Check for updates

The Manchester School

WILEY

The Manchester School

ORIGINAL ARTICLE OPEN ACCESS

Asymmetric Priors in Agency Under Maximum-Entropy

Óscar Gutiérrez

Departament of Business Economics, Universitat Autònoma de Barcelona, Barcelona, Spain

Correspondence: Óscar Gutiérrez (oscar.gutierrez@uab.es)

Received: 19 July 2024 | **Revised:** 14 August 2025 | **Accepted:** 25 August 2025

Keywords: asymmetric priors | first-order approach | maximum-entropy principle | moral hazard | principal-agent model

ABSTRACT

This article analyzes the agency relationship in a setting where principal and agent have asymmetric priors on output. These differing priors stem from distinct information sets regarding the (unknown) output distribution, being assumed that both principal and agent rely on the *maximum-entropy principle* to estimate the respective distributions. The main conclusions obtained are the following. (1) Additional information on output does not necessarily lead to an improved equilibrium. (2) If the principal holds an ex ante piece of private information, it may be optimal for her to hide it and design the agency contract ignoring the additional information. (3) If the agent holds an ex ante piece of information, the contract may fail to be incentive-compatible; in such a case both parties improve if the agent shares his information and the principal ascribes to it, which justifies the common prior assumption.

JEL Classification: D81, D82, D86, J33, M40

1 | Introduction

In the analyses of the agency relationship under moral hazard, it is usual to adopt the *common prior assumption*, according to which both principal and agent agree on the output probability distribution. The role of different prior beliefs on the relevant output distribution and their effects on agency contracts and value creation is thus scarcely considered in agency theory.[1] In this paper, we challenge the common assumption that the principal and agent have perfectly specified and identical output distributions, arguing that is restrictive and unrealistic. We examine the role of asymmetric priors (or different beliefs[2]) due to different ex ante information sets in a context of uncertainty where principal and agent have only partial information about the probability distribution describing output. By applying the maximum-entropy principle, principal and agent can infer the probability distribution that best represents output based on their available information.

The maximum-entropy principle (MEP, henceforth) can be viewed as a generalization of Laplace's principle of indifference, according to which all events must be taken as equally likely if there is no evidence to the contrary. Thus, for a set of possible outcomes $\{x_1, x_2, ..., x_N\}$, the probability associated to x_i is simply $P(X = x_i) = 1/N$. The MEP generalizes this simple principle to encompass situations where some additional information on the probability distribution is available. According to the MEP, the distribution with the largest entropy is the best guess given the information available on a system. When a constraint is inherent to the problem -for example, when a statistical moment of the distribution is known-the MEP gives a natural way to incorporate such information into the mathematical specification of output. Thus, the proposed framework offers a

Óscar Gutiérrez did not receive support from any organization for the submitted work. He gratefully thanks Vicente Salas for his advice and the fruitful discussions, and Fran Ruiz and Joaquín Poblete for their insightful comments and suggestions made on earlier versions of the paper.

This is an open access article under the terms of the Creative Commons Attribution-NonCommercial-NoDerivs License, which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.

© 2025 The Author(s). *The Manchester School* published by The University of Manchester and John Wiley & Sons Ltd.

parsimonious approach to handling uncertainty in agency settings, which are usually modeled in rather complex models.

In our context, the role played by an *additional piece of information* is key: it represents additional knowledge on the output distribution, codified via an additional constraint to the mathematical problem associated with finding the output distribution. The framework proposed is suitable for analyzing the implications of ex ante private information on equilibrium, contracts, and value creation. First, the MEP allows us to construct output distributions based on the information already available for principal and agent. Second, additional pieces of information affect output –and consequently incentives– differently for each player.

Let us briefly review some concepts on decision-making under uncertainty and their relation to the principal-agent model. The *subjective expected utility* framework (SEU, henceforth) applies to decision-making problems when underlying uncertainty can be represented by a known probability distribution, a situation referred to as *risk*. If a probability distribution is not available for the decision-maker, the situation is instead known as *Knightian uncertainty, ambiguity*, or simply *uncertainty*. For decision-making problems under uncertainty, a strand of decision theory has proposed the use of ambiguity-aversion models; Gilboa and Schmeidler (1989), Klibanoff et al. (2005), Hansen and Sargent (2001). In these models, the decision-maker faces a problem that can be described by means of *a set* of probability distributions. By applying a decision criterion, such as max-min, the decision-maker can compare available alternatives (see Gilboa and Schmeidler 1989). The ambiguity-aversion approach to decision-making is appealing as it circumvents the so-called *Ellsberg paradox*, for which SEU models cannot account for (see Ellsberg 1961). On the other hand, some inconsistencies arise in ambiguity-aversion models. Al-Najjar and Weinstein (2009) present a detailed critical analysis showing that ambiguity-aversion models engender more problems than they solve. Ambiguity-aversion models have been applied to agency under moral hazard; see Weinschenk (2010), Lopomo et al. (2011), Kellner (2017) and references therein. A key challenge in applying ambiguity aversion models to agency lies in the extensive information the principal must have about the agent's beliefs and preferences (including the set of distributions describing the output signal and the agent's decision criterion).

In this paper we propose a simpler alternative to agency under uncertainty, where the standard principal-agent model of Mirrlees-Holmstrom (see Holmstrom 1979) is modified to consider an output distribution built upon the (perhaps diverging) pieces of information that the contracting parties possess. It consists of using a unique output distribution (for each contracting party) that is inferred from the already available information on the economic environment, the so-called *maximum-entropy distribution* (MED, henceforth). Such information can typically include some moments of the output distribution, the distribution support, or the events with zero probability. Since any smooth density function can be constructed from its distribution moments using a Taylor expansion, it is reasonable to consider moments as relevant pieces of information about the output distribution. The use of moments as partial information on an unknown distribution is not new (see e.g. Carrasco et al. (2018) and references therein). The present MEP approach to agency accommodates different output distributions derived from diverging information sets. Compared to ambiguity-aversion models, its advantages include simplicity, the compatibility with scarce information on output, limited knowledge regarding agents' beliefs and preferences, and compliance with Savage axioms. However, a key drawback is the assumption of a unique distribution describing the underlying uncertainty, making it inherently incompatible with Ellsberg paradox and other limitations of the SEU framework.

