

Transparency, career concerns, and incentives for acquiring expertise*

Heski Bar-Isaac[†]

NYU

October 2011

Abstract

An agent, who cares about signaling his ability, chooses among different projects that generate observable outcomes. The agent's information about which project delivers a good outcome depends on both his ability and his effort. This paper examines how the agent's incentives for effort change depending on whether or not the agent's project choice is observed. If this choice is publicly observed, the agent's project choice is distorted towards particular types of projects. When the outcomes of these advantaged projects are particularly sensitive to the agent's information, such transparency boosts the agent's information-gathering incentives. However, when public observation of project choice leads the agent to choose information-insensitive projects, then such transparency dampens incentives. This provides a more nuanced view of the implications of action transparency in the literature on career concerns for experts.

Keywords transparency, career concerns, expertise, information acquisition, information-sensitive

JEL: D82, D83, M54, M59

1 Introduction

Consider the relationship between a firm's shareholders and its CEO, a relationship of perennial interest. A theme much discussed by interested parties and in the popular press is the degree of transparency in overseeing the CEO's activities. For example, Harvey Pitt, a former chairman of the Securities and Exchange Commission, has argued that

*Conversations and comments from John Asker, Roland Benabou, Luis Cabral, Catherine de Fontenay, Joyee Deb, Gilat Levy, Meg Meyer, Andrea Prat, Susanna Sollstrom-Matthews, Bauke Visser and audiences at the 6th IIOC, GAMES 2008, NASM 2009, Econometric Society North American Winter Meeting, 2010, Erasmus School of Economics, and, particularly, Ignacio Esponda, Alessandro Gavazza, the Editor, and three anonymous referees, were helpful in the preparation of this paper.

[†]Contact info: Bar-Isaac: heski@nyu.edu; Department of Economics, Stern School of Business, NYU, 44 West 4th street 7-73, NYC, NY 10012 USA

“[t]ransparency is key” (Pitt, 2005). Transparency is also at the heart of numerous popular and academic discussions of politicians, managers, and countless other instances of principal-agent problems, where agents’ incentives arise through career concerns.

A literature on career concerns for experts—one that considers agents endowed with signals who choose between projects (or provide advice)—speaks to the costs of transparency. A key insight in this literature is that agents’ career concerns lead them to distort their project-selection decisions or actions—namely, towards “smart” actions that are relatively likely to be undertaken by more capable agents.¹ Transparency on actions—that is, allowing principals to observe not only the outcome of a project but also the choice of project directly—leads to a greater distortion of the agent’s project selection.² However, whereas this literature takes the ability of an actor as synonymous with the quality of the information he holds and treats it as exogenous, the focus of this paper is to examine the agent’s incentives to undertake effort to improve the quality of his information. This paper shows that transparency on actions also affects an agent’s information-gathering incentives.

While much of the literature focuses on cases in which the agent has a choice between only two actions—one safe and one risky—it is important for the results of this paper that there are more than two actions.³ I consider the minimal case with one safe action and two risky actions. The safe action is information-insensitive in the sense that its outcome is perfectly predictable ex-ante and independently of the agent’s information. Instead, risky actions are information-sensitive, since the agent’s information is useful for distinguishing which of the two risky actions is likely to generate a better outcome.

This distinction between information-sensitive and -insensitive types of action is useful since an agent who anticipates taking a smart action that is information-insensitive will have little reason to gather information. However, if the agent anticipates that career concerns will lead him to choose an action whose outcome depends a great deal on his information, he is more likely to gather the information that will help him make a better choice.

As an illustration of the central mechanism, consider the example of a CEO who can choose between maintaining the status quo (an information-insensitive choice) or restructuring a particular business process. Suppose, further, that shareholders and other potential employers believe that a more capable CEO is relatively more likely to force a restructuring, so that some form of restructuring is seen as smart. This is likely to be

¹This idea is clearly elucidated in Prat (2005), and present in numerous papers, including Harrington (1993), Brandenburger and Polak (1996), Fingleton and Raith (2005), Swank and Visser (2007), and Levy (2007).

