Equilibrium Contracts and Boundedly Rational Expectations*

Heiner Schumacher† Heidi Christina Thysen‡

KU Leuven London School of Economics

Version: October 30, 2019

Abstract

We study a principal-agent framework in which the agent forms beliefs based on a misspecified or simplified model of the principal's project. She fits this model to the objective probability distribution to predict output under alternative actions. Under mild restrictions, the agent has correct beliefs on the equilibrium path, so that the optimal contract is non-exploitative. However, she may overestimate the productivity of her effort. The scope for such biases depends on the agent's role in the organization, i.e., through which channels she influences the output. We obtain new results on the informativeness principle and the risk-incentive trade-off.

Keywords: Bayesian Networks, Principal-Agent Relationship, Moral Hazard
JEL Classification: D03, D82, D86

*We gratefully acknowledge financial support by the ERC Advanced Investigator grant no. 692995. We thank Yair Antler, Felix Bierbrauer, Kfir Eliaz, Florian Englmaier, Guido Friebel, Heiko Karle, Botond Kőszegi, Gilat Levy, Debraj Ray, Ronny Razin, Karl Schlag, Klaus Schmidt, Dirk Sliwka, Balázs Szentes, and Yves Le Yaouanq as well as seminar audiences at Aarhus University, Boston University, University of Cologne, Dalhousie University, London School of Economics, Ludwig-Maximilians University Munich, Tilburg University, the EEA-ESEM 2017 in Lisbon, and the ESSET 2018 Gerzensee for their valuable comments and suggestions. We especially thank Ran Spiegler for his invaluable support of the project. The usual disclaimer applies.

†Corresponding Author. KU Leuven, Department of Economics, Naamsestraat 69, 3000 Leuven, Belgium, ++32 163 74 579, E-mail: heiner.schumacher@kuleuven.be.

‡Department of Economics, London School of Economics, h.c.thysen@lse.ac.uk.

1 Introduction

The canonical principal-agent model of contracting under asymmetric information assumes that the agent knows the probabilistic consequences of all available actions. Formally, these are defined by a production function $p(y \mid a)$, where y is the contractible output and a the agent's action. Given the incentives provided by the contract, the agent chooses an action that – according to this function – maximizes her expected payoff. However, in an organization, $p(y \mid a)$ is typically a complex object. It may reflect information that is unavailable to the agent or that the agent cannot process due to cognitive limitations. Herbert Simon therefore proposed that administrative behavior must be "boundedly rational" (Simon 1947, 1955).

In this paper, we examine a contracting framework in which the agent has boundedly rational beliefs about the production function. She estimates $p(y \mid a)$ based on data generated by the true production process, the implemented action α^*, and a non-parametric subjective model $\mathcal{R}$. A model $\mathcal{R}$ is a collection of variables and causal relationships between these variables. This model may be misspecified. For example, it may be "too simple" relative to the complexity of the organization: empirical regularities that matter for the principal's project may not appear in $\mathcal{R}$. The agent's beliefs about $p(y \mid a)$ will be denoted by $p_{\mathcal{R}}(y \mid a; \alpha^*)$. An equilibrium contract implements action α^* if it is optimal for the agent to choose α^* under this contract given her beliefs $p_{\mathcal{R}}(y \mid a; \alpha^*)$. We study the properties of the optimal equilibrium contract and the implications of misspecifications in $\mathcal{R}$ for the principal-agent relationship.

To capture the agent's limited understanding of her environment, we apply Spiegler's (2016) Bayesian network approach. As an illustration, consider the following example:

> *"Marketer Example."* The agent is a marketer whose job is to increase sales y. One strategy to increase sales is to make cold-calls $a \in \{0, 1\}$, that is, calling potential customers without prior consent. Making cold-calls increases the set of customers $x_1 \in \{0, 1\}$ who know about the firm's product, but also reduces the firm's reputation $x_2 \in \{0, 1\}$ since some customers are annoyed by being cold-called. Both a larger customer set x_1 and a better reputation x_2 increase expected sales. However, when choosing her action, the marketer does not take the firm's reputation into account. The only mechanism on her mind is that making cold-calls enlarges the set of potential customers, and that more potential customers translate into more sales.

The Bayesian network approach roughly works as follows in the marketer example.[1] The setting describes an "extended production function" $p(x_1, x_2, y \mid a)$, i.e., a joint probability

[1] Missing technical details will be explained thoroughly in the next section.

distribution over customer set, reputation and sales for any given action. This function captures the objective model $\mathcal{R}^*$ of the project: $\mathcal{R}^*$ contains all relevant variables, a, x_1, x_2, y, and the causal relationships between these variables. The agent's subjective model $\mathcal{R}$ is a simplified version of $\mathcal{R}^*$ as it only contains the variables a, x_1, y, and their causal relationships. Her beliefs are derived by fitting $\mathcal{R}$ to the objective probability distribution, which is generated by the implemented action α^* and the extended production function $p(x_1, x_2, y \mid a)$. Thus, the different elements in the agent's subjective model $\mathcal{R}$ are quantified using input from the true data-generating process. Combining these elements yields the agent's beliefs $p_{\mathcal{R}}(y \mid a; \alpha^*)$.

If $\mathcal{R}$ differs from $\mathcal{R}^*$, the agent's beliefs about $p(y \mid a)$ may be biased, and both the incentive compatibility and the participation constraint could in principle be affected by the bias. Our first important observation is that a weak restriction on the agent's subjective model guarantees that the participation constraint is not affected. The agent correctly predicts the marginal equilibrium distribution over output if she takes into account the correlation between any two variables in $\mathcal{R}$ that have a joint influence on a third variable in $\mathcal{R}$ (Spiegler 2017). Formally, this is the case if $\mathcal{R}$ is "perfect." When $\mathcal{R}$ is perfect (as in the marketer example), the optimal equilibrium contract does not exploit the agent in the sense that the agent's expected payment is below her reservation utility. In contrast, exploitation would typically occur if one directly assumes biased beliefs, such as overconfidence (e.g., Kőszegi 2014). Importantly, a perfect $\mathcal{R}$ ensures in many cases that there are no informational cues in the data the agent gathers on the equilibrium path that could alert her about the misspecification in $\mathcal{R}$. Therefore, misspecifications in the agent's subjective model allow for sustainable belief biases. Throughout the paper, we assume that the agent's subjective model $\mathcal{R}$ is perfect.

A misspecification in the agent's subjective model $\mathcal{R}$ can change the incentive compatibility constraint even when $\mathcal{R}$ is perfect. In the marketer example, if the principal implements making cold-calls, then, by not taking reputation into account, the agent overestimates the drop in sales after deviation to not making cold-calls (i.e., she is "control optimistic"). This relaxes the incentive compatibility constraint, so that the principal can implement cold-calls with fewer incentives than if the agent had rational expectations. The principal then strictly benefits from the simplification in the agent's model.

However, a misspecification in $\mathcal{R}$ does not always affect the incentive compatibility constraint. We call the agent "behaviorally rational" if she correctly anticipates the production function, or, formally, $p_{\mathcal{R}}(y \mid a; \alpha) = p(y \mid a)$ for all possible a and α, regardless of the parametrization of the extended production function. We characterize a correspondence $H^*(\mathcal{R}^*)$ that indicates for a given objective model $\mathcal{R}^*$ the set of variables the agent must take into account in her simplified subjective model $\mathcal{R}$ so that she is behaviorally rational. We show that $H^*(\mathcal{R}^*)$ is often a strict subset of the variables from the extended production function, and that the difference between a variable $i \in H^*(\mathcal{R}^*)$ and a variable $j \notin H^*(\mathcal{R}^*)$ can be quite nuanced. Here is a very simple example: Consider a version of the marketer example where making cold-calls has no effect on reputation (still, reputation is stochastic and influences sales). The objective model $\mathcal{R}^*$ then has no link between action and reputation. The agent is now behaviorally rational even if she does not take reputation into account. She correctly anticipates the production function even though she ignores the influence of reputation on output.

The characterization of $H^*(\mathcal{R}^*)$ shows which variables matter for the agent's beliefs. An important interpretation of the objective model $\mathcal{R}^*$ is that it captures the agent's job, i.e., through which tasks, interactions, and decision-making powers she influences the final output.[2] In the canonical contracting model, these aspects are immaterial since behavior is governed by the production function $p(y \mid a)$. In our framework, we can have two extended production functions that give rise to the same "reduced-form" production function $p(y \mid a)$, but that differ in their causal model $\mathcal{R}^*$, and hence in the extent to which simplifications affect $p_{\mathcal{R}}(y \mid a; \alpha^*)$.

One application of this finding is that we can examine which organizational features potentially cause the agent to overestimate the productivity of her effort. As the marketer example shows, this happens if the agent does not take into account a partial negative effect of her effort on the output. There are several intuitive reasons why this may happen. Consider an agent in a management position in which her effort influences the behavior of other workers (e.g., a group of telemarketers). If the agent does not understand the difficulties of their job (e.g., that telemarketing has a partial negative effect on sales through reputation), she overestimates her subordinates' – and hence her own – productivity. There are different instances where this could happen: The agent may be a technical expert who is promoted into a management position in which she oversees the actions of workers whose job she does not fully understand. Alternatively, it may be the case that subordinates do not communicate the problems they face to their managers (out of career concerns). These phenomena are usually discussed critically in the management literature, but in our framework they advance the agent's effort motivation and hence benefit the principal.

When the agent's subjective model $\mathcal{R}$ is misspecified so that the incentive compatibility constraint is affected, this naturally impacts on the comparative statics of the optimal equilibrium contract. We revisit two predictions of the canonical contracting framework that received considerable attention in the literature: the informativeness principle (e.g., Holmström 1979, Chaigneau et al. 2019) and the trade-off between risk and incentives (e.g., Prendergast 2002, Corgnet and Hernán-González 2019). In both cases, our framework provides explanations for

[2]For example, one can interpret $\mathcal{R}^*$ as an adjusted depiction of the organizational chart of the principal's project. As the CEO the agent would influence his senior managers who in turn influence their subordinates' behavior and so forth. A simplification in $\mathcal{R}$ then captures that the agent ignores a certain part of the organization.

empirical phenomena that are at odds with these predictions.

First, the informativeness principle states that the optimal contract should condition on an additional variable z if and only if it provides additional information about the agent's action. This is no longer true in our framework. We show that if $\mathcal{R}$ is such that the agent correctly anticipates the joint distribution of z and y, and if in $\mathcal{R}$ variable z is a sufficient statistic of variable y, the optimal equilibrium contract only conditions on z (even when according to the objective model both y and z are informative about the agent's action). For example, if the agent does not take into account the correlation between her output z and the output of others (induced by variations in the state of the economy), the optimal equilibrium contract only conditions on z. Hence, it is incomplete and may reward the agent for windfall gains.

Next, the canonical contracting model suggests that there should be a negative relationship between risk and incentives. If the variance in output increases through a mean-preserving spread in $p(y \mid a)$, the provision of incentives becomes more costly so that all else equal the optimal contract offers fewer incentives. Again, this is no longer necessarily true if the agent's subjective model is misspecified. Instead, she may interpret an increase in risk as an increase in the productivity of her effort. In this case, the incentive compatibility constraint is relaxed when there is more risk so that there can be a positive relationship between risk and incentives.

Related Literature. Our basic model is the principal-agent framework introduced by Holmström (1979) and Grossman and Hart (1983). In this framework, both principal and agent know the production function $p(y \mid a)$. There are different approaches to relax this assumption.

First, several contracting models directly assume that the agent's beliefs about the production function are biased, i.e., $\hat{p}(y \mid a) \neq p(y \mid a)$; see Fang and Moscarini (2005), Van den Steen (2005), Gervais and Goldstein (2007), Santos-Pinto (2008), De la Rosa (2011), Sautmann (2007, 2013), Spinnewijn (2013, 2015). Specifically, this approach is used to model an overconfident agent who overestimates the probability of good states and underestimates the probability of bad states. This typically allows the principal to exploit the agent by paying more after high high output and much less after low output, in which case the agent's expected payoff is below her reservation utility.

Second, a rich literature builds state-space models of "unawareness" (e.g., Dekel et al. 1998, Heifetz et al. 2006, 2013) and applies them to contract theoretical settings. Auster (2013) examines a principal-agent model with an agent who is unaware of some potential output levels y, which again implies that the contract is exploitative. Von Thadden and Zhao (2012, 2014) assume that the agent is unaware of her available actions a, which relaxes incentive compatibility when the principal implements the default action.

Third, in order to justify biased beliefs, several papers assume that the agent knows the link between action and outcomes $p(y \mid a)$, but derives anticipatory utility from optimistic beliefs.

She therefore chooses beliefs $\hat{p}(y \mid a)$ that solve the trade-off between the losses from biased decision-making and the gains from anticipation; see Bénabou and Tirole (2002), Brunnermeier and Parker (2005), and Kőszegi (2006). For an organizational context, Bénabou (2013) shows how the interaction between group members can make the suppression of bad news a strategic complement, so that collective denial of adverse signals ("groupthink") occurs in equilibrium. Immordino et al. (2015) show that if the anticipatory utility is not too important, the principal may provide incentives so that it is optimal for the agent to choose correct beliefs.

Relative to these literatures, our approach to boundedly rational expectations and contracting is more conservative. The agent derives her beliefs from the true data-generating process, as in the canonical model; she just may not take into account all empirical regularities that matter for the principal's project. The misspecification in the agent's subjective model may cause her to overestimate her productivity; however, under a weak technical restriction, she still correctly anticipates the equilibrium distribution over output, which makes the belief bias sustainable. The size of the belief bias does not depend on parameters of the agent's personality, but instead on the specification of the objective production process. Further, whether there is scope for control optimism or not, depends on the agent's function in the organization, that is, on tasks and decision-making rights. Finally, our framework is portable to different settings: For an agent with given subjective model $\mathcal{R}$, it allows us to study how beliefs and equilibrium contract vary in the production process and contracting environment.

We also contribute to the literature on Bayesian networks and directed acyclic graphs (DAGs), which have been used extensively in the artificial intelligence literature. Moreover, they are used as visual inspection tool when choosing explanatory variables, e.g., Shrier and Pratt (2008). In these papers, DAGs are interpreted as a representation of causal relationships. This viewpoint is also promoted by Pearl (2009) who provides a broad introduction to DAGs.[3] In economics, Spiegler (2016, 2017) uses Bayesian networks to model agents with boundedly rational expectations. DAGs provide a general method to capture a variety of different inference errors such as reverse causation and coarseness. We build on these insights and apply them to contracting. Other recent papers use causal models to capture boundedly rational decision makers in monetary policy (Spiegler 2019), political competition (Eliaz and Spiegler 2018), Bayesian persuasion (Eliaz et al. 2018), and decision theory (Schenone 2019).

The remainder of the paper is organized as follows. Section 2 describes our framework. In Section 3, we examine how a misspecification in the agent's subjective model affects the contracting problem. In Section 4, we characterize when a misspecification leads to biased beliefs about the production function, and illustrate the implications of this characterization. In Section 5, we analyze how simplifications in the agent's subjective model affect the informa-

[3]For other general introductions to DAGs see, for example, Koski and Noble (2009).

tiveness principle and the trade-off between risk and incentives. Section 6 concludes. Omitted proofs and further results can be found in the Online Appendix.

2 The Model

We consider a standard principal-agent problem and combine it with the Bayesian network model of boundedly rational beliefs, as introduced in Spiegler (2016).

Basic Framework. The principal proposes a contract $(w(y), a)$, where $w(y) \in W$ is the agent's wage conditional on the output $y \in Y$ and $a \in A$ the action that the principal wishes the agent to choose. Let W be the set of possible incentive schemes, $A \subset \mathbb{R}$ a finite set of actions, and $Y \subset \mathbb{R}$ a finite set of outputs. The agent can reject or accept the contract. If she rejects it, she enjoys the outside option value $\bar{U}$, while the principal earns zero. If she accepts the contract, she chooses an action $a \in A$. Mixed action profiles are denoted by $p(a) \in \Delta(A)$. The agent's personal cost of choosing a is given by a function $c(a)$. The action stochastically influences the project's output. The agent's utility from wage w is given by the utility function $u(w)$, which is weakly concave and exhibits $\lim_{w \to -\infty} u(w) = -\infty$. When the output is y and the agent's action is a, the principal's payoff is $V = y - w(y)$ and the agent's payoff is $U = u(w(y)) - c(a)$.

Causal Structure. We model the causal structure though which the agent's action affects the output y. Let $N^* = \{0, ..., n\}$ be the set of relevant project variables (or nodes). They comprise the agent's action and output, but may also include other variables. A generic realization of variable i is given by $x_i \in X_i$, where X_i is a finite set that contains at least two elements. Node 0 is the agent's action ($x_0 = a$, $X_0 = A$) and node n is the output ($x_n = y$, $X_n = Y$). We use these labels interchangeably. The state is a vector $x^* = (x_0, x_1, ..., x_n)$ and the set of all states is $X^* = \times_{i \in N^*} X_i$. Let x_S be the vector of variables in $S \subset N^*$.

Denote by $p(x_1, \ldots, x_n \mid a)$ the extended production function. For any action $a \in A$, it has full support over $X_1 \times ... \times X_n$. We represent its causal structure by an irreflexive, asymmetric and acyclic binary relation R^* over N^*, and denote it by the DAG $\mathcal{R}^* = (N^*, R^*)$, see the graph on the left in Figure 1 for an example. For two nodes $i, j \in N^*$ one may read iR^*j as "node i impacts on node j." The set of nodes that influence i is defined, with abuse of notation, as $R^*(i) = \{j \in N^* \mid jR^*i\}$. Nothing influences the agent's action, $R^*(0) = \emptyset$. The probability distribution over states $p(x^*)$ then naturally factorizes according to $\mathcal{R}^*$ via the formula

$$p(x^*) = \prod_{i \in N^*} p(x_i \mid x_{R^*(i)}). \tag{1}$$

The "objective model" $\mathcal{R}^*$ is one of the sparsest DAGs so that $p(x^*)$ factorizes according to

$\mathcal{R}^*$. That is, $\mathcal{R}^*$ contains conditional independence assumptions, and the extended production function satisfies all of them.