The analysis builds on Gutiérrez and Salas-Fumás (2021), who explore the implications of applying the MEP to agency contracting in situations where principal and agent only possess partial knowledge of the output probability distribution, limiting their analysis to *symmetric priors*. The authors characterize the second-best agency contract derived from MEDs that reflect the (shared) information of principal and agent. They demonstrate that second-best contracts tend to be more complex as additional information about the output is incorporated to the problem, encompassing many compensation schemes observed in real agency relationships. In contrast, the principal-agent model proposed here is intended to explore the effects of different priors, which stem from (possibly) different pieces of information available to each party. This approach allows us to analyze how different priors on output influence second-best contracts and welfare by examining the impact of an additional piece of information on the agency game.

The model provides several novel contributions to the principal-agent literature. First, an additional piece of information (either shared by the parties or privately held by one) does not necessarily lead to welfare increase (in the Pareto sense). Furthermore, the effect of an additional piece of information can differ depending on which party holds it. Specifically, if the principal holds an additional piece of information, she may be better-off by ignoring it when designing the agency contract, a result connected to *Holmstrom's Informativeness Principle* and *Myerson's Inscrutability Principle*. Conversely, if the agent holds the additional piece of information, the contract designed by the (uninformed) principal may lead to higher effort than the one the principal anticipated. Then, there is a possible welfare improvement for both parties if the agent shares his prior with the principal and she adopts the agent's belief. Under certain conditions, an additional piece of information on output for the agent does not alter the equilibrium outcome of the contract designed according to the principal's priors. Finally, the question of the common prior assumption is revisited, offering different justifications for its use.

Section 2 presents the agency model and its MEP approach. In Section 3, we consider the case of symmetric priors and derive some preliminary results. In Section 4, we consider that the principal has ex ante private information on the output distribution. In Section 5, the agent has ex ante private information. Section 6 discusses and summarizes the paper's results.

14679957, 2026, 1, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/manc.70010 by Spanish Cochrane National Provision (Ministerio de Sanidad), Wiley Online Library on [13/05/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

2 | The Setup

2.1 | Assumptions

First, we present the principal-agent model under moral hazard in its classical formulation, following J. Mirrlees and B. Holmstrom; see Holmstrom (1979). We then introduce the key features of our approach, which represents a variation of the Mirrlees-Holmstrom framework.

2.1.1 | Canonical Agency Setup

A risk-neutral principal (she) hires a risk-averse and effort-averse agent (he) to perform an unobservable task that affects an *ex post* signal X, which we identify with output. The agent's input (say effort) is denoted as a, and entails an opportunity cost $C(a)$, with $C'(a) > 0$ and $C''(a) > 0$. The agent's effort is private information, and both principal and agent observe the realized output x (X represents the random output while x refers to a specific ex-post realization, though both terms can be used interchangeably). The principal offers the agent an output-dependent compensation $s(x)$ for the costly input in order to (partially) solve the moral hazard problem. The net utility of the agent is additively separable, $U(w)-C(a)$, where w represents the payment from the principal corresponding to $s(x)$ in our case, and $U(\cdot)$ is a strictly increasing and strictly concave N-M utility function: $U'(\cdot) > 0$, $U''(\cdot) < 0$ (the agent is risk-averse). To justify the application of the first-order approach to agency problems (FOA henceforth), it is useful to assume the strict concavity of function $U[J(\cdot)]$, with $J(x) \overset{def}{=} (U')^{-1}(1/x)$; see Jewitt (1988, Eqs. (2.9) and (2.12)), Chade and Swinkels (2020).

The optimization problem for the principal is stated as follows:

$$\begin{aligned} &\max_{s(x),a} \int [x - s(x)] f(x,a)dx \\ &s.t. \int U(s(x)) f(x,a)dx - C(a) \geq R \qquad \text{(P1)} \\ &a \in \arg\max_{a'} \int U(s(x)) f(x,a')dx - C(a'), \end{aligned}$$

where $f(x,a)$ represents the output density function for a given effort a, and R denotes the agent's reservation utility. The two constraints in (P1) are respectively known as the *participation constraint* (PC) and the *incentive compatibility constraint* (IC).

The optimization problem (P1) is difficult to solve. In our context, however, the output distribution is assumed to belong to a particular class of functions (the so-called MED functions; see Sec. 2.1.2). This assumption helps validate the FOA, as the MEDs belong to the exponential family of distributions, which, along with other non-stringent conditions, ensure the validity of the FOA; see Chade and Swinkels (2020).[3] According to the FOA, the original constraint (IC), $a \in \arg\max_{a'} \int U(s(x)) f(x,a')dx - C(a')$, is replaced by the relaxed constraint (IC'): $\int U(s(x)) f_a(x,a)dx - C'(a) = 0$, where $f_a \equiv \partial f / \partial a$. When the FOA holds, the optimal sharing rule $s^*(x)$ satisfies the two following conditions:

$$\frac{1}{U'(s^*(x))} = \lambda + \mu \frac{f_a(x,a^*)}{f(x,a^*)} \qquad \text{(1a)}$$

$$\begin{aligned} &\int (x - s^*(x)) f_a(x,a^*)dx \\ &+ \mu \left[\int U(s^*(x)) f_{aa}(x,a^*)dx - C''(a) \right] = 0, \end{aligned} \qquad \text{(1b)}$$

where λ and μ represent the Lagrange multipliers associated to (PC) and (IC'), respectively. The optimal contract is derived from Equation (1a) as a function of λ, μ and a^*, which in turn are obtained from (PC), (IC') and Equation (1b).

2.1.2 | The MEP Approach to Agency

In the MEP approach to agency proposed here, the output distributions are not explicitly known but are inferred from the information available to the contracting parties, and this inference is carried out by invoking the Maximum-Entropy Principle. The approach is based on the two following assumptions.

A. *The expected output is driven by the agent's effort. Specifically, and for the sake of simplicity, we assume $E(X) = a$. This information is common knowledge for both principal and agent.*

B. *Principal and agent appeal to the* Maximum-Entropy Principle *to guess an output distribution. The information-constrained problems from which they infer the output distributions depend on the information set available for each one, so the* MEDs *of principal and agent can diverge.*

2.1.2.1 | Discussion of (A1) and (A2). According to (A1), principal and agent ignore the output distribution for a given a (or the implicit probability values associated with the states of nature). The minimum information on the output distribution shared by principal and agent at the start of the relation is that the expected output depends on the agent's effort, say $E(X) = a$ for simplicity. The equality $E(X) = a$ can be justified by the minimum information necessary for the principal to consider hiring the agent, that is, for the agency relationship to be meaningful. The agreement between principal and agent on $E(X) = a$ is relatively undemanding, especially when compared to the assumption that principal and agent agree on the entire output distribution.