²See, in particular, Prat (2005).

³This difference in the setup provides an important contrast with the model of Suurmond, Swank and Visser (2004), who do allow the expert to exert effort to improve his information.

the case if, for example, the firm is facing perceived difficulties. In this case, if the CEO’s course of action is publicly observed, then he may be more likely to choose to restructure (even if his information suggests that this is unwarranted). Anticipating this, he might, therefore, exert more effort in exploring different restructuring alternatives. This ex-ante boost to his incentives for information-gathering could overcome any inefficiency in his ex-post decision on the course of action.

This paper presents a simple model to explore this intuition and inform a more complete account of the effects of transparency.

Related Literature

This paper brings together two strands in the career-concerns literature. Interestingly, both have their origins in Holmström (1982/99). In the model presented in Section 2 of Holmström (1982/99)—similar to traditional reputation models (for example, the seminal work of Kreps and Wilson, 1982 and Milgrom and Roberts, 1982)—principals have no doubts as to the appropriate action (simply to exert more effort) but do not know whether the agent has taken such effort. Career concerns for displaying productive capability and the conflation of ability and effort in generating outcomes, lead the agent to undertake such costly actions (see Section 5 of Bar-Isaac and Tadelis, 2008, for a brief review). In contrast, in the literature on career concerns for expertise, depending on the realized state of the world, any of the actions could turn out to be optimal ex-post (as is the case in Section 3 of Holmström, 1982/99), and so the reputational concern is for the appropriate use of judgement.^{4,5}

As pointed out above, this literature typically supposes that the precision of the agent’s information is exogenous; instead, this paper supposes that the agent can take actions that affect the quality of his information. Career concerns—specifically the concern to demonstrate that one has the productive capability to effectively gather useful information—provide incentives for information-gathering, and this information is then used to make decisions, as in the experts literature. In order to highlight endogenous information-gathering, I simplify an aspect that is considered at length in the experts literature by treating the benefit associated with the public believing that the agent has taken a particular “smart” action as exogenous.⁶ I return to this simplifying assumption in Section 4 below.

⁴Another stream of the literature—for example, Benabou and Laroque (1992)—considers the case in which the quality of the signal is fixed, but the actors may differ in the extent to which their preferences are aligned with the principal.

⁵In most of this literature, different choices have no current costs associated with them (costs are purely future reputational costs). However, adding costs to exercise judgement by taking one action rather than another, as in Ely and Välimäki (2003), does not change the argument.

⁶Much of the literature on career concerns for experts is concerned with endogenizing the “smart” actions and their rewards. This is often in the context of competition among experts. Examples include Scharfstein

There is a small literature that, like this paper, considers career concerns for experts who can gather information and, thereby, endogenize expertise. These papers differ a little in their setups and, more substantively, in the questions that they address.⁷

Milbourn, Shockley, and Thakor (2001) do not consider the agent taking an action directly. Instead, they suppose that the principal, rather than the agent, makes this decision but does so after directly observing the agent’s information.⁸ They show that the agent’s incentives to acquire information are higher than first best when he has career concerns regarding his ability to generate good projects. In this paper, though, agents’ career concerns are for their expertise—the quality of their information or, rather, their ability to generate useful information. In addition to making investments in the quality of their information, agents also make decisions about what projects to implement on the basis of that information. Indeed, our focus is on the interaction of the effort-choice and the decision of which project to implement.

Most closely related is Suurmond, Swank, and Visser (2004), which considers an agent (who may or may not know his own ability) who exerts costly effort to become better informed about the state of the world.⁹ In their model, an agent chooses whether to implement a particular project or maintain the status quo. If he implements the project, the state of the world becomes known; if he chooses not to, the state remains unknown. The market observes the decision on the project, and observes the state only when the project is implemented. Thus, as in this paper, one action (status quo) is “safe” in that it ensures that no more information is generated. In their environment, an agent who does not know his own type will make socially optimal project-selection decisions. It is worth noting that in their environment, depending on its realization, a more-informative signal can lead the agent to take the safe or the risky action, in contrast with the environment in this paper, where information helps determine which risky action is the better one. Suurmond, Swank, and Visser characterize equilibrium efforts and decisions, and contrast what happens if the

and Stein (1994), Zwiebel, (1995), Prendergast and Stole (1996), Ottaviani and Sørensen (2006a,b,c), and Levy (2004).

