RATIONALITY & COMPETITION
CRC TRR 190

Pay Transparency in Organizations

Amir Habibi (HU Berlin)

Discussion Paper No. 395

May 10, 2023

Collaborative Research Center Transregio 190 | www.rationality-and-competition.de
Ludwig-Maximilians-Universität München | Humboldt-Universität zu Berlin
Spokesperson: Prof. Georg Weizsäcker, Ph.D., Humboldt University Berlin, 10117 Berlin, Germany
info@rationality-and-competition.de

Pay Transparency in Organizations

Amir Habibi*

March 15, 2023

Abstract

I study when a firm prefers to be transparent about pay using a simple multidimensional signaling model. Pay transparency within the firm means that a worker can learn about his own worker-firm match from another worker's pay. This can either encourage or discourage workers—which affects retention—and so creates a trade-off for the firm when it commits to a level of transparency. The model predicts that when few workers have a high worker-firm match, transparency is always preferred by the firm and becomes more favorable as the value of retaining these 'star' workers increases. This prediction is consistent with the firms in the field that choose to be internally transparent about pay. The model also predicts that transparency leads to pay compression, again consistent with evidence from the field.

Keywords: pay transparency, bonus pay, multidimensional signaling, relative pay

JEL Classification: D82, D86, J30, M52

1 Introduction

Private organizations are increasingly committing to pay transparency—making each worker's pay observable to all others within the firm.[1] Pay transparency is particularly

*Humboldt University of Berlin. Email: amir.habibi@hu-berlin.de. I am indebted to Martin Cripps and Konrad Mierendorff for their encouragement and supervision, to Vasiliki Skreta for guidance and support, and to Peter Spittal for numerous insightful conversations. I would also like to thank Guo Bai, Anna Becker, Helmut Bester, Antonio Cabrales, Hugo Freeman, Antonio Guarino, Nathan Hancart, Philippe Jehiel, Ian Jewitt, Attila Lindner, Albert Marcet, James Maynard, Fabien Postel-Vinay, Ian Preston, Nikita Roketskiy, Anja Schöttner, two anonymous referees and audience members at UCL, Carlos III de Madrid, ENTER Jamboree 2019, RES Symposium of Junior Researchers 2019, Oxford, and Humboldt University of Berlin for their helpful comments. I am also grateful for financial support from UCL through the Ricardo Scholarship and from the German Research Foundation (DFG) through CRC TRR 190 (Project number 280092119).

[1] In a survey of 715 UK firms, 18% reported an increase in disclosure of pay outcomes between 2015 and 2017, while almost none reported a decrease (CIPD (2017)).

prevalent in industries in which workers' performance is subjective and there is a high degree of heterogeneity in the quality of the worker-firm match—this includes the technology industry and creative industries such as advertising. In a recent survey of 400 US advertising and marketing hiring managers, 77% reported their organization offers some level of salary transparency, with 34% reporting full transparency.[2] One rationale for choosing to be transparent is that a firm wants to help manage the retention of its best workers. For example, the social analytics platform SumAll has committed to be transparent about the pay of its workforce internally with the following rationale:

> *When Dane Atkinson started [...] SumAll [...], he too was looking for a way to attract and retain talented people. Informed by two decades of experience as a serial entrepreneur, board member, advisor and executive, he was also trying to mitigate several factors that contributed to high turnover of staff at other companies.*[3]

This is echoed in a recent Harvard Business Review article discussing the costs and benefits of pay transparency:

> *[Pay Transparency] facilitates attracting and retaining talent.*[4]

A benefit of committing to transparency is that it gives an employer credibility, enabling her to demonstrate that she is treating everyone in the same way. If a worker sees that not only did he not receive a bonus (or equivalently, a pay rise), but also that none of his peers received a bonus, he will infer that it was more likely that the employer was not able to pay anyone a bonus because funds were not available. On the other hand, a downside of transparency is that it enables workers to compare themselves to others. In particular, when a worker sees that he was paid less than a peer, he will infer that his employer values him less, meaning he will become discouraged and potentially leave the firm or exert less effort. The contribution of this paper is to propose a model that formalises this trade-off and to use it to determine when transparency or no transparency is more favorable for a firm. Understanding what tradeoffs a firm faces when making decisions about transparency is important both to inform managers, and also to inform policy makers who mandate pay transparency rules.[5]

[2] www.roberthalf.com/blog/management-tips/how-much-do-you-make. Last accessed 14/09/2020.

[3] www.theguardian.com/business/2015/jul/10/salary-wage-glassdoor-payscale-buffer-sumall. Last accessed 08/27/2019.

[4] Lam et al. (2022).

[5] There are a number of other factors that may affect a firm and policy makers' decisions on pay transparency that I do not analyse. These include discrimination and the gender pay gap (Baker et al. (2019)), and public aversion to high pay in the public sector (Mas (2017)). Related to this, in the United States there have been recent mandates that firms must disclose salary ranges in job advertisements. https://www.cnbc.com/2023/01/03/where-us-companies-have-to-share-salary-ranges-with-workers-by-law.html Last accessed 03/07/2023.

In the model a principal employs two agents. At the start of the game, the principal commits to make bonuses either transparent or not transparent. After this, the principal privately learns each agent's match quality with the firm (from now on I refer to this as 'productivity'). The principal wants to encourage all agents to stay at the firm rather than take an outside offer. A more productive agent produces a greater surplus if he stays at the firm, and the additional surplus is shared such that both the principal and the agent enjoy greater benefits.[6] The principal also privately learns whether she is able to pay bonuses—variation in this may be due to a lack of funds or the opportunity cost of investing elsewhere. The principal uses discretionary bonuses to signal to more productive agents that they have good prospects at the firm and that they will benefit from staying.

The principal's choice of transparency affects what an agent learns (or does not learn) about his own productivity from the other agent. Under no transparency, in equilibrium, the principal pays a bonus only when she is able to do so and when the agent has high productivity.[7] The agent has just one piece of information when updating his belief about his productivity. If he is paid a bonus he learns for sure he has high productivity, while if he is not paid a bonus he becomes pessimistic. Under transparency, in equilibrium, the principal uses different strategies depending on the parameters of the model. The key tension occurs when one agent has high productivity and the other low. If the principal pays only the high productivity agent a bonus, the low productivity agent becomes pessimistic. This is not so damaging if either the value of retaining the high productivity agent is much higher than retaining the low productivity agent, or if paying neither agent a bonus leaves both agents relatively pessimistic. When these are satisfied, then in equilibrium, the principal pays just one agent a bonus. Transparency is always optimal for the principal. Furthermore, as the value of retaining a high productivity agent increases, the difference in the value of transparency and no transparency increases. When the conditions above are not met, the equilibrium is qualitatively different. The principal only pays bonuses when *both* agents have high productivity. Intuitively when only one agent has high productivity, making low productivity agents pessimistic outweighs the benefit of making the high productivity agent more optimistic by paying him a bonus. Now it is possible for no transparency to be optimal. This happens when either the value of retaining high productivity agents is relatively high (but not too high to revert to the previous case), or the likelihood of an agent having high productivity (the prior) is high.

To summarise, transparency is the optimal choice for the principal when the number of 'star' workers with high productivity is relatively low or the value of retaining these workers is high (compared to the value of retaining less productive workers). This is consistent with the high levels of transparency reported in creative industries discussed

[6]The model is static and so these continuation values exogenously take this form. In Section 4.3, I discuss a dynamic version of the model in which these values arise endogenously.

[7]As described more precisely below I focus on the (unique) equilibrium that is selected by the relevant selection criteria.

above—where star workers are both valuable and scarce. Furthermore, these industries match other key features of the model such as firms having an informational advantage over workers about their productivity (subjective evaluations).[8]

The model also predicts that there will always be pay compression under transparency. More precisely given the productivity of the agents and the principal's cost shocks, bonuses will be lower compared to under no transparency. Previous empirical work has documented pay compression resulting from both transparency and how workers compare themselves to peers. I discuss the related literature in Section 5.