The *entropy* H of the discrete probability distribution $\{p_1, ..., p_N\} \equiv \{P(x_1), ..., P(x_N)\}$ gives a measure of its uncertainty, and it is defined as $H(p_1, ..., p_N) \overset{def}{=} -\sum_{i=1}^{N} p_i \ln(p_i)$. This mathematical specification ensures that if the event x_i (with probability p_i) is the realized one, then an amount of information $-\ln(p_i)$ is gained, and hence, the information carried by a distribution will be equal to $-\sum_{i=1}^{N} p_i \ln(p_i)$.

Information Theory (Shannon 1948) proposes measuring the uncertainty faced by a decision-maker through the *entropy* of

14679957, 2026, 1, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/manc.70010 by Spanish Cochrane National Provision (Ministerio de Sanidad), Wiley Online Library on [18/05/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

the probability distribution $(p_1, ..., p_N)$. In the agency model used in this paper, the output signal is continuous, and entropy is then defined as $H(X) \overset{def}{=} -\int f(x)\ln[f(x)]dx$. In our context, the basic information required for the agency relationship to take place is $E(X) = \int xf(x)dx = a$.

As has been mentioned, the MEP generalizes the *Laplace's principle of indifference* (aka *Principle of insufficient reason*). The MEP also prescribes to choose the probability distribution consistently with the information available (and no more than that). In fact, the distribution with maximum entropy assures that no arbitrary assumptions are unconsciously introduced in the inference.[4] Therefore, if principal and agent reason under the information theory principles from different information sets (thus forming different priors), the two will separately choose different probability distributions to describe output. Consequently, the expectations for principal and agent involved in the agency problem will be calculated from the corresponding output distributions.

2.2 | Additional Information

In the MEP approach to agency, the role played by an additional piece of information is crucial. If an additional piece of information is incorporated into the problem (of any contracting party), the information-constrained problem becomes:

$$\begin{aligned} &\max_{f(x)} H(X) \\ &s.t. \int f(x)dx = 1 \\ &\int xf(x)dx = a \\ &\int \varphi(x)f(x)dx = b \end{aligned} \tag{P2}$$

where b is a model parameter and $\varphi(x) \neq x$ represents a smooth function. It must be noted that the right-hand side of the last constraint is party specific (i.e., b_A for the agent and b_P for the principal); so, the agent's and the principal's prior distributions ($f^A(x)$ and $f^P(x)$, respectively) are obtained by solving different versions of the problem (P2).

The action (or effort) a is not given ex ante but is chosen from a continuous range of values, so a in (P2) is treated as a model parameter. In the simplest case, see (A1), the constraint $\int \varphi(x)f(x)dx = b$ drops from (P2); in such a case, the only information available on output is $E(X) = a$, and the output density is given by $f_{ME}(x) = (1/a)\exp(-x/a)$. In general, the maximum-entropy density corresponding to (P2) is (see Zellner and Highfield 1988):

$$f_{ME}(x) = (1/\Omega)\exp[-Cx - D\varphi(x)]$$

where C and D denote the multipliers associated with the two last constraints in (P2) and the constant $\Omega = \int \exp[-Cx - D\varphi(x)]dx$ ensures that the mass of probability is 1.

Under different (but public) priors, the expectations involved in the agency problem are calculated according to the corresponding prior distributions of principal and agent. The agency problem is then written as:

$$\begin{aligned} &\max_{s(x),a} E^P[x - s(x)] = \int [x - s(x)]f^P(x, a)dx \\ &s.t.\ E^A[U(s(x))] - C(a) = \int U(s(x))f^A(x, a)dx - C(a) \geq R \\ &\frac{d}{da}\{E^A[U(s(x))]\} - C'(a) = \int U(s(x))f^A_a(x, a)dx - C'(a) = 0, \end{aligned} \tag{P3}$$

In (P3), the superscripts denote the contracting party according to which the expectations are calculated, and the corresponding output distributions are obtained from the information available for each party by appealing to the MEP. It describes a scenario where the parties have asymmetric (but known) priors, meaning they *agree to disagree*. In contrast, this paper assumes that principal and agent do not agree to disagree; instead, one party possesses hidden information about the output distribution, which may or may not be communicated to the uninformed party (except in the next section, where priors are symmetric).

3 | Symmetric Priors

In this section we explore the implications of an additional piece of information on the output distribution assuming identical prior beliefs. The additional piece of information (known for principal and agent) must be incorporated through an additional constraint to problem (P2), whose solution is the output MED (identical for principal and agent in this particular case). Perhaps counterintuitively, an additional piece of information does not *necessarily* reduce the entropy of the MED, $H(X)$, which can increase or decrease the additional information.[5] In our context, the "additional piece of information" refers to the third constraint in (P2) but can be generalized to an arbitrary number of constraints.

More importantly, from the economic perspective, an additional piece of information does not *necessarily* increase welfare in the Pareto sense (Proposition 1). This result is linked to the fact that an additional piece of information can lead to an output MED that is more disperse than the one built without considering the additional constraint. Specifically, if the additional information leads to an output distribution with higher variance, which is mathematically possible, then the output distribution will lead to a welfare reduction, regardless of whether the entropy of the MED increases or decreases with the additional information. However, in most cases, an additional piece of information does increase welfare.

Proposition 1. *Consider the agency model of Section 2. Assume that principal and agent have identical beliefs on the output distribution. Then, an additional piece of information in (P2) for principal and agent does not necessarily improve welfare in the Pareto sense (even if the entropy of the MED decreases).*

Proof. See Appendix. □

14679957, 2026, 1, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/manc.70010 by Spanish Cochrane National Provision (Ministerio de Sanidad), Wiley Online Library on [13/05/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

4 | The Effects of Ex-Ante Private Information: The Principal

Traditional agency does not consider situations where one party has ex ante private information on output. In Sections 4 and 5 we consider that either principal or agent has a piece of information that the other party ignores.

In Section 4 the principal has the informational advantage. We show that the principal may optimally ignore the private information when designing the agency contract. In fact, in some circumstances the principal will not make use of her private information on output since the corresponding contract would reveal such information to the agent. A situation where the principal has ex ante private information on output may correspond to a principal who is able to carry out the tasks entrusted to the agent, that is, a principal with knowledge of the economic environment (e.g., the principal is the CEO of a multinational company and the agent is a national manager). In Section 5, it is the agent who has private information on the output distribution. This situation may correspond to a principal without the skills or knowledge necessary to complete the tasks entrusted to the agent (e.g., the principal is the shareholders of the company).