⁷There is also a related literature that considers “delegated expertise” in a contractual environment, where outcome-contingent contracts aim to give incentives both for information-gathering and for decisions that use the information appropriately. The seminal paper in this literature is Demski and Sappington (1987); more recent work includes Szalay (2005), and Malcomson (2009). In particular, this literature highlights that contracts might deliberately manipulate and distort decisions in order to boost incentives for information-gathering. In this paper, the interpretation is that such distortions do not arise by design contractually, but instead through some reputational effects that can boost or dampen incentives for information-acquisition.

⁸I discuss such a case in the last paragraph of Section 4 below. There, I argue that in our setup direct observation of the information by the principal can lead to information-gathering incentives that can be higher or lower than other benchmarks that I consider.

⁹Also related is Swank and Visser (2008), which compares two agents who take actions sequentially.

agent knows his own type with what happens if he does not. They find that when the agent knows his own type, a low-type agent makes unconventional project choices to try to mimic the high type and this, in turn, gives a high-type agent stronger incentives to gather information and separate from the low type. Here, instead, I consider only the setting in which the agent does not know his own type and the focus is on transparency—comparing the case in which the agent’s action is observed to the case in which it is not.¹⁰ The key result is that the effect of such transparency on incentives for effort depends on whether or not the “smart” actions that transparency encourages are information-sensitive.

2 Model

An agent must make two strategic decisions: first, whether or not to exert effort in improving the quality of the signal he observes and, second, following observation of this signal, which action to take. The action, in turn, leads to some outcome that is observed by everyone.

There are three actions that the agent can take. One is a safe action and yields an outcome 0 with probability 1. The other two actions are risky and are ex-ante indistinguishable; of these, one is good and the other bad. The good action delivers 0 with probability ρ and $G > 0$ with probability $1 - \rho$, and the bad action delivers 0 with probability ρ and $-B$ with probability $1 - \rho$, where $B > 0$.¹¹

Before choosing which action to take, the agent observes a signal that indicates which risky action is the good one. This signal depends on the both agent’s type and his effort choice.

Both the public and the agent assign prior belief λ to the agent being a high type and $1 - \lambda$ to the agent being a low type. The agent chooses effort $e \in \{0, 1 - h\}$ to improve the quality of the signal that he generates, which, in turn, may inform his choice of action. Choosing $e = 0$ is costless and choosing $e = 1 - h$ comes at a cost c . If the agent is the low type the signal is pure noise, but if the agent is the high type the signal is accurate with probability $h + e$, where $h \geq \frac{1}{2}$.¹²

The agent is risk-neutral. The payoff that he seeks to maximize through his choice of effort and action consists of several components. First, the agent values the outcome of the

¹⁰In the core model, Suurmond, Swank, and Visser (2004) outcomes are not observed, and so if actions are also not observed, then there is no updating at all. They consider, in Section 7, an extension where outcomes are observed, but then, since the status quo and the project always generate different outcomes, even if the action is not observed, it could be inferred. Thus, it is impossible to address our central mechanism in their model without altering it.

¹¹It is convenient to suppose that the good and bad projects deliver 0 with the same probability, but this is not critical for the results of the paper.

¹²This implies that a high-type agent who exerts effort can perfectly identify the good risky action. This can be weakened without affecting the qualitative results.

project.¹³ Next, the agent seeks to maximize his posterior reputation less the cost of any effort that he undertakes. This is a reduced form for the standard career-concern model, such as Holmström (1982/99), where different members of the public compete in a second period for the services of the agent. Finally, there is a payoff R associated with the public believing that the agent has taken the risky action. This is interpreted as a reputational benefit (or cost) that is enjoyed if the agent is directly observed as taking a risky action or if it is thought that the agent is taking a risky action. We say that the smart action is safe if $R < 0$ and risky if $R > 0$.