2 Model

In this section I begin by describing the model. After describing the model, I discuss the modelling assumptions and their connection to the literature.

2.1 Set-up

Players. There is a principal (she) and two agents (he), indexed by $i = 1, 2$.

Information. Each agent's productivity (or worker-firm match quality) is given by $\theta_i \in \{H, L\}$. The productivity of each agent is independently drawn with the probability of high productivity given by $\Pr[\theta_i = H] = p_0 \in (0, 1)$. The principal faces uncertainty on his ability to pay bonuses, this is given by $B \in \{0, \bar{B}\}$, where $B = 0$ is the 'high cost state' in which paying bonuses is too costly. The prior probability of the 'low cost state' is given by $\Pr[B = \bar{B}] = q \in (0, 1)$. I make the assumption that $\frac{1}{2}\bar{B} \geq g_H^P$ (where g_H^P is defined below).[9] Each agent receives an outside option—this is drawn independently from $u_i \sim U[0, 1]$. At the start of the game the productivity of each agent, the marginal cost of paying bonuses, and the outside option of each agent is drawn independently and unknown to all players who share a common prior.

Actions and timing.

1. The principal decides on a level of transparency. Denote this decision by $a^P \in \{N, T\}$ where N and T represent no transparency and full transparency.

2. The principal *privately* learns the productivity of the agents (θ_1, θ_2) and her ability to pay bonuses (B).

[8] The key features of the model and testable implications are discussed in Section 6.

[9] As will become clear below, this means that in the low cost state the principal can always pay bonuses that are sufficiently high to support all possible types of signaling equilibria.

3. The principal chooses whether or not to pay each agent a bonus, b_i subject to her budget constraint $b_1 + b_2 \leq B$. If $a^P = N$ (no transparency) agent i only learns b_i, while if $a^P = T$ (transparency) agents learn both b_1 and b_2.
4. The agents learn their outside options u_i.[10]
5. The agents simultaneously choose whether to stay at the firm or to quit. Denote this decision by $a_i^A \in \{S, Q\}$.
6. The players receive their payoffs that are given below.

Beliefs and strategies. The principal's strategy is to choose a level of transparency $a^P \in \{N, T\}$ in the initial node. If she (privately) learns that $B = 0$, her choice of action is degenerate: she chooses $b_1 = b_2 = 0$. If she (privately) learns that $B = \bar{B}$ and that (θ_1, θ_2), given her choice of transparency, she chooses a distribution over bonuses

$$\sigma : \{N, T\} \times \{H, L\}^2 \to \Delta(\mathcal{B}).$$

The set of possible bonuses that the principal can pay is assumed to be $b_i \in 0 \cup [\epsilon, \bar{B}]$ with $\epsilon > 0$. I consider the limit as $\epsilon \to 0$. This is a technical assumption to rule out strictly positive bonuses being arbitrarily small. As will become clear below this ensures existence of an equilibrium that survives the selection criterion that I use (otherwise the set of equilibrium surviving the selection criterion is empty). $\mathcal{B}$ is the set of bonuses the principal can pay the two agents subject to the budget constraint:

$$\mathcal{B} \equiv \left\{(b_1, b_2) : b_1, b_2 \in 0 \cup [\epsilon, \bar{B}] \text{ and } b_1 + b_2 \leq \bar{B}\right\}.$$

Agent i updates his belief about his productivity (θ_i) and the principal's costs (B) following his own bonus b_i, and in the case of $a^P = T$, the other agent's bonus b_j.[11] He then chooses a quitting decision formally given by

$$a_i^A : \mathbb{R}_+ \times [0, 1] \to \{S, Q\},$$
$$a_i^A : \mathbb{R}_+^2 \times [0, 1] \to \{S, Q\},$$

in the case of no transparency and transparency respectively.[12]

Payoffs. The principal's payoff is given by

[10] The analysis does not rely on this being privately learned.

[11] Note that agent i may also make inferences about θ_j. However, this will never be relevant for his quitting decision.

[12] I assume that in the case of indifference the agent chooses to stay and so the agent will never play a mixed strategy. This will be without loss in equilibrium due the distribution of u_i having no atoms.

$$V = \sum_i -b_i + \mathbb{1}[a_i^A = S] g^P_{\theta_i},$$

where $g^P_{\theta_i}$ is the expected future surplus that the principal will earn from an agent with productivity θ_i. Assume that $g^P_H > g^P_L = 1$. This means that the principal wants to retain all agents, but prefers to retain agents with high productivity.

Agent i's payoff is given by

$$U_i = b_i + \mathbb{1}[a_i^A = S] g^A_{\theta_i} + \mathbb{1}[a_i^A = Q] u_i,$$

where $g^A_{\theta_i}$ is the expected future surplus if an agent of productivity θ_i stays at the firm. I assume $1 = g^A_H > g^A_L = 0$. This means that an agent prefers to stay if he has high productivity.

Equilibrium. The equilibrium concept is perfect Bayesian equilibrium.[13] In equilibrium, upon observing the bonuses, each agent i updates his belief about θ_i and B given the principal's strategy using Bayes rule. They then best respond given this updated belief and their outside options. The principal chooses a strategy $a^P \in \{N, T\}$ to maximise her expected payoff for the rest of the game given her strategy b_1, b_2 (once she learns θ_1, θ_2 and B) and the agents' best responses. After learning θ_1, θ_2 and B, the principal chooses (a distribution over) a pair of bonuses (b_1, b_2). This maximises her expected payoff given her choice of transparency, the beliefs this induces for the agents and the agents' corresponding best responses.

I also make some natural restrictions to the set of equilibria that I focus on. Following a^P and $B = \bar{B}$,[14] the strategy of the principal is a mapping from a pair of productivities (θ_1, θ_2) to (a distribution over) a pair of bonuses (b_1, b_2)

$$\{H, L\}^2 \to \Delta(\mathcal{B}).$$

Let

$$\sigma^{b_i b_j}_{\theta_i \theta_j} \equiv \Pr[b_1 = b_i, b_2 = b_j | \theta_1 = \theta_i, \theta_2 = \theta_j]$$

denote the (mixed) strategy of the principal and let σ be the vector of the principal's entire strategy. In order to provide sufficient conditions for the uniqueness result below, I restrict attention to *symmetric equilibria* so that each agent is treated in the same way (in expectation) given their productivity.

[13]Note that this does not pin down off-path beliefs and so there will possibly be multiple equilibria. I discuss later which equilibria I choose to focus on among those that are possible.

[14]$B = 0$ leads to a degenerate choice of action $(b_1, b_2) = (0, 0)$

Definition 1. *A* ***symmetric equilibrium*** *restricts* σ *so that*

$$\sigma_{\theta_i \theta_j}^{b_i b_j} = \sigma_{\theta_j \theta_i}^{b_j b_i},$$

for any realisations of θ_i*,* θ_j*, and choice of* b_i *and* b_j*.*

I also restrict equilibria so that if one agent is of high productivity and the other agent is of low productivity, the principal must not pay a higher bonus to the low productivity agent.

Definition 2. *A* ***monotonic equilibrium****, is one in which* $b_i \geq b_j$ *if* $b_i = H$ *and* $b_j = L$*.*

In Section 3, I elaborate on which equilibria this rules out, and I argue why these equilibria are unrealistic.

2.2 Discussion of the model

Principal's informational advantage. A key feature of the model is that the principal has a better knowledge of the agents' productivity than the agents have themselves. In many organizational settings—for example, in professional services such as law or consultancy or in a technology start-up—junior employees have little experience and so are not able to evaluate their own ability as accurately as more senior and experienced employees. This assumption is in line with subjective (or private) evaluations, the relevance of this assumption is discussed extensively in the survey by Prendergast (1999). The model makes an extreme assumption that agents are completely uninformed. The results do not qualitatively change if the agents get noisy signals about their own productivity.