Consider first that the principal holds an additional piece of information. In our approach, an additional piece of information often improves welfare (in the Pareto sense); see the proof of Proposition one in the Appendix. If this is the case, the principal will take into account the additional information when designing the agency contract. However, the opposite conclusion is also possible: taking into account an additional piece of information to design the agency contract may diminish the principal's expected profit, and then the principal would optimally ignore it when designing the contract. Consider that output follows a gamma distribution of parameters p and a/p, with $p \neq 1$. The gamma distribution is the MED for which $E(X)$ and $E[\ln(X)]$ are constrained to be known, and it is frequently used in agency settings (see Gutiérrez and Salas-Fumás 2021); its density function is $f(x) = \frac{1}{(a/p)^p \Gamma(p)} x^{p-1} e^{-(p/a)x}$. Routine calculations yield $E[\ln(X)] = \ln(a/p) + \Gamma'(p)/\Gamma(p)$ with $\Gamma(\cdot)$ denoting the gamma function (a generalization of the factorial function: $\Gamma(n) = (n-1)!$ for integer numbers). We can observe that the additional constraint provides information on a, so one would expect that taking into account such information when designing the contract should improve welfare. However, if $p < 1$, the use of the additional piece of information "$E[\ln(X)] = \ln(a/p) + \Gamma'(p)/\Gamma(p) = \text{constant}$" can entail a welfare reduction; see Appendix. In such a case the private information will be strategically ignored by the principal in the design of the agency contract: from the contract issued by the principal, the agent could infer the principal's prior since the correspondence $s^*(x) \mapsto p$ is one-to-one (recall that $E(X) = a$ is common knowledge). The result is related to the fact that the additional piece of information (inferred from the contract) would inform the agent that the output MED has a higher variance than he anticipated (according to his prior). Since the agent is risk-averse, the principal finds suboptimal to design an agency contract that reveals such information.[6]

The following result points out the existence of setups where the principal strategically hides the additional piece of information on output, ignoring it when designing the contract. This conclusion resembles some previous results, as Weinschenk (2010); see Sec. 6 for a discussion. Weinschenk (2010) contradicts Holmstrom's *Informativeness Principle*, which establishes that any informative signal must be incorporated into the agency contract. However, Proposition 2 does not contradict it, since (ex ante) *pieces of information* are not the same as (ex-post) *signals*, and additional pieces of information can be "non-informative" in Holmstrom's sense (*informative* is equivalent to *valuable*, and valuable means that principal and agent can be strictly better-off incorporating the signal to the agency contract; see Holmstrom (*op. cit*, p. 83)).

Proposition 2. *Consider the MEP agency setup of Section 2. Assume that the principal has an additional piece of information. The net utility function of the agent is given by* $u(s, a) = \ln(s) - ka^2/2$, *with* k *a positive constant. Then, there exist mathematical specifications under which the principal improves by leaving aside the additional piece of information when designing the agency contract.*

Proof. In the Appendix we demonstrate the existence of parameter configurations validating the proposition. □

According to this result, it may be optimal for the principal to use a contract that does not reveal her private information, which relates it to the literature on mechanism design by an informed principal: the agency game becomes a signaling game in which the principal may reveal something about her private information through the agency contract. In a more general setting, Myerson proved his *Principle of Inscrutability*, which establishes the optimality of proposing inscrutable mechanisms to prevent the agent from learning anything about the principal's private information; see Myerson (1983). The result of Proposition 2 is related to the inscrutability principle, as when designing the agency contract the principal will avoid revealing a piece of information if it reduces her expected profits (from the contract issued by the principal, the agent could infer the principal's prior). If, instead, the principal can improve by revealing her private information (e.g., output follows a gamma with $p > 1$), she will communicate it to the agent (via cheap talk communication). It must be recalled that, according to the inscrutability principle, the principal does not use the contract to send messages to the agent; see Myerson (1983) and Galperti (2015, Sec. 1).[7] In our context, a contract which reveals that the principal's prior follows a gamma with $p < 1$ induces the agent to exert a lower effort (than if the additional information is ignored, which corresponds to an exponential output distribution). As has been mentioned in Section 3, this situation is related to the fact that the agent is risk-averse and the principal's private information leads to an output distribution more disperse than the exponential one.

5 | The Effects of Ex-Ante Private Information: The Agent

In this section, we consider the case where the agent holds an additional piece of information that the principal ignores. This assumption is line with Holmstrom (1979, Sec. 6), who argues

14679957, 2026, 1, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/manc.70010 by Spanish Cochrane National Provision (Ministerio de Sanidad), Wiley Online Library on [13/05/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

that an "unrealistic feature is the assumption that the agent chooses his action having the same information as the principal". First, we analyze an agency model where the principal can adopt the agent's priors as her own. In this setup, the principal-agent game remains incentive-compatible, and equilibrium is achieved even though the principal lacks the agent's ex ante private information about output (Proposition 3). Therefore, the additional piece of information has no consequences on the contract design, which provides a justification for the common prior assumption usual in agency. Secondly, we give conditions under which the agent maximizes his expected utility by exerting an effort above the one anticipated by the principal (Proposition 4). However, such situation is Pareto-dominated by the equilibrium in which the agent shares his information on output with the principal: as the agent is better-off by sharing his prior, he optimally sends cheap talk signals to the principal, who also improves by adhering to the agent's prior.

Proposition 3. *Consider the MEP agency setup of Section 2. Assume that the agent has an additional piece of information that sets* $E(\ln X) =$ constant (so, the agent's output MED is a gamma distribution). The utility function of the agent is $U(s,a) = 2\sqrt{s} - ka^2/2$, with $k > 0$. Then, the contract designed by the principal is second-best optimal, and the agent exerts the effort anticipated by the principal.