Figure 1: An objective model $\mathcal{R}^*$ (left) and the agent's subjective model $\mathcal{R}$ (right).

Beliefs, Personal Equilibrium, and Equilibrium Contract. The agent has her own subjective model $\mathcal{R} = (N, R)$, as, for example, the graph on the right in Figure 1. We assume that N is a subset of N^* and contains at least action a and output y, with $R(0) = \emptyset$ (the agent knows that she does not receive any information prior to choosing her action). If $\mathcal{R}$ differs from $\mathcal{R}^*$, we say that the agent's subjective model is misspecified. A simplification is a misspecification where nodes from N^* are missing in $\mathcal{R}$, but the rest is unchanged, so that iRj for $i, j \in N$ if and only if iR^*j (as in model $\mathcal{R}$ of Figure 1). We will mostly discuss the effects of simplifications, but many results hold for general misspecifications. Denote by $x = (x_i)_{i \in N}$ the state vector for the agent's subjective model and $X = \times_{i \in N} X_i$. The agent fits her subjective model to the data generated by p, so her beliefs factorize according to the formula

$$p_{\mathcal{R}}(x) = \prod_{i \in N} p(x_i \mid x_{R(i)}). \tag{2}$$

She chooses the prescribed action from the contract only if it maximizes her expected utility given the wage scheme $w(y)$ and her subjective beliefs $p_{\mathcal{R}}(x)$. Since her action potentially influences her beliefs, we formalize the agent's action choice as a personal equilibrium. For this, define by $p_{\mathcal{R}}(y \mid a'; p(a))$ the agent's belief about the distribution over output after choosing action a' when her subjective model is $\mathcal{R}$ and her personal equilibrium is $p(a)$.

Definition 1. *The action $p(a)$ is a personal equilibrium at $\mathcal{R}$ and $w(y)$ if for all actions $a \in A$ in the support of $p(a)$ we have*

$$a \in \arg\max_{a'} \sum_{y \in Y} p_{\mathcal{R}}(y \mid a'; p(a)) u(w(y)) - c(a),$$

where $p_{\mathcal{R}}(y \mid a'; p(a)) = \lim_{k \to \infty} p_{\mathcal{R}}(y \mid a'; p^k(a))$ for all actions $a' \in A$ and a sequence $p^k(a) \to p(a)$ of fully mixed action profiles.

With the full support assumption, a fully mixed action profile ensures that all conditional probabilities are well-defined, in particular, those at variables in $\mathcal{R}$ that are directly influenced by a (such as variable 1 in the subjective model $\mathcal{R}$ of Figure 1). The definition requires that equilibrium beliefs are the limit of a sequence of fully mixed profiles. A personal equilibrium always exists in our framework; see Online Appendix A.1. We call a contract $(w(y), p(a))$ an "equilibrium contract" if $p(a)$ is a personal equilibrium at $\mathcal{R}$ and $w(y)$. An optimal equilibrium contract is an equilibrium contract that maximizes the principal's expected payoff. For convenience, we denote beliefs by $p_{\mathcal{R}}(y \mid a; a^*)$ when a pure action a^* is implemented, and $p_{\mathcal{R}}(y \mid a; \alpha)$ with $p(a = 1) = \alpha$ when we have a binary action set $A = \{0, 1\}$.

The proposed solution concept is static. The agent's beliefs are derived from a probability distribution that could be influenced by the action that the equilibrium contract implements. One interpretation is that the agent is experienced and thus has data on how her action impacts on the variables in her subjective model. An alternative interpretation is that there are (or have been) many other agents in the organization who exchange data with their new colleague to which she can fit her subjective model.

3 The Optimal Equilibrium Contract

In this section, we study the properties of the optimal equilibrium contract for a given extended production function $p(x_1, ..., x_n \mid a)$ and subjective model $\mathcal{R}$. If $(w^*(y), p^*(a))$ is an optimal equilibrium contract, then $w^*(y), p^*(a)$ solve the maximization problem

$$\max_{w(y)\in W, p(a)\in\Delta(A)} \sum_{a\in A} \sum_{y\in Y} p(a)p(y \mid a)(y - w(y)) \tag{3}$$

subject to the constraints

$$p(a) \in \Delta(A) \;\; \textit{is a personal equilibrium at } \mathcal{R} \textit{ and } w(y), \tag{IC}$$

$$\sum_{a'\in A} \sum_{y\in Y} p(a')[p_{\mathcal{R}}(y \mid a'; p(a))u(w(y)) - c(a)] \geq \bar{U}. \tag{PC}$$

If the agent's subjective model $\mathcal{R}$ equals the objective model $\mathcal{R}^*$, the problem collapses to the canonical principal-agent problem and can be solved as suggested by Grossman and Hart (1983). We first find for each pure action $a \in A$ the wage scheme $w(y)$ that implements this action at lowest possible cost. Then we choose the action-incentive scheme combination that maximizes the principal's profit. If the agent's subjective model $\mathcal{R}$ differs from the objective model $\mathcal{R}^*$, we find the optimal equilibrium contract by applying the same procedure. However, since the agent's beliefs $p_{\mathcal{R}}(y \mid a; p(a))$ may depend on the implemented action $p(a)$, the first step has to be done for all pure and mixed actions $p(a) \in \Delta(A)$.

Suppose the agent is risk-averse with unlimited liability, and the principal implements a (possibly mixed) action $p(a)$. The Kuhn-Tucker conditions for the principal's problem are then necessary and sufficient for an optimum. Choose any action a in the support of $p(a)$. The optimal incentive scheme is then characterized by the first-order condition

$$\frac{1}{u'(w(y))} = \frac{p_{\mathcal{R}}(y; p(a))}{p(y)} \left[\mu + \sum_{a' \in A} \lambda_{a'} \frac{p_{\mathcal{R}}(y \mid a; p(a)) - p_{\mathcal{R}}(y \mid a'; p(a))}{p_{\mathcal{R}}(y; p(a))} \right] \tag{4}$$

for all $y \in Y$, where μ and $\lambda_{a'}$ are the usual Lagrange multipliers for the participation and incentive compatibility constraint, respectively; $p(y) = \sum_{a \in A} p(a)p(y \mid a)$ is the distribution over output and $p_{\mathcal{R}}(y; p(a)) = \sum_{a \in A} p(a)p_{\mathcal{R}}(y \mid a; p(a))$ is the agent's belief about the distribution over output when her subjective model is $\mathcal{R}$ and the equilibrium action is $p(a)$.

Equation (4) allows us to disentangle how a misspecification in $\mathcal{R}$ may change the contracting problem. First, the *PC* is affected when the agent holds biased beliefs about the equilibrium distribution over outcomes; see the first term on the right of equation (4). In Subsection 3.1, we state a sufficient condition on $\mathcal{R}$ so that this belief is unbiased. Second, the *IC* may be affected. Suppose the principal implements a pure action a and $p_{\mathcal{R}}(y; a) = p(y)$. The ratio in the squared brackets then becomes $1 - \frac{p_{\mathcal{R}}(y|a';a)}{p_{\mathcal{R}}(y|a;a)}$, in which case the optimal incentive scheme depends on a likelihood ratio as in the canonical framework. Any difference between the contracts under the subjective model $\mathcal{R}$ and the objective model $\mathcal{R}^*$ is then driven by differences between the objective and subjective likelihood ratios. In Subsection 3.2, we examine how these differences may affect the optimal equilibrium contract.

3.1 Correct Expectations on the Equilibrium Path

We use a Bayesian network result from Spiegler (2017) that characterizes under what circumstances the agent's beliefs over the output distribution are identical to the equilibrium output distribution, so that $p_{\mathcal{R}}(y; p(a)) = p(y)$ for all $p(a) \in \Delta(A)$. To this end, we introduce a few definitions. A v-collider is a triple of nodes (i, j, k) such that iRj, kRj and there is no link between i and k (neither iRk nor kRi is in R). The set of v-colliders of a DAG is called its v-structure. A DAG is called perfect if it has an empty v-structure. A subset of nodes $S \subset N$ is a clique in $\mathcal{R} = (N, R)$ if iRj or jRi for any two nodes $i, j \in S$. For example, in the DAG $\mathcal{R}^*$ from Figure 1, the set $S = \{1, 3, 4\}$ is a clique, while the set $S' = \{2, 3, 4\}$ is not. Each node is a clique in itself, so the output node n is a clique. As for the output y we define, for any clique $S \subset N$, $p(x_S) = \sum_{a \in A} p(a)p(x_S \mid a)$, and $p_{\mathcal{R}}(x_S; p(a)) = \sum_{a \in A} p(a)p_{\mathcal{R}}(x_S \mid a; p(a))$. The following result essentially restates Proposition 2 from Spiegler (2017).

Proposition 1 (Equilibrium Beliefs). *If the agent's model $\mathcal{R} = (R, N)$ is perfect, her equilibrium beliefs satisfy $p_{\mathcal{R}}(x_S; p(a)) = p(x_S)$ for all $p(a) \in \Delta(A)$ and any clique $S \subset N$.*

If the agent's subjective model $\mathcal{R}$ is perfect, then in a personal equilibrium, the agent correctly anticipates the marginal distribution over each variable in her model, and also the joint distribution over variables in cliques. The intuition behind this result is that perfectness excludes biased estimates due to neglect of correlation. Imagine two variables i, j that influence a third variable k. Suppose that i and j are correlated, and that the agent treats them as uncorrelated. Through the application of the factorization formula (2), the agent may then obtain a biased estimate of the marginal distribution $p(x_k)$. Perfectness implies that the agent always checks for correlations between two variables i, j when, according to her subjective model, they influence a third variable k. We obtain two useful corollaries from Proposition 1.

Corollary 1. *If the agent's model $\mathcal{R} = (R, N)$ is perfect and her equilibrium action is a pure action a^*, her equilibrium beliefs satisfy $p_{\mathcal{R}}(x_S \mid a^*; a^*) = p(x_S \mid a^*)$ for every clique $S \subset N$.*

If the equilibrium contract implements a pure action a^*, the agent's belief over the joint distribution of any clique S conditional on her equilibrium action is correct. Corollary 1 is in general not true if the equilibrium contract implements a mixed action $p^*(a)$. While the agent still gets the marginal equilibrium distribution over each variable right, her beliefs may also exhibit $p_{\mathcal{R}}(x_i \mid a'; p^*(a)) \neq p(x_i \mid a')$ for an action a' in the support of $p^*(a)$. Thus, the agent's expected utility conditional on a' may be biased, $\mathbb{E}_{\mathcal{R}}[u(w(y)) \mid a'; p^*(a)] \neq \mathbb{E}[u(w(y)) \mid a']$. The second direct implication of Proposition 1 is the following result.

Corollary 2. *Suppose $(w(y), p(a))$ is an equilibrium contract. If $\mathcal{R} = (R, N)$ is perfect, the PC is satisfied at this contract if and only if this is also the case under the objective model $\mathcal{R}^*$.*

If $\mathcal{R}$ is perfect, the incentive scheme has to satisfy the same participation constraint as under the objective model. Thus, an agent with a misspecified – but perfect – model cannot be exploited. Throughout the paper, we will assume that $\mathcal{R}$ is perfect. As we see next, this does not imply that the principal cannot benefit from the agent's misperception.

3.2 Incentive Effects

We examine how a misspecification in the agent's subjective model $\mathcal{R}$ can change the equilibrium contract when $\mathcal{R}$ is perfect. By Corollary 2, only the incentive compatibility constraint could then be affected by the misspecification. We examine a simple setting with two effort

Figure 2: Objective model $\mathcal{R}^*$ (left) and subjective model $\mathcal{R}$ (right) in the marketer example.

levels $a \in \{0, 1\}$, two output levels $y \in \{y_L, y_H\}$ with $y_H > y_L$, and cost $c(1) = c > c(0) = 0$. The probability of output y_H increases in the agent's effort.

Consider the marketer example from the introduction. Figure 2 shows the objective model $\mathcal{R}^*$ and the agent's subjective model $\mathcal{R}$. Node 1 is the set of customers who are informed about the firm's product. It can be small ($x_1 = 0$) or large ($x_1 = 1$). Node 2 is the firm's reputation, which can be bad ($x_2 = 0$) or good ($x_2 = 1$). The subjective model $\mathcal{R}$ captures that the agent does not take reputation into account. For the objective probability distribution, we use the parametrization $p(x_i = 1 \mid x_{R(i)}) = \beta_i + \sum_{j \in R(i)} \beta_{ji} x_j$ for $i \in \{1, 2\}$ and $p(y_H \mid x_1, x_2) = \beta_3 + \beta_{13} x_1 + \beta_{23} x_2$. Making cold-calls increases the customer set, $\beta_{01} > 0$, and decreases reputation, $\beta_{02} < 0$; the customer set x_1 and reputation x_2 have a positive influence on sales y, $\beta_{13} > 0$ and $\beta_{23} > 0$. We obtain the following result.

Proposition 2 (Marketer Example)**.** *Consider the marketer example of this subsection.*

(a) The simplification in the agent's subjective model $\mathcal{R}$ relaxes the IC for $\alpha = 1$.

(b) The optimal equilibrium contract implements $\alpha \in \{0, 1\}$. If and only if effort costs c are small enough, the optimal equilibrium contract implements $\alpha = 1$ and the principal strictly benefits from the simplification in the agent's subjective model $\mathcal{R}$.

Before we prove this result, we explain the intuition behind it and its implications. First, consider statement (a). When the principal implements $\alpha = 1$, the agent overestimates the drop in expected output when she exerts low instead of high effort. According to her subjective model $\mathcal{R}$, the only effect of her action on the output occurs through the size of the customer set x_1; she does not take into account that a deviation to low effort would also have a positive effect on expected reputation x_2, which translates into a positive effect on expected output. Formally, the *IC* under the objective model $\mathcal{R}^*$ is

$$[\beta_{01}\beta_{13} + \beta_{02}\beta_{23}]\,(u(w(y_H)) - u(w(y_L))) - c \geq 0. \tag{5}$$

The term in squared brackets is the effect of effort on output and contains the customer set channel $\beta_{01}\beta_{13}$ and the reputation channel $\beta_{02}\beta_{23}$. Under the subjective model $\mathcal{R}$, this second channel is missing, so that the IC that implements $\alpha = 1$ becomes

$$\beta_{01}\beta_{13}\,(u(w(y_H)) - u(w(y_L))) \geq c. \tag{6}$$

Since the effect of effort on reputation β_{02} is negative, the simplification in $\mathcal{R}$ relaxes the IC. As long as $\alpha \in (0, 1)$, the reputation effect is partly reflected in $p(y_H \mid x_1)$; the extent of this depends on α since α affects the correlation between customer set and reputation.

Next, consider statement (b). The observation that the principal will implement a pure action would be trivial in the canonical framework with rational expectations. This is not the case here as the agent's perceived effect of effort on output $p_{\mathcal{R}}(y_H \mid a = 1; \alpha) - p_{\mathcal{R}}(y_H \mid a = 0; \alpha)$ may vary non-monotonically in α. In the present setting, the perceived effect of effort on output is maximal at $\alpha = 1$, so that there is no reason for the principal to implement a mixed action. At the end of this subsection, we present an example where the unique optimal equilibrium contract indeed implements a mixed action $\alpha \in (0, 1)$.

Importantly, if the agent chooses a pure action, then, by Corollary 1 and the fact that $\mathcal{R}$ is perfect, she correctly anticipates the joint distribution over all variables in $\mathcal{R}$ conditional on her equilibrium action. Thus, in the data that the agent gets under the optimal equilibrium contract, there are no informational cues which could alarm her about a misspecification in her subjective model. This is a crucial difference between the present framework and models where beliefs over outcomes are biased for equilibrium actions.

Finally, the last part of statement (b) spells out that the principal strictly benefits from the simplification in $\mathcal{R}$ when effort costs are small enough so that it is profitable to implement high effort. The principal would have no incentive to correct the agent's view on the production process (if this were possible). This is of course not true in general. For example, if the agent's action has a positive effect on reputation, $\beta_{02} > 0$, the simplification in $\mathcal{R}$ tightens the IC for $\alpha = 1$ as the agent does not take all positive effects of her action on output into account.

