Competition and Coming Clean

Cary Deck and Dustin Tracy

June 21, 2020

Abstract

Conflicts of interest arise in many principal agent settings. Conventional wisdom suggests that the bias such incentives generate can be reduced with disclosure requirements, but at least some studies have found disclosure can exacerbate the level of bias. We consider an alternative force to reduce bias - market competition between agents and consider the effect of disclosure in the presence of market pressure. In each round of our experimental task, principals face a choice between two options and select among two agents from whom to seek advice. We test three treatments: one, in which the principals rate agents after receiving advice and other principals can condition which agent they select on those ratings; a second, which is similar to the first treatment but where the conflict of interest is disclosed; and a third in which market pressure is removed. Our results demonstrate that without disclosure there is no bias when there is competitive pressure while introducing disclosure generates bias. As expected, removing competitive pressure also leads to biased behavior.

Keywords: Conflict of Interest, Principal-Agent, Disclosure, Market Competition

JEL-Codes: *D82, G14, G18, I18, K12, L14, M55*

1 Introduction

For many interactions the parties have conflicting motivations. In standard principal agent problems, the benefits of the agent's effort accrue to the principal and thus they disagree on the agent's optimal effort. However, in practice the principal and agent recognize this conflict and understand that they ultimately have to compromise in order to both benefit. They can agree to some form of monitoring so principals can then construct incentive contracts to align the agent's incentives with their own. However when the principal does not full understand the agent's motivation, it can undermine the ability to reach mutually beneficial arrangements. For example, someone might not want to hire a lawyer to sue a company if they knew the lawyer's spouse worked for that company. Someone might be less willing to invest in a start-up if he knew the financial advisor promoting the company was the parent of the start-up's CEO. A patient might question a doctor more thoroughly if she knows the doctor receives a kickback from the pharmaceutical rep for the prescribed drug.

Lawyers, financial advisors, and doctors are said to have conflicts of interest if their personal motivation is at odds with their professional responsibilities to serve their clients (Boatright, 2000). Survey literature reveals conflicts are common and undermine trust in professionals. Scholl and Hung (2018) find there is generally low trust in finical professional, and while 51% of consumers would prefer their adviser did not have a conflict of interest, consumers lacked awareness of advisers' responsibilities and how they were compensated. Such conflicts of interest are particularly common among physicians; 94% of doctors a financial relationship with a pharmaceutical company, most commonly (83%) receiving gifts (meals), though 35% had received funding for continuing education or a conference and 28% had received payment for speaking or research (Campbell et al. (2007)). Patients viewed personal gifts to physician less favorably than office use gifts, believing they could impact the cost of care (Mainous (1995)). Patients who believe physicians accept gifts from the pharmaceutical industry trust doctors significantly less than those patients who do

not believe physicians accept gifts (Grande et al., 2012). Unsurprisingly, physicians believe that gifts from the pharmaceutical industry are less influential and more acceptable compared to patient attitudes Gibbons et al. (1998).

The conventional wisdom is that if a principal is aware of an agent's conflict of interest then the principal can take appropriate counteraction. As stated by Cain et al. (2005, p. 2), "Common sense suggests that recipients of advice will benefit from the being more fully informed when they are made aware of an advisor's conflict of interest." This notion has led to widespread requirements for disclosure of conflicts of interest, particularly in markets for "credence goods" where the client may not be able to determine the quality of service provided even after the fact.[1] The propensity for expert dishonesty in these markets has been established in economics through both lab and field experiment–see Kerschbamer and Sutter (2017) for a review of this literature.

There is general support for disclosure of conflicts although there is debate about the potential impact. A majority of patients would want to know if their doctor received a gift over $100 from a pharmaceutical company. A majority also say it would undermine their trust in their doctor if they knew their doctor had received such a gift (Green et al., 2012). Disclosing physicians' financial relations within the healthcare industry improved patient awareness of the relationship but had no impact on the likelihood of attending an appointment or stated trust and confidence in physicians (Rose et al., 2019). In Massachusetts, a disclosure law reduced prescriptions for both branded and generic antidepressants, antipsychotics and statins (Guo and Sriram, 2017). However, the 2012 Federal sunshine law had little impact on average payments from pharmaceutical companies to prescribes, but instead led to higher payments to fewer doctors (Guo et al., 2017).

In a simple experiment Cain et al. (2005) show that disclosure actually can exacerbate the problem introduced by a conflict of interest rather than reducing it. Specifically they find that in the absence of disclosure, agents provide biased estimates that encourage principals to take actions in the interest of the agent rather than the principal, but disclosing the conflicts of interest actually leads to substantially *increased* bias by agents. While the principals anticipate agent bias, they greatly underestimate its magnitude and thus disclosure actually harms the principals. While the counterintuitive finding of Cain et al. (2005) suggests that it may be optimal to not disclose conflicts of interest, their setting is one in which there is no competitive pressure on agents to perform well. The dishonest (criminal) behavior has little to no cost, so per Becker (1968) is expected to be prevalent. Reputation and the fear of lost future opportunity make dishonesty more costly. For example, Darby and Karni (1973) discuss how market forces and reputation can limit fraud for credence goods and Wolinsky (1993) presents a theoretical proof of how reputation and consumer search can discipline experts.