Uncertainty on the principal's ability to pay bonuses. There is uncertainty on the ability of the principal to pay bonuses and this is privately known by the principal. Uncertainty on the ability to pay bonuses is equivalent to uncertainty in the marginal cost of paying a bonus—I have assumed that in the high cost state, the marginal cost is very high so the principal cannot pay bonuses. All firms will have some uncertainty on the opportunity costs of paying bonuses. This will be particularly pertinent in smaller organizations or start-ups that are more likely to be cash constrained due to a shock or be in a position where they need to prioritise new projects—this corresponds to the marginal costs being so high that the firm cannot pay bonuses.[15] In terms of information, more junior employees won't necessarily have good knowledge of how the firm is performing, and consequently what funds are available to pay bonuses. Even if the firms accounts were made public, the management (principal) will almost certainly hold some private

[15]As discussed on p.59 in Bewley (1999), the most common way that firms react to financial distress is to freeze wages or reduce bonuses or raises.

information about possible future investment opportunities that will affect the opportunity cost of paying bonuses today. In a different setting, Li and Matouschek (2013) make similar assumptions on the uncertainty of the marginal cost of paying bonuses.[16]

Peer effects. Agents' productivities are independent by assumption. In other papers where agents learn about themselves through peers, correlation in agents' types are used to drive results. For example, this is the case in Battaglini et al. (2005) and Halac et al. (2017). Papers that include a contest where types are heterogenous among contestants (such as Ederer (2010)) may not have explicit correlation between peers, but still have the effect that the marginal benefit of effort depends on the type of peers. Although these assumption may be valid in some organizational settings, I have made the independence assumption in order to not obfuscate the channel in which agents learn in my model.

Future share of surplus. The parameters g_θ^P and g_θ^A are the expected future share of surplus received by the principal and agent when an agent with productivity θ stays at the firm. The assumption is that both parties are better off in the future when the agent is more productive which makes sense in any organizational setting. This is a key assumption in my model that allows for bonuses to act as signals. I also discuss a dynamic version of the model in order to provide micro-foundations for the assumptions made (see Section 4.3). The assumption $g_L^P > 0$ means that the principal always wants an agent to stay at the firm regardless of his productivity. This makes sense particularly if the cost of hiring a new worker is high. Setting $g_L^P = 1$ simplifies the analysis, but result but results do not qualitatively change if $g_L^P > 0$ and $g_L^P \neq 1$. Again, the assumptions that $g_H^A = 1$ and $g_L^A = 0$ simplify the analysis, but results do not qualitatively change if $1 > g_H^A > g_L^A > 0$.

Communication between agents. In the model agents do not share information with each other about their pay. This is consistent with evidence from the field—Cullen and Perez-Truglia (2022) find in a large commercial bank that although employees have a high willingness to pay for accurate information about the salary of their peers, they are not able to report them accurately. Obviously, in my model agents would benefit from learning the other agent's bonus. But since they do not benefit from sharing their own bonus, agents do not have an incentive to do this.[17]

[16]They study relational contracts in which the principal privately learns the cost of paying a bonus to the agent. They motivate uncertainty on the cost of paying workers from a well known case study describing the situation Lincoln's Electric faced following financial difficulty after expanding to foreign markets—Hastings (1999).

[17]If there is a small preference for privacy, agents would have a strict preference for not sharing their bonus.

3 Analysis

In this section I analyse the model and determine the optimal level of transparency for the principal highlighting the tradeoff between transparency and no transparency. I start by providing separate analysis for the two possible transparency decisions the principal can make, $a^P \in \{N, T\}$. With slight abuse of terminology, I will refer to these as the *no transparency game* and the *transparency game* throughout the rest of the paper. As is typical in signaling games, in each 'game' there will potentially be multiple equilibria. For the main part of the analysis I focus on the (unique) equilibrium that is selected by the appropriate equilibrium selection criterion. I use the D1 criterion (Cho and Kreps (1987)).[18] In the final part of this section I also discuss the full set of (non-selected) equilibria. I characterise which equilibrium is selected in each game selected for all parameters. I then compare payoffs between the transparency and no transparency games and provide comparative statics.

3.1 No transparency game

Definition 3. *In the no transparency game a **separating equilibrium** has the principal pay the following bonuses b_i to agent i:*

- $b_i = b_N > \epsilon$ *if* $\theta_i = H$ *and* $B = \bar{B}$,
- $b_i = 0$ *otherwise.*

The beliefs off the equilibrium path are as follows: if $b_i \in (\epsilon, b_N)$ then agent i has a belief that $\theta_i = L$ with probability 1, and if $b_i > b_N$ the agent i can have any belief.

In a separating equilibria, the principal uses the bonus $b_i = b_N$ to signal to an agent that he has high productivity. The agents have beliefs consistent with this strategy. An agent is more likely to stay at the firm following a bonus since he has higher beliefs. This then makes it worthwhile for the principal to only pay a bonus to a high productivity agent.

The beliefs of the agents are as follows. Taking the strategy of the principal as given above, the updated beliefs of agent i following realisations of b_i (on the equilibrium path) are

$$\Pr[\theta_i = H | b_i = b_N] = 1,$$

$$\Pr[\theta_i = H | b_i = 0] = p_N,$$

[18]For a text book definition see Fudenberg and Tirole (1991). The D1 criterion rules out equilibria with 'unrealistic' off-path beliefs. It is similar in spirit, but slightly weaker than the widely used intuitive criterion (also in Cho and Kreps (1987)). But as will become clear below, it is necessary to use the D1 criterion in my setting since the intuitive criterion has no bite.

where $p_N \equiv \frac{p_0(1-q)}{1-p_0q}$ is the belief of the agent following no bonus. These beliefs are illustrated in Figure 3 (in Section 3.3 where I compare beliefs in the no transparency game to those in the transparency game).

Next I consider the agent's best response given these beliefs and the principal's incentive constraints. Agent i's best response (on the equilibrium path) is given by

$$a_i^A = \begin{cases} S & \text{if } b_i = b_N\text{; or } b_i = 0 \text{ and } u_i \leq p_N, \\ Q & \text{otherwise.} \end{cases}$$

Note that the best response does not depend on the bonus payment of agent $j \neq i$ since agent i does not see this under no transparency—this means that I can consider the incentives for the principal to pay a bonus to each agent separately. The incentive constraints of the principal to ensure that she pays a bonus when $\theta_i = H$ and $B = \bar{B}$ and does not pay a bonus when $\theta_i = L$ and $B = \bar{B}$ are given by

$$-b_N + g_H^P \geq -0 + p_N g_H^P, \tag{3.1}$$

$$-0 + p_N \geq -b_N + 1, \tag{3.2}$$

respectively. Combining these gives[19]

$$1 - p_N \leq b_N \leq (1 - p_N)g_H^P. \tag{3.3}$$

Notice that if these constraints are satisfied, the principal not only has incentives to choose the appropriate bonuses from those on the equilibrium path, but also has no incentive to deviate to a bonus that is off the equilibrium path.

Define the lowest bonus that satisfies 3.3 as $b_N^* = 1 - p_N$.