Proof of Proposition 2. To illustrate our approach, we present the proof of Proposition 2. We first derive $p_{\mathcal{R}}(y_H \mid a; \alpha)$ for a given mixed equilibrium action $\alpha \in (0, 1)$. The agent's equilibrium beliefs about the joint probability distribution of the variables in $\mathcal{R}$ is given by $p_{\mathcal{R}}(a, x_1, y) = p(a)p(x_1 \mid a)p(y \mid x_1)$. Since node 0 and node 1 form an ancestral clique, $p(x_1 \mid a)$ is independent of α and we have $p(x_1 = 1 \mid a) = \beta_1 + \beta_{01}a$. However, $p(y \mid x_1)$ depends on α since the distribution over y also depends on x_2. To get $p(y \mid x_1)$, we first derive $p(x_2 = 1 \mid x_1)$, i.e., the probability that $x_2 = 1$ given that value x_1 is observed at node 1 when the agent's equilibrium action is α. We get

$$p(x_2 = 1 \mid x_1 = 1) = \frac{\alpha(\beta_1 + \beta_{01})(\beta_2 + \beta_{02}) + (1-\alpha)\beta_1\beta_2}{\beta_1 + \alpha\beta_{01}}, \tag{7}$$

$$p(x_2 = 1 \mid x_1 = 0) = \frac{\alpha(1 - \beta_1 - \beta_{01})(\beta_2 + \beta_{02}) + (1-\alpha)(1-\beta_1)\beta_2}{1 - \beta_1 - \alpha\beta_{01}}. \tag{8}$$

With this we can calculate the equilibrium probability that output y_H realizes after observing $x_1 = 1$ and $x_1 = 0$, respectively:

$$p(y_H \mid x_1 = 1) = \beta_3 + \beta_{13} + \frac{\alpha(\beta_1 + \beta_{01})(\beta_2 + \beta_{02}) + (1-\alpha)\beta_1\beta_2}{\beta_1 + \alpha\beta_{01}}\beta_{23}, \tag{9}$$

$$p(y_H \mid x_1 = 0) = \beta_3 + \frac{\alpha(1 - \beta_1 - \beta_{01})(\beta_2 + \beta_{02}) + (1-\alpha)(1-\beta_1)\beta_2}{1 - \beta_1 - \alpha\beta_{01}}\beta_{23}. \tag{10}$$

From $p_{\mathcal{R}}(a, x_1, y)$ we can now calculate the agent's subjective probability of a high output after high and low effort, respectively:

$$p_{\mathcal{R}}(y_H \mid a = 1; \alpha) = (\beta_1 + \beta_{01})p(y_H \mid x_1 = 1) + (1 - \beta_1 - \beta_{01})p(y_H \mid x_1 = 0), \tag{11}$$

$$p_{\mathcal{R}}(y_H \mid a = 0; \alpha) = \beta_1 p(y_H \mid x_1 = 1) + (1 - \beta_1)p(y_H \mid x_1 = 0). \tag{12}$$

We then use these terms to compute the *IC* for $\alpha \in (0, 1)$,

$$[p_{\mathcal{R}}(y_H \mid a = 1; \alpha) - p_{\mathcal{R}}(y_H \mid a = 0; \alpha)]\,(u(w(y_H)) - u(w(y_L))) = 0. \tag{13}$$

Applying Definition 1, we obtain the *IC* for $\alpha = 1$, which is the inequality in (6). Since $\beta_{02} < 0$, this completes the proof of statement (a). To prove statement (b), note first that both *IC* and *PC* must be binding at the optimal equilibrium contract. Simple calculations show that $\beta_{01}, \beta_{13}, \beta_{23} > 0$ and $\beta_{02} < 0$ imply

$$p_{\mathcal{R}}(y_H \mid a = 1; \alpha) - p_{\mathcal{R}}(y_H \mid a = 0; \alpha) \leq \beta_{01}\beta_{13}, \tag{14}$$

for all $\alpha \in (0, 1]$; that is, when the agent exerts high effort with positive probability, her perceived effect of effort on output is largest at $\alpha = 1$. The principal then cannot gain from implementing a mixed action. Finally, given that the optimal equilibrium contract implements either $\alpha = 0$ or $\alpha = 1$, the last part of statement (b) follows from a simple comparison of expected profits under the equilibrium contracts that implement these two actions. □

Mixed action example. We show by example that it is not always optimal for the principal to implement a pure action. Consider again the marketer example. Assume that the agent is risk-neutral, protected by limited liability so that $w(y) \geq 0$, her outside option value is zero, and

$y_L = 0$. Suppose payoff parameters are such that the principal optimally implements $\alpha > 0$. Standard arguments show that $w(y_L) = 0$, and that $w(y_H)$ is chosen so that the IC in (13) is satisfied. The principal's expected payoff from this contract is then

$$\mathbb{E}[V] = [\alpha p(y_H \mid a = 1) + (1 - \alpha)p(y_H \mid a = 0)] \left(y_H - \frac{c}{\Delta_{\mathcal{R}}(\alpha)} \right), \tag{15}$$

where $\Delta_{\mathcal{R}}(\alpha) = p_{\mathcal{R}}(y_H \mid a = 1; \alpha) - p_{\mathcal{R}}(y_H \mid a = 0; \alpha)$ is the agent's perceived effect of effort on output. The slope of $\Delta_{\mathcal{R}}(\alpha)$ at $\alpha = 1$ is

$$\left. \frac{d\Delta_{\mathcal{R}}(\alpha)}{d\alpha} \right|_{\alpha=1} = \beta_{01}\beta_{02}\beta_{23} \left(\frac{\beta_1}{\beta_1 + \beta_{01}} - \frac{1 - \beta_1}{1 - \beta_1 - \beta_{01}} \right). \tag{16}$$

Let the agent's action have a positive impact on both customer set and reputation, $\beta_{01} > 0$ and $\beta_{02} > 0$. Then for $\beta_{01} \to 1 - \beta_1$ the slope converges to minus infinity. Thus, if all else equal β_{01} is sufficiently close to $1 - \beta_1$, then, starting from $\alpha = 1$, a small reduction in α reduces $w(y_H)$, and in terms of profits, this reduction overcompensates the smaller probability of high output. The optimal equilibrium contract then implements $\alpha \in (0, 1)$. Thus, when the agent is induced to switch between periods of working hard and periods of shirking, her effort appears to her as particularly important for the final output.

4 Behavioral Rationality

We learned in the previous section that a simplification in the agent's subjective model may bias her beliefs about the production function, so that the incentive compatibility constraint is affected. However, does a simplification in $\mathcal{R}$ automatically imply that the agent's beliefs are biased? In this section, we find that the answer to it is negative. The agent may correctly anticipate the true production function, i.e., $p_{\mathcal{R}}(y \mid a; p(a)) = p(y \mid a)$ for all $a \in A$ and $p(a) \in \Delta(A)$, even when her subjective model $\mathcal{R}$ misses out variables from $\mathcal{R}^*$. Specifically, this statement can hold independent of the parametrization of the extended production function as long as it factorizes according to the objective causal structure $\mathcal{R}^*$. We then say that the agent is "behaviorally rational."

We characterize for given objective model $\mathcal{R}^*$ when the agent is behaviorally rational, and when her beliefs about the production function change if an (additional) node from N^* is dropped from her subjective model $\mathcal{R}$. We will see that two extended production functions – which involve the same set of nodes N^* and may give rise to the same $p(y \mid a)$ – can differ in the extent to which simplifications affect the agent's beliefs about $p(y \mid a)$. This extent depends on the "channels" in $\mathcal{R}^*$ through which the agent's action affects the output. Intuitively, they describe the agent's role in the organization, that is, which components or behaviors of others the agent affects directly or indirectly through her action. This allows us to identify several processes in an organization that potentially cause the agent to overestimate the productivity of her effort. We proceed as follows. In Subsection 4.1, we extend our marketer example to illustrate the influence of the agent's job on the scope for biased beliefs and control optimism. In Subsection 4.2, we characterize when the agent is behaviorally rational and generalize the main findings from Subsection 4.1.

4.1 The Agent's Job and the Scope for Control Optimism

We examine the interaction between the agent's job, model misspecification, and incentives. Let the agent first work as an ordinary marketer whose job is to increase sales. This time, making cold-calls is not part of her job. Her effort only has a (positive) effect on the customer set. However, there is a group of employees engaged in telemarketing. Their effort – making cold-calls – impacts on the customer set and the firm's reputation in the usual manner. The objective model $\mathcal{R}^*$ on the left of Figure 3a represents the causal structure of this extended production function. Throughout, we use our parametrization with binary outcomes at all variables $i \in N^*$ and $p(x_i = 1 \mid x_{R(i)}) = \beta_i + \sum_{j \in R(i)} \beta_{ji} x_j$. The telemarketers either conduct cold-calls or not, $\beta_1 \in \{0, 1\}$; cold-calls have a negative effect on reputation, $\beta_{13} < 0$; the customer set has a positive effect on reputation, $\beta_{23} > 0$.[4] All formal proofs of this subsection are in Online Appendix A.2.

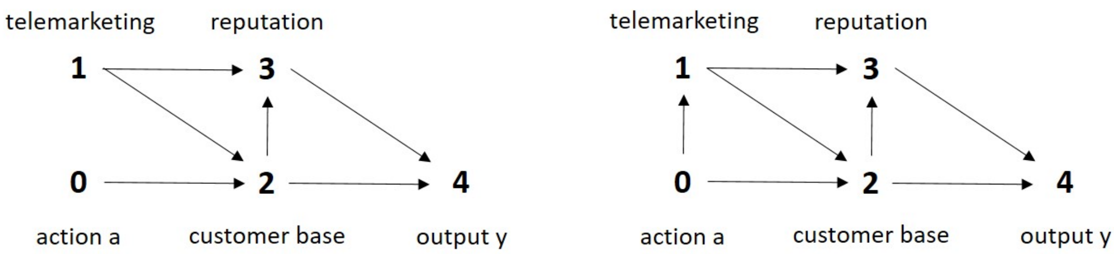

Figure 3a: Objective model $\mathcal{R}^*$ (left) when the agent works as ordinary marketer, and objective model $\mathcal{R}^{**}$ (right) when the agent works as "head of marketing."

Imagine that the marketer neither takes into account the telemarketers' operation nor the firm's reputation so that her subjective model is given by $\mathcal{R}$ on the upper-left of Figure 3b. When choosing effort, she only considers how her action impacts sales through the size of the customer set. Does this misspecification change incentives? The answer is negative. We can show – using the results from the next subsection – that the agent's subjective beliefs about the

[4]To simplify the discussion, we introduce the link between customer set and reputation. Moreover, we violate our full support assumption by assuming $p(x_1 = 1) \in \{0, 1\}$. Formally, we consider the limit case of $p(x_1 = 1) \in \{\varepsilon, 1 - \varepsilon\}$ when $\varepsilon \to 0$.

production function are correct, so that $p_{\mathcal{R}}(y_H \mid a; \alpha) = p(y_H \mid a)$ for all $a \in \{0, 1\}$ and $\alpha \in [0, 1]$. Thus, given her role in the principal's project (as captured by $\mathcal{R}^*$), the subjective model $\mathcal{R}$ is rich enough to produce correct predictions for off-equilibrium actions. The agent may ignore important parts of the project and still act as if she were fully rational. The equilibrium contract is then the same as in the canonical model.

Importantly, telemarketing still matters for the principal. The probability distribution over sales depends on whether cold-calls are made or not. It is just not essential for the agent to take the telemarketers' activity into account when estimating how her action influences the output.

Is there any simplification that makes the agent overestimate the effectiveness of her effort, such that the principal benefits from it? Again, the answer is negative. If the agent does not take node 2 into account, she believes that her action has no consequences for the output. It would then be impossible implement high effort. If only node 1 or only node 3 were missing in her subjective model, the agent would again have correct beliefs about the production function. Thus, there is no scope for control optimism when the agent works as ordinary marketer.

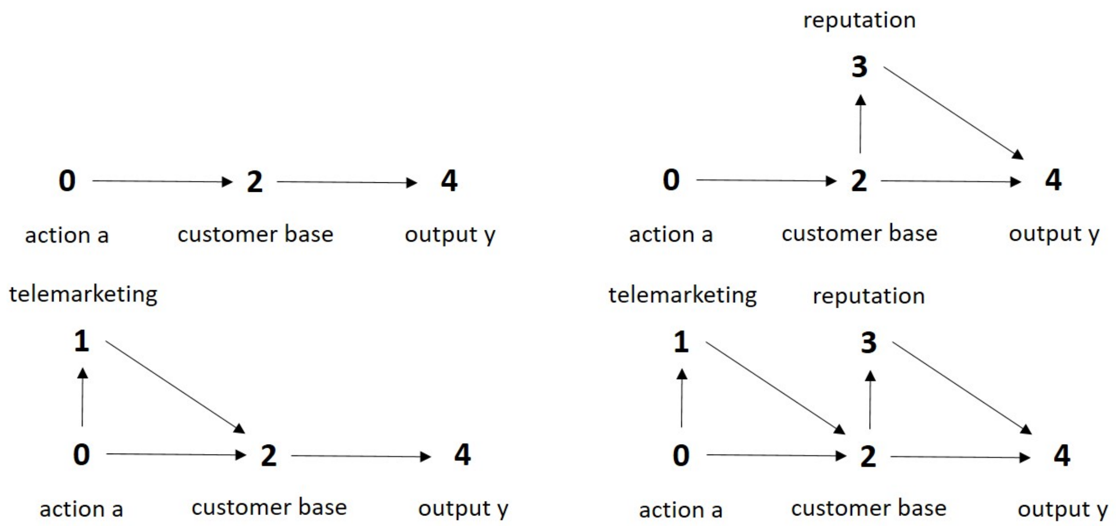

Figure 3b: Subjective models $\mathcal{R}$ (upper-left), $\mathcal{R}_1$ (upper-right), $\mathcal{R}_2$ (lower-left), $\mathcal{R}_3$ (lower-right).

Next, we alter the agent's job by promoting her to "head of marketing." Her action now influences the telemarketers' effort, for example, by motivating or inspiring the telemarketers. Instead of $p(x_1 = 1) = \beta_1$, we now have $p(x_1 = 1 \mid a) = \beta_1 + \beta_{01}a$. To keep things as close as possible to the previous case, we assume $\beta_1 = 0$ and $\beta_{01} = 1$.[5] Hence, the agent needs to act in order to get the telemarketers going. The objective model of the extended production function is given by $\mathcal{R}^{**}$ on the right of Figure 3a. How does a misspecification in the agent's subjective model now affect equilibrium beliefs and incentives in this environment?

[5] We show in the proofs for this subsection that our results do not depend on this assumption.

Let us first assume that the agent has the same subjective model $\mathcal{R}$ as before. She neglects both the telemarketer's activity and the firm's reputation. This of course is not realistic since as "head of marketing" the agent should be aware of her subordinates' basic activities; so we will relax this assumption below. The misspecification now affects incentives. Under the objective model $\mathcal{R}^{**}$ the IC that implements high effort would be

$$[(\beta_{02} + \beta_{01}\beta_{12})(\beta_{24} + \beta_{23}\beta_{34}) + \beta_{01}\beta_{13}\beta_{34}](u(w(y_H)) - u(w(y_L))) \geq c. \tag{17}$$

In contrast, under the subjective model $\mathcal{R}$ this IC becomes

$$(\beta_{02} + \beta_{01}\beta_{12})(\beta_{24} + \beta_{23}\beta_{34})(u(w(y_H)) - u(w(y_L))) \geq c, \tag{18}$$

Thus, the simplification relaxes the IC. This occurs since the agent does not take the negative influence of cold-calls on reputation into account. However, through the estimate of the link between the agent's action and the customer set, she implicitly takes into account her positive influence on the telemarketers' effort which in turn positively affects the customer set. Therefore, by being promoted to a job where the agent also influences telemarketing, she starts to overestimate her productivity. The principal benefits from this since the misspecification again reduces the need to provide incentives.

The assumption that the agent does not include the telemarketer's activity in her subjective model seems a bit odd, given that she is the head of marketing. Therefore, let her subjective model be given by $\mathcal{R}_2$ on the lower-left of Figure 3b. She now takes into account her influence on the telemarketers, and that the telemarketers increase the customer base when exerting effort. Does this inclusion correct, at least partly, the agent's beliefs? It turns out that this is not the case. The models $\mathcal{R}$ and $\mathcal{R}_2$ produce the same beliefs about the effectiveness of effort, i.e., $p_{\mathcal{R}}(y_H \mid a;\alpha) = p_{\mathcal{R}_2}(y_H \mid a;\alpha)$ for all $a \in \{0,1\}$ and $\alpha \in [0,1]$. Thus, including more variables does not necessarily make the agent more rational. This also holds for the models $\mathcal{R}_1$ and $\mathcal{R}_3$ in Figure 3b. Note that $\mathcal{R}_3$ is almost equal to the objective model $\mathcal{R}^{**}$. Only the link between telemarketing and reputation is missing. Yet, all subjective models in this figure produce the same beliefs. Thus, a small misspecification in the agent's subjective model can render several important variables as inessential for estimating the production function.

Proposition 3 (Scope for Control Optimism)**.** *Consider the job examples of this subsection.*

(a) If the agent works as ordinary marketer (objective model $\mathcal{R}^{*}$), the misspecification in $\mathcal{R}$ has no effect on the IC and the optimal equilibrium contract is the same as under the canonical model. There is no simplification that generates control optimism.

(b) If the agent works as "head of marketing" (objective model $\mathcal{R}^{**}$), the misspecification in

$\mathcal{R}$ generates control optimism and relaxes the IC; the subjective models $\mathcal{R}$, $\mathcal{R}_1$, $\mathcal{R}_2$, and $\mathcal{R}_3$ generate the same beliefs about the production function.

Proposition 3 illustrates how the agent's job matters for optimal incentives. The two jobs with objective models $\mathcal{R}^*$ and $\mathcal{R}^{**}$ may give rise to the same production function $p(y \mid a)$,[6] so that incentives would be identical under rational expectations. However, effort motivation is larger under the job with objective model $\mathcal{R}^{**}$ when the agent's subjective model is simplified in a way that benefits the principal. The crucial difference between the jobs are the sets of channels through which the action affects the output. In the next subsection, we will formally define these channels.

The findings in Proposition 3 allow for several new interpretations. First, parts (a) and (b) combined demonstrate that an agent's degree of control optimism may be determined by the nature of her job. In the example, the agent with misspecified model $\mathcal{R}$ was behaviorally rational in her job as ordinary marketer, but overestimated the importance of her effort after being promoted to "head of marketing" where she influences the actions of others. Thus, in our framework, the agent's control optimism is not caused by certain features of her personality, but it is a consequence of her environment when her subjective model does not capture all empirical regularities of this environment.