Thus we ask 1) if market pressure can eliminate bias on the part of agents and 2) what the effect of disclosure is in the presences of competitive pressure. The literature does not offer a clear prediction to either question. In credence markets, consumers often struggle to select an expert when there are many to select from because reputation can be noisy and it is not clear on which criteria the choice should be made. This may be particularly true in healthcare markets. Pope (2009) finds patients select hospitals based on rankings, but not underlying quality measures. There has been debate about the impact of rankings from the state Department of Public Health on mortality rates for coronary artery bypass graft surgery. Hannan et al. (1994) argue providing rankings improves quality and decreases mortality rates. Mukamel and Mushlin (1998) find the rankings increased market share for low mortality surgeons and hospitals leading to increased prices for the more highly demanded providers. Dranove et al. (2003) claim that the rankings reduce consumer welfare, at least in the short term as doctors attempting to improve their grades elect to operate on healthier patients. They argue this leads to increased costs, resulting in no improvement in the outcomes of the healthier patients, but worse outcomes for less healthy patients. Cain et al.'s results may well carry over to situations with reputational concerns if people fail to realize the bias agents exhibit. Carl (2008)

[1] Our good, in some respects, is an "Experience Goods", the principals eventually learn the true value, just as someone who buys a bottle of wine eventually tastes it. However, the distinction of whether experience leads to knowledge of the true value is not the key to assessing is one made the right decision. In our task, what is relevant is not value of the option, but its value net opportunity cost–the difference between selected option and the alternative. Just as our wine drinker might wonder: was the bottle not selected better? Our principals, only get a noisy signal of relative value so never are certain they made the choice.

actually find that disclosure of payments *increases* an agent's credibility when engaged in word of mouth product promotion in a setting with endogenous disclosure. More generally, one could argue that disclosure, particularly if there are many things that are being disclosed such as in housing mortgages, actually creates confusion and thus cost ultimately doing more harm than good. People may not know what to do with the disclosed information and can only pay attention to so many things, so it may be best to keep them focused on key features of the situation (Ben-shahar and Schneider (2011)).

We examine the effects of competitive pressure on agent bias using a similar task to Cain et al. (2005). Our principals are forced to rely on an agent's assessment of two options, similar to a patient whose physician is recommending between two courses of action. The principal receives a benefit based on the selected option that is independent of the agent's behavior. The agent has a financial incentive to encourage the principal to select a particular option, but critically the agent also has an incentive for principals to choose to solicit the agent's assessments. We compare treatments with and without the disclosure of the agent's conflicts of interest. In the absence of disclosure, while under the forces of market competition, we find that agents do not provide systematically biased assessments. Like Cain et al. (2005) we do find that disclosure creates a problem in that assessments become biased and principals are unable to fully incorporate that bias when making decisions. Further, we verify that removing the competitive pressure increases bias in assessments. Our results are consistent with a related experiment by Church and Kuang (2009), who show that the ability for the principal to sanction the agent can improve outcomes. However, the sanctions they introduce are meant to model actions such malpractice claims against doctors and thus are extreme and costly. In reality, there are many instances in which expert advice may not optimize consumer outcomes but the advice falls within the bounds of reasonable professional judgment or is otherwise unlikely to lead to formal sanctions. For example, a doctor may prescribe a medication that resulted in a worse outcome for the patient or exposes them to more potential side effects, but still has an FDA indication for the diagnosed ailment. Thus, our results show that market pressure can serve as a substitute for direct sanctions while being more broadly applicable.

2 Experimental Design

We first provide an overview of the main experimental task and then describe the various treatments. The task involves a principal ("investor" in the subject interface) that must select between two options labeled A and B. The options are jars of coins and the principal's value for an option equals the value of the coins in the associated jar. The principal cannot observe the jars directly, but selects an agent ("expert") who can observe a picture of the two jars. The agent reports an estimate of the value of the coins in each jar to the principal. The rationale for using jars of coins is that the agent only receives a noisy signal of the value of each option and thus always has plausible deniability, even to the researcher, that the reported estimates match the agent's actual beliefs.[2] The agent receives a fixed compensation of \$0.75 for each principal served plus \$0.50 if a served agent selects option A. After selecting an option, the value of the selected option was revealed to the principal as was a noisy signal of the value of the option that was not selected. The noisy signal was drawn from a uniform distribution centered at the true value with a range of \$4.00. The agents did not receive feedback on the quality of their estimates. They were only informed if they were selected and if they earned a bonus. Figure 1 provides a graphical depiction.

For each session, a group of 12 subjects entered the lab. Each subject was randomly assigned to either the role of an agent or to the role of a principal. Subjects maintained their roles throughout the experiment, which lasted for an undisclosed number of periods. Within a period, two principals and two agents were randomly matched. The principals then individually selected the agent with whom they wanted to interact that period.[3] Agents provided estimations for both options before learning how many principals would observe those estimations so that all agents made two estimates every period, akin to using the strategy method. Agents and principals were anonymous to minimize issues of reciprocity and other repeated play incentives.

[2]This process is similar to Cain et al. (2005) who used jars of coins for the same reason.

[3]Both principals could select the same agent from the pair.