Payoffs when actions are observed

If actions are publicly observed, then the agent earns

$$0 - c1_e + \lambda \tag{1}$$

if taking the safe action, where 1_e is an indicator function that takes the value 1 if the agent exerts effort and 0 otherwise. The posterior on taking the safe action is necessarily λ since the agent does not know his type (and taking efforts and observing signals provides him no additional information on his own type), so that the decision to take a safe or a risky action cannot, in itself, be informative of type.¹⁴

It is convenient to introduce notation λ_g and λ_b for the posteriors following a good or bad outcome.¹⁵

If actions are observed and the agent takes a risky action, the agent obtains R . In addition, he believes that he is an H -type with probability λ . Consequently, he anticipates choosing the good risky action with probability $\lambda(1_e + h(1 - 1_e)) + \frac{1-\lambda}{2}$. When he does so, with probability $(1 - \rho)$, it generates a payoff of G and the public posterior will shift to λ_g . With probability ρ , a risky project (whether good or bad) delivers an outcome of 0, and nothing can be inferred from this outcome so that posterior remains at λ . Finally, the project may be bad (with probability $1 - \lambda(1_e + h(1 - 1_e)) - \frac{1-\lambda}{2}$), generate a bad outcome (with probability $1 - \rho$), and then provide the agent a direct project payoff $-B$ and a reputational payoff of λ_b . Therefore, the overall value of taking the preferred risky

¹³Since G and B are not normalized, we can vary the weight attached to this component by scaling G and B up or down.

¹⁴This is necessarily the case both on- and off-equilibrium.

¹⁵The values λ_g and λ_b will, of course, depend on the public (on- or off-) equilibrium beliefs about the agent's choice of effort.

action can be written as:¹⁶

$$R - c1_e + \rho\lambda + (1-\rho)(\lambda(1_e + h(1-1_e)) + \frac{1-\lambda}{2})(G + \lambda_g) + (1-\rho)(1-\lambda(1_e + h(1-1_e)) - \frac{1-\lambda}{2})(\lambda_b - B). \quad (2)$$

Payoffs when actions are not observed

If actions are not publicly observed, the values of taking the safe or the risky action are a little more involved; in particular, the payoff on taking the safe action depends on the public's equilibrium expectations, and can be written as

$$R1_R - c1_e + \lambda, \quad (3)$$

where 1_R is an indicator function that takes the value 1 if the public expects the agent to take the risky action in equilibrium and 0 otherwise. Similarly, the value of taking the preferred risky action can be written as:

$$\rho R1_R + (1-\rho)R - c1_e + \rho\lambda + (1-\rho)(\lambda(1_e + h(1-1_e)) + \frac{1-\lambda}{2})(G + \lambda_g) + (1-\rho)(1-\lambda(1_e + h(1-1_e)) - \frac{1-\lambda}{2})(\lambda_b - B). \quad (4)$$

Note that, when a good or bad outcome is observed, the public can then be certain that the agent has taken a risky action (accounting for the term $(1-\rho)R$ in this expression). Instead, if an outcome of 0 is observed, then the public belief that the agent has taken the risky action is 1_R . Otherwise, expression (4) is similar to (2).

Summary

Summarizing the timing of the model: (a) the agent chooses $e \in \{0, 1-h\}$; (b) the agent observes a signal indicating which risky action is the good one; (c) the agent chooses whether to take the risky action that he believes is more likely to be good or the safe action;¹⁷ and (d) the project delivers an outcome of 0 if the agent chose the safe action, or G , $-B$, or 0 if he chose a risky action.