Second, part (b) offers a new perspective on the phenomenon that managers often do not completely understand the difficulties that their rank-and-file workers face (e.g., Porter and Nohria 2018). Specifically, this can happen when an individual worked as specialist in her previous position, but then was promoted to a management position where she influences the activity of individuals whose jobs she often does not fully understand. This lack of knowledge is typically regarded as a problem since it may lead to conflicts or inefficient managerial decisions. However, as our example shows, it also can have positive effects on effort motivation, in particular, when the agent does not take into account a partial negative effect of her subordinates' behavior on the final output, and she motivates this behavior through her action.

Third, what leads to control optimism in part (b) is the agent's ignorance of the partially negative consequences of her subordinates' activity for the final output. Our framework does not provide an explanation for why a certain node is in the agent's subjective model or not. However, in an organizational context, there can be good reasons why the agent only takes into account the positive aspects of her subordinates' activity. For example, subordinates may have an incentive to communicate why their effort is effective, and at the same time be reluctant to

[6] Specifically, when we denote parameters for the job with objective model $\mathcal{R}^*$ ($\mathcal{R}^{**}$) with "$*$" ("$**$") we only have to select parameters so that $\beta^*_{02}(\beta^*_{24} + \beta^*_{23}\beta^*_{34}) = (\beta^{**}_{02} + \beta^{**}_{01}\beta^{**}_{12})(\beta^{**}_{24} + \beta^{**}_{23}\beta^{**}_{34}) + \beta^{**}_{01}\beta^{**}_{13}\beta^{**}_{34}$.

communicate the disadvantages of their activity.[7] In terms of our example, the telemarketers may know that cold-calls displeases some customers. However, they may not want to make the agent aware of this, e.g., when having career concerns. Indeed, it is difficult for CEOs to obtain unbiased information about what their subordinates to. Porter et al. (2004) find that "[all] information coming to the top is filtered [...] Receiving solid information becomes even more difficult because immediately upon appointment, the CEO's relationships change. Former peers and subordinates who used to constitute an informal channel [...] go on their guard. Even those the CEO was closest to are wary of delivering bad news."

4.2 A General Result on Behavioral Rationality

To obtain a general result on behavioral rationality, we first assume that the objective model $\mathcal{R}^*$ is perfect, and that the agent's subjective model $\mathcal{R}$ is a simplification of $\mathcal{R}^*$. Note that $\mathcal{R}$ will then be perfect. No v-structure emerges if we take out nodes from a perfect $\mathcal{R}^*$ and all links attached to them. The assumptions on $\mathcal{R}^*$ and $\mathcal{R}$ are not overly restrictive: Note that any probability distribution $p(x^*)$ factorizes according to some perfect DAG $\mathcal{R}^*$. The assumption on $\mathcal{R}$ is satisfied by almost all subjective models we consider in this paper. Below, we (partially) extend our behavioral rationality result to imperfect objective models. All formal proofs for this subsection are in Online Appendix A.3.

In the following, we characterize for any perfect $\mathcal{R}^*$ the subset of nodes the agent needs to have in her subjective model $\mathcal{R}$ so that she acts as if she had fully rational beliefs about the production function. We formally define behavioral rationality.

Definition 2. *An agent with subjective DAG $\mathcal{R}$ is behaviorally rational if, at any probability distribution $p(x) \in \Delta(X)$ that factorizes according to $\mathcal{R}^*$, we have $p_{\mathcal{R}}(y \mid a; p(a)) = p(y \mid a)$ for all $a \in A$ and $p(a) \in \Delta(A)$.*

We use the following definitions and results from the Bayesian network literature. Consider any DAG $\mathcal{R} = (N, R)$. Its skeleton $(N, \tilde{R})$ is obtained by making the DAG undirected. We have $i\tilde{R}j$ if and only if iRj or jRi.

Definition 3. *Two DAGs $\mathcal{R}$ and $\mathcal{G}$ are equivalent if $p_{\mathcal{R}}(x) \equiv p_{\mathcal{G}}(x)$ for every $p(x) \in \Delta(X)$.*

Proposition 4 (Verma and Pearl 1991)**.** *Two DAGs $\mathcal{R}$ and $\mathcal{G}$ are equivalent if and only if they have the same skeleton and v-structure.*

[7] A large literature in organizational economics studies strategic information transmission in organizations (e.g., Aghion and Tirole 1997). The models in this literature are built on the common prior assumption, i.e., all parties have a correct prior of what other parties may know. This is not the case in our framework. The crucial point here is that strategic communication may directly influence how the agent perceives the production process.

Two different models produce the same beliefs if they share the same skeleton and the same set of v-colliders. To illustrate, consider the two models in Figure 1. The DAGs $\mathcal{R}^*$ and $\mathcal{R}$ are not equivalent since they have different skeletons. Next, consider a DAG $\mathcal{G}$ that only differs from $\mathcal{R}$ in Figure 1 in that the link between the nodes 1 and 4 is reversed. $\mathcal{R}$ and $\mathcal{G}$ then have the same skeleton, but a different v-structure, so that they are not equivalent.

We need a few more definitions. A subset of nodes $M \subset N$ is called ancestral in $\mathcal{R}$ if for all nodes $i \in M$ we have $R(i) \subset M$. A path τ of length d from node i to node j is a sequence of nodes $\tau_0, \tau_1, ..., \tau_d$ so that $\tau_0 = i$, $\tau_d = j$, and $\tau_{h-1}\tilde{R}\tau_h$ for all $h \in \{1, ..., d\}$. The length of the shortest path between i and j is called the distance between these nodes and denoted by $d(i, j)$. A path of length d is active if there is no $h \in \{1, ..., d-1\}$ so that $\tau_{h-1}R\tau_h$ and $\tau_{h+1}R\tau_h$.

Define by $\mathcal{E}$ the set of DAGs in the equivalence class of $\mathcal{R}^*$ in which the action node 0 is ancestral (nothing influences the agent's action). In each of these DAGs, all active paths between the action node 0 and any node i point towards i. Thus, the assumption that node 0 is ancestral pins down the direction of many links in a perfect DAG. We call such links "fundamental links." There is a close connection between fundamental links and the set of nodes that can be removed while maintaining behavioral rationality.

Definition 4. *Consider two nodes $i, j \in N^*$. If iGj for all $\mathcal{G} = (G, N^*) \in \mathcal{E}$, then the link iGj is called fundamental link and denoted by iEj.*

An intuition for fundamental links is that they capture empirically relevant directions of causality (given agreement on the ancestral node). Specifically, they describe how the agent's action impacts on other variables. Consider $\mathcal{R}^*$ from Figure 1. Since the action node is ancestral, the links pointing from node 0 to other nodes are fundamental ($0R^*1$, $0R^*2$, and $0R^*3$). Thus, the two links pointing into the output node ($1R^*4$ and $3R^*4$) also must be fundamental. If we would turn around one of them, we would create a v-collider since there is no link between node 0 and node 4. The remaining links $1R^*2$, $1R^*3$, and $2R^*3$ are not fundamental. Below, we present an algorithm that identifies all fundamental links in any perfect DAG $\mathcal{R}^*$. For now, we go a step further and consider sequences of fundamental links.

Definition 5. *Let τ be an active path in $\mathcal{R}^*$. Then τ is a fundamental active path if all the links between neighboring nodes in τ are fundamental.*

Fundamental active paths are what we so far called "channels." Consider again $\mathcal{R}^*$ from Figure 1. The path $\tau = \{0, 1, 4\}$ is a fundamental active path since both links $0R^*1$ and $1R^*4$ are fundamental. In contrast, the active path $\tau' = \{0, 2, 3, 4\}$ is not fundamental since the link $2R^*3$ is not fundamental. We define the set of nodes that are part of at least one fundamental active path between the action and the output by

$$H^*(\mathcal{R}^*) := \{i \in N^* \mid i \text{ is part of a fundamental active path between 0 and } n \text{ in } \mathcal{R}^*\}.$$

It turns out that the nodes in $H^*(\mathcal{R}^*)$ are exactly those nodes the agent needs to have in her subjective DAG in order to be behaviorally rational, provided that her subjective DAG is a simplification of $\mathcal{R}^*$. We can prove this by finding a DAG $\mathcal{G}$ that is equivalent to $\mathcal{R}^*$ and in which there are no links pointing from nodes in $N^* \setminus H^*(\mathcal{R}^*)$ to nodes in $H^*(\mathcal{R}^*)$. In this DAG, the nodes that are not in $H^*(\mathcal{R}^*)$ have no influence on the output, so the agent can safely ignore them. By Proposition 4, the agent correctly anticipates the production function if $H^*(\mathcal{R}^*) \subseteq N$.

Proposition 5 (Behavioral Rationality)**.** *Let $\mathcal{R}^*$ be a perfect DAG and let the agent's subjective DAG $\mathcal{R}$ be a simplification of $\mathcal{R}^*$. The agent is behaviorally rational if and only if $\mathcal{R}$ contains all nodes from $H^*(\mathcal{R}^*)$.*

Proposition 5 implies that the agent does not necessarily have to take into account all variables of her (potentially) complex environment in order to be behaviorally rational. In particular, this holds independent of the parametrization of the extended production function. For example, when $p(x_1, ..., x_4 \mid a)$ factorizes according to $\mathcal{R}^*$ in Figure 1, the agent can ignore node 2 and still would behave as in the contracting model with common priors. The intuition is that when $H^*(\mathcal{R}^*) \subseteq N$, then the information captured through the variables in $H^*(\mathcal{R}^*)$ already includes the probabilistic information from variables outside $H^*(\mathcal{R}^*)$. Conversely, if the agent's subjective model does not include all variables from $H^*(\mathcal{R}^*)$, she is not behaviorally rational. In this case, we can find a parametrization of $p(x_1, ..., x_n \mid a)$ such that the incentive compatibility constraint is affected by the simplification in the agent's subjective model $\mathcal{R}$.

Next, Proposition 5 also shows that different misspecifications can have the same effect on incentives. Consider the two models $\mathcal{R}_1$ and $\mathcal{R}_2$ from the job example in Figure 3b. The set of nodes on fundamental active paths is the same for these two models, $H^*(\mathcal{R}_1) = H^*(\mathcal{R}_2) = \{0, 2, 4\}$. This implies that the agent's beliefs under these models are identical. Thus, it does not matter for the equilibrium contract whether the agent ignores node 1, node 3, or both nodes. Therefore, the ignorance about one channel of causality may render another channel unimportant. A further interpretation is that two agents with different subjective models may still have the same beliefs about the production function. We capture this result in a general statement. Consider a DAG $\mathcal{R} = (N, R)$ and a subset $\tilde{N} \subset N$. Denote by $\mathcal{R}^{[\tilde{N}]} = (\tilde{N}, \tilde{R})$ with $\tilde{R} = (\tilde{N} \times \tilde{N}) \cap R$ the DAG $\mathcal{R}$ restricted on $\tilde{N}$.

Corollary 3. *Let $\mathcal{R}_1 = (N_1, R_1)$ and $\mathcal{R}_2 = (N_2, R_2)$ be two perfect DAGs. Suppose there* *exists a DAG* $\mathcal{R}_3$ *so that* $\mathcal{R}_3^{[N_1]} = \mathcal{R}_1$ *and* $\mathcal{R}_3^{[N_2]} = \mathcal{R}_2$*. If* $H^*(\mathcal{R}_1) = H^*(\mathcal{R}_2)$*, then we have that* $p_{\mathcal{R}_1}(y \mid a; p(a)) = p_{\mathcal{R}_2}(y \mid a; p(a))$ *for all* $a \in A$ *and* $p(a) \in \Delta(A)$*.*

Identification of fundamental links. We provide an algorithm that identifies $H^*(\mathcal{R}^*)$ in perfect DAGs. Nodes that are connected by fundamental links in perfect DAGs exhibit characteristics that are easy to identify.

Proposition 6 (Fundamental Links)**.** *Let* $\mathcal{R}^*$ *be a perfect DAG and consider two adjacent nodes* $i, j \in N^*$*. The link* iR^*j *is fundamental if and only if at least one of the following conditions is satisfied:*

(a) we have $d(0, i) = d(0, j) - 1$*;*

(b) there exists a node $k \in N^*$ *such that* kEi *and* $k \notin R^*(j)$*.*

From this result we can derive an algorithm that finds all fundamental links in a perfect DAG $\mathcal{R}^*$. Let the topological ordering of $\mathcal{R}^*$ be such that every link is directed from an earlier to a later node. Then find for each node i the distance to the action node, $d(0, i)$. Links between nodes of differing distance are fundamental links. Next, check the links between nodes i, j that are of equal distance to the action node. Let N_d be the nodes that are at distance d to the action node. Consider the smallest element of N_d, say i, and any $j \in N_d$ with iR^*j. A link iR^*j is fundamental if and only if there exists a node k so that there is a fundamental link from k to i, but no link from k to j. Continue in this manner to evaluate all links between nodes in N_d, going sequentially from the smallest to the largest node in N_d. Do this for all distances $d > 0$.

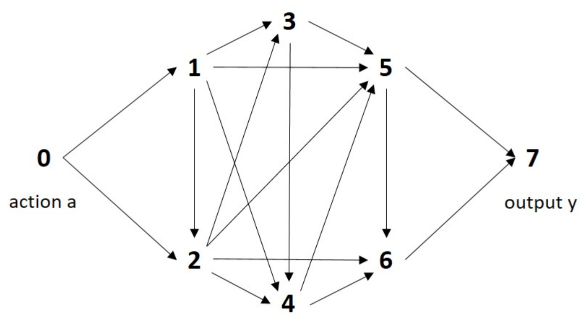

Figure 4: Example model $\mathcal{R}^*$.

It is not always simple to spot the nodes that are not in $H^*(\mathcal{R}^*)$. In this case, Proposition 6 is helpful. Consider the perfect DAG $\mathcal{R}^*$ in Figure 4. Condition (a) from Proposition 6 implies that all links which connect nodes of different distances to the action node are fundamental. The remaining links are $1R^*2$, $3R^*4$, $3R^*5$, $4R^*5$, $4R^*6$, and $5R^*6$. Condition (b) from Proposition 6 then implies that $4R^*6$ and $5R^*6$ are fundamental links, while the remaining links are non-fundamental. We therefore get $H^*(\mathcal{R}^*) = N^* \setminus \{3\}$.

Imperfect objective models. In several applications, the objective model $\mathcal{R}^*$ is imperfect. Nevertheless, we can apply Proposition 5 to these models to detect nodes that can be dropped from the agent's subjective model while preserving behavioral rationality. Note that one can make any imperfect DAG perfect by adding links between nodes that create v-colliders. If p is consistent with $\mathcal{R}^*$, it is consistent with any DAG that adds links to $\mathcal{R}^*$. Consider a perfect DAG $\hat{\mathcal{R}}$ that is identical to the imperfect DAG $\mathcal{R}^*$ except that it has additional links. Suppose all these additional links disappear when we take out the nodes that are not in the agent's subjective model $\mathcal{R} = (N, R)$. Then from Proposition 5 we immediately get that the agent is behaviorally rational if N contains $H^*(\hat{\mathcal{R}})$. We state this result formally.

Corollary 4. *Let $\mathcal{R}^* = (N^*, R^*)$ be the (possibly imperfect) objective DAG and $\mathcal{R} = (N, R)$ the agent's subjective DAG. The agent is behaviorally rational if there is a perfect DAG $\hat{\mathcal{R}} = (N^*, \hat{R})$ with $\hat{\mathcal{R}}^{[N]} = \mathcal{R}$ and $R^* \subseteq \hat{R}$, so that $\mathcal{R}$ contains all nodes from $H^*(\hat{\mathcal{R}})$.*

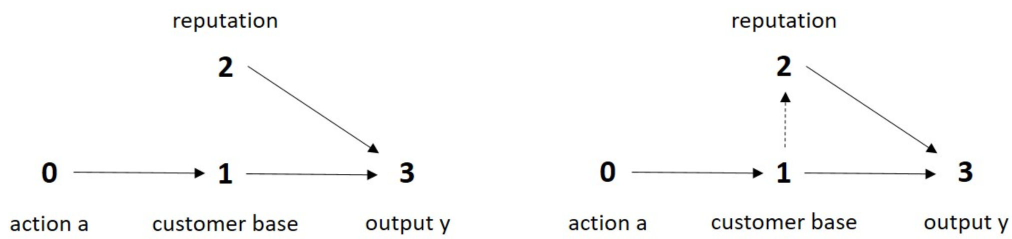

Figure 5: Imperfect model $\mathcal{R}^*$ and perfect model $\hat{\mathcal{R}}$.

As an illustration, consider the marketer example from Subsection 3.2 when the agent's effort has no impact on reputation, $\beta_{02} = 0$. The causal structure of this production function is then given by the imperfect DAG $\mathcal{R}^*$ on the left of Figure 5. The perfect DAG $\hat{\mathcal{R}}$ on the right is identical, except that it has an additional link $1\hat{R}2$. In this model, node 2 is not on a fundamental active path. Hence, the agent is behaviorally rational if her subjective model does not take this node into account.

5 Comparative Statics

The canonical contracting framework produces some clear predictions about the shape of the optimal contract: Holmström's informativeness principle states which variables should be used in an optimal contract; risk-aversion on the side of the agent implies that there should be a negative relationship between risk and incentive provision. However, the empirical evidence on these comparative statics often rejects these predictions. In this section, we consider the informativeness principle (Subsection 5.1) and the risk-incentive trade-off (Subsection 5.2) in our framework, and examine to what extent misspecifications in the agent's subjective model can explain the empirical evidence.[8]

5.1 Observability and Incentives

An important question in contract theory is on which information the principal should condition the agent's wage. Consider a setting with risk-averse agent and unlimited liability. For this setting, the informativeness principle states that the optimal contract conditions on an additional variable z if and only if it is informative about the agent's effort, i.e., if and only if the likelihood ratio $\frac{p(y,z|a')}{p(y,z|a)}$ varies in z for some y. However, relative to this benchmark, many actual compensation contracts seem to use too little information. For example, they do not make use of peer-performance and therefore reward CEOs for windfall gains (Bebchuk and Fried 2004).