Figure 1: Round Design

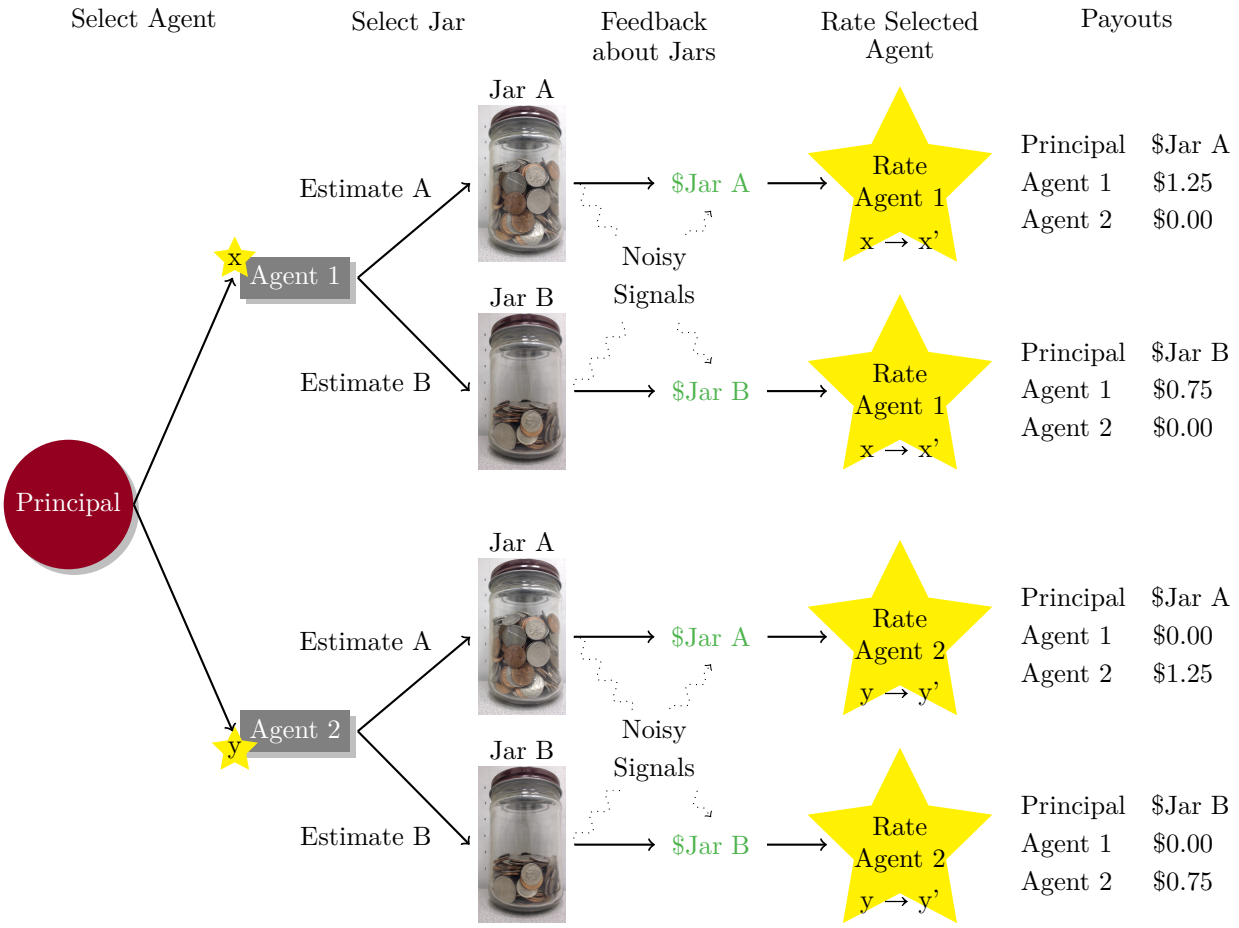

The main portion of our study consists of three between-subjects treatments. In the Market Treatment, the agent's conflict of interest was not disclosed to the principals. Additionally, principals rated the agent on a scale of one to five stars after the task was completed. Periodically during each session, the experiment was restarted by resetting all agent ratings.[4] The set of periods over which reputations persisted was referred to as a block. Blocks lasted at least five periods, but starting in the 5th period the block continued for another period with a 75% chance. We used the same realized random draw for block duration for all sessions; there were a total of 29 rounds across four blocks in each session. After the first period in each block, principals could observe the average rating of the two agents and select which one to use.[5]

The Disclosure Treatment was the same as the Market Treatment except that the conflict of interest was disclosed to principals in the instructions. Thus, the Disclosure Treatment continues to have the disciplining market forces as agents competed for future principals based on reputation. In the No Reputation Treatment, the conflict of interest was disclosed, but principals did not rate agents. Because this removes the dimension along which agents can compete for principals, choices of agents effectively became random.[6]

The jars that were used contained mixtures of quarters, dimes, nickels and pennies. The values of the contents ranged form \$13.32 to \$22.43 with a mean value of \$17.77. Figure 2 shows a sample image of a jar containing \$17.73. All subjects were informed that jars contained between \$10 and \$25. To control for jar specific effects, each jar of coins was observed twice by each subject at some point during the 29 periods. Further, all jars served as both option A and as option B.[7]

Figure 2: Sample Image

Guessing the value of the coins is challenging. Thus, we conducted a fourth treatment, which we refer to as the Guess Treatment, where a new group of subjects observed the jars used in the main experiment.

[4]This was done out of concern that an agent might develop a low rating early in the experiment while she was still learning how the process worked and then never being able to recover.

[5]Principals were assigned an agent in the first period of every block to ensure that every agent would have a rating in the second period.

[6]The two agents observed by two principals were listed in different orders in case principals were more likely to select the left or right one ceteris paribus.

[7]Due to a programming error, not all jars are observed as both Option A and Option B by all subjects. Subsequent analysis takes this into account.

These subjects were asked to report their best guess of the value of the coins in each of the 30 jars. A standard quadratic scoring rule was used to incentivize the subjects.

3 Results

For each of the three main treatments, we conducted three sessions with 12 subjects apiece for a total of 108 subjects. These sessions lasted up to 90 minutes and subjects earned an average of \$22.90 including a \$7 participation payment. Twenty four subjects participated in the Guess treatment in a single session. This session lasted one half hour and the average earnings were \$14.58 including the \$7 participation payment.

The results are presented in three subsections. The first subsection reports the estimates of jar values from the Guess Treatment as these serve as a reference point for estimates reported in the three the main treatments. The second subsection focuses on the behavior of agents in the No Reputation, Market, and Disclosure treatments while the third subsection reports the behavior of principals in those treatments.

3.1 Incentivized Estimates

The mean guess is \$14.61, while the mean of the true values was \$17.78.[8] Figure 3 shows a boxplot of the estimates from the Guess Treatment for each jar, ordered by jar value. The ◆'s show the actual values of the jars. The "wisdom of crowds" does not prevail.[9] Our subjects generally underestimate the value of the coins–the median estimate fell below the true value for 27 out 30 jars. A sign test rejects that errors are unbiased (p-value $< .001$). The mean error in estimation is \$3.17, while the mean absolute error is \$4.58. This leads to Finding 1.

Finding 1: Subjects systematically underestimate the value of jars, in contrast to the "wisdom of crowds".

Because Finding 1 shows that incentivized estimates are systematically different from true values, we use the estimates from the Guess Treatment as the basis of comparison for identifying agent bias in estimates from the main treatments.