Proposition 1. *In the no transparency game, the unique equilibrium surviving the D1 criterion is the pure strategy separating equilibrium with $b_i = b_N^*$.*

Omitted proofs can be found in Appendix A. In the proof I show that apart from the 'least cost' separating equilibrium with $b = b_N^*$, all equilibria do not satisfy the D1 criterion. The intuition for why this is the case is similar to the reason that only the least cost separating equilibrium is uniquely selected by the D1 criterion in the canonical Spence education signaling model (Spence (1973)). Pooling with pessimistic off-path beliefs is ruled out because a deviation to a higher bonus would be more likely to be for a 'high type' (principal observing a high productivity agent). This in turn means that 'low types' are ruled out and given this, the deviation becomes profitable for the high type meaning this equilibrium violates the selection criterion. Furthermore, a separating equilibrium

[19] The assumption that $\frac{1}{2}\bar{B} \geq g_H^P$ ensures that the principal can always pay bonuses in this range when $B = \bar{B}$.

with 'too high' a bonus is also ruled out. A lower bonus could only be profitable for a high type, thus ruling out off-path beliefs that support such an equilibrium.[20]

3.2 Transparency game

The transparency game is more complex to analyse than the no transparency game, and so to ease exposition I relegate some of the analysis to Appendix A.2. Next I state two assumptions. Whether or not the these are satisfied will be critical in terms of determining the type of equilibrium that is selected—and ultimately whether transparency or no transparency is optimal.

Assumption 1. $p_T \equiv \frac{p_0(1-q)}{1-2p_0q+p_0^2q} < \frac{1}{2}$.

Assumption 2. $g_H^P \geq \frac{p_T}{1-p_T}$.

Assumption 1 is satisfied if either p_0 is sufficiently low, so high productivity agents are uncommon, or q is sufficiently high, so high cost shocks are uncommon. Assumption 2 is satisfied if given p_0 and q (and consequently p_T), g_H^P is above a given cutoff (which is increasing in p_T). The different parameters that satisfy these assumptions are illustrated graphically in Figure 1 that depicts the results of Proposition 2 below.

I define a separating equilibrium as before. As will become clear below this only exists for some parameters.

Definition 4. *In the transparency game a* ***separating equilibrium*** *has the principal pay the following combination of bonuses:*[21]

- $(b_i, b_j) = (\bar{b}_T, \bar{b}_T)$ *(with* $\bar{b}_T > \epsilon$*) if* $\theta_i = \theta_j = H$ *and* $B = \bar{B}$,
- $(b_i, b_j) = (\underline{b}_T, 0)$ *(with* $\underline{b}_T > \epsilon$*) if* $\theta_i = H$, $\theta_j = L$ *and* $B = \bar{B}$,
- $(b_i, b_j) = (0, 0)$ *otherwise.*

The beliefs off the equilibrium path are as follows: if (b_i, b_j) *is such that* $\min\{b_i, b_j\} < \max\{\underline{b}_T, \bar{b}_T\}$, *then agents* i *and* j *both have beliefs that* $\theta_i = \theta_j = L$ *with probability 1, and otherwise agents can have any belief.*

[20] In fact, in the no transparency game, the least cost separating equilibrium would be selected even if the weaker refinement, the intuitive criterion (Cho and Kreps (1987)) was used. However, since in the transparency game the D1 criterion is required to make a meaningful selection, I use this for both games for consistency. The reason for this difference is as in the application of these refinements to the canonical education signaling models: with two types, the intuitive criterion is sufficient, while with more types, the D1 criterion is required (this is detailed on p.212 in Cho and Kreps (1987)).

[21] Note that it is possible to have such that $(b_i, b_j) = (0, \underline{b}_T)$ when $\theta_i = H$, $\theta_j = L$, however this is ruled out by the monotonicity assumption (Definition 2). Also note that as discussed in an earlier version of the paper Habibi (2020), using increases in wages in place of bonuses rules out such an equilibrium. Furthermore, note that this strategy cannot be an equilibrium in the no transparency game where agents don't observe each other's bonuses.

These equilibria are similar to the separating equilibria in the no transparency game—the principal pays an agent a positive bonus if and only if the agent has high productivity and the principal is able to do so. However notice that the bonus levels may depend on the productivity (and bonus) of the other agent.

Now, as in the no transparency analysis, I describe the beliefs of the agents. Taking the strategy of the principal as given above, the updated beliefs of agent i following realisations of (b_i, b_j) (on the equilibrium path) are

$$\Pr[\theta_i = H | b_i \in \{\underline{b}_T, \bar{b}_T\}] = 1 \text{ for any } b_j,$$

$$\Pr[\theta_i = H | b_i = 0, b_j = 0] = p_T,$$

$$\Pr[\theta_i = H | b_i = 0, b_j = \underline{b}_T] = 0.$$

Recall that $p_T \equiv \frac{p_0(1-q)}{1-2p_0q+p_0^2q}$, which is the belief of the agent following no bonus and having observed that the other agent also received no bonus. These beliefs are illustrated in Figure 4 (in Section 3.3). The principal's incentive constraints are detailed in Appendix A.2. Using these constraints, the lowest bonus levels that satisfy all constraints are[22]

$$\underline{b}_T^* = 1 - 2p_T,$$
$$\bar{b}_T^* = 1 - p_T.$$

Next I define a different class of equilibria that will be selected when all separating equilibria are ruled out by the D1 criterion.

Definition 5. *In the transparency game a* ***partially separating equilibrium*** *has the principal pay the following combination of bonuses:*

- $(b_i, b_j) = (\bar{b}'_T, \bar{b}'_T)$ *(with* $\bar{b}'_T > \epsilon$*) if* $\theta_i = \theta_j = H$ *and* $B = \bar{B}$*,*
- $(b_i, b_j) = (0, 0)$ *otherwise.*

The beliefs off the equilibrium path are as follows: if (b_i, b_j) *is such that* $\min\{b_i, b_j\} < \bar{b}'_T$*, then agents* i *and* j *both have beliefs that* $\theta_i = \theta_j = L$ *with probability 1, and otherwise agents can have any belief.*

In these equilibria, the principal only pays a bonus if *both* agents have high productivity and the principal has a low cost shock. This means that when only one agent has high productivity, the principal pays both agents no bonus. The reason is that by separating and revealing the productivity of each agent perfectly, the principal incurs such a great cost by making the low productivity agent pessimistic, she would prefer for both agents to become slightly pessimistic by paying no bonus.

[22]Note if Assumption 1 is satisfied it must be that $\underline{b}_T^* > 0$.

As before, I describe the beliefs of the agents. Taking the strategy of the principal as given above, the updated beliefs of agent i following realisations of (b_i, b_j) (on the equilibrium path) are

$$\Pr[\theta_i = H | b_i = \bar{b}'_T] = 1,$$

$$\Pr[\theta_i = H | b_i = 0] = p'_T,$$

where $p'_T \equiv \frac{1-qp_0^2}{(1-q)p_0+q(1-p_0^2)}$. The lowest cost bonus satisfying the principal's constraints is $\bar{b}'^*_T = \frac{1}{2}(1-p'_T)(g_H^P+1)$ (details are again in Appendix A.2).

Now I provide a result that characterises the equilibrium uniquely selected by the D1 criterion. The D1 criterion needs to be adapted since there are multiple 'receivers' (agents) as the refinement is only defined for games with a single sender and a single receiver. I provide a formal definition of the 'multi-receiver D1 criterion' adapted for my setting in Appendix B and argue why I think it is appropriate.