The informativeness principle no longer holds when the agent's subjective model $\mathcal{R}$ is misspecified. Suppose that $\mathcal{R}$ is such that the participation constraint is undistorted. Then it does not matter whether z is informative about the agent's effort, but it matters whether the agent believes that variable z contains additional information about her effort. As an illustration, consider the marketer example from Subsection 3.2 and assume that the principal can also condition the agent's wage on the size of the customers set x_1. How does this affect the optimal equilibrium contract? If the agent had rational expectations, the optimal wage scheme would condition both on the customer set x_1 and sales y since neither variable is a sufficient statistic of the other. In contrast, according to model $\mathcal{R}$, sales y are just a noisy signal of the size of the customer set x_1. Recall that the agent overestimates the productivity of her effort when the contract conditions on y. However, x_1 is more informative about the agent's effort than y. The latter aspect turns out to dominate the former. Therefore, the optimal equilibrium contract only conditions on x_1 and appears as "incomplete." This result can be generalized.

Proposition 7 (Observability and Incentives)**.** *Suppose the agent is risk-averse with unlimited liability. Let z and y be two contractible variables. The optimal equilibrium contract does not condition on y if the agent's subjective model $\mathcal{R}$ satisfies the following conditions:*

(a) $p_{\mathcal{R}}(z, y; p(a)) = p(z, y)$ *for all* $p(a) \in \Delta(A)$*, and*

(b) $a \perp_{\mathcal{R}} y \mid z$.

Below we provide the detailed proof. Condition (a) requires that the agent has correct beliefs about the joint equilibrium distribution of y and z. By Proposition 1, it is satisfied if $\mathcal{R}$

[8] Additionally, in Online Appendix A.5, we analyze the effectiveness of team incentives when the size of the team becomes large, and the agent neglects the contributions of others to the output.

is perfect, and y and z form a clique. This condition ensures that the participation constraint is unaffected by the misspecification, so that optimal incentives only depend on likelihood ratios. Condition (b) requests that, according to the agent's model $\mathcal{R}$, the output y is independent of the agent's action a, conditional on outcome z. This condition directly implies that, in the agent's mind, z is a sufficient statistic of y. In the marketer example, this is the case for y and x_1 since $R(y) = \{x_1\}$.

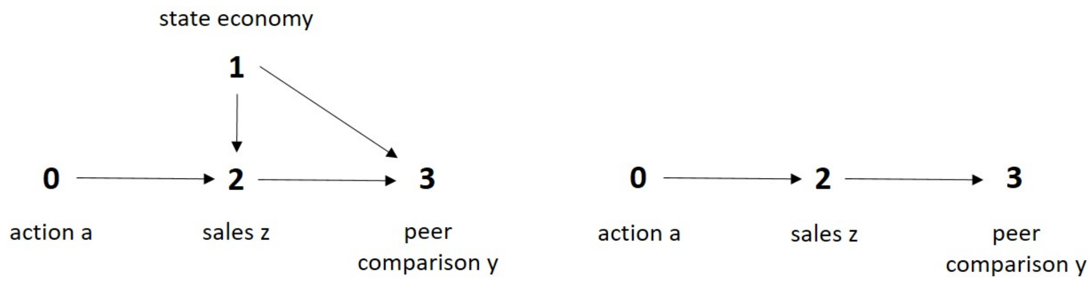

Figure 6: Objective model $\mathcal{R}^*$ (left) and subjective model $\mathcal{R}$ in the peer-comparison example.

As an application, consider the objective model $\mathcal{R}^*$ on the left in Figure 6, and the agent's subjective model $\mathcal{R}$. According to the objective model, the agent's effort directly influences sales z; the state of the economy x_1 also influences sales z, and a "peer-comparison variable" y which captures how good the agent's sales z are relative to those of others. That is, the state of the economy induces a correlation between the agent's and others' sales, which the optimal contract under $\mathcal{R}^*$ would exploit to better tailor the agent's wage to her effort. However, by Proposition 7, the optimal equilibrium contract under $\mathcal{R}$ does not make use of peer-information as the agent does not take into account the confounding factor. Using peer-information in the contract would only increase the risk-premium that the agent requires to accept the contract, and would not have additional incentive effects. The equilibrium contract therefore is incomplete and may reward the agent for windfall gains that come from good states of the economy.[9]

Proof of Proposition 7. Suppose the principal wishes to implement $p(a)$. For any action a in the support of $p(a)$ the optimal incentive scheme is characterized by the first-order condition

$$\frac{1}{u'(w(y,z))} = \frac{p_{\mathcal{R}}(y,z;p(a))}{p(y,z)}\left[\mu + \sum_{a'\in A}\lambda_{a'}\frac{p_{\mathcal{R}}(y,z\mid a;p(a)) - p_{\mathcal{R}}(y,z\mid a';p(a))}{p_{\mathcal{R}}(y,z;p(a))}\right]. \tag{19}$$

[9]Using the results from Section 4, one can further examine which misspecifications lead to incomplete contracts. For example, consider the following extension of this setting: The state of the economy x_1 influences sales z and rival sales x_3. Sales and rival sales determine the peer-comparison variable y; the rest remains the same. We obtain the same result as above if the agent's subjective model $\mathcal{R}$ is identical to the objective model, except that the link between state of economy x_1 and rival sales x_3 is missing (to see this, we only have to apply Corollary 4 to $\mathcal{R}$). Hence, the agent may take into account that her sales are partially determined by the state of the economy. Still, the optimal equilibrium contract is incomplete if she does not take into account that the state of the economy also affects rival sales.

From assumption (a) we get $p_{\mathcal{R}}(y, z; p(a)) = p(y, z)$, and from assumption ($b$) we obtain that for all $a \in A$ we have $p_{\mathcal{R}}(y, z \mid a; p(a)) = p_{\mathcal{R}}(z \mid a; p(a))p(y \mid z)$. Hence, we get

$$p_{\mathcal{R}}(y, z \mid a; p(a)) - p_{\mathcal{R}}(y, z \mid a'; p(a)) = \frac{p_{\mathcal{R}}(y, z; p(a))}{p_{\mathcal{R}}(z; p(a))}[p_{\mathcal{R}}(z \mid a; p(a)) - p_{\mathcal{R}}(z \mid a'; p(a))]. \quad (20)$$

The first-order condition therefore simplifies to

$$\frac{1}{u'(w(y, z))} = \mu + \sum_{a' \in A} \lambda_{a'} \frac{p_{\mathcal{R}}(z \mid a; p(a)) - p_{\mathcal{R}}(z \mid a'; p(a))}{p_{\mathcal{R}}(z; p(a))}. \quad (21)$$

Since the right-hand side of this first-order equation is independent of y, the optimal incentive scheme does not condition on y, which completes the proof. □

5.2 Risk and Incentives

Another important implication of the canonical contracting model is a trade-off between risk and incentives. A risk-averse agent demands a risk premium for accepting a wage schedule with uncertain wage payments. Thus, an increase in risk drives up the costs of providing incentives. Consequently, the provision of effort incentives should decrease in the riskiness of the environment. However, empirically this relationship does not hold in general (e.g., Prendergast 2002). Field evidence on the relationship between risk an incentives for CEO compensation is mixed, and for other domains, such as franchising, a positive relationship can be observed. In contrast, a negative relationship is obtained in lab experiments where subjects know the true production function (e.g., Corgnet and Hernán-González 2019). In this subsection, we continue our marketer example to show how the relationship between risk and incentives may become positive when the agent has a simplified model of the project.

To study the relationship between risk and incentives, the literature typically uses a setting with continuous actions, normally distributed output and exponential utility (so that the optimal contract is linear). To properly apply our framework, we consider a setting with discrete actions and outputs that captures the negative relationship between risk and incentives. Let there be a binary action $a \in \{0, 1\}$ and three equidistant output levels, y_L, y_M, y_H with $y_L < y_M < y_H$. The level of risk is indexed by a variable $\xi \in [0, \bar{\xi}]$. The production function is $p(y_L \mid a) = \beta_L(\xi) - \beta a$, $p(y_M \mid a) = \beta_M(\xi)$, and $p(y_H \mid a) = \beta_H(\xi) + \beta a$, where $\beta_L(\xi) = \beta_H(\xi)$ for all ξ. An increase in risk ξ shifts probability mass from the medium output y_M to the extreme outputs y_L and y_H, i.e., $\beta_L'(\xi) = \beta_H'(\xi) = \varepsilon$ for some $\varepsilon > 0$ and $\beta_M'(\xi) = -2\varepsilon$. The agent has a piecewise linear utility function $u(w) = w$ for $w \geq 0$ and $u(w) = \lambda w$ with $\lambda > 1$ for $w < 0$; her reservation utility is $\bar{U} = 0$.

We now fit the marketer example from Subsection 3.2 to the present setting. The objective causal model is given by $\mathcal{R}^*$ on the left of Figure 2, while the agent's subjective model is given by $\mathcal{R}$ on the right of this figure. We use our usual binary notation, except for the output. The probability of low, middle and high output conditional on x_1 and x_2 is given by

$$p(y_H \mid x_1, x_2) = \beta_3^H(\xi) + \beta_{13}(\xi)x_1 + \beta_{23}(\xi)x_2, \tag{22}$$

$$p(y_M \mid x_1, x_2) = \beta_3^M(\xi), \tag{23}$$

$$p(y_L \mid x_1, x_2) = \beta_3^L(\xi) - \beta_{13}(\xi)x_1 - \beta_{23}(\xi)x_2. \tag{24}$$

We assume that the level of risk ξ changes the importance of the size of the customer set and the firm's reputation for the final output. The larger the risk, the more important are these two factors to obtain a high rather than a small output. We capture this by assuming

$$\beta_{13}(\xi) = \bar{\beta}_{13}\left(1 + \frac{\xi}{\beta_{01}\bar{\beta}_{13}}\right) \text{ and } \beta_{23}(\xi) = \bar{\beta}_{23}\left(1 + \frac{\xi}{\mid \beta_{02} \mid \bar{\beta}_{23}}\right) \tag{25}$$

for two values $\bar{\beta}_{13}, \bar{\beta}_{23} > 0$ with $\beta_{01}\bar{\beta}_{13} + \beta_{02}\bar{\beta}_{23} = \beta$. We choose the functions $\beta_3^H(\xi)$, $\beta_3^M(\xi)$ and $\beta_3^L(\xi)$ so that the objective probability model generates the production function from above.[10]

Proposition 8 (Risk and Incentives)**.** *Consider the marketer example of this subsection.*

(a) Under the objective model $\mathcal{R}^*$ the expected wage that implements $a = 1$ increases in risk ξ. Thus, there exists an interval $[c_L, c_H]$ so that if $c \in (c_L, c_H)$, then for some $\xi^* \in (0, \bar{\xi})$ the optimal equilibrium contract implements $a = 1$ if $\xi < \xi^*$ and $a = 0$ if $\xi > \xi^*$.

(b) Under the subjective model $\mathcal{R}$ the expected wage that implements $a = 1$ decreases in risk ξ if the slope $\beta_L'(\xi) = \beta_H'(\xi) = \varepsilon$ is small enough. In this case, there exists an interval $[c_L, c_H]$ so that if $c \in (c_L, c_H)$, then for some $\xi^* \in (0, \bar{\xi})$ the optimal equilibrium contract implements $a = 0$ if $\xi < \xi^*$ and $a = 1$ if $\xi > \xi^*$.

The proof is in Online Appendix A.4. We explain why part (a) holds. When the agent has rational expectations, the incentive compatibility constraint that ensures high effort equals

$$\beta(u(w_H) - u(w_L)) \geq c. \tag{26}$$

The optimal wage schedule that implements high effort is given by

$$w(y_L) = -\frac{1}{2\lambda\beta}c, \; w(y_M) = 0, \text{ and } w(y_H) = \frac{1}{2\beta}c. \tag{27}$$

[10]Specifically, we derive $\beta_3^H(\xi)$ and $\beta_3^L(\xi)$ from $\beta_H(\xi) = \beta_3^H(\xi) + \beta_1\beta_{13}(\xi) + \beta_2\beta_{23}(\xi)$ and $\beta_L(\xi) = \beta_3^L(\xi) - \beta_1\beta_{13}(\xi) - \beta_2\beta_{23}(\xi)$. Since $\beta_H(\xi) = \beta_L(\xi)$ for all ξ, we have $\beta_3^M(\xi) = 1 - 2[\beta_3^H(\xi) + \beta_1\beta_{13}(\xi) + \beta_2\beta_{23}(\xi)]$.

Note that a change in risk ξ affects neither the optimal wage schedule, nor the incentive compatibility constraint in (26). In terms of incentives, the effect of risk on the importance of the customer set and reputation cancel each other out. However, an increase in risk exposes the agent to more variation in her wage, so that she requires a higher risk-premium. Hence, when the principal implements high effort, his expected payment to the agent increases in risk. Thus, for an interval of cost levels $[c_L, c_H]$, if $c \in (c_L, c_H)$, the optimal equilibrium contract implements high effort if and only if the level of risk is sufficiently small, so that we obtain a negative relationship between risk and incentives.

Next, consider part (b). If the agent does not take the reputation channel into account, an increase in risk appears to her as an increase in the productivity of her effort, as the association between customer set and sales becomes stronger. The incentive compatibility constraint that ensures high effort now equals

$$\beta_{01}\beta_{13}(\xi)(u(w_H) - u(w_L)) \geq c. \tag{28}$$

Recall that $\beta_{13}(\xi)$ increases in ξ, so an increase in risk ξ relaxes this IC. The optimal wage schedule that implements $\alpha = 1$ is given by

$$w(y_L) = -\frac{\beta_H(\xi) + \beta - \beta_{01}\beta_{13}(\xi)}{\lambda(\beta_H(\xi) + \beta_L(\xi))\beta_{01}\beta_{13}(\xi)}c, \; w(y_M) = 0, \text{ and } w(y_H) = \frac{\beta_L(\xi) - \beta + \beta_{01}\beta_{13}(\xi)}{(\beta_H(\xi) + \beta_L(\xi))\beta_{01}\beta_{13}(\xi)}c. \tag{29}$$

A change in risk now has two countervailing effects on the expected payment when the principal implements high effort. It again increases the risk premium that the agent requires, but it also relaxes the incentive compatibility constraint. Which effect dominates depends on the probability model and the utility function. If the slope $\beta'_L(\xi) = \beta'_H(\xi) = \varepsilon$ is small enough, an increase in risk reduces the expected payment to the agent at all $\xi \in [0, \bar{\xi}]$. We the obtain a positive relationship between risk and incentives: For an interval of cost levels $[c_L, c_H]$, if $c \in (c_L, c_H)$, the optimal equilibrium contract implements high effort if the level of risk is sufficiently large, and low effort through a fixed wage otherwise.

6 Conclusion

In this paper, we applied Spiegler's (2016) Bayesian network framework to analyze optimal contracting when the agent forms beliefs about the production function based on a misspecified model of the principal's project. The objective causal model may be very complex. It potentially contains empirical regularities that the agent does not consider due to cognitive limitations or because they are never brought to the agent's attention. An important feature of misspecifications is that they can bias beliefs while also being sustainable, in the sense that the agent cannot detect them from the data she gathers on the equilibrium path. The principal benefits from misspecifications if the agent does not take into account a partial negative effect of her effort on the output. The agent then overestimates the productivity of her effort, which in turn relaxes incentive compatibility. Whether such simplifications exist, depends on the agent's role in the organization, that is, through which channels her effort influences outcomes. For example, a front-line worker may not fully understand the workings of the organization around her, but still act as if she were fully rational. In contrast, a high-ranking manager, who affects the output by influencing the behavior of many subordinates, overestimates her own productivity if she does not take into account the challenges that her subordinates face in their routines. Thus, when the agent has a simplified model of the project, being control optimistic can be a consequence of the agent's job (and not necessarily of her personality).

The Bayesian network approach allows us to systematically study which variables matter for the agent's beliefs in a particular setting. Moreover, it is portable so that we can examine how the optimal equilibrium contract varies in features of the contracting environment. As an example, we revisited the informativeness principle and the risk-incentive trade-off. When the agent has a misspecified model of the principal's project, the predictions of the canonical contracting model about the comparative statics in these domains may be violated: The agent may interpret an informative variable as uninformative, in which case the optimal equilibrium contract is incomplete. She may also interpret an increase in risk as an increase in the productivity of her effort, in which case there can be a positive relationship between risk and incentives.

Throughout the paper, we focused on a simple contracting framework so that we can identify precisely how misspecifications in the agent's model affect incentive contracts. Future research can extend the framework by considering more advanced topics in organizational economics, such as team incentives, relational contracts, strategic communication and delegation. The Bayesian network approach offers a very disciplined tool to study the effects of bounded rationality on organizations, and we think that our results are useful in this respect.

References

AGHION, PHILIPPE, AND JEAN TIROLE (1997): "Formal and Real Authority in Organizations," *Journal of Political Economy*, 105(1), 1–29.

AUSTER, SARAH (2013): "Asymmetric awareness and moral hazard," *Games and Economic Behavior*, 82, 503–521.

BEBCHUK, LUCIAN ARYE, AND JESSE FRIED (2004): *Pay Without Performance: The Unfulfilled Promise of Executive Compensation*, Harvard University Press, Cambridge.