Proof. See Appendix. □

Hence, if the agent's utility is $u(s,a) = 2\sqrt{s} - ka^2/2$, output follows an exponential distribution according to the principal and a gamma distribution accordingto the agent (both MEDs with $E(X) = a$), then the common prior assumption is innocuous (the same conclusion holds if the principal believes that output follows any other gamma MED with $E(X) = a$). The reason lies in the fact that the private information on output of the agent (p) has no effect on (IC'), which implies that the agent chooses a level of effort that coincides with the effort anticipated by the principal. Specifically, the best-response function of the agent to the contract $s(x) = (\alpha+\beta x)^2$ turns out to be $a_A{}^* = 2\beta/k$ (with k the constant of effort-aversion), independent of parameter p; see Appendix. This best-response function is notably simple for an agency model, and resembles the one of the Holmstrom and Milgrom (1987) linear model.

However, the last result does not hold for most specifications of the income utility, for example $U(s) = \log(s)$. In general, the equilibrium of the game depends on the agent's prior, and the principal cannot act as if the agent's prior coincides with her own prior belief. In such circumstances, the agent may choose a level of effort different from the one anticipated by the principal. Next, we provide a model where the agent chooses a level of effort higher than the one anticipated by the principal, implying that both parties could improve if the agent shared his information on output. The fact that the principal can be better-off is a consequence of that, when the FOA holds, given the sharing rule the principal would like to see the effort increase; see Holmstrom (1979, Prop. 1) or Jewitt (1988, 1180). The following result summarizes the point.

Proposition 4. *Consider the MEP agency setup of Section 2. Assume that the agent has an additional piece of information that sets* $E(\ln X) =$ constant (so, the agent's output MED is a gamma distribution of parameters p and a/p). The agent's utility function is given by $U(s,a) = \ln(s) - ka^2/2$, with $k > 0$.

Then, if $p > 1$, the agent exerts a higher level of effort than anticipated by the principal, resulting in an improvement over the reservation utility.

Proof. See Appendix. □

The outcome obtained in Proposition 4 is suboptimal, as principal and agent could improve if the agent discloses his private information on output. In return, he demands a higher reservation utility, at least equal to the utility reached when the additional piece of information is not revealed. A corollary of Proposition 4 is the following.

Corollary 1. *Consider that the agent can renegotiate his reservation utility under the conditions of Proposition 4. Then, the agent is better-off by sharing his private information on output (through cheap talk communication), and the principal is better-off by adhering to the agent's beliefs.* □

It is worth observing that Corollary 1 contradicts the way in which Karni (2008) justifies the common prior assumption. According to Karni (2008), in static settings the principal can ascribe to the agent her own set of priors (since the utilities and subjective probabilities may be confounded; see Karni (2008, Sec. 5)). We provide an alternative justification for the common prior assumption, but now the principal must accommodate her priors to the agent's.

5.1 | Illustration

Take $k = 0.02$ and $R = 0.8$ (recall that $U(s,a) = \ln(s) - a^2/2$). The second-best contract from the principal's point of view (i.e., with no information on parameter p) is $s(x) = 1.002 + 0.498x$, which induces an effort of 5.458; the expected profit $E^P[x\text{-}s^*(x)]$ turns out to be equal to 1.750. Consider now that the agent has an additional piece of information, according to which $p = 2$; given that such information is not available for the principal, the contract issued by the principal remains, $s(x) = 1.002 + 0.498x$. Then, as pointed out by Proposition 4, the agent's best choice of effort increases, from 5.458 to 5.788, which is the agent's optimal choice (the effort which satisfies constraint (IC')); the agent's expected utility increases from $R = 0.8$ to 0.897. Finally, according to the agent's prior, the principal's expected profit is $E^A[x - s^*(x)] = 1.904$ ($> E^P[x - s^*(x)] = 1.750$; due to moral hazard, from realized output the principal is unable to infer that the agent has exerted an effort above 5.458). However, as pointed out by Corollary 1, the agent can improve by sharing his private information on output ($p = 2$), since then the principal will issue a contract which allows both parties to improve (whenever the agent can bargain on his reservation utility, which must be above the expected utility obtained if he hides his prior, 0.897).

14679957, 2026, 1, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/manc.70010 by Spanish Cochrane National Provision (Ministerio de Sanidad), Wiley Online Library on [18/05/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

The optimal contract when $p = 2$ is shared information turns out to be $s(x) = 0.648 + 0.584x$; we can observe that the new contract has stronger incentives to exert effort (the variable part of the salary is higher). The agent's best choice of effort increases to 6.271, reaching the minimum utility that he can accept is he shares his private information on output, 0.897 (obviously, the renegotiated reservation utility may be some higher in order to ensure that the agent is strictly better-off than hiding his information). Finally, the principal's expected profit improves to $E^P[x - s^*(x)] = 1.960$. This example illustrates that, if the agent can bargain on his reservation utility, a cheap talk communication can help attain the Pareto equilibrium in the agency game; then both parties improve over the situation in which the agent hides his private information on output and exert an effort above the one anticipated by the principal.

6 | Discussion and Conclusions

The analyses of the agency relationship under moral hazard are often confined to situations where principal and agent agree on the output distribution parametrized by the agent's effort. In rare cases, principal and agent may agree to disagree on the output distributions, and in even less, one of the parties form her/his beliefs using private information on output. In this paper we deal with this question using a novel approach to agency, where the output distribution -assumed to be known in the principal-agent model of Mirrlees-Holmstrom-is substituted by an output distribution built from the pieces of information that principal and agent have. In order to solve the problem derived from the lack of an output distribution, we appeal to the maximum-entropy principle to elucidate "the best" output distribution that can be inferred from the available information.

Information Theory Justifies the maximum-entropy distribution as the maximally noncommittal distribution of missing information (see Shannon 1948; Jaynes 1957). In the case that a constraint can be taken as inherent to the problem (e.g. the output variance or another moment is known), the MEP gives a natural way to incorporate such information into the mathematical specification of output. Using an output distribution other than a MED implies that either available information is being ignored, or non-available information is being used. As principal and agent can have different information or prior beliefs on output, they can disagree on the output distribution.

In the analysis presented, we obtain that diverging information sets for principal and agent can lead to non-standard conclusions in agency. In first place, additional information linking output to effort does not necessarily improve welfare (Proposition 1). As a consequence, if the principal has information on the output distribution that the agent ignores, she may find it optimal to hide it and design the agency contract ignoring the additional information (Proposition 2). In a different context, Weinschenk (2010) derives a similar conclusion, concluding that "the famous Informativeness Principle does not hold: it can be optimal for the principal to ignore an informative but ambiguous performance measure and instead provide incentives solely via other measures". According to Weinschenk (2010), the implicit assumption that the agent knows precisely the statistical properties of the performance measure is too strong (and unrealistic). However, and contrary to Weinschenk (2010), we must point out that Proposition 2 does not contradict Holmstrom's Informativeness Principle, since an additional piece of information is different from a "valuable signal". In the present approach, an additional piece of information can inform the principal that output is more dispersed (than it is when ignoring the additional information), which penalizes the welfare creation due to the agent's risk-aversion. Then, the principal could be better-off by ignoring such information when designing the agency contract. This result relates to the Myerson's inscrutability principle.