Proposition 2. *The multi-receiver D1 criterion uniquely selects the following equilibria:*

- *if Assumption 1 is satisfied, the separating equilibrium with $\underline{b}_T = \underline{b}^*_T$ and $\bar{b}_T = \bar{b}^*_T$ (case 1),*
- *if neither Assumption 1 nor Assumption 2 are satisfied, the partially separating equilibrium with $\bar{b}'_T = \bar{b}'^*_T$ (case 2),*
- *otherwise, the separating equilibrium with $\underline{b}_T = \epsilon$ and $\bar{b}_T = \bar{b}^*_T$ (case 3).*

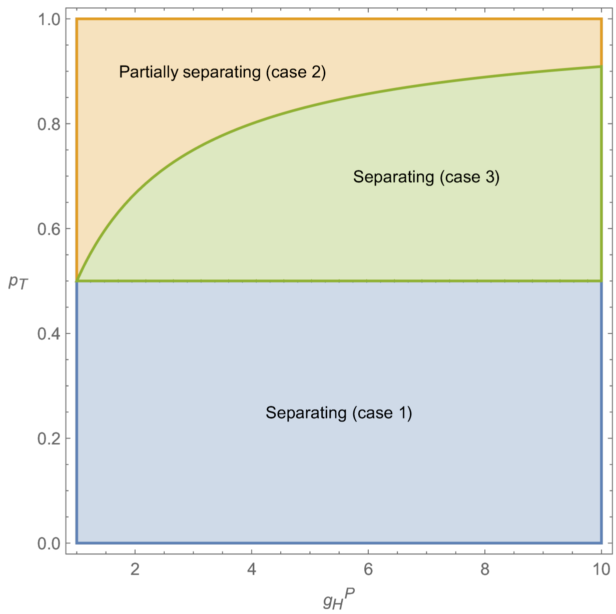

Figure 1: Selected equilibria under transparency.

Figure 1 illustrates the equilibria in each of the three cases above. In case 1, when Assumption 1 is satisfied, the principal is able to separate and pay positive bonuses to high productivity agents. As in the no transparency game, the D1 criterion selects the lowest cost bonus. A feature of this equilibrium is that the bonus is lower when only one agent is paid a bonus. This is because in this case the principal incurs the cost of making the low productivity agent more pessimistic and so does not need to incur such a high signaling cost through a bonus.

In case 2, when only one agent has high productivity, the principal cannot separate from when both agents have low productivity. Suppose there was such a separating equilibrium (as in case 1 above), then there would be a profitable deviation for the principal to pay no bonus and induce beliefs p_T for both agents. This is the case because the payoff from retaining the high productivity agent is sufficiently low (Assumption 2).

In case 3, the principal can separate when only one agent has high productivity. However, compared to case 1, the bonus paid to the high productivity agent here now has $\underline{b}_T^* < 0$. Thus to separate the principal must pay a strictly positive bonus, the lowest possible bonus is $b_i = \epsilon$. Recall that I consider the limit as $\epsilon \to 0$ so in effect the principal incurs no signaling cost, but each agent does learn their productivity rather than remain at the prior. This is the equilibria that requires the technical assumption of having a lowest strictly positive bonus ϵ. If this were not the case so the principal could pay a bonus $b_i \in (0, \epsilon)$, although paying a bonus $b_i = \epsilon$ would still be an equilibrium, because the principal could pay a lower bonus, this equilibrium would be eliminated by the D1 criterion by considering a deviation to $b_i \in (0, \epsilon)$.

3.3 Optimal choice of transparency

Now I turn to the economic question of interest: the principal's optimal choice of transparency, a^P. I provide two sets of results. First, I compare the absolute value of transparency and no transparency (Proposition 3). Second, I provide comparative statics on the difference between the value of transparency and no transparency (Proposition 4).

Proposition 3. *If either Assumption 1 or 2 is satisfied (so a separating equilibrium is selected in the transparency game) then transparency is optimal.*

If both Assumption 1 and 2 are not satisfied (so a partially separating equilibrium is selected in the transparency game), then under some parameters (g_H^P, p_0 and q), transparency is optimal for the principal and for others no transparency is optimal.

Define the difference between the expected value of the transparency game and the no transparency game for the principal as

$$D_{TN} \equiv \mathbb{E}V(T) - \mathbb{E}V(N).$$

Proposition 4. *If either Assumption 1 or 2 is satisfied, increasing the value of retaining a high productivity agent (g_H^P) increases D_{TN}.*

If both Assumption 1 and 2 are not satisfied, increasing the value of retaining a high productivity agent (g_H^P) decreases D_{TN}.

I graphically illustrate Proposition 3 in Figure 2 for the prior on the low cost state fixed at $q = .5$. The region where no transparency is optimal is characterised in Appendix **??**. Roughly speaking this has high values of p_0 (so most agents have high productivity) and low values of g_H^P (so retaining high productivity agents is not so valuable for the principal).

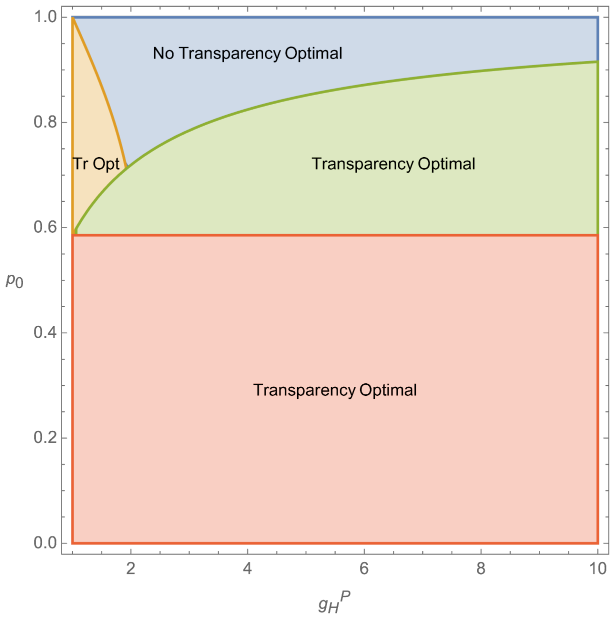

Figure 2: Optimal choice of transparency for prior probability of high productivity agent p_0 and value of retaining high productivity agent g_H^P. Probability of low cost state is fixed at $q = .5$.

I start by discussing the first parts of Propositions 3 and 4 where there is a separating equilibrium selected in each game. It is instructive to begin by comparing how the distribution of posterior beliefs for different productivity agents differs in each game, these are depicted in Figures 3 and 4. The arrows indicate the distribution of posterior beliefs of the agents. In Figure 3, for the no transparency game, an agent either perfectly learns he has high productivity or becomes pessimistic with belief p_N. In Figure 4, for the transparency game, the distribution of beliefs depends on the bonus given to the other agent. The agent still perfectly learns that he has high productivity after a bonus, but now how pessimistic he becomes after no bonus depends on whether the other agent is paid a bonus. A key difference between transparency and no transparency is what happens when neither agent is paid a bonus. In the transparency game, an agent that is not paid a bonus and observes the other agent is also not paid a bonus becomes *less*

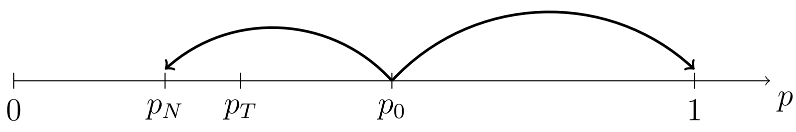

Figure 3: **Posterior beliefs of agent i under no transparency**. *Posterior beliefs do not depend on the other agent's bonus.*

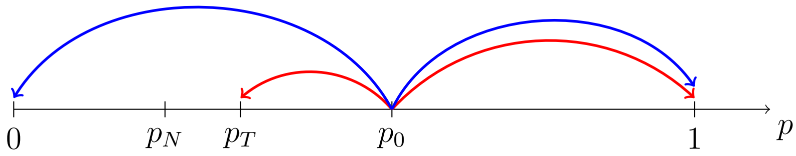

Figure 4: **Posterior beliefs of agent i under transparency in a separating equilibrium**. *Posterior beliefs depend on the other agent's bonus: the blue line represents when $b_j > 0$ and the red line represents when $b_j = 0$.*

pessimistic under transparency than when not paid a bonus under no transparency—the agent understands that it is more likely that the firm had a high cost shock ($B = 0$) in this scenario under transparency. This 'trust effect' is formally captured by the fact $p_T > p_N$, and means that when the principal cannot pay agents a bonus ($B = 0$), they are less likely to quit the firm.