BÉNABOU, ROLAND (2013): "Groupthink: Collective Delusions in Organizations and Markets," *Review of Economic Studies*, 80(2), 429–462.

BÉNABOU, ROLAND, AND JEAN TIROLE (2002): "Self-Confidence and Personal Motivation," *Quarterly Journal of Economics*, 117(23), 871–915.

BRUNNERMEIER, MARKUS, AND JONATHAN PARKER (2005): "Optimal Expectations," *American Economic Review*, 95(14), 1092–1118.

CHAIGNEAU, PIERRE, ALEX EDMANS, AND DANIEL GOTTLIEB (2019): "The informativeness principle without the first-order approach," *Games and Economic Behavior*, 113, 743–755.

CORGNET, BRICE, AND ROBERTO HERNÁN-GONZÁLEZ (2019): "Revisiting the Trade-off Between Risk and Incentives: The Shocking Effect of Random Shocks?," *Management Science*, 65(3), 1096–1114.

DE LA ROSA, ENRIQUE (2011): "Overconfidence and Moral Hazard," *Games and Economic Behavior*, 73(2), 429–451.

DEKEL, EDDIE, BARTON LIPMAN, AND ALDO RUSTICHINI (1998): "Standard State-Space Models Preclude Unawareness," *Econometrica*, 66(1), 159–173.

ELIAZ, KFIR, AND RANI SPIEGLER (2018): "A Model of Competing Narratives," CEPR Discussion Paper No. DP13319.

ELIAZ, KFIR, RANI SPIEGLER, AND HEIDI THYSEN (2018): "Strategic Interpretations," Working Paper.

FANG, HANMING, AND GIUSEPPE MOSCARINI (2005): "Morale Hazard," *Journal of Monetary Economics*, 52(4), 749–777.

GERVAIS, SIMON, AND ITAY GOLDSTEIN (2007): "The Positive Effects of Biased Self-Perceptions in Firms," *Review of Finance*, 11(3), 453–496.

GROSSMAN, SANFORD, AND OLIVER HART (1983): "An Analysis of the Principal-Agent Problem," *Econometrica*, 51(1), 7–45.

HEIFETZ, AVIAD, MARTIN MEIER, AND BURKHARD SCHIPPER (2006): "Interactive unawareness," *Journal of Economic Theory*, 130(1), 78–94.

HEIFETZ, AVIAD, MARTIN MEIER, AND BURKHARD SCHIPPER (2013): "Unawareness, beliefs, and speculative trade," *Games and Economic Behavior*, 77(1), 100–121.

HOLMSTRÖM, BENGT (1979): "Moral Hazard and Observability," *Bell Journal of Economics*, 10(1), 74–91.

IMMORDINO, GIOVANNI, ANNA MARIA MENICHINI, MARIA GRAZIA ROMANO (2015): "Contracts with Wishful Thinkers," *Journal of Economics and Management Strategy*, 24(4), 863–886.

KOSKI, TIMO, AND JOHN NOBLE (2009): *Bayesian Networks: An Introduction*, Wiley Series in Probability, Wiley.

KŐSZEGI, BOTOND (2006): "Ego Utility, Overconfidence, and Task Choice," *Journal of the European Economic Association*, 4(4), 673– 707.

KŐSZEGI, BOTOND (2014): "Behavioral Contract Theory," *Journal of Economic Literature*, 52(4), 1075–1118.

PEARL, JUDEA (2009): *Causality: Models, Reasoning, and Inference*, Cambridge University Press.

PORTER, MICHAEL, JAY LORSCH, AND NITIN NOHRIA (2004): "Seven surprises for new CEOs," *Harvard Business Review*, 82(10), 62–72.

PORTER, MICHAEL, AND NITIN NOHRIA (2018): "How CEOs manage time," *Harvard Business Review*, 96(4), 42–51.

PRENDERGAST, CANICE (2002): "The Tenuous Trade-Off between Risk and Incentives," *Journal of Political Economy*, 110(5), 1071–1102.

SANTOS-PINTO, LUÍS (2008): "Positive Self-Image and Incentives in Organisations," *Economic Journal*, 118(531), 1315–1332.

SAUTMANN, ANJA (2007): "Self-Confidence in a Principal-Agent Relationship," unpublished manuscript, Brown University.

SAUTMANN, ANJA (2013): "Contracts for Agents with Biased Beliefs: Some Theory and an Experiment," *American Economic Journal: Microeconomics*, 5(3), 124–156.

SCHENONE, PABLO (2013): "Causality: A Decision Theoretic Framework," Working Paper, California Institute of Technology.

SIMON, HERBERT (1947): *Administrative Behavior*, Macmillan, London.

SIMON, HERBERT (1955): "A Behavioral Model of Rational Choice," *Quarterly Journal of Economics*, 69(1), 99–118.

SHRIER, IAN, AND ROBERT PLATT (2008): "Reducing bias through directed acyclic graphs," *BMC Medical Research Methodology*, 8(70).

SPIEGLER, RAN (2016): "Bayesian Networks and Boundedly Rational Expectations," *Quarterly Journal of Economics*, 131(3), 1243–1290.

SPIEGLER, RAN (2017): "Data Monkeys: A Procedural Model of Extrapolation from Partial Statistics," *Review of Economic Studies*, 84(4), 1818–1841.

SPIEGLER, RAN (2019): "Can Agents with Causal Misperceptions be Systematically Fooled?," *Journal of the European Economic Association*, forthcoming.

SPINNEWIJN, JOHANNES (2013): "Insurance and Perceptions: How to Screen Optimists and Pessimists," *Economic Journal*, 123(569), 606–633.

SPINNEWIJN, JOHANNES (2015): "Unemployed but Optimistic: Optimal Insurance Design with Biased Beliefs," *Journal of the European Economic Association*, 13(1), 130–167.

VON THADDEN, ERNST-LUDWIG, AND XIAOJIAN ZHAO (2012): "Incentives for Unaware Agents," *Review of Economic Studies*, 79(3), 1151–1174.

VON THADDEN, ERNST-LUDWIG, AND XIAOJIAN ZHAO (2014): "Multitask agency with unawareness," *Theory and Decision*, 77(2), 197–222.

VAN DEN STEEN, ERIC (2005): "Organizational Beliefs and Managerial Vision," *Journal of Law, Economics, and Organization*, 21(1), 256–283.

VERMA, THOMAS, AND JUDEA PEARL (1991): "Equivalence and Synthesis of Causal Models," *Uncertainty in Artificial Intelligence*, 6, 255–268.

A Online Appendix

A.1 Existence of a Personal Equilibrium

We show that a personal equilibrium exists at any admissible $\mathcal{R}$ and $w(y) \in W$. Note that $\Delta(A)$ is non-empty, compact, and convex. Define the best-response correspondence $BR : \Delta(A) \to \Delta(A)$ by

$$BR(p(a)) = \arg \max_{p'(a') \in \Delta(A)} \sum_{a' \in A} \sum_{y \in Y} p'(a')[p_{\mathcal{R}}(y \mid a'; p(a))u(w(y)) - c(a)]. \tag{30}$$

For every $p(a) \in \Delta(A)$ we have that $BR(p(a))$ is non-empty and convex. The latter statement follows since any convex combination of pure actions that are optimal for the agent is an element of $BR(p(a))$. Definition 1 and the factorization formula in (2) imply that the agent's beliefs $p_{\mathcal{R}}(y \mid a'; p(a))$ are continuous in $p(a)$. Therefore, we also must have that $\sum_{a' \in A} \sum_{y \in Y} p'(a')[p_{\mathcal{R}}(y \mid a'; p(a))u(w(y)) - c(a)]$ is continuous in $p(a)$. Hence, $BR(p(a))$ is upper hemi-continuous. The existence of a personal equilibrium then follows from Kakutani's theorem.

A.2 Omitted Proofs from Subsection 4.1

We first derive the IC under the objective model $\mathcal{R}^*$. The probability of high output after high and low effort, respectively, are given by

$$\begin{aligned} p(y_H \mid a = 1) &= \beta_4 + [\beta_2 + \beta_{02} + (\beta_1 + \beta_{01})\beta_{12}]\beta_{24} \\ &\quad + [\beta_3 + (\beta_1 + \beta_{01})\beta_{13} + (\beta_2 + \beta_{02} + (\beta_1 + \beta_{01})\beta_{12})\beta_{23}]\beta_{34}, \quad (31) \\ p(y_H \mid a = 0) &= \beta_4 + [\beta_2 + \beta_1\beta_{12}]\beta_{24} + [\beta_3 + \beta_1\beta_{13} + (\beta_2 + \beta_1\beta_{12})\beta_{23}]\beta_{34}, \quad (32) \end{aligned}$$

so that the effect of effort on the probability of high output equals

$$p(y_H \mid a = 1) - p(y_H \mid a = 0) = (\beta_{02} + \beta_{01}\beta_{12})(\beta_{24} + \beta_{23}\beta_{34}) + \beta_{01}\beta_{13}\beta_{34}. \tag{33}$$

Next, we drive the IC under the subjective model $\mathcal{R}$ when the equilibrium action is $\alpha \in [0, 1]$. We calculate

$$\begin{aligned} p(x_1 = 1 \mid x_2 = 1; \alpha) &= \frac{\alpha(\beta_1 + \beta_{01})(\beta_2 + \beta_{02} + \beta_{12}) + (1 - \alpha)\beta_1(\beta_2 + \beta_{12})}{\beta_2 + \beta_1\beta_{12} + \alpha(\beta_{02} + \beta_{01}\beta_{12})}, \quad (34) \\ p(x_1 = 1 \mid x_2 = 0; \alpha) &= \frac{\alpha(\beta_1 + \beta_{01})(1 - \beta_2 - \beta_{02} - \beta_{12}) + (1 - \alpha)\beta_1(1 - \beta_2 - \beta_{12})}{1 - \beta_2 - \beta_1\beta_{12} - \alpha(\beta_{02} + \beta_{01}\beta_{12})}, \quad (35) \end{aligned}$$

and

$$p(x_3 = 1 \mid x_2 = 1; \alpha) = \beta_3 + p(x_1 = 1 \mid x_2 = 1; \alpha)\beta_{13} + \beta_{23}, \tag{36}$$

$$p(x_3 = 1 \mid x_2 = 0; \alpha) = \beta_3 + p(x_1 = 1 \mid x_2 = 0; \alpha)\beta_{13}. \tag{37}$$

The agent's belief about the probability of high output after $x_2 = 1$ and $x_2 = 0$, respectively, is therefore given by

$$p_{\mathcal{R}}(y_H \mid x_2 = 1; \alpha) = \beta_4 + \beta_{24} + [\beta_3 + p(x_1 = 1 \mid x_2 = 1; \alpha)\beta_{13} + \beta_{23}]\beta_{34}, \tag{38}$$

$$p_{\mathcal{R}}(y_H \mid x_2 = 0; \alpha) = \beta_4 + [\beta_3 + p(x_1 = 1 \mid x_2 = 0; \alpha)\beta_{13}]\beta_{34}. \tag{39}$$

Since $p_{\mathcal{R}}(x_2 \mid a; \alpha) = p(x_2 \mid a)$ for all $a \in \{0, 1\}$ and $\alpha \in [0, 1]$, the agent's belief about the effect of effort on the probability of high output under $\mathcal{R}$ equals

$$\begin{aligned} p_{\mathcal{R}}(y_H \mid a = 1; \alpha) - p_{\mathcal{R}}(y_H \mid a = 0; \alpha) = {} & (\beta_{02} + \beta_{01}\beta_{12})(\beta_{24} + \beta_{23}\beta_{34}) + (\beta_{02} + \beta_{01}\beta_{12})\beta_{13}\beta_{34} \\ & \times [p(x_1 = 1 \mid x_2 = 1; \alpha) - p(x_1 = 1 \mid x_2 = 0; \alpha)]. \end{aligned} \tag{40}$$

Recall that $\beta_{13} < 0$. By comparing (33) and (40) we get that at $\alpha = 1$ the misspecification in $\mathcal{R}$ relaxes the IC if and only if

$$\beta_{01} > \frac{\beta_{12}(\beta_1 + \beta_{01})(1 - \beta_1 - \beta_{01})(\beta_{02} + \beta_{01}\beta_{12})}{(1 - \beta_2 - \beta_{02} - \beta_{12}(\beta_1 + \beta_{01}))(\beta_2 + \beta_{02} + \beta_{12}(\beta_1 + \beta_{01}))}, \tag{41}$$

which implies the statement in the main text.

Proof of Proposition 3. We prove the statement in (a). Since $\beta_1 \in \{0, 1\}$, we can rewrite the probability model without variable 1. The corresponding objective model $\tilde{\mathcal{R}}^*$ equals $\mathcal{R}^*$ in Figure 3a without node 1. We now apply Propositions 5 and 6. In model $\tilde{\mathcal{R}}^*$, node 3 is not on a fundamental active path. Hence, the agent with subjective model $\mathcal{R}$ is behaviorally rational which yields the result. We prove the statements in (b). The first statement is shown in the text. The second statement follows from Corollary 3. Note that in all models the set of nodes on fundamental active paths is identical. □

A.3 Omitted Proofs from Section 4.2

We first prove Proposition 6 and then Proposition 5. To this end, we prove several intermediate results. We first note that in a perfect DAG $\mathcal{R}^*$ the link $iR^* j$ is fundamental if the nodes i and j

differ in their distance to the action node 0.

Lemma 1. *Let $i, j \in N^*$ be adjacent nodes in $\mathcal{R}^*$. If $d(0,i) = d(0,j) - 1$, then iEj.*

Proof. First, suppose $d(0,i) = 0$ so that $i = 0$. Since node 0 is ancestral, we must have iGj in every DAG $\mathcal{G} \in \mathcal{E}$. Next, suppose $d(0,i) = d > 0$. Since $\mathcal{R}^*$ is perfect and node 0 is ancestral, there exists an active path of length d from node 0 to node i. Denote by k the direct ancestor of i on this path. There cannot exist a link between k and j, otherwise we would have $d(0,i) = d(0,k)$, a contradiction. Thus, we must have iGk in every DAG $\mathcal{G} \in \mathcal{E}$, otherwise we would have a v-collider at node i. □

Lemma 2. *Let $i, j \in N^*$ and iR^*j. If there exists a node $k \in N^*$ such that kEi and $k \notin R^*(j)$, then iEj.*

Proof. If there is a fundamental link from node k to node i, then iR^*j implies that we cannot have jR^*k. Otherwise, we would have a directed cycle. Node j and node k are therefore not adjacent. Hence, if jGi in some DAG $\mathcal{G} \in \mathcal{E}$, there would be a v-collider at i, a contradiction. □

The "if"-statement of Proposition 6 follows directly from Lemma 1 and Lemma 2. For the "only if"-statement we need two more results. The first one provides a condition under which a link is not fundamental.

Lemma 3. *Let $i, j \in N^* \setminus \{0\}$ and iR^*j. If $R^*(i) \subset R^*(j)$, then the link between i and j is not fundamental.*

Proof. Consider the DAG $\mathcal{G} = (G, N^*)$ that is identical to $\mathcal{R}^*$ except that it reverses the link between i and j. The assumption $R^*(i) \subset R^*(j)$ rules out that there are v-colliders in $\mathcal{G}$. Assume that there is a cycle in $\mathcal{G}$. Since $\mathcal{R}^*$ is acyclic, the cycle must contain jGi. Further, there must exists a node k and a link kGj which is part of the cycle. Since $\mathcal{R}^*$ is perfect, we must have $k\tilde{R}^*i$. Assume first that we have kR^*i. Then jGi implies that kGi is not part of the cycle. Thus, there must exist an active path τ of some length d so that $\tau_0 = i$ and $\tau_d = k$. But then there is a cycle consisting of the link kGi and τ. This cycle also exists in $\mathcal{R}^*$, a contradiction. Next, assume that we have iR^*k. Since $i \neq 0$ and $R^*(i) \subset R^*(j)$, there exists a node l with lR^*i and lR^*j. Since $\mathcal{R}^*$ is perfect, we also must have $l\tilde{R}^*k$. The same applies to all $l' \in R^*(i)$. Hence, starting from $\mathcal{R}^*$, we can reverse the links between i and j as well as between i and k and obtain a DAG $\mathcal{G}' \in \mathcal{E}$. □

The second result for the proof of the "only if"-statement of Proposition 6 demonstrates that for each node i in a perfect DAG $\mathcal{R}^*$ there exists a DAG $\mathcal{G} \in \mathcal{E}$ in which there is no non-fundamental link that points to i.