3.2 Behavior of Agents

Table 1 reports the estimates made by agents for all jars by treatment, and by whether or not there was a bonus to the agent if the principal chose that jar. The first row reports the estimates for the jars when they are bonus earning and the second row reports estimates of jar value when those jars are non-bonus earning.

Table 1: Means by Treatment

	No Reputation	Baseline	Disclosure	Actual Value
Estimate Jar A	16.82	16.15	16.28	17.81
Estimate Jar B	16.65	16.36	15.88	17.81

Table 2 reports similar information to Table 1 but for only the cases in which agents saw the identical jar as both Jar A and Jar B.[10] However rather than raw estimates, this table reports deviations from estimates

[8] This number differs from the mean actual value reported in the next subsection because in the other treatments the majority of jars were observed twice while two were observed only once, given that there were an odd number of periods, as described in the previous section. When calculating the mean actual value for the main treatments we weight by the frequency that jars were seen.

[9] In his book, Surowiecki (2005) provides several examples where the mean of estimates from a group of individuals accurately reflect the true value (e.g. weight of an ox, or number of jelly beans in a jar).

[10] There was a logic error in the experiment's software which meant that even though all agents observed all jars twice, sometimes an agent saw a particular jar as Jar a twice or Jar B twice. The intention was that each agent would see all jars once as Jar A and once as Jar B.

Figure 3: Boxplots of Incentivized Estimates and Actual Values

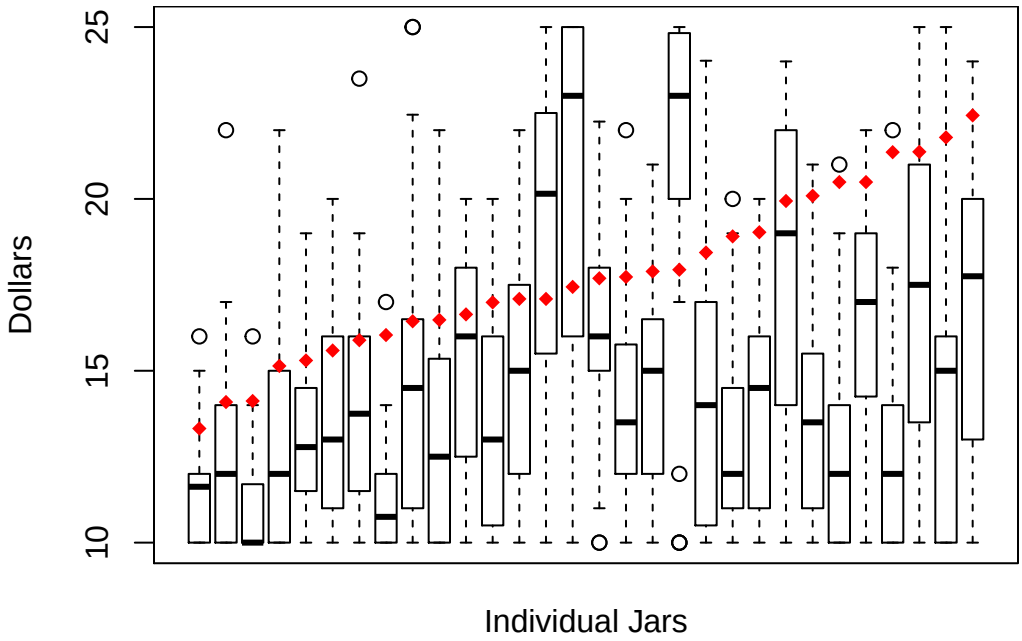

in the Guess Treatment. Additionally, Table 2 report values of (two-sided) paired t-tests of each agent's estimates for individual jars when the jar is and is not bonus eligible. For the No Reputation Treatment, the difference in the estimated value of the jar is statistical significant. That is, Agents estimate the value of a jar to be \$0.47 higher when it is bonus eligible than when that same jar is not bonus eligible. In the Market Treatment, when agents have reputation concerns, the difference is not statistically different from zero indicating that agents are not systematically favoring the bonus eligible jar in the Market Treatment. In the Disclosure Treatment, the difference is similar to that in the No Reputation Treatment suggesting that disclosure introduces bias into the estimation in the presence of competitive forces.

Table 2: Deviations from Incentivized Estimates for Cases in which each Jar was both A and B

	No Reputation	Baseline	Disclosure
Estimate Jar A	2.40	1.77	2.14
Estimate Jar B	1.92	1.89	1.63
Difference	0.47	-0.11	0.52
# Jar Pairings	324	300	300
T-test p value	0.01	0.51	0.01

Table 3 reports regression results, in which the dependent variable is the agents' estimate of jar value. In all specifications, a dummy variable indicating whether the jar was bonus eligible (Jar A) or not (Jar B) is interacted with a categorical variable indicating treatment. Specifications 1-4 use the actual jar value as a control while the specifications in Columns 5 and 6 use the jar specific mean estimates from the Guess Treatment. Specifications 2, 4 and 6 include dummy variables for each jar as controls. Specifications 3, 4 and 6 cluster errors by agents. The interaction of Jar B and Baseline Treatment is the omitted combination, therefore the "Jar A and Baseline" coefficient in each specification indicate how much higher estimates for Jar A are than Jar B in the Baseline. The results show Jar A estimates are actually 11.2 cents lower than

Jar B estimates. This is the opposite direction than would happen if agents were systematically pushing Jar A, but the difference is not statistically significant in any model. To get estimates of the differences between reported values of Jar A and B in the other treatments, we take the difference in the coefficients. Significance tests for the effect of being bonus eligible or not are reported at the bottom of Table 3 for the Disclosure and No Reputation treatments. Agent estimates for Jar A are 51.7 cents higher than agent estimates for Jar B in the Disclose Treatment. For the No Reputation Treatment, the difference is 47.3 cents. The reported p-values are for two-sided tests, thus if one expects agents to be biased due to the conflicts of interest then the p-values for the appropriate one-sided tests would be half of those reported in the table.

Together, the evidence in Table 2 and Table 3 support Finding 2.

Finding 2: In the absence of disclosure, agents do not exhibit bias when there is market pressure. However, in the presence of market pressure requiring disclosure leads to biased behavior by agents.