¹⁶Here and below, I write expected values when the agent takes the risky action that his information suggests is likely to be the good one. In principle, the agent could choose to take the action that he thinks is bad, but it is easy to show that if there is an equilibrium where the agent exerts effort and takes the risky action that he thinks is bad, then there is an equilibrium where the agent exerts effort and takes the risky action that he thinks is good. This seems the more reasonable case to focus on.

¹⁷I could also consider the case in which the agent can also choose to take the action he believes is more likely to be bad. However, given the symmetry in the model, if there is an equilibrium where the agent takes the more-likely-to-be-bad risky action, there is also an equilibrium where the agent takes the more-likely-to-be-good risky action. However, the reverse is not true since the agent cares about the outcome of the project. Given our focus on equilibrium incentives for effort (that, as I argue, require the agent taking the risky action), it is reasonable to focus on the case in which the agent is relatively likely to take such an action.

It is assumed throughout that project outcomes are publicly observed. The focus of the paper is on comparing the agent's effort incentives when his choice of action can be observed to his effort when it cannot. The agent's strategy consists of a choice of effort, and then a choice of whether or not to choose a risky action (where this choice might depend on his effort choice).

3 Analysis

When actions are observed, the agent has a greater incentive to take a risky action if $R > 0$ and to take a safe action if $R < 0$. Since effort is of no value if the agent intends to take the safe action, it is intuitive that when $R > 0$, effort is more likely when actions are observed, and when $R < 0$, effort is more likely when actions are not observed.

We can formalize this intuition through the following results.

Proposition 1 *(i) Suppose that the smart action is risky ($R \geq 0$) and that there is an equilibrium where the agent exerts effort when the agent's choice of action is not observed; then, there is an equilibrium where the agent exerts effort when the choice of action is observed;*

(ii) Suppose that the smart action is safe ($R \leq 0$) and that there is an equilibrium where the agent exerts effort when the agent's choice of action is observed; then, there is an equilibrium where the agent exerts effort when the choice of action is not observed.

Proof. First notice, following (1) and (3) that there can be no equilibrium where the agent takes the safe action and exerts effort.

(i) Intuitively, in this case, the expected payoff on the equilibrium path (where equilibrium involves effort and the risky action) is the same whether or not the action is observed, but observability decreases the deviation payoff (to deviating to exerting no effort and taking the safe action), so that if there is an equilibrium with effort when the action is not observed, then there should also be such an equilibrium when effort is observed.

Formally, we can write down the payoffs associated with different strategies and demonstrate this explicitly.

If there is an equilibrium where actions are not observed and the agent exerts effort, then the equilibrium strategy gives a payoff as defined in (4), where $1_R = 1$ and $1_e = 1$ and $\lambda_g = \frac{\lambda}{\lambda + (1-\lambda)^{\frac{1}{2}}}$ and $\lambda_b = 0$. The agent's expected payoff in such an equilibrium is given by the following expression.

$$R - c + \rho\lambda + (1 - \rho)\left(\lambda + \frac{1 - \lambda}{2}\right)(G + \lambda_g) + (1 - \rho)\left(1 - \lambda - \frac{1 - \lambda}{2}\right)(\lambda_b - B). \quad (5)$$

If there is, indeed, such an equilibrium when actions are not observed—that is, one

where the agent exerts effort and takes the risky action—the expected payoff in equilibrium, (5), must be a higher payoff than the expected payoff from deviating to take the safe action (with no effort), which, following (3) with $1_R = 1$ and $1_e = 0$, is equal to $R + \lambda$. The payoff (5) must also be higher than the expected payoff from deviating to take the risky action but exerting no effort, which, following (4) with $1_R = 1$ and $1_e = 0$, is given by

$$R + \rho\lambda + (1 - \rho)(\lambda h + \frac{1 - \lambda}{2})(G + \lambda_g) + (1 - \rho)(1 - \lambda h - \frac{1 - \lambda}{2})(\lambda_b - B). \quad (6)$$