Propositions 1 and 2 have that the principal pays higher bonuses under no transparency than under transparency when both agents are paid a bonus. Furthermore, under transparency there is a lower bonus if only one agent is paid a bonus. The intuition for the differences in bonus are as follows. Under transparency when the principal pays only one agent a bonus, because by discouraging the agent not paid a bonus, she incurs an additional cost compared to no transparency. This means she does not need to incur such a high signaling cost through the bonus to signal to the high productivity agent. Even when paying a bonus to both agents, the principal does not need to incur such a high signaling cost compared to no transparency. The reason is that if she deviates and does not pay a bonus to either agent then the agents are not as discouraged as under no transparency (because $p_T > p_N$).

To understand why transparency is always optimal consider different states of the world. When only one agent has high productivity (and the principal can pay bonuses), under transparency, the low productivity agent becomes more pessimistic. However, the principal's lower cost of bonuses under transparency offset the downside of the low productivity agent being more likely to quit. Transparency is advantageous in all other states of the world: First, when both agents have high productivity (and the principal can pay bonuses), the agents are equally optimistic under transparency and no transparency, however the principal has lower bonus costs under transparency; and second, when the principal does not pay bonuses (either because she cannot or both agents have low productivity), the agents are less pessimistic under transparency again benefiting the principal.

When there is a separating equilibrium, the comparative statics on the value of retaining high productivity agents provides a sharp prediction—transparency becomes relatively more favourable as the value of retaining high productivity agents goes up (g_H^P). The reason for this when the principal cannot pay high productivity workers a bonus ($B = 0$), they are more likely to be retained under transparency, and this probability only depends on the priors. Increasing the value of retaining them thus makes transparency more beneficial. Meanwhile, the high productivity agents are equally likely to be retained under both transparency and no transparency when they are paid a bonus.

I now move onto discussing the second parts of Propositions 3 and 4 where a partially separating equilibrium selected in the transparency game. It is now possible that no transparency is optimal. The reason is that when only one agent has high productivity (and the principal can pay bonuses), the principal cannot signal to the high productivity agent by paying him a positive bonus. This means under transparency, the high productivity agent becomes more pessimistic compared to no transparency where he is paid a bonus—and learns his productivity. This benefit is greater when the value of retaining the high productivity agent increases.[23] This is why the region where no transparency is optimal has higher values of g_H^P within the region where the partially separating equilibrium is selected. A similar reasoning is behind the comparative statics in this region—increasing g_H^P in this region leads to transparency becoming relatively more favorable.

3.4 Other equilibria

For the analysis above I have focused on the equilibrium selected in each game by the appropriate equilibrium refinement. In order to make a comparison of payoffs across two different games, it is problematic to compare outcomes across a set of equilibria. Furthermore the full set of equilibria, in particular in the transparency game, is very large. However, for completeness, I now discuss other equilibria that can arise in each game and how this affects the tradeoff between transparency and no transparency.

First, for the class of equilibria I focus on (separating in the no transparency game, and both separating and partially separating in the transparency game depending on the parameters), there is a continuum of other equilibria. These have higher bonuses being paid, with pessimistic beliefs if a lower bonus is paid. These are eliminated by the D1 criterion, but also note that they yield a lower payoff for the principal—there are higher

[23]However, for a given p_T, there will be threshold where increasing g_H^P will lead to a separating equilibrium existing and being selected.

bonus costs but the agents' distribution of beliefs is the same as before.

Another point to note is that when the partially separating equilibrium is selected in the transparency game (when Assumptions 1 and 2 are both not satisfied), then there does not exist a separating equilibrium in this game—it is not that it is not selected.

In both games there are also mixed strategy (or semi pooling equilibria). These are all eliminated by the D1 criterion (as described in the proof of Propositions 1 and 2). These equilibria are also much more complicated to implement compared to the straightforward pure strategies in the equilibria in the main analysis.

Finally, in both games there is a pooling equilibrium. This must have that no bonus is always paid and is supported by pessimistic off-path beliefs when a (positive) bonus is paid. If this equilibrium is played in both games then transparency plays no role—there is nothing to learn from the other agent's bonus since it is already known. Thus, transparency and no transparency always yield the same payoff to the principal.

4 Extensions

In this section I consider a number of extensions of the model. These illustrate the driving force of the key insights as well as provide justifications for the modelling assumptions.

4.1 The role of commitment

I have assumed that the principal has the ability to commit to (full) transparency or no transparency about the agents' bonuses—two very simple information structures. For a firm these are realistic policies it could commit to, since deviations can be detected and result in a reputational cost.[24] I do not allow the principal to commit to more general mechanisms that may partially reveal the bonus of the other agent, or reveal information about the principal's costs (B). In practice it might be difficult for a firm to commit to these types of information structure since deviations would be hard to detect. In the case of disclosing information about costs, it might be possible to commit to disclose information about their balance sheet, however, outside opportunities that affect the opportunity cost of paying bonuses today are likely to be the private information of the management (principal) and cannot be credibly disclosed. In addition, even if there is a more general mechanism that is optimal, since in some circumstances (full) transparency is preferred over no transparency, my model shows that no transparency can be suboptimal.

To illustrate the importance of the commitment assumption in the model, I consider what happens if the principal cannot commit to transparency. This exercise illustrates the

[24] The assumption of commitment to a disclosure policy in an organizational setting is also made in Jehiel (2015). Here the disclosure is not about the pay of other workers, but about other unknown features of the environment—for example, the monitoring technology.

key role this assumption plays in the model. The result is reminiscent of the unravelling result in Milgrom (1981): the principal is forced to be transparent in all cases since not being transparent means that the principal is choosing to hide bad news. In this case 'bad news' is that there are bonuses available and that the principal chose not to pay one of the agents a bonus.

To relax commitment in what I refer to as the 'game without commitment', consider the same set up but with the timing is changed so that the principal chooses the level of transparency after learning (θ_1, θ_2) and B.[25] Furthermore, I restrict attention to when Assumption 1 is satisfied ('case 1' above). Recall that this meant that under both transparency and no transparency when only one agent has high productivity, the principal pays a (positive) bonus to the high productivity agent. To avoid complications with equilibrium selection, I also restrict the principal's bonuses to be $b_i \in \{0, b_N^*\}$ following $a^P = N$ and $b_i \in \{0, \underline{b}_T^*, \bar{b}_T^*\}$ following $a^P = T$. I also assume that following any choice of a^P the principal pays bonuses as in the equilibria described in Propositions 1 and 2.[26] Finally in the case of indifference, I assume that the principal chooses $a^P = T$ rather than $a^P = N$.

Proposition 5. *In the game without commitment, if for each choice of a^P the principal plays the respective equilibrium described in Propositions 1 and 2, then in any equilibrium the principal chooses transparency ($a^P = T$) for any realisation of θ_1, θ_2 and B with probability 1.*

The intuition is as follows. An agent now updates his beliefs twice: After observing a^P, the agent updates his beliefs given the strategy played by the principal in equilibrium; and then the agent updates his beliefs again after observing his bonus (and potentially the other agent's bonus). In the case that $B = \bar{B}$ and $\theta_1 = \theta_2 = H$, the principal pays both agents a bonus, and so the principal strictly prefers to choose $a^P = T$ and pay the lower bonus $\bar{b}_T^* < b_N^*$ and still induce high beliefs. There are two other sets of 'states' to consider: First, when the principal would definitely not pay a bonus to either agent (for any choice of a^P), which happens when either $B = 0$ or $B = \bar{B}$ and $\theta_1 = \theta_2 = L$; and second, when $B = \bar{B}$ and $\theta_1 \neq \theta_2$, where the principal would pay a bonus only to one agent. In the proof I show that in both cases the principal must choose $a^P = T$ due to an 'unraveling effect'. For example, suppose she chose $a^P = T$ only in the first case, and 'obfuscated' in the second case when only one agent had high productivity by choosing $a^P = N$—so the low productivity agent did not learn that the other agent was paid a bonus. Now, in the second case, the agent not paid a bonus would update his belief to $\Pr[\theta_i = H] = 0$, and not to p_N as he does in the game with commitment. This is because he knows that following $a^P = N$, only one agent has high productivity and that agent

[25] Formally (1) and (2) in the 'Actions and timing section' are switched.

