Costly transparency

Justin Fox [a,1], Richard Van Weelden [b,*]

[a] Department of Political Science, Yale University, 77 Prospect St., New Haven, CT 06520, United States
[b] Department of Economics, University of Chicago, 1126 E. 59th St., Chicago, IL 60637, United States

ARTICLE INFO

ABSTRACT

We consider whether a career-minded expert would make better decisions if the principal could observe the consequences of the expert's action. The previous literature has found that this "transparency of consequence" can only improve the efficacy of the expert's decision making. We show, however, that this conclusion is very sensitive to the specified cost structure: if learning the consequences of the expert's action makes the expert more likely to choose the action most likely to correspond to the true state of the world, when costs are asymmetric, this can be associated with a decrease in the principal's expected welfare. In addition, we show that, when the prior on the state of the world is sufficiently strong, if the principal benefits from learning the consequences of the expert's action, her utility is higher if she observes only the consequences and not the action taken. For such priors, the optimal transparency regime will involve either the principal observing only the expert's action or only the consequences of the expert's action: it will never be optimal to observe both. We illustrate these results with examples from finance and public policymaking.

1. Introduction

In many situations, individuals delegate authority to an informed expert: investors give money to a fund manager in the hopes that it will be invested wisely; voters elect politicians to make decisions for them; judges are charged with interpreting the law and upholding the constitution on behalf of the people. Of course, whenever the principal delegates to an expert, there is always the concern that the expert may not act in the principal's interest. It is usually argued that such concerns are lessened, however, if the principal can observe the consequences of the expert's action: learning the consequences of the action increases the incentive for the expert to take actions which are likely to benefit the principal.[2] While observing the consequences of the expert's action is certainly beneficial in many situations, we identify conditions under which it can decrease the principal's welfare.

We analyze a model in which the principal's preferred course of action varies with the underlying state of the world. There is a strong prior about this state, but some mistakes may be more costly than others. We show that, when the expert's primary concern is to promote his own career, and it is very costly if the expert takes the

incorrect action in the less likely state of the world, principal welfare is decreased by observing the consequences of the expert's action. There are many such environments in which the ex-ante less likely event is associated with a greater potential cost: stock market crashes are relatively rare events but can lead to large losses or even bankruptcy; the defendant in a trial is more likely to be guilty but the costs of convicting the innocent are greater than the costs of acquitting the guilty; a Senator who believed that Iraq more likely than not had weapons of mass destruction might also believe the costs of invading if they didn't have weapons were greater than the costs of not invading if they did. Our results then show that observing the consequences of the expert's action can be socially harmful in many economically relevant environments.

Consider an expert, who may either be the high type, in which case he observes a perfectly accurate signal of the state, or the low type, in which case he observes a noisy but informative signal, who makes a decision on behalf of a principal. Further, assume that rather than seeking to promote the principal's welfare, the expert's objective is to maximize the principal's ex-post belief that the expert is the high type. Now suppose there is a strong prior about the state of world, and consider the decision of a low type expert who observes a signal contrary to that prior. If only the high type expert were to ever go against the prior, and the state of the world is not learned, the low type expert could mimic the high type expert by choosing the ex-ante less-likely alternative. So, in equilibrium, some low type experts will follow their signal. Now, suppose the state of the world will be learned with certainty before the expert's reputation is assigned. The low type expert can no longer mimic the high type by going against the prior — if he does, and the state matches the prior, the

expert will reveal himself to be the low type. Hence, if the prior is high enough, the low type expert will ignore his private information and take the action favored by the prior — even if it's in the principal's interest for the expert to follow his signal.

So, when the prior on the state of the world is sufficiently strong, transparency about the consequences of the expert's action increases the probability a low type expert takes the action recommended by the prior rather than his signal, which increases the probability the expert's action matches the state. Whether this is socially harmful or beneficial depends on the costs associated with different mistakes. When costs are symmetric — that is, the costs of different mistakes are the same — transparency of consquence can only increase principal welfare. When costs are asymmetric, however, it may be socially optimal for the expert to go against the prior even if, conditional on his private information, the state of the world more likely than not matches the prior. The principal's welfare will then be decreased if she observes the consequences of the expert's action.

One extremely natural application of our model is to the decision faced by a fund-manager. First, the manager's reputation for competence is supremely important, so it makes sense to model the manager's objective as maximizing his reputation for competence. Second, a market crash is a relatively rare event but carries with it extremely large costs for investors. Third, there is extremely fast feedback about the consequences of the manager's decisions. Finally, there is much concern in the financial literature about herd behavior (e.g. Scharfstein and Stein, 1990). Our results then raise the possibility that there is so much herding precisely *because* there is so much information available about the quality of the manager's choice: a manager who observes a signal that a crash is possible, but not likely, would be reluctant to act on such information for fear of being proven wrong.[3] So the career concerns of the manager lead him to expose investors to excessive risk of catastrophic loss — a distortion which is heightened when the manager's performance is more easily evaluated.

