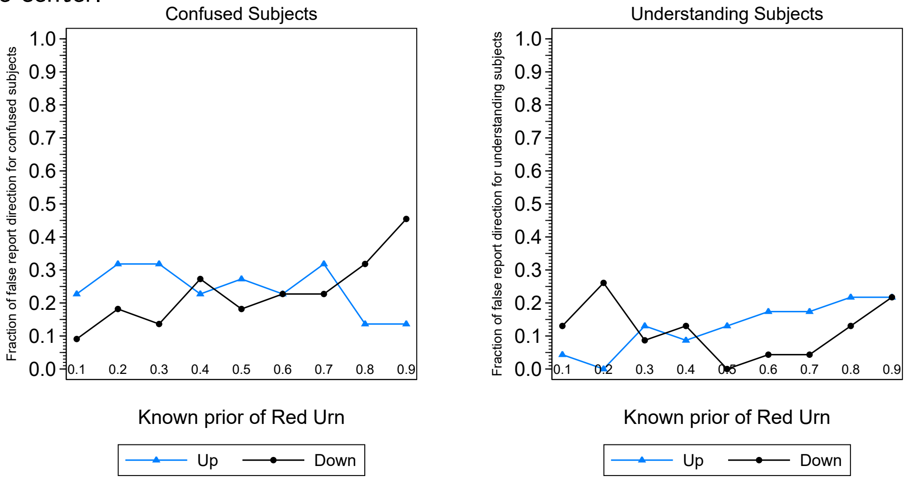

Results: Priors

Most belief elicitations are more complicated than scenarios with endowed priors

Result 4: Revealing incentives did not induce more effort in prior reports

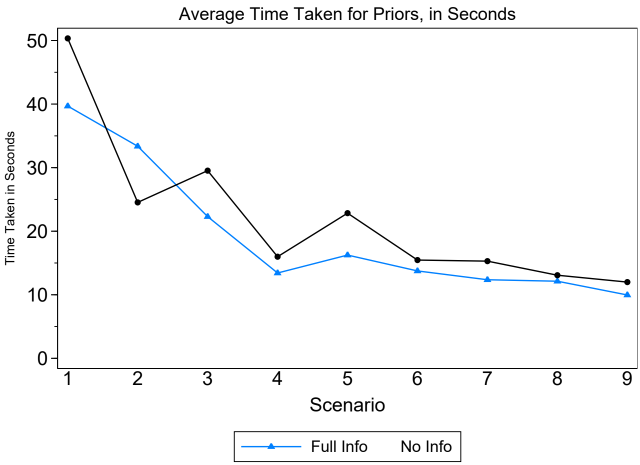

Results: Priors

Do subjects with a better understanding of probability perform better?

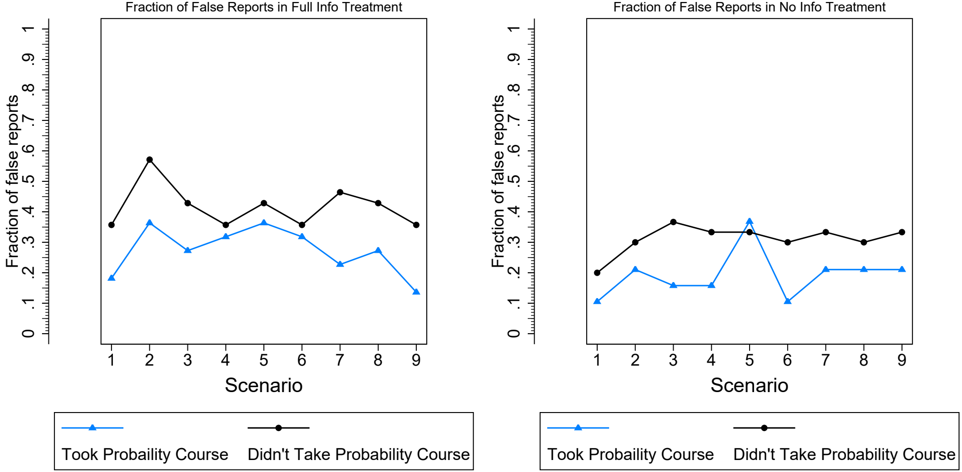

Results: Updates

- Incentives did not affect the quality of updates (deviation from Bayes, correct direction)
- Probability coursework did not affect the quality of updates
- Understanding incentives did not affect updates
- None of these factors influenced effort

Results: Updates

Q: Did incentives induce more effort on the more complicated task of belief updating?