Lemma 4. *For all nodes $i \in N^*$ there exists a DAG $\mathcal{G} \in \mathcal{E}$ in which all non-fundamental links adjacent to node i point away from i.*

Proof. Let N_d be the set of nodes that have distance $d > 0$ to the action node 0. Denote by $N_d^{[\kappa]}$, $\kappa = 1, 2, ...$, the maximal subset of nodes that (i) are at distance $d > 0$ from the action node 0 and (ii) are connected through non-fundamental links (i.e., for any two nodes $i, j \in N_d^{[\kappa]}$ there exists a path between i and j consisting of non-fundamental links). **Step 1.** We show that all nodes in a given set $N_d^{[\kappa]}$ have the same parents outside of $N_d^{[\kappa]}$. Consider two nodes $i, j \in N_d^{[\kappa]}$ that are connected through the non-fundamental link iR^*j. By definition kEi for each $k \in R^*(i) \setminus N_d^{[\kappa]}$ for each $i \in N_d^{[\kappa]}$. Since $\mathcal{R}^*$ is perfect, this implies that $R^*(j) \setminus N_d^{[\kappa]} \subset R^*(i) \setminus N_d^{[\kappa]}$. Since iR^*j is non-fundamental, we also must have $R^*(i) \setminus N_d^{[\kappa]} \subset R^*(j) \setminus N_d^{[\kappa]}$ so that $R^*(i) \setminus N_d^{[\kappa]} = R^*(j) \setminus N_d^{[\kappa]}$. The result follows from the fact that, by assumption, all nodes in $N_d^{[\kappa]}$ are connected through non-fundamental links. **Step 2.** Consider two links $i \in N_d^{[\kappa]}$ and $i' \in N_d^{[\kappa']}$ with $\kappa \neq \kappa'$ that are adjacent. Assume w.l.o.g. that iR^*i'. By definition, iR^*i' is a fundamental link. Step 1 then implies that iEj' for all $j' \in N_d^{[\kappa']}$. Thus, there cannot exist nodes $j \in N_d^{[\kappa]}$ and $j' \in N_d^{[\kappa']}$ so that $j'R^*j$. Otherwise, we would have $j'Ej$ and $j'Ei$ for all $i \in N_d^{[\kappa]}$, a contradiction. Thus, there cannot exist nodes $i, j \in N_d^{[\kappa]}$ and $i', j' \in N_d^{[\kappa']}$ such that iR^*i' and $j'R^*j$. **Step 3.** Note that, since $\mathcal{R}^*$ is perfect, by Lemma 1 all links between N_d and N_{d+1} point away from the nodes in N_d. **Step 4.** We now can prove Lemma 4. Take any node $i \in N^*$ and assume w.l.o.g. that $i \in N_d^{[\kappa]}$. Consider the DAG $\mathcal{G}^{[\kappa]} = (N_d^{[\kappa]}, G^{[\kappa]})$ where $G^{[\kappa]}$ is identical to R^* restricted on $N_d^{[\kappa]}$. Since $\mathcal{R}^*$ is perfect, $\mathcal{G}^{[\kappa]}$ also must be perfect. Corollary 1 from Spiegler (2019) implies that there exists a DAG $\mathcal{Q}^{[\kappa]}$ in which node i is ancestral and that is equivalent to $\mathcal{G}^{[\kappa]}$. Choose such a $\mathcal{Q}^{[\kappa]}$ and replace $\mathcal{G}^{[\kappa]}$ in the original DAG $\mathcal{R}^*$ by $\mathcal{Q}^{[\kappa]}$. Call the resulting DAG $\mathcal{Q}^*$. Step 1 implies that there are no v-colliders in $\mathcal{Q}^*$, and Step 2 and 3 imply that there are no cycles in $\mathcal{Q}^*$, which proves the result. □

Proof of Proposition 6. The "if"-statement follows from Lemma 1 and Lemma 2. We prove the "only if"-statement. Consider any two adjacent nodes $i, j \in N^*$ with $iR^*j$and $d(0, i) = d(0, j)$. Suppose that for any node $k \in R^*(i)$ with a fundamental link kR^*i we also have $k \in R^*(j)$. By Lemma 4, we can find a DAG $\mathcal{G} \in \mathcal{E}$ in which all non-fundamental links are turned away from node i. In this DAG, we have $G(i) \subset G(j)$. From Lemma 3 it then follows that the link iR^*j is not fundamental. This completes the proof. □

Before we can prove Proposition 5, we need two more results. We will use the following definitions. A path τ of length d is directed if for any $h \in \{1, ..., d\}$ we have $\tau_{h-1}R\tau_h$ on this path. For any DAG, the topological ordering is a sequence of nodes such that every link is directed from an earlier to a later node in the sequence.

Lemma 5. *Let $M \subset N^* \setminus H^*$ be a set of nodes connected through non-fundamental links. Suppose there are two nodes $i, j \in H^*$ with non-fundamental links to nodes in M. Then i and j are adjacent.*

Proof. Assume w.l.o.g. that i, j are on a fundamental active path between 0 and n (the argument for $n+1$ is identical). As in the proof of Lemma 4, let N_d be the set of nodes that have distance $d > 0$ to the action node 0. Let $E(i)$ be the set of nodes k with kEi. By Lemma 1, there is a $d > 0$ so that $i, j \in N_d$ and $M \subset N_d$. By Lemma 2, we must have $E(i) = E(j)$ since these nodes are connected through non-fundamental links. Choose any node $k \in N_{d-1}$ with $k \in H^*$ and kR^*i. By Lemma 2, we also have kR^*j. We can now choose two fundamental active paths $\tau^{[i]}, \tau^{[j]}$ from node 0 to node n so that (i) $k \in \tau^{[i]}$ and $k \in \tau^{[j]}$, (ii) $i \in \tau^{[i]}$ and $j \in \tau^{[j]}$, (iii) all nodes on $\tau^{[i]}$ and $\tau^{[j]}$ before k are identical, and (iv) there is not any node on $\tau^{[i]}$ ($\tau^{[j]}$) between k and i (k and j). Since $i, j \in H^*$ this is possible. Now define by $m_1^{[i]}$ ($m_1^{[j]}$) the last node on $\tau^{[i]}$ ($\tau^{[j]}$) before node n; by $m_2^{[i]}$ ($m_2^{[j]}$) the penultimate node on $\tau^{[i]}$ ($\tau^{[j]}$) before node n, and so forth. Since $\mathcal{R}^*$ is perfect, $m_1^{[i]}$ and $m_1^{[j]}$ must be adjacent. Since $m_1^{[i]}$ and $m_1^{[j]}$ are adjacent and $\mathcal{R}^*$ is perfect, $m_2^{[i]}$ and $m_2^{[j]}$ must be adjacent, and so forth. If nodes i and j are both the t'th node from n in $\tau^{[i]}$ ($\tau^{[j]}$), we are done. Assume that this is not the case, and that w.l.o.g. node i is the t'th node from n while node j is the t''th node from n, with $t' > t$. Then i is adjacent to $m_t^{[j]}$, and also to all nodes on $\tau^{[j]}$ between $m_t^{[j]}$ and j (including j) through non-fundamental links, otherwise there would be a contradiction to $E(i) = E(j)$. □

The next result is crucial for the proof of Proposition 5. It shows that all nodes that are not on a fundamental active path between action and outcome nodes can be made "unimportant" in the sense that they have no impact on outcomes. Formally, this means that we can find a DAG in $\mathcal{E}$ in which all links between one node in H^* and one node in $N^* \setminus H^*$ point towards the node in $N^* \setminus H^*$.

Lemma 6. *There exists a DAG $\mathcal{G}^* \in \mathcal{E}$ such that in $\mathcal{G}^*$ all links with one end in H^* and the other in $N^* \setminus H^*$ point from H^* to $N^* \setminus H^*$.*

Proof. The proof proceeds by steps. **Step 1.** Consider any maximal set $M \subset N^* \setminus H^*$ of nodes connected through non-fundamental links and let $M^+ \subset H^*$ be the set of nodes that have non-fundamental links to nodes in M. By Lemma 1, there is a $d > 0$ so that $M, M^+ \subset N_d$. Denote by M^{++} the set of nodes in $N_d \cap H^*$ with fundamental links into M. Since the nodes in M are connected through non-fundamental links, there is a fundamental link from any node

$i \in M^{++}$ to any node in M. Thus, any node in M^{++} must also be adjacent to any node in M^{+}, so $M^{+} \cup M^{++}$ is a clique. **Step 2.** Consider the DAG $\bar{\mathcal{G}} = (N, \bar{G})$, where $N = M \cup M^{+} \cup M^{++}$ and $\bar{G}$ is identical to R^{*} restricted on N. By construction, this DAG is perfect. Hence, Corollary 1 from Spiegler (2019) implies that there exists a DAG $\bar{\mathcal{G}}^{+}$ in which the clique $M^{+} \cup M^{++}$ is ancestral and that is equivalent to $\bar{\mathcal{G}}$. We choose such a $\bar{\mathcal{G}}^{+}$ with the property that the ordering of the nodes in $M^{+} \cup M^{++}$ is the same as in $\bar{\mathcal{G}}$ (this is possible since $M^{+} \cup M^{++}$ is a clique, and all links between nodes $M^{+} \cup M^{++}$ and nodes in M point towards the latter one). Consider now the DAG $\mathcal{G}$ that is identical to $\mathcal{R}^{*}$ except that $\bar{\mathcal{G}}$ is replaced by $\bar{\mathcal{G}}^{+}$. We show that there are no cycles or v-colliders in $\mathcal{G}$ so that it is equivalent to $\mathcal{R}^{*}$. Consider any node $i \in N_{d-1} \cup N_d$ that is outside $M \cup M^{+} \cup M^{++}$ and that has a fundamental link into a node in M. Since the nodes in M are connected through non-fundamental links, node i has a fundamental link into every node in M (otherwise, i would belong to M, a contradiction). This rules out v-colliders. Any link between a node in N_d and a node in N_{d+1} points into the latter one. Hence, by construction, there cannot be cycles or v-colliders in $\mathcal{G}$. We obtain $\mathcal{G}^{*}$ by performing the same changes for any maximal set $M \subset N^{*} \setminus H^{*}$ of nodes connected by non-fundamental links in $\mathcal{R}^{*}$. □

Proof of Proposition 5. First, we show the "if"-statement. Assume that the agent's subjective DAG $\mathcal{R}$ is aware of all the nodes in H^{*}. Consider the DAG $\mathcal{G}^{*} \in \mathcal{E}$ in which all links with one end in H^{*} and the other in $N^{*} \setminus H^{*}$ point from H^{*} to $N^{*} \setminus H^{*}$. By Lemma 6, this DAG exists. From Proposition 4 it follows that $p_{\mathcal{G}^{*}}(x_{H^{*}}) = p(x_{H^{*}})$ for all distributions $p \in \Delta(X)$. Consider the subgraph $\mathcal{G} = (G, N)$ where G equals G^{*} restricted on N. Since none of the nodes in $N \setminus H^{*}$ impacts on any node in H^{*}, we have $p_{\mathcal{G}}(x_{H^{*}}) = p_{\mathcal{G}^{*}}(x_{H^{*}})$ for all $p \in \Delta(X)$. By construction, the DAGs $\mathcal{R}$ and $\mathcal{G}$ are equivalent so that we have $p_{\mathcal{R}}(x_{H^{*}}) = p_{\mathcal{G}}(x_{H^{*}}) = p_{\mathcal{G}^{*}}(x_{H^{*}}) = p(x_{H^{*}})$ for all distributions $p \in \Delta(X)$, which proves the "if"-statement. Next, we show the "only if"-statement. Assume that there is one node $i \in H^{*}$ that is not in the agent's subjective model. Assume w.l.o.g. that i is on a fundamental active path τ between the action node 0 and the output node n. We find a probability distribution $p \in \Delta(X)$ so that $p_{\mathcal{R}}(x_n \mid x_0) \neq p(x_n \mid x_0)$. Let k be the k'th node in τ. Consider a probability distribution with the following properties: $p(x_j \mid x_{R^{*}(j)}) = p(x_j)$ for all nodes $j \notin \tau$ that are between the nodes 0 and n, and $p(x_k \mid x_{R^{*}(k)}) = p(x_k \mid x_{k-1})$. Clearly, such a distribution can have the desired property. □

Proof of Corollary 3. Denote $H^{*}(\mathcal{R}_1) = H^{*}(\mathcal{R}_2) = H$. By Proposition 5 there is a DAG $\mathcal{R}_1^{[1]}$ that is equivalent to $\mathcal{R}_1$ and in which all links between a node $i \in H$ and node $j \in N_1 \setminus H$ is turned away from i. Thus, we have

$$p_{\mathcal{R}_1}(x_H) = \sum_{x_{N_1 \setminus H} \in X_{N_1 \setminus H}} p_{\mathcal{R}_1}(x_{N_1}) = \sum_{x_{N_1 \setminus H} \in X_{N_1 \setminus H}} p_{\mathcal{R}_1^{[1]}}(x_{N_1}) = p_{\mathcal{R}_1^{[1]}}(x_H). \tag{42}$$

Note that foll all $i \in H$ we have that $R_1^{[1]}(i) \subset H$. Consider the restriction of $R_1^{[1]}$ on H, $R_1^{[H]}$. We then have

$$p_{\mathcal{R}_1^{[1]}}(x_H) = \prod_{i \in H} p(x_i \mid x_{R_1^{[1]}(i)}) = \prod_{i \in H} p(x_i \mid x_{R_1^{[H]}(i)}) = p_{\mathcal{R}_1^{[H]}}(x_H). \tag{43}$$

Define $\mathcal{R}_2^{[1]}$ and $\mathcal{R}_2^{[H]}$ just like $\mathcal{R}_1^{[1]}$ and $\mathcal{R}_1^{[H]}$. By assumption, we have $iR_1^{[H]}j \in R_1^{[H]}$ if and only if $iR_2^{[H]}j \in R_2^{[H]}$ or $jR_2^{[H]}i \in R_2^{[H]}$. Thus, $\mathcal{R}_1^{[H]}$ and $\mathcal{R}_2^{[H]}$ have the same skeleton. Since $\mathcal{R}_1$ and $\mathcal{R}_2$ are perfect, so are $\mathcal{R}_1^{[H]}$ and $\mathcal{R}_2^{[H]}$. Hence $\mathcal{R}_1^{[H]}$ and $\mathcal{R}_2^{[H]}$ are equivalent, so that

$$p_{\mathcal{R}_1^{[H]}}(x_H) = p_{\mathcal{R}_2^{[H]}}(x_H). \tag{44}$$

From the equations (42) to (44), we get $p_{\mathcal{R}_1}(x_H) = p_{\mathcal{R}_2}(x_H)$, which implies the result. □

A.4 Omitted Proofs from Subsection 5.2

Proof of Proposition 8. We first prove statement (a). To this end, we derive the optimal contract under the objective model $\mathcal{R}^*$ that implements high effort. For convenience, we abbreviate $w_H = w(y_H)$, $w_M = w(y_M)$, and $w_L = w(y_L)$. Standard arguments show that both *IC* and *PC* must be binding at the optimal contract, and that $w_L < 0$ and $w_H > 0$ at the optimum. Assume for the moment that $w_M \geq 0$ under the optimal contract. The *IC* is then

$$\beta(w_H - \lambda w_L) = c, \tag{45}$$

and the *PC* equals

$$(\beta_H(\xi) + \beta)w_H + \beta_M(\xi)w_M + (\beta_L(\xi) - \beta)\lambda w_L = 0. \tag{46}$$

From the *IC* we get

$$w_H = \frac{c}{\beta} + \lambda w_L, \tag{47}$$

We plug this into the *PC*, solve for w_M, and get

$$w_M = -\frac{\beta_H(\xi)}{\beta_M(\xi)\beta}c - \frac{\beta_L(\xi) + \beta_H(\xi)}{\beta_M(\xi)}\lambda w_L. \tag{48}$$

The expected wage payment of the principal when he implements $a = 1$ equals

$$\mathbb{E}[w \mid a = 1] = (\beta_H(\xi) + \beta)w_H + \beta_M(\xi)w_M + (\beta_L(\xi) - \beta)w_L. \tag{49}$$

Using the results from above, the expected wage payment simplifies to

$$\mathbb{E}[w \mid a = 1] = c - (\beta_L(\xi) - \beta)(\lambda - 1)w_L. \tag{50}$$

The optimal wage w_L minimizes this term subject to the constraint that w_M in (48) remains weakly positive. The solution implies that $w_M = 0$, and $w(y_L) = -\frac{1}{2\lambda\beta}c$ as well as $w(y_H) = \frac{1}{2\beta}c$. We obtain the same result when we go through the same steps while assuming $w_M \leq 0$. With this we can compose the expected wage payment $\mathbb{E}[w \mid a = 1]$ and obtain

$$\frac{\partial \mathbb{E}[w \mid a = 1]}{\partial \xi} = \frac{\varepsilon}{2\beta}c - \frac{\varepsilon}{2\lambda\beta}c > 0. \tag{51}$$

Hence, the expected wage payment to implement $a = 1$ strictly increases in risk. The expected wage payment to implement $a = 0$ is zero for all risk levels. This yields us statement (a). Next, we prove statement (b). We first derive the agent's beliefs about the production function at $\alpha = 1$. As in the proof of Proposition 2, we find $p(x_2 = 1 \mid x_1 = 1)$ and $p(x_2 = 1 \mid x_1 = 0)$. At $\alpha = 1$, we have $p(x_2 = 1 \mid x_1 = 1) = p(x_2 = 1 \mid x_1 = 0) = \beta_2 + \beta_{02}$, and thus

$$p(y_H \mid x_1 = 1) = \beta_3^H(\xi) + \beta_{13}(\xi) + (\beta_2 + \beta_{02})\beta_{23}(\xi), \tag{52}$$
$$p(y_H \mid x_1 = 0) = \beta_3^H(\xi) + (\beta_2 + \beta_{02})\beta_{23}(\xi), \tag{53}$$
$$p(y_M \mid x_1 = 1) = \beta_3^M(\xi), \tag{54}$$
$$p(y_M \mid x_1 = 0) = \beta_3^M(\xi), \tag{55}$$
$$p(y_L \mid x_1 = 1) = \beta_3^L(\xi) - \beta_{13}(\xi) - (\beta_2 + \beta_{02})\beta_{23}(\xi), \tag{56}$$
$$p(y_L \mid x_1 = 0) = \beta_3^L(\xi) - (\beta_2 + \beta_{02})\beta_{23}(\xi). \tag{57}$$