Table 3: Regression Estimates for Agents' Estimates of Jars

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)
Constant	9.110***		9.110***		2.631***	
	(0.701)		(0.722)		(0.587)	
Value	0.376***	0.899***	0.376***	0.899***		
	(0.0309)	(0.0266)	(0.0377)	(0.0218)		
Mean Guess					0.937***	1.019***
					(0.0296)	(0.0247)
Jar B & Disc	-0.498	-0.508	-0.498	-0.508	-0.507	-0.508
	(0.628)	(0.573)	(0.561)	(0.527)	(0.591)	(0.527)
Jar B & No Rep.	0.479	0.0320	0.479	0.0320	0.213	0.0320
	(0.622)	(0.570)	(0.517)	(0.505)	(0.587)	(0.505)
Jar A & Baseline	-0.112	-0.112	-0.112	-0.112	-0.112	-0.112
	(0.266)	(0.212)	(0.188)	(0.189)	(0.222)	(0.189)
Jar A & Disc	0.0197	0.00887	0.0197	0.00887	0.00989	0.00887
	(0.628)	(0.573)	(0.710)	(0.651)	(0.591)	(0.651)
Jar A & No Rep.	0.952	0.505	0.952*	0.505	0.686	0.505
	(0.622)	(0.570)	(0.502)	(0.487)	(0.587)	(0.487)
Observations	1,848	1,848	1,848	1,848	1,848	1,848
Number of Agents	54	54	54	54	54	54
Overall R2	0.0728	0.368	0.0728	0.368	0.313	0.368
Jar # Controls	N	Y	N	Y	N	Y
Clustering	N	N	Y	Y	N	Y
Tests of difference between Jar A and Jar B						
JarA - JarB for Disc	0.517	0.517	0.517	0.517	0.517	0.517
SE	(0.266)	(0.204)	(0.308)	(0.310)	(0.213)	(0.321)
P>|t|	0.0518	0.0146	0.137	0.140	0.0265	0.0950
JarA - JarB for No Rep	0.473	0.473	0.473	0.473	0.473	0.473
SE	(0.256)	(0.212)	(0.319)	(0.321)	(0.222)	(0.310)
P>|t|	0.0643	0.0201	0.0927	0.0950	0.0197	0.140

Standard errors in parentheses *** p<0.01, ** p<0.05, * p<0.1

3.3 Behavior of Principals

Principals take up to three sequential actions: selecting an agent, choosing a jar based on the selected agent's estimates, and rating the agent. In this subsection we analyze these three activities in term.

3.3.1 Selecting an Agent

Table 4 reports marginal effects from panel data probit estimation of the likelihood that an agent is chosen given the star ratings of the two agents observed by the principal.[11] Each additional star in an agent's rating over and above the rating of the other agent rating increases the likelihood of the higher rated being selected by about 32%. Specification 2 interacts a dummy variable for the Disclosure Treatment with the star difference. The coefficient estimate of the interaction term is statistically indistinguishable from zero, meaning that the impact of stars does not vary across treatments. These results indicate that the rating system is applying competitive pressure on agents as principals rely upon it for selection. This is stated as Finding 3.

Finding 3: Principals rely on the rating system when selecting an agent and thus the rating system adds competitive pressure to the process.

Table 4: Estimated Marginal Effects of Likelihood of Choosing Agent 2

VARIABLES	(1)	(2)
Stars Agent 2 - Stars Agent 1	0.318***	0.318***
	(0.0169)	(0.0438)
Disc * (Stars Agent 2 - Stars Agent 1)		-4.25e-05
		(0.0868)
Observations	900	900
Log Likelihood	-349.2	-349.2

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

3.3.2 Choosing a Jar

Table 5 reports estimated marginal effects from probit regressions predicting that principals will select Jar B, which is the non-bonus generating jar. The interpretation of the marginal effect for the difference in an agent's estimates for the jars is that for every dollar that the estimate of Jar A exceeds Jar B there is a 7% decrease in the chance that Jar B is picked. The negative sign (in all specifications) is what one would expect if principals believe agent estimations convey information. According to specification 2, principals in the Disclosure Treatment are actually less likely to choose Jar B, all else equal, by 11.1% as compared to the Market Treatment. This effect is the opposite of what one would observe if principals were wary of choosing Jar A in the Disclosure Treatment, but it is consistent with a desire to help the agents earn a larger profit. No Reputation does not impact the likelihood of choosing Jar B relative to Market, conditional on the estimated difference in jar values. Specification 3 shows that the impact of the jar estimate differential does not vary between the Market and Disclosure treatments, but the effect of the estimate differential is marginally smaller in No Reputation than in the Market. This leads to Finding 4.

Finding 4: Subjects rely on the agents estimates of the jar value to determine which option to select in all three treatments. However, Disclosure actually increases the likelihood that principal will select the option that benefits the agent ceteris paribus.

[11] Note it was necessary to exclude the No Reputation Treatment from this analysis because the are no ratings in that treatment.

Table 5: Estimated Marginal Effects of Likelihood of Choosing Jar B

VARIABLES	(1)	(2)	(3)
Estimate Jar A - Estimate Jar B	-0.0722***	-0.0718***	-0.0918***
	(0.00556)	(0.00563)	(0.0120)
Disclosure		-0.111***	-0.108**
		(0.0423)	(0.0427)
No Reputation		0.0155	0.0161
		(0.0302)	(0.0301)
Disclosure * (Est. Jar A - Est. Jar B)			0.0178
			(0.0215)
No Reputation * (Est. Jar A - Est. Jar B)			0.0376*
			(0.0197)
Observations	1,566	1,566	1,566
Base Rate	0.423	0.466	0.462
Log Likelihood	-810.7	-804.8	-795.3
Standard errors in parentheses	*** p<0.01, ** p<0.05, * p<0.1		

3.3.3 Rating the Expert

Figure 4 plots how often each star rating was given for both the Market and Disclosure treatments. The modal response was four stars. The mean number of stars in the Disclosure Treatment (3.397) was slightly lower than the mean in the No Disclosure Treatment (3.487), but the difference is not statistically significant (p-value = 0.22).