Another situation in which our model could be applied is to a legislator who must decide whether to approve an executive's proposal. Consider, for example, a U.S. Senator asked to vote on whether to authorize the Bush administration to use force against Iraq. Suppose the Senator (or more accurately the Senator's constituents) felt that war was justified if weapons of mass destruction existed, but not justified if they didn't. As, in the lead up to the war, the prior that weapons of mass destruction existed was high, our results indicate that if the Senator had weak information that weapons of mass destruction did not exist, he would have been more likely to oppose the war if it would not have been learned whether they existed.[4] Further, if the costs of an unjustified war were sufficiently high, this opposition would have been in the voters' interest.

As the above examples indicate, in environments with a sufficiently unbalanced prior, learning the consequences of an expert's action only increases the incentive to herd on the prior. As such, efforts to improve transparency, and thereby increase the speed with which the consequences of the expert's action are revealed, though frequently beneficial, can, in some cases, make the principal worse off. Consequently, increasing the frequency of disclosure for financial institutions may have the counter-productive effect of increasing the manager's incentive to herd. Similarly, increasing the effectiveness of the media, which makes it more likely that the consequences of a politician's actions will be learned before the next election, can make politicians more reluctant to go against the conventional wisdom, even when doing so would benefit the public.

While for most agency problems, the more information the principal has, the better the agent's behavior, the literature provides some examples where more information can be harmful.[5] One paper which is especially relevant for our work, Prat (2005), argues that observing the action a career-minded expert takes may lead to socially harmful distortions. That said, in the established results in the literature (e.g. Canes-Wrone et al., 2001; Maskin and Tirole, 2004; Prat, 2005), it is never harmful to observe the consequences of the expert's action. In particular, Prat (2005, 863) describes "the main contribution of [Prat's] paper is to show that, while transparency on consequences is beneficial, transparency on action can have detrimental effects." When the prior on the state of the world is sufficiently strong, and the costs are sufficiently asymmetric, however, we show this is reversed: the principal is made better off observing the action ("transparency of action") but worse off observing the consequences of that action ("transparency of consequence").

To understand how principal welfare varies with the transparency regime, and how the standard welfare results can be reversed in an asymmetric cost environment, consider the expert's incentives when the prior is sufficiently unbalanced that, even if a low type expert observes a contrary signal, the state is more likely than not to match the prior. As previously noted, in settings in which the principal observes the expert's action, it is more difficult for a low type expert to mimic the high type by going against the prior when the principal also observes the consequences of that action. The fear of being proven wrong then makes the low type expert more reticent about acting on a signal contrary to the prior when both the consequences and the action are observed than when only the action is. However, when some low types follow the prior rather than their signal, if the principal sees the expert go against the prior, this reveals positive information about the expert's type. So, when the consequences are observed, also observing the action taken causes the expert to go against the prior more often. Because transparency of consequence and transparency of action push the expert's behavior in opposite directions, they have an opposite effect on principal welfare.

This means that we can order the likelihood of the low type following his signal across the three transparency regimes: when the principal observes only the action taken but not the consequences, when the principal observes only the consequences but not the action, and when both the action and the consequences are observed. When the prior is sufficiently unbalanced, and the expert knows his type, the expert is most likely to herd on the prior when only the consequences are observed and least likely to when only the action is. Which transparency regime results in the highest welfare depends on the cost structure. When costs are symmetric, so the first-best decision rule involves the expert taking the action most likely to match the state, it is optimal to have transparency over consequence, but not action. However, when the costs are sufficiently asymmetric, so the first-best decision rule involves the expert matching his action to his signal, even if it goes against the prior, it is optimal to have transparency over action, but not consequences. As transparency of consequence is almost always harmful when transparency of action is beneficial, generically, some form of non-transparency increases welfare.

The paper is organized as follows: Section 2 describes the model, Section 3 presents our results, and Section 4 concludes. The proofs of our results are in Appendix A.

2. Model

A privately informed expert makes a decision on behalf of a principal.[6] We assume that there are two states of the world, $\omega \in \{0,1\}$, with prior probability that the state is 1, $Pr(\omega=1)=\pi>1/2$. There

---

[3] Robert Rodriquez, CEO of First Pacific Advisors, a hedge fund which divested its portfolio of subprime mortgages well before the financial crises, argues that, when it comes to risky but widely held assets, managers feel compelled to "be fully invested for fear of underperforming" (Rodriquez, 2009).

[4] This assumes that the objective of the Senator was to signal competence rather than to signal ideology or toughness. It also assumes that the Senator's vote would not affect the probability of learning the state.

[5] See Prat (2005) for a discussion.

[6] We use male pronouns for the expert and female pronouns for the principal.