[26] This simplifies the exposition.

will be paid a bonus.

4.2 Correlation in agent's productivity and outside options and agents receiving informative signals

A simplify assumption of the model is that the agents' outside options are not correlated with their productivities. In reality, a more productive agent is likely to receive better outside offers from other firms. In Appendix C, I show that the results do not qualitatively change when this is the case (under some mild parameter restrictions).

There are two effects when more productive agents receive better outside offers. First, agents (of any productivity) will learns about their productivity from their outside offer. Second, high productivity agents are more likely to quit since they are more likely to receive better outside offers. Together these mean that the principal needs to incur a lower signaling cost. In order to be able to signal to high productivity agents in a separating equilibrium, the value of retaining high productivity agents (g_H^P) needs to be sufficiently high (this is the required parameter restriction). If this is satisfied then the selected equilibria are qualitatively the same as before and equivalent results can be derived.

4.3 Dynamic model: Endogenising continuation values

The static model discussed takes the continuation values of the players as exogenously given and is somewhat 'reduced-form'. In this section I discuss a dynamic extension of the model in which the continuation payoffs become *endogenous*. The full description of the model and formal analysis is provided in Habibi (2020), here I just provide a summary of the model and results.

The baseline version of the dynamic model considers a principal interacting with a single agent over an infinite horizon. Each period the agent gets an i.i.d. outside offer and needs to decide whether to stay at the firm or take the outside offer—the game ends once she takes the outside offer. Each period the principal also learns whether she is able to pay bonuses or not that period—this is also i.i.d. across periods. I assume that the principal can either pay no bonus or a fixed bonus. The agent's productivity at any point in time is either high or low, and it varies stochastically with positive persistence from one period to the next. This means that if an agent has high productivity today, he is likely to have high productivity in the near future. As before, it is assumed that only the principal learns the agent's productivity perfectly. At the start of every period the agent has a belief about his productivity—this is the 'state' variable. The principal is assumed to get a higher flow payoff from a high productivity agent. This provides the incentives for her to pay a bonus if and only if agent has high productivity and the the principal is able to pay bonuses. I prove that there is a Markov perfect equilibrium where this strategy is played by the principal in every period. Furthermore, I prove that the continuation values for both player's are increasing in the agents' productivities—this provides a micro-foundation for the assumptions made on parameters representing the continuation values in my static model (g_θ^P and g_θ^A).

5 Related literature

There has been relatively little theoretical work on pay transparency—however, a number of recent papers explore different mechanisms to the one in my model. Both Long and Nasiry (2020) and Halac et al. (2021) consider static models of a single firm with a small number of workers. Long and Nasiry (2020) consider the effect of transparency in a model with moral hazard. In contrast to my paper, they assume that in addition to a standard utility function, workers have a disutility of being paid less than others. They provide conditions under which transparency is optimal for the firm—the intuition for why transparency is beneficial is that workers dislike being paid less than others and so exert additional effort in order to avoid this. Halac et al. (2021) consider a team production model where the principal provides contracts to workers that makes a payment dependent on the team's output. Transparency in their model is about whether a worker can see how their contract compares to the contract that other workers are offered—this is their 'rank' in the firm. They find that less transparency—'rank uncertainty'—means that the principal can provide incentives at a lower (expected) cost. Both these papers consider moral hazard, in contrast, my model considers retention of workers—whether they choose to stay at the firm and forgo their outside option.

Cullen and Pakzad-Hurson (2021) explore an alternative implication of pay transparency in a dynamic model with many workers. Their workers have homogeneous and observable productivity—this is in contrast to my model where the driving force is heterogeneity of the workers and the informational advantage or subjective evaluations of the firm. In their model, workers arrive over time, and the effect of increased transparency is that it commits the firm to negotiate more aggressively with workers in future because it does not want to be seen by other workers to pay high wages. They find that as long as the workers have some bargaining power, it is optimal for a firm to commit to full transparency. This contrasts with my model where, depending on the setting, either transparency or no transparency can be an optimal choice for the firm—and this is consistent with what is observed in the real world. Although it is likely both mechanisms are relevant in practice, my model is particularly applicable for settings in which evaluations of workers are subjective, and there is a lot of heterogeneity in the worker-firm match quality.[27]

There is a growing empirical literature that studies the effect of pay transparency policies. However, these primarily analyse the introduction of disclosure rules mandated by policy makers, such as publishing aggregate statistics on the gender pay gap (Baker et al. (2019)). These are important questions, but are somewhat different to the focus of my paper—whether a firm wants to make an active choice to internally disclose individual worker pay. One paper that does study disclosure of individual pay in a workplace is Blanes i Vidal and Nossol (2011). In their study workers receive feedback on their relative performance (and pay, which is piece rate). They find that anticipating this, workers exert more effort. This provides an alternative rationale for pay transparency that is complimentary to my model. Whilst my model focuses on the effect of transparency on workers learning about their own ability, their rationale focuses on providing direct incentives in an environment with moral hazard.

There are a number of other papers with 'bonus-as-signal' models—where an informed principal who uses bonuses to signal her private information to an agent. Bénabou and Tirole (2003) and Fuchs (2015) both study settings with a single agent and the principal's private information is only in a single dimension. Having a principal who is informed about an agent's productivity is also related to the literature on subjective (or private) evaluations—see, for example, MacLeod (2003).

From a theoretical point of view, I analyse a multidimensional signaling model where the sender (principal) has multiple dimensions of private information (her ability to pay bonuses and the agents' productivities) and the signaling (bonuses) is in a single dimension. As is typical in such models, in some instances it is not possible for the receiver (agent) to attribute the signal to the type of the sender, leading to a signal extraction problem. These types of model have been studied in other contexts. For example, Frankel and Kartik (2019) study a multidimensional signaling model in which the focus is on comparative statics with respect to the informativeness of the set of equilibria as the 'stakes' of the game change—the stakes can be thought of as a larger audience of receivers.[28] In all existing papers in this literature, the sender wants to induce a higher belief about the *same* dimension of the state for *all* members of the audience. In contrast, in my model, different parts of the audience (agents) are interested in different dimensions of the state—their own productivity and the ability of the principal to pay a bonus (and not the productivity of other agents). A novelty of my model is the trade-off the principal faces when designing the informational environment in which the signaling game, with multiple receivers, takes place. In publicly revealing the signal sent to each agent, she

[27] Note that Cullen and Pakzad-Hurson (2021) do extend their model to have heterogeneous productivity in workers with the firm having an informational advantage, however, the driving force of the model—the commitment to not bargain with future workers—remains unchanged.

[28] There are many other papers that study models of multidimensional signaling, a prominent example is Bénabou and Tirole (2006).

potentially reveals information about the state to other agents, and depending on the state, this might be beneficial or detrimental to her.

Finally, pay compression as a result of transparency has been documented empirically in Obloj and Zenger (2022) and Mas (2017). Obloj and Zenger (2022) analyses a large data set of academics from 1997 to 2017 and shows that transparency leads to pay compression. Mas (2017) studies the effect of public sector managers' pay being disclosed to the public. The paper argues that the findings are consistent with public aversion to high pay which is not something considered in my model. In a theoretical contribution, Cabrales et al. (2008) consider an equilibrium labor market model where workers have social preferences. Among their findings is that social preferences lead to pay compression—the broad intuition is that lower paid workers get more discouraged under social preferences so firms have an incentive to reduce the gap in pay.