On the other hand, if the agent has an additional piece of information over the principal, the contract might not be incentive-compatible: the agent could improve by exerting a higher effort than the second-best effort anticipated by the principal according to her beliefs, thereby reaching an expected total utility above the agent's reservation utility (Proposition 4). As this outcome is Pareto-dominated, the agent has incentives to use a cheap talk to communicate the principal his prior on output (whenever he can renegotiate his reservation utility). Then the principal will ascribe to the agent's prior, since an agency contract designed according to the agent's prior allows her to improve (Corollary 1). In the particular case where the agent holds additional information about output, according to which it follows a gamma distribution, and his utility is $u(s, a) = 2\sqrt{s} - ka^2/2$, the additional piece of information has not effect on the game outcome. In such case the private information on output does not alter the incentives constraint, so the agent chooses the level of effort anticipated by the principal (Proposition 3). So, just as Proposition 3 provides a particular agency model in which different priors on output have no consequences on the second-best equilibrium, supporting Karni (2008), Proposition 4 contradicts Karni's result in the sense that it is the principal who ascribes to the agent's priors.

Endnotes

[1] Karni (2008, Sec. 5) represents an exception. He provides a rationale for the foundations of the parameterized distribution formulation of agency theory.

[2] Differences in beliefs are used in contract theory to endogenize different modes of organization; see Van den Steen (2005), Giat et al. (2010), Jung and Subramanian (2014), Alonso and Câmara (2016), Gutiérrez and Ortín-Ángel (2016), and others. In the principal-agent model, the roles of employer and employee are pre-assigned (appealing, e.g., to wealth constraints).

[3] When the agent's reservation utility is sufficiently low, the FOA can fail; see Gutiérrez (2012, 2024) and Moroni and Swinkels (2014). We rule out such situations.

[4] Information Theory justifies the maximum-entropy distribution as the least biased (or maximally noncommittal) distribution given missing information; see Shannon (1948) and Jaynes (1957).

[5] The *entropy* $H(\cdot)$ of a distribution X can be seen as a measure of the information gained when uncertainty is unraveled. $H(X) = E[I(x)]$, with $I(x)$: $= -x \cdot \ln(x)$ denoting the information conveyed by x. As *entropy* can be understood as "lack of information", it may sound odd that the entropy of a distribution increases with an additional piece of information.However, is not difficult to find examples where an additional constraint leads to a MED with higher entropy than the

14679957, 2026, 1, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/manc.70010 by Spanish Cochrane National Provision (Ministerio de Sanidad), Wiley Online Library on [13/05/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

exponential distribution (which is the MED obtained by ignoring the third constraint). For example, the entropy of the half-normal distribution is higher than the entropy of the exponential distribution if the standard deviation of the first one is sufficiently large (the half-normal distribution is the MED obtained by imposing that the two first distribution moments are known). In most situations such "anomaly" does not arise, and an additional piece of information does decrease entropy (the gamma distribution provides an example).

[6] We can prove such a claim for the gamma distribution: if output follows a gamma distribution with parameters p and a/p, the expected output is a and the variance is $p(a/p)^2 = a^2/p$. Then, the lower the shape parameter p the higher the variance (i.e., the output is more dispersed). The effect of a lower p on output is assesed through (IC'): $\int_0^{+\infty} U(s(x)) f_a(x;a)dx - C'(a) = \int_0^{+\infty} U(s(x))\left(p/a^2\right)(x-a)f(x;a)dx - a = 0$ (taking $C(a) = a^2/2$): a lower p (so, a more dispersed output) implies a lower effort (for the (IC') to hold), which potentially reduces welfare (in the Pareto sense).

[7] I thank an anonymous referee for suggesting this interpretation.

References

Al-Najjar, N. I., and J. Weinstein. 2009. "The Ambiguity Aversion Literature: A Critical Assessment." *Economics and Philosophy* 25, no. 3: 249–284. https://doi.org/10.1017/s026626710999023x.

Alonso, R., and O. Câmara. 2016. "Bayesian Persuasion With Heterogeneous Priors." *Journal of Economic Theory* 165: 672–706. https://doi.org/10.1016/j.jet.2016.07.006.

Carrasco, V., V. F. Luz, N. Kos, M. Messner, P. Monteiro, and H. Moreira. 2018. "Optimal Selling Mechanisms Under Moment Conditions." *Journal of Economic Theory* 177: 245–279. https://doi.org/10.1016/j.jet.2018.05.005.

Chade, H., and J. Swinkels. 2020. "The No-Upward-Crossing Condition, Comparative Statics, and the Moral Hazard Problem." *Theoretical Economics* 15, no. 2: 445–476. https://doi.org/10.3982/te2937.

Ellsberg, D. 1961. "Risk, Ambiguity, and the Savage Axioms." *Quarterly Journal of Economics* 75, no. 4: 643–669. https://doi.org/10.2307/1884324.

Galperti, S. 2015. "Common Agency With Informed Principals: Menus and Signals." *Journal of Economic Theory* 157: 648–667. https://doi.org/10.1016/j.jet.2015.01.018.

Giat, Y., S. T. Hackman, and A. Subramanian. 2010. "Investment under Uncertainty, Heterogeneous Beliefs, and Agency Conflicts." *Review of Financial Studies* 23, no. 4: 1360–1404. https://doi.org/10.1093/rfs/hhp096.

Gilboa, I., and D. Schmeidler. 1989. "Maxmin Expected Utility With Non-Unique Prior." *Journal of Mathematical Economics* 18, no. 2: 141–153. https://doi.org/10.1016/0304-4068(89)90018-9.

Gutiérrez, Ó. 2012. "On the Consistency of the First-Order-Approach to Principal-Agent Problems." *Theoretical Economics Letters* 2, no. 2: 157–161. https://doi.org/10.4236/tel.2012.22028.

Gutiérrez, Ó. 2024. "On the Range of Applicability of the First-Order Approach to Agency Problems." *Annals of Economics and Finance* 25, no. 2: 649–662. http://down.aefweb.net/AefArticles/aef250208Gutierrez.pdf.