Figure 4: Histogram of Stars (Ratings) of Agents

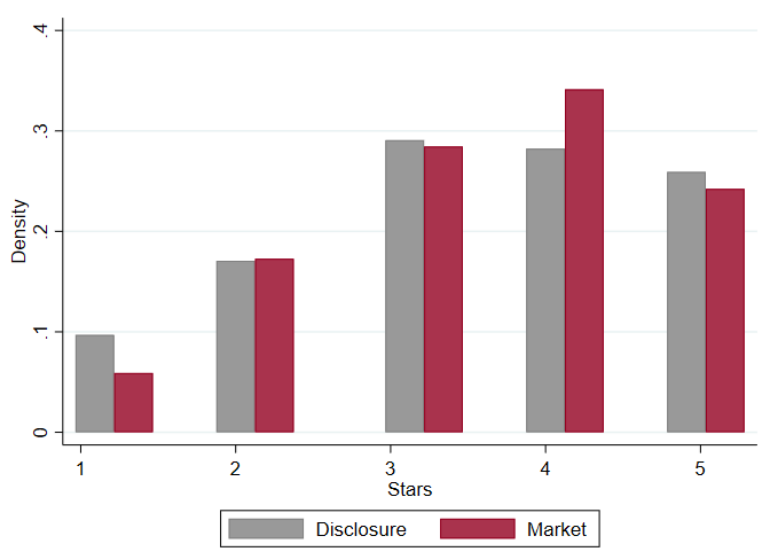

Table 6 reports regression estimates for the number of stars principals awarded agents after learning about how much was actually in the jar selected and receiving the noisy signal of the value of the other jar. *DIRECTACC* is a binary variable indicating that the ordering the agent provided is consistent with

ordering based on the information subsequently revealed to the principal. That is, this variable equals one if the agent estimated Jar X had a greater value and the ex-post information of the principal also indicated Jar X had the greater value. This coefficient is positive and statistically significant meaning that principals rate agents higher when the agent's rank order Jars correctly. *OverEstPick* identifies how each dollar the agent overestimated the jar selected by the principal impacts the number of stars the principal awards the agent. Principals do indeed punish agents for overestimating the value of the selected jar. Similarly, *UnderEstPick* identifies how each dollar the agent underestimated the jar selected by the principal impacts the number of stars the principal gives the agent. Including both variables in the regression allows for over- and underestimation to have different effects. Ultimately, the data reveal that overestimation is punished more harshly than underestimation of the selected jar. That is, the coefficients for *OverEstPick* and *UnderEstPick* are statistically different from each other (p-value < 0.001). *OverEstOther* and *UnderEstOther* examine the effect of over- and underestimation of the jar not selected relative to the noisy signal of its value. As one would expect errors in the estimates of the jar not selected lead to lower ratings, but the reduction is only marginally different for over- and underestimation. While the punishment for underestimation is similar for the two jars, overestimation of the selected jar is viewed as worthy of a greater rating reduction than overestimation of the jar not selected. Specifications 2 and 3 of Table 6 allow for variation in rating by treatment. While some of the additional terms are marginally significant the overall result remain effectively unchanged. In the appendix, we provide additional evidence of robustness using an ordered probit estimation to account for the discrete nature of the star rating. Again, the overall implications are unchanged. Table 6, as well as the material in the appendix, provides the support for Finding 5.

Finding 5: Principals reward agents with higher ratings when the agent appears to correctly identify the best option and punish agents with lower ratings for misestimation. The punishment is harshest for overestimating the value of the option that was actually selected.

Table 6: Estimated Impacts on Rating Given to Agent

	(1)	(2)	(3)
DIRECTACC	0.517***	0.517***	0.501***
	(0.113)	(0.113)	(0.111)
OverEstPick	-0.244***	-0.244***	-0.274***
	(0.027)	(0.027)	(0.031)
UnderEstPick	-0.108***	-0.108***	-0.161***
	(0.031)	(0.031)	(0.033)
OverEstOther	-0.050**	-0.050**	-0.087***
	(0.024)	(0.025)	(0.024)
UnderEstOther	-0.106***	-0.106***	-0.068***
	(0.022)	(0.022)	(0.017)
DISC		-0.010	-0.127
		(0.144)	(0.272)
UnderEstPickDisc			0.110*
			(0.062)
OverEstPickDisc			0.058
			(0.053)
UnderEstOtherDisc			-0.081*
			(0.043)
OverEstOtherDisc			0.056
			(0.047)
DirectAccDISC			0.045
			(0.226)
Constant	4.011***	4.017***	4.068***
	(0.139)	(0.164)	(0.179)
Observations	1,044	1,044	1,044
Number of UniqueID	36	36	36
Tests of symmetry			
OverEstPick = UnderEstPick	< 0.001	< 0.001	< 0.001
OverEstOther = UnderEstOther	0.063	0.065	0.476
OverEstPick = OverEstOther	< 0.001	< 0.001	< 0.001
UnderEstPick= UnderEstOther	0.958	0.961	0.006
Robust standard errors in parentheses	*** p<0.01, ** p<0.05, * p<0.1		

4 Discussion

Conventional wisdom suggests that conflict of interest should be disclosed. However, this need not be the case if disclosure emboldens those with conflicts and encourages them to act in ways not anticipated by others. This is the surprising pattern reported in Cain et al. (2005). Our paper extends the basic set-up of Cain et al. (2005) to introduce competition between agents who have a conflicts of interest in the pursuit of serving principals. The goal of our paper is two-fold. First, we ask if the disciplining effect of market pressure can effectively offset the harmful incentives created by conflicts of interest. Second we ask whether disclosure has negative effects in the presence of such competitive pressure