Now, suppose that actions are observed and that the public believes that the agent is exerting effort; then, doing so and taking the risky action yields the same expected payoff as under no observation—that is, equal to (5). The deviation to taking the risky action and exerting no effort yields the same payoff as in the case in which the action is not observed—that is, an expected payoff equal to (6). However, deviating and taking the safe action (with no effort) yields a lower payoff—namely, following (1), it yields λ rather than $R + \lambda$ in the case in which the action is unobserved—this is necessarily lower since, by assumption, $R \geq 0$. Thus, the “on-equilibrium” payoffs are the same, and the “off-equilibrium” deviations are either the same (in the case of taking the risky action) or lower (if taking the safe action) when actions are observable rather than unobservable. It follows immediately that there is also an equilibrium where the agent exerts effort when the choice of action is observed.

(ii) If there is an equilibrium where actions are observed and the agent exerts effort, then the equilibrium strategy gives a payoff as defined in (2), where $1_e = 1$ and $\lambda_g = \frac{\lambda}{\lambda + (1 - \lambda)\frac{1}{2}}$ and $\lambda_b = 0$, which is equal to (5). For this to be an equilibrium, this must be a higher payoff than deviating to take the safe action (with no effort), which, following (1), is λ , and it must be higher than the expected payoff is in deviating to take the risky action but taking no effort, which, following (2), with $1_e = 0$ is equal to (6). Thus, $(5) \geq \max\{(6), \lambda\}$

Instead, if actions are not observed and the public expects the agent to be taking the risky action, then we can set $1_R = 1$ in (4) and (3) to obtain that the value of taking the risky action and exerting effort is equal to (5); the value to taking the risky action and exerting no effort is equal to (6); and the value of taking the safe action is $R + \lambda$. Therefore, it is an equilibrium to exert effort and take the risky action as long as $(5) \geq \max\{(6), R + \lambda\}$, which is implied by $(5) \geq \max\{(6), \lambda\}$ and $R \leq 0$. ■

Corollary 1 *(i) Suppose that the smart action is risky ($R \geq 0$) and that there is no equilibrium where the agent exerts effort when the agent’s choice of action is observed; then, there is no equilibrium where the agent exerts effort when the choice of action is not observed;*

(ii) Suppose that the smart action is safe ($R < 0$) and that there is no equilibrium where the agent exerts effort when the agent's choice of action is not observed; then, there is no equilibrium where the agent exerts effort when the choice of action is observed.

4 Discussion

The model aims to illustrate as simply and succinctly as possible that when public observation of actions affects the choice of action, this, in turn, affects an agent's incentives for gathering information. Specifically, if transparency encourages actions that are information-insensitive, it blunts information-gathering incentives, but augments them if it encourages actions that are information-sensitive.

Although this intuition appears fairly straightforward, it requires several moving parts (albeit familiar ones from various associated literatures). For a career concern to operate, there must be uncertainty regarding the agent's type—here, I have taken the minimal case in which there are only two possible types. For incentives for information-gathering to be relevant, there must be a decision to be made about information-gathering—here, this is a simple binary decision. Moreover, the information must have the potential to change the action chosen; I allow for this by supposing that there are two risky actions and that the information indicates which is the good one. In terms of outcomes, there must be scope to learn from observation whether the information was accurate or not and to update on the agent's type. Finally, for there to be a genuine distinction between observing actions and not observing them, outcomes cannot perfectly reveal the choice of action—here, since the risky action can lead to the same outcome as the safe action (with probability ρ), the public is, in some circumstances, unable to distinguish between them if not observing the actions directly.

Thus, all these elements of the model are necessary, but they do lead to some notational burden. In this paper, I have simplified with respect to the remaining necessary element—that reputational concerns might drive the agent towards one action rather than another if actions are observed—by taking the strength of this concern as exogenous. One justification for treating it this way is to suppose that it arises from social approval or disapproval for particular actions, as considered in Daughety and Reinganum (2010). More broadly, one might suppose that it is a reputational cost or benefit for an orthogonal component of reputation. For example, if there is some probability that the agent might be a type that simply does not have the operational capability to undertake the risky action or to undertake a safe action, then R can be considered as reflecting a reputational premium/cost to appearing to be one of those commitment types. Finally, another approach, that is more in keeping with the literature on career concerns for experts, is explored in Bar-Isaac (2008).