Gutiérrez, O., and P. Ortín-Ángel. 2016. "Entrepreneurship and the Legal Form of Businesses: The Role of Differences in Beliefs." *Review of Law & Economics* 12, no. 1: 119–151. https://doi.org/10.1515/rle-2014-0027.

Gutiérrez, Ó., and V. Salas-Fumás. 2021. "Agency Contracts Under Maximum-Entropy." *Entropy* 23, no. 8: 957. https://doi.org/10.3390/e23080957.

Hansen, L. P., and T. J. Sargent. 2001. "Robust Control and Model Uncertainty." *American Economic Review* 91, no. 2: 60–66. https://doi.org/10.1257/aer.91.2.60.

Holmstrom, B. 1979. "Moral Hazard and Observability." *Bell Journal of Economics* 10, no. 1: 74–91. https://doi.org/10.2307/3003320.

Holmstrom, B., and P. Milgrom. 1987. "Aggregation and Linearity in the Provision of Intertemporal Incentives." *Econometrica* 55, no. 2: 303–328. https://doi.org/10.2307/1913238.

Jaynes, E. T. 1957. "Information Theory and Statistical Mechanics." *Physical Review* 106, no. 4: 620–630. https://doi.org/10.1103/physrev.106.620.

Jewitt, I. 1988. "Justifying the First-Order Approach to Principal-Agent Problems." *Econometrica* 56, no. 5: 1177–1190. https://doi.org/10.2307/1911363.

Jung, H. W., and A. Subramanian. 2014. "Capital Structure Under Heterogeneous Beliefs." *Review of Finance* 18, no. 5: 1617–1681. https://doi.org/10.1093/rof/rft042.

Karni, E. 2008. "Agency Theory: Choice-Based Behavioral Foundations of the Parametrized Distribution Formulation." *Economic Theory* 36, no. 3: 337–351. https://doi.org/10.1007/s00199-007-0277-9.

Kellner, C. 2017. "The Principal-Agent Problem With Smooth Ambiguity." *Review of Economic Design* 21, no. 2: 83–119. https://doi.org/10.1007/s10058-017-0198-4.

Klibanoff, P., M. Marinacci, and S. Mujerki. 2005. "A Smooth Model of Decision Making Under Ambiguity." *Econometrica* 73, no. 6: 1849–1892. https://doi.org/10.1111/j.1468-0262.2005.00640.x.

Lopomo, G., L. Rigotti, and C. Shannon. 2011. "Knightian Uncertainty and Moral Hazard." *Journal of Economic Theory* 146, no. 3: 1148–1172. https://doi.org/10.1016/j.jet.2011.03.018.

Moroni, S., and J. Swinkels. 2014. "Existence and Non-Existence in the Moral Hazard Problem." *Journal of Economic Theory* 150: 668–682. https://doi.org/10.1016/j.jet.2013.09.011.

Myerson, R. B. 1983. "Mechanism Design by an Informed Principal." *Econometrica* 51, no. 6: 1767–1797. https://doi.org/10.2307/1912116.

Shannon, C. E. 1948. "A Mathematical Theory of Communication." *Bell System Technical Journal* 27, no. 3: 379–423. https://doi.org/10.1002/j.1538-7305.1948.tb01338.x.

Van den Steen, E. 2005. "Organizational Beliefs and Managerial Vision." *Journal of Law, Economics, and Organization* 21, no. 1: 256–283. https://doi.org/10.1093/jleo/ewi011.

Weinschenk, P. 2010. "Moral Hazard and Ambiguity." *MPI Collective Goods Preprint, (2010/39)*.

Zellner, A., and R. A. Highfield. 1988. "Calculation of Maximum Entropy Distributions and Approximation of Marginal Posterior Distributions." *Journal of Econometrics* 37, no. 2: 195–209. https://doi.org/10.1016/0304-4076(88)90002-4.

Appendix

Proof of Proposition 1: We utilize the gamma distribution to provide examples where the introduction of additional information (coded through the third constraint in the information-constrained problem (P2)) reduces welfare in the Pareto sense (i.e., the additional information leads to a Pareto-dominated equilibrium). Depending on the nature of the information, however, the opposite can also occur. Here, it is assumed that principal and agent share identical prior beliefs (so the output MEDs for principal and agent are identical).

Consider a gamma distribution with parameters p and a/p (if p is an integer number, output may consist of the aggregation of a number p of i. i.d. exponential random variables whose expected values are a/p). This distribution is the solution to (P2) with $E[\ln(X)]$ constrained to be a

14679957, 2026, 1, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/manc.70010 by Spanish Cochrane National Provision (Ministerio de Sanidad), Wiley Online Library on [13/05/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

constant (which may be dependent on p and a). The density is $f(x) = \frac{1}{(a/p)^p \Gamma(p)} x^{p-1} e^{-(p/a)x}$; the parameterization ensures that $E(X) = a$, and the likelihood ratio is $f_a/f = -p/a + (p/a^2)x$. Assume that the agent's utility is $u(s, a) = 2\sqrt{s} - a^2/2$ with reservation utility $R = 1.2$. Consider first that the only available information on output is $E(X) = a$ (i.e., output is exponentially distributed, which corresponds to $p = 1$). In this case the second-best contract is $s(x) = (0.456 + 0.379x)^2$, the second-best effort is $a^* = 0.758$, and the principal's expected profit is $E^P[x-s(x)] = 0.123$. Consider instead that $p = 1.5$ (in this case the density is inversely U-shaped); as $p \neq 1$, additional information on X is now available. The second-best contract is now $s(x) = (0.431 + 0.411x)^2$, it induces a second-best effort $a^* = 0.821$, and the principal's expected profit is $E^P[x\text{-}s(x)] = 0.155$. So, the new equilibrium improves the former one, where no additional information on X was available. This illustration shows the typical case where additional information contributes to enhance welfare.

Consider now that $p = 0.9$. The optimal agency contract is now $s(x) = (0.463 + 0.371x)^2$, the induced level of effort is $a^* = 0.742$, and the principal's expected profit is $E^P[x-s(x)] = 0.114$. Therefore, an additional piece of information (which in this case also entails a reduction of the distribution entropy, see Section 3) is compatible with a welfare loss.