In a treatment with No Reputation, the estimates provided by the agents are biased in a self-serving manner. This result is unsurprising given that the situation lacked the disciplining effect that comes with concern for reputation when one has to compete for future principals. By contrast, in the Market Treatment, when principals rated agents based on performance and subsequently principals were allowed to select agents based upon past ratings then agent behavior did not exhibit a self-serving bias. This is despite the fact that the agent's conflict of interest was not disclosed to the principal. That is, the effect of competition was sufficient to discipline the agent's behavior. In a third treatment which combined disclosure with competitive pressure, self serving behavior was again observed, consistent with the surprising results of Cain et al. (2005). Beyond examining the effect of competitive pressure and disclosure on conflicts of interest, our results provide two additional insights. One has to do with the manner in which people rate those who have provided service. Specifically, we find that misestimation of an option, whether it is selected or not, leads to a reduction in the awarded rating. However, overestimating the benefit from a selected option is viewed as more worthy of scorn than underestimating the value of an option not selected despite these errors having similar implications for the principal. It is not clear if punishment was less severe because the subjects did ot realize the implications or because the signal was noisy. Our experiment was not designed to differentiate between the two channels. A second insight is that the subjects in our experiment did not exhibit a wisdom of crowds. As part of the experiment, subjects had to guess the value of jars of coins and they systematically underestimated those values.

Market competition's success at preventing self serving behavior despite a conflict of interest was dependent upon two key features of our experimental environment, which has important implications for future research and policy applications. First our principals, had a good but not perfect signal of the true value of the good selected and its associated opportunity cost. Second, principals had free access to ratings which were good indicator of true agent performance. While our results indicate that mandated disclosure of conflicts of interest are likely to be detrimental when these conditions are met, they leave open the questions as to how to promote these conditions and what is the best policy when these conditions cannot be met.

References

Becker, G. S. (1968, April). Crime and Punishment: An Economic Approach. *Journal of Political Economy 76*(2), 169.

Ben-shahar, O. and C. E. Schneider (2011). The Failure of Mandated Disclosure. *University of Pennsylvania Law Review 159*(3), 647. Publisher: University of Pennsylvania Law School.

Boatright, J. R. (2000). Conflicts of Interest in Financial Services. *Business and Society Review 105*(2), 201–219. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/0045-3609.00078.

Cain, D. M., G. Loewenstein, and D. A. Moore (2005). The dirt on coming clean: Perverse effects of disclosing conflicts of interest. *The Journal of Legal Studies 34*(1), 1–25.

Campbell, E. G., R. L. Gruen, J. Mountford, L. G. Miller, P. D. Cleary, and D. Blumenthal (2007, April). A National Survey of Physician–Industry Relationships. *New England Journal of Medicine 356*(17), 1742–1750. Publisher: Massachusetts Medical Society _eprint: https://doi.org/10.1056/NEJMsa064508.

Carl, W. (2008, July). The role of disclosure in organized word-of-mouth marketing programs. *Journal of Marketing Communications 14*(3), 225–241. Publisher: Routledge.

Church, B. and X. Kuang (2009, June). Conflicts of Interest, Disclosure, and (Costly) Sanctions: Experimental Evidence. *The Journal of Legal Studies 38*(2), 505–532.

Darby, M. R. and E. Karni (1973). Free competition and the optimal amount of fraud. *The Journal of law and economics 16*(1), 67–88.

Dranove, D., D. Kessler, M. McClellan, and M. Satterthwaite (2003). Is More Information Better? The Effects of "Report Cards" on Health Care Providers. *Journal of Political Economy 111*(3), 555. Publisher: The University of Chicago Press.

Gibbons, R. V., F. J. Landry, D. L. Blouch, D. L. Jones, F. K. Williams, C. R. Lucey, and K. Kroenke (1998, March). A comparison of physicians' and patients' attitudes toward pharmaceutical industry gifts. *Journal of General Internal Medicine 13*(3), 151–154.

Grande, D., J. A. Shea, and K. Armstrong (2012, March). Pharmaceutical Industry Gifts to Physicians: Patient Beliefs and Trust in Physicians and the Health Care System. *Journal of General Internal Medicine 27*(3), 274–279.

Green, M. J., R. Masters, B. James, B. Simmons, and E. Lehman (2012, May). Do Gifts From the Pharmaceutical Industry Affect Trust in Physicians? *FAMILY MEDICINE 44*(5), 7.

Guo, T. and S. Sriram (2017). 'Let the Sun Shine in': The Impact of Industry Payment Disclosure on Physician Prescription Behavior. *SSRN Electronic Journal*.

Guo, T., S. Sriram, and P. Manchanda (2017). The Effect of Information Disclosure on Industry Payments to Physicians. *SSRN Electronic Journal*.

Hannan, E. L., H. Kilburn, M. Racz, E. Shields, and M. R. Chassin (1994, March). Improving the Outcomes of Coronary Artery Bypass Surgery in New York State. *JAMA 271*(10), 761–766. Publisher: American Medical Association.

Kerschbamer, R. and M. Sutter (2017). The economics of credence goods–a survey of recent lab and field experiments. *CESifo Economic Studies 63*(1), 1–23.

Mainous, A. G. (1995, April). Patient perceptions of physician acceptance of gifts from the pharmaceutical industry. *Archives of Family Medicine 4*(4), 335–339.

Mukamel, D. B. and A. I. Mushlin (1998). Quality of Care Information Makes a Difference: An Analysis of Market Share and Price Changes after Publication of the New York State Cardiac Surgery Mortality Reports. *Medical Care 36*(7), 945–954. Publisher: Lippincott Williams & Wilkins.

Pope, D. G. (2009). Reacting to rankings: evidence from "America's Best Hospitals". *Journal of health economics 28*(6), 1154–1165.

Rose, S. L., S. Sah, R. Dweik, C. Schmidt, M. Mercer, A. Mitchum, M. Kattan, M. Karafa, and C. Robertson (2019, March). Patient responses to physician disclosures of industry conflicts of interest: A randomized field experiment. *Organizational Behavior and Human Decision Processes*.

Scholl, B. and A. A. Hung (2018). The Retail Market for Investment Advice. Technical report, Secuity and Exchange Commission, Washington, D.C.