From this, we can derive the agent's beliefs about the production function at $\alpha = 1$ as

$$p_{\mathcal{R}}(y_H \mid a = 1; \alpha = 1) = \beta_3^H(\xi) + (\beta_1 + \beta_{01})\beta_{13}(\xi) + (\beta_2 + \beta_{02})\beta_{23}(\xi), \tag{58}$$
$$p_{\mathcal{R}}(y_M \mid a = 1; \alpha = 1) = \beta_3^M(\xi), \tag{59}$$
$$p_{\mathcal{R}}(y_L \mid a = 1; \alpha = 1) = \beta_3^L(\xi) - (\beta_1 + \beta_{01})\beta_{13}(\xi) - (\beta_2 + \beta_{02})\beta_{23}(\xi), \tag{60}$$

and

$$p_{\mathcal{R}}(y_H \mid a = 0; \alpha = 1) = \beta_3^H(\xi) + \beta_1\beta_{13}(\xi) + (\beta_2 + \beta_{02})\beta_{23}(\xi), \tag{61}$$
$$p_{\mathcal{R}}(y_M \mid a = 0; \alpha = 1) = \beta_3^M(\xi), \tag{62}$$
$$p_{\mathcal{R}}(y_L \mid a = 0; \alpha = 1) = \beta_3^L(\xi) - \beta_1\beta_{13}(\xi) - (\beta_2 + \beta_{02})\beta_{23}(\xi). \tag{63}$$

At $\alpha = 1$, the IC is therefore given by

$$\beta_{01}\beta_{13}(\xi)(u(w_H) - u(w_L)) \geq c. \tag{64}$$

The rest of the proof proceeds as in the proof of statement (a). Assume that at the optimal equilibrium contract, we have $w_M \geq 0$. From the IC, we then get that at the optimal equilibrium contract, we must have

$$w_H = \frac{c}{\beta_{01}\beta_{13}(\xi)} + \lambda w_L, \tag{65}$$

and from the PC we then get that

$$w_M = -\frac{\beta_H(\xi) + \beta - \beta_{01}\beta_{13}(\xi)}{\beta_M(\xi)\beta_{01}\beta_{13}(\xi)} - \frac{\beta_L(\xi) + \beta_H(\xi)}{\beta_M(\xi)}\lambda w_L. \tag{66}$$

With this, we can calculate the expected wage payment under the optimal equilibrium contract that implements $\alpha = 1$ as

$$\mathbb{E}[w \mid a = 1; \mathcal{R}] = c - (\beta_L(\xi) - \beta)(\lambda - 1)w_L. \tag{67}$$

The optimal wage w_L minimizes this term subject to the constraint that w_M in (66) remains weakly positive. The solution implies that $w_M = 0$ as well as

$$w_L = -\frac{\beta_H(\xi) + \beta - \beta_{01}\beta_{13}(\xi)}{\lambda(\beta_H(\xi) + \beta_L(\xi))\beta_{01}\beta_{13}(\xi)}c \text{ and } w_H = \frac{\beta_L(\xi) - \beta + \beta_{01}\beta_{13}(\xi)}{(\beta_H(\xi) + \beta_L(\xi))\beta_{01}\beta_{13}(\xi)}c. \tag{68}$$

We obtain the same result when we go through the same steps while assuming $w_M \leq 0$. With this we can compose the expected wage payment at the optimal equilibrium contract as

$$\mathbb{E}[w \mid a = 1; \mathcal{R}] = \frac{(\lambda - 1)(\beta_H(\xi) + \beta)(\beta_L(\xi) - \beta) + (\lambda + 1)\beta_{01}\beta_{13}(\xi)}{\lambda(\beta_H(\xi) + \beta_L(\xi))\beta_{01}\beta_{13}(\xi)}. \tag{69}$$

We differentiate this expression with respect to risk ξ and find

$$\lim_{\varepsilon \to 0} \frac{\partial \mathbb{E}[w \mid a = 1; \mathcal{R}]}{\partial \xi} = -\frac{\lambda(\lambda - 1)(\beta_H(\xi) + \beta_L(\xi))(\beta_H(\xi) + \beta)(\beta_L(\xi) - \beta)}{[\lambda(\beta_H(\xi) + \beta_L(\xi))\beta_{01}\beta_{13}(\xi)]^2} < 0. \tag{70}$$

Hence, if ε is sufficiently small, the expected wage payment needed to implement $\alpha = 1$ decreases in risk ξ. The rest of the proof of statement (b) proceeds in the same way as for statement (a). □

A.5 Team Size and Incentives

In a team incentive problem, effort incentives are provided by tying each team member's payoff to the joint output y. The effectiveness of team incentives is constrained by the size of the team. When an agent's relative contribution to the output becomes small, it is typically no longer profitable for the principal to condition her pay on y, as the incentive effect would be outweighed by the costs of incentive provision (e.g., Kandel and Lazear 1992). An important implication of this result is that stock-options should be granted only to those employees whose actions significantly move the stock price. However, many firms grant stock options also to non-executive employees, and there is evidence that these have positive incentive effects (e.g., Hochberg and Lindsey 2010). In the following, we provide a belief-based explanation for this phenomenon. Specifically, we show that output-based incentives can remain effective in large teams when agents do not take into account the contributions of others to the final output.

Team incentives and optimal project size. We consider a simple team setting in which the principal chooses both incentives and the size of the team. Let there be m identical agents who can choose between high and low effort $a \in \{0, 1\}$. We suppress notation for individual agents. For convenience, we assume that agents are risk-neutral and protected by limited liability, so that $w(y) \geq \bar{w} > 0$ for all $y \in Y$. The project output is either large ($y = y_H$) or small ($y = y_L$). The team size m scales these payoffs. We have $y_H = m^\theta \bar{y}_H$ and $y_L = m^\theta \bar{y}_L$, for some $\theta \in (0, 1)$, and normalize $\bar{y}_L = 0$. If the share k of agents exerts high effort, the probability of a high output is $kB + D$, where B, D are positive constants with $B + D < 1$. Thus, as the team becomes large, the relative influence of a single agent on the final output becomes small. The cost of high effort for the individual agent is c and the cost of low effort is 0.

The principal chooses both team size m and agents' incentives $w(y)$. If he wishes to implement high effort from m agents, the optimal wage scheme is a bonus scheme with $w(y_H) = \bar{w} + \frac{mc}{B}$ and $w(y_L) = \bar{w}$. The principal's profit is then given by

$$(B + D)\left(m^\theta \bar{y}_H - \frac{m^2 c}{B}\right) - m\bar{w}. \tag{71}$$

Observe that $w(y_H)$ converges to infinity for $m \to \infty$. As project size increases, it becomes prohibitively costly to provide effort incentives, as an individual agent's influence on the final outcome becomes small. If the principal wishes to implement low effort from m agents, the optimal wage scheme is a fixed-wage $w(y_H) = w(y_L) = \bar{w}$ and the corresponding profit is $Dm^\theta \bar{y}_H - m\bar{w}$. Denote by $m^{[a]}(c)$ the optimal team size if the principal implements action $a \in \{0, 1\}$ and effort costs are given by c. This value is uniquely defined. We then get the following result: There is a $c^* > 0$ such that the principal optimally implements high effort with project size $m^{[1]}(c)$ if $c \le c^*$, and low effort with project size $m^{[0]}(c)$ if $c > c^*$.

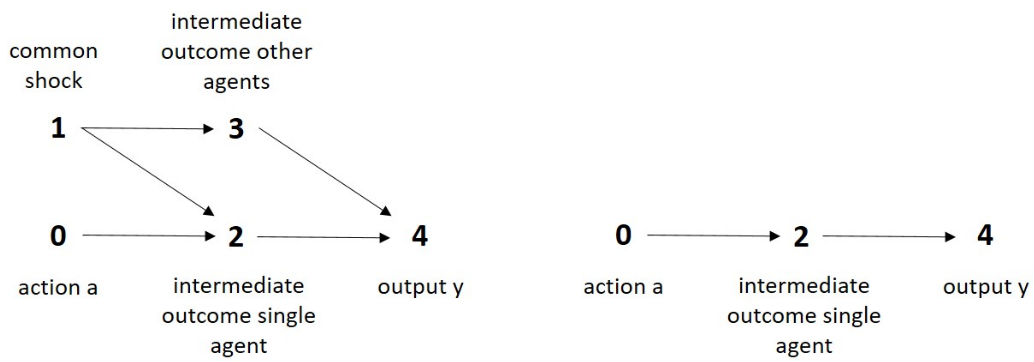

Figure 7: The objective model $\mathcal{R}^*$ and the agent's subjective model $\mathcal{R}$ in the team size example.

Team incentives and optimal project size with misspecified model. We now consider an extended production function that is consistent with the production function indicated above, and that allows us to study how team size and incentives change when agents do not take their colleagues' effort into account. Consider the objective model $\mathcal{R}^*$ on the left of Figure 6. Node 0 is the effort of a single agent. We use the fact that all agents are symmetric, and assume that each of the other $m-1$ agents exerts high effort with probability α^o. Through her effort, the single agent affects an intermediate outcome $x_2 \in \{0,1\}$. Denote by $x_3 \in \{0,1\}^{m-1}$ the $m-1$-dimensional vector of intermediate outcomes of all other agents. The probability of high output increases linearly in the number of high intermediate outcomes, $p(y_H \mid x_2, x_3) = \beta_{24}x_2 + \beta_{24} \parallel x_3 \parallel$, where $\parallel . \parallel$ is the sum of entries in a vector. There is a common shock $x_1 \in \{0,1\}$ that occurs with probability $p(x_1 = 1) = \beta_1$. It positively affects each agent's intermediate outcome, $p(x_2 = 1 \mid a, x_1) = \beta_{02}a + \beta_{12}x_1$, where $\beta_{02} + \beta_{12} < 1$ and $\beta_1\beta_{12} \ge \frac{1}{2}$; for any other agent, the probability of a high intermediate outcome is $\beta_{02}\alpha^o$ if $x_1 = 0$ and $\beta_{02}\alpha^o + \beta_{12}$ if $x_1 = 1$. We define $B \equiv \beta_{02}\beta_{24}$ with $\beta_{24} = \frac{\bar{\beta}_{24}}{m}$ for some $\bar{\beta}_{24}$, and $D \equiv \beta_1\beta_{12}\bar{\beta}_{24}$. Thus, the production function is the same as above; optimal team size and incentives would remain unchanged if the agents' subjective model would be given by $\mathcal{R}^*$. We assume now that agents ignore the contributions of others. Let an agent's subjective model be given by $\mathcal{R}$ on the right of Figure 6. We then obtain the following result.

Proposition 9 (Team Size and Incentives)**.** *Consider the team size example of this section.*

(a) Under the objective model $\mathcal{R}^*$ there is a unique $c^* > 0$ such that the principal optimally implements high effort with team size $m^{[1]}(c)$ if $c \le c^*$, and low effort with team size $m^{[0]}(c)$ if $c > c^*$.

(b) Under the subjective model $\mathcal{R}$ there is a unique $c^{**} > c^*$ such that the principal optimally

implements high effort with team size $m^{[1]}_{\mathcal{R}}(c) > m^{[1]}(c)$ *if* $c \leq c^{**}$*, and low effort with team size* $m^{[0]}(c)$ *if* $c > c^{**}$.

Thus, whenever effort costs are small enough so that $c \leq c^{**}$, the principal chooses a team size that is "too large" for also tying the agents' pay to the output. It then appears as if incentives are provided to too many employees. However, as we show next, the simplification in the agents' subjective model causes them to overestimate the importance of their effort for the final output, so that granting these incentives remains profitable for the principal.

We explain the intuition behind Proposition 9. According to $\mathcal{R}$, if $p(a = 1) = \alpha$, the agent's belief about how her intermediate outcome affects the final output equals

$$p(y_H \mid x_2 = 1) - p(y_H \mid x_2 = 0) = \beta_{24}[1 + \xi(\alpha, \beta_1, \beta_{02}, \beta_{12})(m - 1)], \tag{72}$$

where $\xi(\alpha, \beta_1, \beta_{02}, \beta_{12}) = \frac{\beta_1(1-\beta_1)\beta_{12}^2}{(\beta_{12}+\alpha\beta_{02})(1-\beta_{12}-\alpha\beta_{02})} \in (0, 1)$. Note that under the objective model, the value in (72) would be equal to β_{24} and therefore vanish as the project size m becomes large. Thus, under the subjective model, the agent overestimates the importance of her intermediate outcome for the output. The reason is that a high outcome indicates a positive common shock, which also increases the chance of high intermediate outcomes for all other agents. Under $\mathcal{R}$ the agent falsely attributes the corresponding increase in the probability of a high output y_H to the significance of her intermediate outcome x_2. We show below that her perception of the significance of her intermediate outcome decreases in team size, but converges against a positive constant for $m \to \infty$. Thus, the agent maintains a certain belief in the importance of her effort even when her true impact on the final output vanishes.

When $\mathcal{R}$ is the subjective model of all agents, the principal's profit from implementing high effort at team size m with the optimal incentive scheme is given by

$$(B + D)\left(m^{\theta}\bar{y}_H - \frac{m^2 c}{B}\frac{1}{1 + \xi(1, \beta_1, \beta_{02}, \beta_{12})(m - 1)}\right) - m\bar{w}. \tag{73}$$

From this we can derive the optimal project size $m^{[1]}_{\mathcal{R}}(c)$ at cost c. The profit from implementing low effort from m agents remains the same as under the objective model. Proposition 9 then follows from a comparison of the profit levels in (71) and (73).

Mathematical Details. We fit the agent's subjective model $\mathcal{R}$ to the probability distribution, taking α and α^o as given. First, we calculate the probabilities that there is a common shock, given that $x_2 = 1$ and $x_2 = 0$, respectively. We get

$$p(x_1 = 1 \mid x_2 = 1) = \frac{\beta_1\beta_{12} + \alpha\beta_1\beta_{02}}{\beta_1\beta_{12} + \alpha\beta_{02}}, \tag{74}$$

$$p(x_1 = 1 \mid x_2 = 0) = \frac{\beta_1(1 - \beta_{12} - \alpha\beta_{02})}{1 - \beta_1\beta_{12} - \alpha\beta_{02}}. \tag{75}$$

The agent's subjective probability of a high output y_H after a high intermediate outcome $x_2 = 1$ is then given by

$$\begin{aligned} p(y_H \mid x_2 = 1) = {} & \frac{\beta_1\beta_{12} + \alpha\beta_1\beta_{02}}{\beta_1\beta_{12} + \alpha\beta_{02}} \left[1 + \sum_{k=0}^{m-1} \binom{m-1}{k} (\beta_{12} + \alpha\beta_{02})^k (1 - \beta_{12} - \alpha\beta_{02})^{m-1-k} k\right] \beta_{24} \\ & + \left(1 - \frac{\beta_1\beta_{12} + \alpha\beta_1\beta_{02}}{\beta_1\beta_{12} + \alpha\beta_{02}}\right) \left[1 + \sum_{k=0}^{m-1} \binom{m-1}{k} (\alpha\beta_{02})^k (1 - \alpha\beta_{02})^{m-1-k} k\right] \beta_{24}. \end{aligned} \tag{76}$$

Using $\binom{m}{k} p^k (1-p)^{m-k} k = mp$ we get

$$p(y_H \mid x_2 = 1) = \beta_{24}(1 + \alpha^o\beta_{02}(m-1)) + \frac{\beta_1\beta_{12} + \alpha\beta_1\beta_{02}}{\beta_1\beta_{12} + \alpha\beta_{02}} \beta_{24}\beta_{12}(m-1). \tag{77}$$

Similarly, we get

$$p(y_H \mid x_2 = 0) = \beta_{24}\alpha^o\beta_{02}(m-1) + \frac{\beta_1(1 - \beta_{12} - \alpha\beta_{02})}{1 - \beta_1\beta_{12} - \alpha\beta_{02}} \beta_{24}\beta_{12}(m-1). \tag{78}$$

From equations (77) and (78) we can then derive $p(y_H \mid x_2 = 1) - p(y_H \mid x_2 = 0)$ and the incentive compatibility constraint. From this IC we can derive that if the principal wishes to implement high effort from m agents, then the optimal incentive scheme is

$$w(y_H) = \bar{w} + \frac{cm}{B[1 + \xi(1, \beta_1, \beta_{02}, \beta_{12})(m-1)]} \quad \text{and} \quad w(y_L) = \bar{w}. \tag{79}$$

From this the principal's profit in equation (73) follows. Note that $\beta_1\beta_{12} \geq \frac{1}{2}$ implies that $\xi(\alpha, \beta_1, \beta_{02}, \beta_{12})$ is maximal at $\alpha = 1$. Thus, the principal cannot gain by implementing high effort with probability $\alpha \in (0, 1)$.

Finally, we show that $\xi(\alpha, \beta_1, \beta_{02}, \beta_{12}) < 1$ for all admissible values $\alpha, \beta_1, \beta_{02}, \beta_{12}$. This inequality is identical to $\beta_1\beta_{12}(1 - \beta_{12}) + \alpha\beta_{02}(1 - 2\beta_1\beta_{12} - \alpha\beta_{02}) > 0$. Since $1 > \beta_{02} + \beta_{12}$ and $\alpha \leq 1$, this inequality is implied by $\beta_1\beta_{12}(1 - \beta_{12}) - \beta_1\beta_{12}\beta_{02} = \beta_1\beta_{12}(1 - \beta_{02} - \beta_{12}) > 0$, which implies the statement above.

References

Hochberg, Yael, and Laura Lindsey (2010): "Incentives, Targeting, and Firm Performance: An Analysis of Non-executive Stock Options," *Review of Financial Studies*, 23(11), 4148–4186.

Kandel, Eugene, and Edward Lazear (1992): "Peer Pressure and Partnerships," *Journal of Political Economy*, 100(4), 801–817.

Spiegler, Ran (2019): "Can Agents with Causal Misperceptions be Systematically Fooled?," *Journal of the European Economic Association*, forthcoming.