A: No.

Result 6: Revealing full incentives did not induce more effort on the updating task.

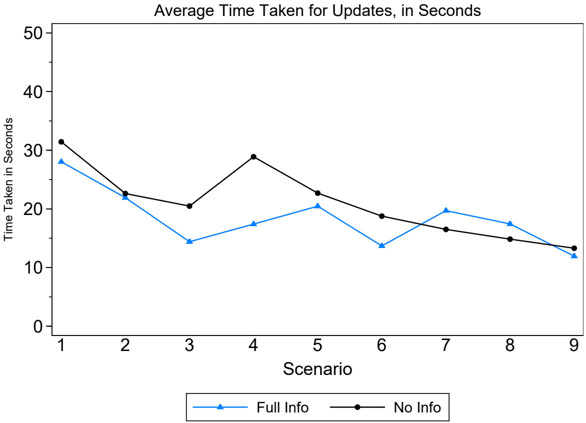

Results: Updates

Most subjects are conservative in their updates

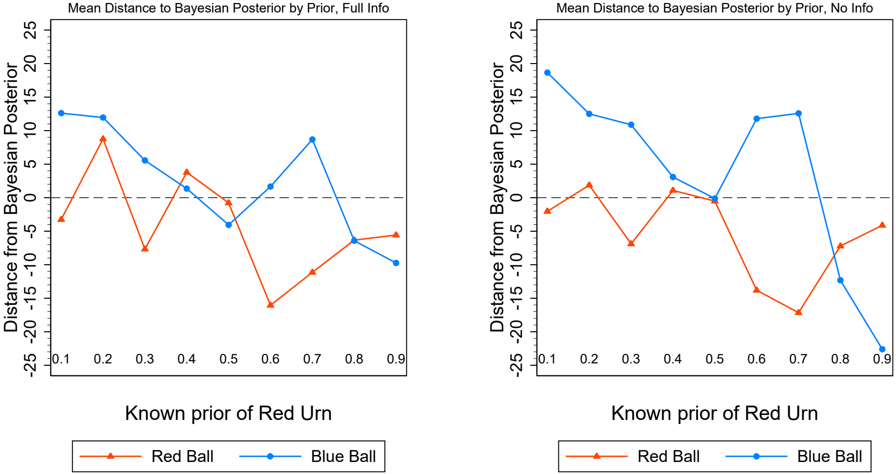

Results: Updates

Incentives did not affect mean absolute deviations from Bayesian posterior

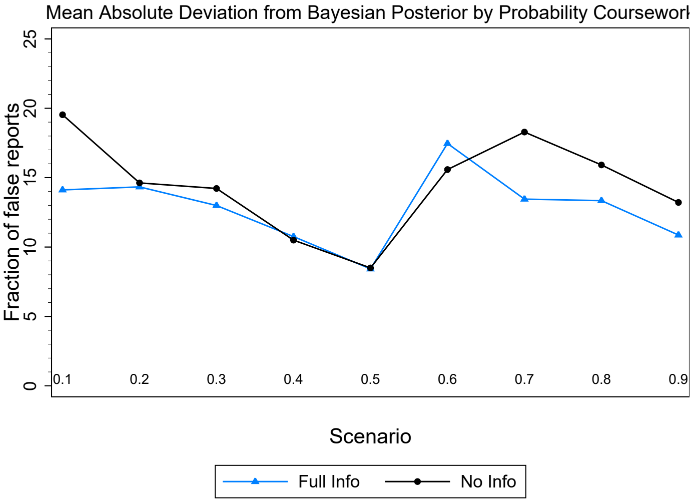

Results: Updates

Incentives did not impact whether subjects updated in the correct direction.