Surowiecki, J. (2005). *The wisdom of crowds* (1st Anchor books ed ed.). New York: Anchor Books. OCLC: 61254310.

Wolinsky, A. (1993). Competition in a Market for Informed Experts' Services. *The RAND Journal of Economics 24*(3), 380–398.

Appendix

Tables A.1 through A.3 present results from analysis complementary to that reported in Table 6. However, rather than treating the number of stars awarded to an agent as a continuous variable, this supplementary analysis treats the ratings as discrete ranked options and thus uses ordered probit regressions. Each table reports the marginal effects of the probit regression in parallel to one of the specifications in Table 6. For these tables, each column is the effects from a single independent variable and each row is the change in the likelihood of an agent receiving that number of stars for a one unit change in the variable. The results are consistent with those from Table 6.

Table A.1: Order Probit of Impacts on Rating given to Agent, Model 1

VARIABLES	(1) DIRECTACC	(2) OverEstPick	(3) UnderEstPick	(4) OverEstOther	(5) UnderEstOther
1 Star	-0.0592***	0.0274***	0.0122***	0.00496*	0.0117***
	(0.0156)	(0.00486)	(0.00353)	(0.00273)	(0.00334)
2 Star	-0.0728***	0.0337***	0.0150***	0.00610*	0.0144***
	(0.0156)	(0.00452)	(0.00461)	(0.00350)	(0.00320)
3 Star	-0.0485***	0.0224***	0.00999***	0.00406	0.00962***
	(0.0116)	(0.00463)	(0.00363)	(0.00251)	(0.00223)
4 Star	0.0404***	-0.0187***	-0.00831***	-0.00338*	-0.00800***
	(0.00918)	(0.00361)	(0.00284)	(0.00205)	(0.00227)
5 Star	0.140***	-0.0649***	-0.0289***	-0.0117*	-0.0278***
	(0.0297)	(0.00753)	(0.00835)	(0.00656)	(0.00569)
Observations	1,044	1,044	1,044	1,044	1,044

Standard errors in parentheses *** p<0.01, ** p<0.05, * p<0.1

Table A.2: Order Probit of Impacts on Rating given to Agent, Model 2

VARIABLES	(1) DIRECTACC	(2) OverEstPick	(3) UnderEstPick	(4) OverEstOther	(5) UnderEstOther	(6) DISC
1 Star	-0.0593***	0.0274***	0.0122***	0.00495*	0.0117***	0.00197
	(0.0156)	(0.00487)	(0.00354)	(0.00275)	(0.00334)	(0.0161)
2 Star	-0.0728***	0.0337***	0.0150***	0.00608*	0.0144***	0.00242
	(0.0157)	(0.00453)	(0.00463)	(0.00353)	(0.00319)	(0.0197)
3 Star	-0.0485***	0.0224***	0.00997***	0.00405	0.00960***	0.00161
	(0.0115)	(0.00458)	(0.00364)	(0.00252)	(0.00221)	(0.0132)
4 Star	0.0403***	-0.0187***	-0.00830***	-0.00337	-0.00800***	-0.00134
	(0.00926)	(0.00365)	(0.00287)	(0.00208)	(0.00227)	(0.0109)
5 Star	0.140***	-0.0649***	-0.0288***	-0.0117*	-0.0278***	-0.00466
	(0.0297)	(0.00749)	(0.00837)	(0.00660)	(0.00569)	(0.0381)
Observations	1,044	1,044	1,044	1,044	1,044	1,044

Standard errors in parentheses *** p<0.01, ** p<0.05, * p<0.1

Table A.3: Order Probit of Impacts on Rating given to Agent, Model 3

VARIABLES	(1) DIRECTACC	(2) OverEstPick	(3) UnderEstPick	(4) OverEstOther	(5) UnderEstOther
1 Star	-0.0623***	0.0305***	0.0184***	0.00868***	0.00697***
	(0.0132)	(0.00511)	(0.00454)	(0.00312)	(0.00235)
2 Star	-0.0767***	0.0376***	0.0227***	0.0107***	0.00859***
	(0.0151)	(0.00493)	(0.00447)	(0.00375)	(0.00284)
3 Star	-0.0513***	0.0251***	0.0152***	0.00715***	0.00575***
	(0.0123)	(0.00573)	(0.00419)	(0.00263)	(0.00186)
4 Star	0.0422***	-0.0206***	-0.0125***	-0.00588**	-0.00472**
	(0.0103)	(0.00434)	(0.00329)	(0.00232)	(0.00183)
5 Star	0.148***	-0.0725***	-0.0438***	-0.0206***	-0.0166***
	(0.0251)	(0.00821)	(0.00863)	(0.00675)	(0.00482)
Observations	1,044	1,044	1,044	1,044	1,044

VARIABLES	(6) DirectAccDISC	(7) UnderEstPickDisc	(8) OverEstPickDisc	(9) UnderEstOtherDisc	(10) OverEstOtherDisc
1 Star	0.00428	-0.0130*	-0.00627	0.00976*	-0.00525
	(0.0197)	(0.00726)	(0.00531)	(0.00550)	(0.00516)
2 Star	0.00527	-0.0160**	-0.00773	0.0120**	-0.00647
	(0.0245)	(0.00778)	(0.00668)	(0.00601)	(0.00616)
3 Star	0.00353	-0.0107*	-0.00517	0.00804*	-0.00433
	(0.0165)	(0.00566)	(0.00487)	(0.00423)	(0.00394)
4 Star	-0.00290	0.00878*	0.00425	-0.00661*	0.00356
	(0.0136)	(0.00467)	(0.00390)	(0.00342)	(0.00337)
5 Star	-0.0102	0.0309**	0.0149	-0.0232*	0.0125
	(0.0471)	(0.0157)	(0.0129)	(0.0120)	(0.0118)
Observations	1,044	1,044	1,044	1,044	1,044

Standard errors in parentheses *** p<0.01, ** p<0.05, * p<0.1