Specifically, in that earlier version of the paper, I present a slightly different version of the model with a more involved signal structure, where, in effect, an agent gets a signal that not only tells him which risky action is most likely to be the good one, but does so with more or less confidence. In that case, different types of agents can face different distributions of signals, so that the signal received (which might encourage particular action-choices) can be informative about type, and so, in turn, tend to favor the choice of one action rather than another. Different signal structures can then endogenously lead the safe action to be “smart” (or, in the language of this paper, for $R < 0$) or for the risky actions to be “smart” ($R \geq 0$).

In analyzing only the payoffs of the agent, the model does not allow for a full welfare analysis; however, a natural benchmark case to consider is the action choices that maximize the expected payoff of the project less the cost of any effort incurred (that is, the surplus generated in the current period, ignoring any private or social value from reputation). With this “efficiency” benchmark, since we can consider different values of (λ, c) while separately changing G and B , it is easy to construct examples in which incentives for effort are too low or too high when there are reputation concerns. While career concerns and transparency on action lead to biased decision-making (towards smart actions), which is clearly bad for welfare, this paper highlights that this bias also has effort consequences. If the bias is towards information-sensitive actions that will encourage the agent to exert effort in information gathering, this can be welfare-enhancing (if the starting point is weak incentives). Instead, if the smart action is information-insensitive, then welfare is damaged both through biased decisions and weak information-gathering incentives.

Finally, another case that could be considered is the case in which the public observes the agent’s information directly and can force action on the basis of that information (corresponding to an active principal who does not delegate, but insists on reviewing all reports that are hard information). Again, since G and B are exogenous parameters of the model, it is easy to find parameters where the public would always force the agent to take the safe action (blunting all incentives for effort) or always take the risky action (generating significant incentives for effort). Instead, individual incentives for effort, as analyzed in Section 3, depend also on R and λ . Consequently, there is no clear comparison of whether direct observation of information should lead to greater or diminished incentives for effort as compared to the cases analyzed above, in which only outcomes are observed or in which action choices are also observed. This holds equally for smart actions that are information-sensitive and -insensitive.

5 Conclusions

This paper has aimed to provide a nuanced view of action-transparency. In particular, it has considered the distinction between smart actions that reveal information and smart actions that conceal information. When the smart action conceals information, transparency on actions reduces incentives for effort. This is intuitive: Returning to the earlier example, when a CEO who is pushed by reputational concerns to maintain the status quo (and following this course reveals nothing more about him), there is little reason for him to exert efforts in exploring whether it makes sense to pursue alternatives. Consider, instead, a CEO who is pushed towards more-interventionist courses of action (that is, intervention is the “smart” course of action); since the outcome will reveal a great deal about the CEO’s decision-making, the CEO, pre-disposed towards such actions, will exert considerable efforts to figure out whether this really makes sense.

In many applications, it is more plausible that smart actions are revealing rather than concealing. For example, an agent who pursues the status quo might be perceived as devoid of ideas. Transparency on actions, which might encourage new initiatives, can be beneficial inasmuch as it might encourage efficient effort from an agent in determining which new initiative is best.