Proof of Proposition 2: Let us create an example showing the result. We recall that the expectations for principal and agent are calculated from the corresponding output distributions, which are built upon the corresponding priors. Consider that the information constraints in the principal's problem (P2) includes $E^P[\ln(X)] = \text{constant}$ (the superscript denotes that the expectation is calculated according to the principal's prior). Hence, and according to principal's beliefs, output follows a gamma distribution. The output density is $f^P(x) = \frac{1}{(a/p)^p \Gamma(p)} x^{p-1} e^{-(p/a)x}$, which leads to $E^P(X) = a$ and $f_a^P/f^P(x; a) = (p/a^2)(x - a)$. The net agent's utility is $u(s, a) = \ln(s) - ka^2/2$, so the second-best optimal agency contract is $s^*(x) = \alpha + \beta x$ (obtained from the FOA).

Next, we construct an example where the principal strategically ignores information when designing the contract. Take $p = 0.9$, $k = 0.02$ and $R = 0.8$. For $p < 1$ the gamma density is strictly decreasing and unbounded from above (at 0). The optimal agency contract is $s(x) = 1.073 + 0.489x$, the induced level of effort is $a = 5.301$, and the principal's expected profit is $E^P[x\text{-}s(x)] = 1.659$. If the principal leaves aside her information on p when designing the contract (so the third constraint is removed from (P2), and the only information available to both principal and agent is $E^P(X) = E^A(X) = a$), the corresponding output MED becomes the exponential distribution (gamma with $p = 1$), with density function $f(x) = (1/a)e^{-x/a}$. In this case the principal would issue the agency contract $s(x) = 1.02 + 0.498x$, which induces an effort of $a = 5.458$ (> 5.301). The principal's expected profit is calculated with the contract actually issued, $s(x) = 1.02 + 0.498x$, but using the gamma distribution with $p = 0.9$ (principal's private information) and the effort $a = 5.458$ (since the agent ignores the piece of information $p = 0.9$, so he chooses the effort for an exponentially distributed output). We obtain $E^P[x-s(x)] = 1.750$ (> 1.659). The expected profit thus improves by hiding the information on output, $p = 0.9$. Hence, sharing an additional piece of information is compatible with a welfare loss, and the principal would then act optimally by hiding it from the agent.

Proof of Proposition 3: The output density according to agent's beliefs is $f^A(x) = \frac{1}{(a/p)^p \Gamma(p)} x^{p-1} e^{-(p/a)x}$. For the principal, output follows an exponential distribution (by the MEP); if $u(s, a) = 2\sqrt{s} - ka^2/2$ and the FOA holds, the second-best agency contract is quadratic in output, $s(x) = (\alpha + \beta x)^2$. The (IC') is: $\int_0^{+\infty} 2\sqrt{(\alpha + \beta x)^2} f_a(x; a)dx - ka = 2\beta \int_0^{+\infty} x f_a(x; a)dx - ka = 0$ (superscripts are omitted when no confussion is possible). The constant parameter α has disappeared from (IC'); this occurs because $\int_0^{+\infty} f_a(x; a)dx = (\partial/\partial a) \int_0^{+\infty} f(x; a)dx = 0$ (the mass of probability of the output MED must be one, so it does not depend on a).

Taking into account that the first two moments of a gamma distribution of parameters p and a/p are $E(X) = a$ and $E(X^2) = a^2/p + a^2$, simple operations lead to:

$$\int_0^{+\infty} x f_a(x; a)dx = \int_0^{+\infty} x(f_a/f) f(x; a)dx = \int_0^{+\infty} x(p/a^2)(x - a) f(x; a)dx$$
$$= (p/a^2)E^A(X^2) - (p/a^2)aE^A(X) = 1.$$

Hence (IC') reduces to $2\beta\text{-}ka = 0$, with β independent of parameter p as it is chosen by the principal (who ignores it). So the ex ante private information of the agent (p) has no effect on (IC'), which implies that the agent chooses a level of effort that coincides with the effort anticipated by the principal, $a_A{}^* = 2\beta/k$.

Now we have to show that the participation constraint (PC) in problem (P1) holds. The PC is: $\int_0^{+\infty} 2\sqrt{(\alpha + \beta x)^2} f(x; a)dx - ka^2/2 \geq R$; operating it can be rewritten as $2\alpha + 2\beta a - ka^2/2 \geq R$, which coincides with the (PC) used by the principal in the contract design (observe that parameter p is absent from the inequality). Consequently, the contract designed by the principal without information on p ensures that the agent reaches his reservation utility: (PC) holds. We conclude that the agent's private information on the output distribution has no effect on the equilibrium of the agency game. □

Proof of Proposition 4: According to the principal the output MED follows an exponential distribution; according to this belief, the optimal contract corresponding to utility $u(s, a) = \ln(s) - a^2/2$ adopts the form $s^*(x) = \alpha + \beta x$, with α and β constants that depend on the multipliers λ and μ (see Section 2); the agent's effort anticipated by the principal is denoted by a_P. According to the agent, the output MED follows a gamma distribution (of parameters p and a/p). Condition (IC') implies:

$$\int_0^{+\infty} \ln(s(x)) f_a^A(x; a)dx - C'(a)$$
$$= \int_0^{+\infty} \ln(\alpha + \beta x)(p/a^2)(x - a) f^A(x; a)dx - ka = 0$$

(recall that the likelihood ratio is $f_a/f = (p/a^2)(x - a)$, and that superscripts are omitted if no confussion is possible). Take $p > 1$; then the agent improve by exerting an effort a_A above the level of effort anticipated by the principal a_P (which corresponds to $p = 1$). Inequality $a_A > a_P$ is due to that the agent must exert an effort higher than a_P for condition (IC') to hold, since the integral $\int_0^{+\infty} \ln(\alpha + \beta x)(p/a^2)(x - a) f(x; a)dx$ increases with p, so that (IC') holds for an effort value $a = a_A$ higher than a_P. So, the equilibrium of the game anticipated by the principal is not realized. Given that the agent's expected utility equals the reservation utility R if $p = 1$ and $a = a_P$, inequalities $a_A > a_P$ and $p > 1$ imply that his net utility will exceed R. Finally, the effort $a_A > a_P$ also entails a higher principal's expected profit $E[x-s(x)]$ because if the FOA holds, the principal would prefer to see the effort increase (see Holmstrom (1979, Prop. 1) or Jewitt (1988, 1180)).

14679957, 2026, 1, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/manc.70010 by Spanish Cochrane National Provision (Ministerio de Sanidad), Wiley Online Library on [13/05/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License