References

- [1] Bar-Isaac, H. 2008. Which kind of transparency and when? Career Concerns, and Incentives for Acquiring Expertise, Working Paper EC-08-23, Stern School of Business, N.Y.U. at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1174502
- [2] Bar-Isaac, H. and S. Tadelis, 2008. Seller Reputation. *Foundations and Trends in Microeconomics*. 4, 273–351.
- [3] Benabou, R. and G. Laroque, 1992. Using Privileged Information to Manipulate Markets: Insiders, Gurus, and Credibility. *Quarterly Journal of Economics*. 107, 921–58.
- [4] Brandenburger, A. and B. Polak, 1996. When Managers Cover Their Posteriors: Making the Decisions the Market Wants to See. *RAND Journal of Economics*. 27, 523–41.
- [5] Daughety, A. F. and J. F. Reinganum, 2010. Public Goods, Social Pressure, and the Choice Between Privacy and Publicity. *American Economic Journal: Microeconomics*. 2, 191–221.
- [6] Demski, J. and D. Sappington, 1987. Delegated Expertise, *Journal of Accounting Research*, 25, 68–89.
- [7] Ely, J. C. & J. Välimäki, 2003. Bad Reputation, *Quarterly Journal of Economics*, 118(3), 785–814.
- [8] Fingleton, J. and M. Raith, 2005. Career Concerns of Bargainers. *Journal of Law*,

- Economics, and Organization.* 2, 179–204.
- [9] Harrington, J. E. Jr., 1993. Economic Policy, Economic Performance, and Elections. *American Economic Review.* 83, 27–42.
 - [10] Holmström, B., 1982/99. Managerial Incentive Problems: A Dynamic Perspective, *Review of Economic Studies*, 66, 169–182, 1999. (originally published in 1982 in *Essays in Honor of Professor Lars Wahlback*).
 - [11] Kreps, D., and R. Wilson, 1982. Reputation and Imperfect Information, *Journal of Economic Theory*, 27, 253–279.
 - [12] Levy, G., 2004. Anti-Herding and Strategic Consultation. *European Economic Review.* 48, 503–525.
 - [13] Levy, G., 2007. Decision Making in Committees: Transparency, Reputation and Voting Rules. *American Economic Review.* 97, 150–168.
 - [14] Malcomson, J. M., 2009. Principal and Expert Agent. *The B.E. Journal of Theoretical Economics.* 9, Article 17.
 - [15] Milbourn, T. T., R. L. Shockley, and A. V. Thakor, 2001. Managerial career concerns and investments in information. *RAND Journal of Economics.* 32, 334–51.
 - [16] Milgrom, P. R. and D. J. Roberts, 1982. Predation, Reputation and Entry Deterrence, *Journal of Economic Theory*, 27, 280–312.
 - [17] Ottaviani, M. and P. Sorensen, 2006. Professional Advice. *Journal of Economic Theory.* 126, 120–42.
 - [18] Ottaviani, M. and P. Sorensen, 2006b. Reputational Cheap Talk. *RAND Journal of Economics.* 37, 155–175.
 - [19] Ottaviani, M. and P. Sorensen, 2006c. The Strategy of Professional Forecasting. *Journal of Financial Economics.* 81, 441–66.
 - [20] Pitt, H. quoted in “Corporate Governance Must Stress Transparency, Pitt Says” at http://www.law.virginia.edu/html/news/2005_spr/pitt.htm
 - [21] Prat, A., 2005. The Wrong Kind of Transparency. *American Economic Review.* 95, 862–77.
 - [22] Prendergast, C. and L. Stole, 1996. Impetuous youngsters and jaded oldtimers. *Journal of Political Economy.* 104, 1105–34.
 - [23] Scharfstein, D. and J. Stein, 1990. Herd behavior and investment. *American Economic Review.* 80, 465–79.
 - [24] Swank, O. and B. Visser, 2007. On Committees of Experts. *Quarterly Journal of Economics.* 122, 337–72.
 - [25] Swank, O. and B. Visser, 2008. The consequences of endogenizing information for the performance of a sequential decision procedure. *Journal of Economic Behavior &*

- Organization*. 65, 667–81.
- [26] Suurmond, G., O. H. Swank, and B. Visser, 2004. On the bad reputation of reputational concerns. *Journal of Public Economics*. 88, 2817–2838.
- [27] Szalay, D., 2005. The Economics of Extreme Options and Clear Advice. *Review of Economic Studies*, 72, 1173–98.
- [28] Zwiebel, J., 1995. Corporate conservatism and relative compensation. *Journal of Political Economy*, 103, 1–25.