6 Discussion and conclusion

In this paper, I propose a model that allows me to analyse what features of a firm make pay transparency more favorable. The model predicts that when there are relatively few high productivity 'star' workers, increasing the firm's value of retaining these workers makes transparency more favorable.

It should be noted that the findings rely on the assumptions of the model, and that these assumptions are more relevant to certain industries and types of workers. The first key assumption is that bonuses (or wage increases) act as a signal to workers from a firm that has better information about their worker-firm match quality. The second key assumption is that there is uncertainty on the firm's ability to pay bonuses or increase wages and that this is the private information of the firm. As already discussed, it is almost certain that firms have some private information about this, however the results are more pronounced when the firm has a greater informational advantage since signaling is 'stronger'. This is likely to be the case in smaller firms or start-ups.

The key predictions of the model could be tested with appropriate data. One would need a dataset from industries that satisfy the key assumptions above and also where there are relatively few high productivity workers. This could be for example, junior employees in law, consulting or advertising. The data would need to include firms' pay transparency policy as well as some measure of the heterogeneity in the worker-firm matches. Further data on retention could also be used to verify that workers receiving higher bonuses or wage increases are more likely to stay at the firm—as is predicted by the model. The first prediction that could be tested is that as the heterogeneity in the worker-firm matches increases, transparency becomes more prevalent. The second prediction that could be tested is that there is pay compression within firms that are transparent. Ideally data would allow a comparison across time at the individual firm level for firms that transition to a more transparent regime over time.

When looking at the wider labor market, it should also be noted that the model has also abstracted from how a decision to commit to transparency affects the composition of workers within a firm. In a well known empirical study, Lazear (2000) finds that performance pay has a significant effect on the sorting into a firm—the introduction of performance pay means that the firm's workforce becomes more productive. A natural (theoretical) question is: what effect does transparency have on the sorting of workers into firms? The theoretical model in Cullen and Pakzad-Hurson (2021) addresses this question within a simple framework with homogeneous workers, but it is not clear what will happen with heterogeneous workers within the signaling framework I have developed.

Appendix A Proofs and omitted analysis

A.1 Proof of Proposition 1

Consider the bonus paid to agent i. In all equilibria must have that $b_i = 0$ if $B = 0$. I will go through possible equilibria and rule them out using the intuitive criterion.[29]

First, there is a pooling equilibrium. This must have $b_i = 0$ for both types when $B = \bar{B}$. If not then the principal could deviate to $b_i = 0$ and save on the bonus inducing the same belief. The equilibrium is supported by pessimistic off-path beliefs, so if $b_i > \epsilon$ agent i believes that $\theta_i = H$ with probability p_0. This equilibrium can be ruled out by the intuitive criterion in a similar way to pooling equilibria being ruled out in the canonical Spence education model. Consider a deviation to $b_i \in \big((1-p_0)g_H^P, (1-p_N)g_H^P\big)$ and assume that this induces the best possible belief—that $\theta_i = H$ with probability 1. It is straightforward to show that such a deviation is only profitable for a type $\theta_i = H$ (and not a type $\theta_i = L$). Type $\theta_i = H$ has a payoff of

$$V_H^* = p_0 g_H^P.$$

The deviation (and given belief) yields a payoff of

$$V_H = g_H^P - b_i > V_H^*.$$

Meanwhile for type $\theta_i = L$ the deviation yields a payoff

$$V_L = 1 - b_i < V_L^*.$$

Thus the intuitive criterion is violated since the only type that benefits from the deviation is θ_H and it is having this type with probability 1 that induces a profitable deviation.

Second, there are separating equilibria where $b_i \in (b_N^*, (1-p_N)g_H^P]$, a higher bonus than when $b_i = b_N^*$. Such an equilibrium can be ruled out by considering a deviation as before to $b_i = b_N^*$. Only type $\theta_i = H$ can find this deviation profitable thus meaning that this does not satisfy the intuitive criterion.

Finally, there are semi-separating equilibria where type $\theta_i = L$ mixes between $b_i = 0$ and $b_i = \bar{b} > 0$. Again this can be ruled out by the intuitive criterion by considering a deviation to $b_i \in \big((1-p_0)g_H^P, (1-p_N)g_H^P\big)$ which can only be profitable for a type $\theta_i = H$.

[29] As explained in the main text this result holds for the intuitive criterion which is a stronger refinement than the D1 criterion stated in the result. The proof is for the intuitive criterion.

A.2 Analysis of transparency game not in main text

I start by providing some more details omitted in the main text for the transparency game.

Separating equilibrium. Agent i's best response (on the equilibrium path) is given by

$$a_i^A = \begin{cases} S & \text{if } b_i > \epsilon; \text{ or } b_i = b_j = 0 \text{ and } u_i \leq p_T, \\ Q & \text{otherwise.} \end{cases}$$

Note that, unlike in the no transparency case, this *does* depend on the bonus payment of agent $j \neq i$ since agent i revises his beliefs of B (and hence his productivity θ_i) based on b_j.

There are six incentive constraints that ensure: type[30] (L, L) doesn't deviate to (H, H) and vis-versa, type (H, L) doesn't deviate to (H, H) and vis-versa, and type (L, L) doesn't deviate to (H, L) and vis-versa.[31] The constraints can be simplified to[32]

$$\begin{aligned} 1 - p_T \leq \bar{b}_T & \leq (1 - p_T) g_H^P, \\ 1 \leq 2\bar{b}_T - \underline{b}_T & \leq g_H^P, \\ 1 - 2p_T \leq \underline{b}_T & \leq (1 - p_T) g_H^P - p_T, \end{aligned}$$

The expressions for $\bar{b}_T^*$ and $\underline{b}_T^*$ used in Proposition 2 follow from the lower bound of the first and third inequalities.

Partially separating equilibrium. Agent i's best response (on the equilibrium path) is given by

$$a_i^A = \begin{cases} S & \text{if } b_i > \epsilon; \text{ or } b_i = 0 \text{ and } u_i \leq p'_T, \\ Q & \text{otherwise.} \end{cases}$$

Unlike the separating equilibrium above, this means that (on the equilibrium path) the agent does not learn from the other agent's bonus.

The principal has just two incentive constraints the first to ensure that types (H, L) does not deviate to (H, H) (this also ensures that type (L, L) does not deviate to (H, H), and the second to ensure that type (H, H) does not deviate to (H, L) (or equivalently (L, L)). These can be simplified to

$$\tfrac{1}{2}(1 + g_H^P)(1 - p'_T) \leq \bar{b}'_T \leq (1 - p'_T) g_H^P. \tag{A.1}$$

[30] Here 'type' refers to the productivity of agents that the principal observes (θ_1, θ_2).

[31] There are obviously identical constraints for (L, H) but I omit these.

[32] Again, notice that the assumption that $\frac{1}{2}\bar{B} \geq g_H^P$ ensures that the principal can always pay bonuses in this range when $B = \bar{B}$.

The expression for $\bar{b}_T'^*$ used in Proposition 2 follows from the lower bound of this inequality.

A.3 Proof of Proposition 2

In all equilibria must have that $(b_i, b_j) = (0, 0)$ if $B = 0$. As in the proof for no transparency (Proposition 1), I will go through possible equilibria and rule them out using the D1 criterion. Recall the notation introduced in the main text to denote mixed strategies in the transparency game (when $B = \bar{B}$):

$$\sigma_{\theta_i \theta_j}^{b_i b_j} \equiv \Pr[b_1 = b_i, b_2 = b_j | \theta_1 = \theta_i, \theta_2 = \theta_j].$$

Assumptions 1 and/or 2 are satisfied (case 1 and 3 in text). Order the principal's 'types', (θ_1, θ_2) when $B = \bar{B}$, from low to high as (L, L), (H, L), (H, H).[33]

Suppose there is an equilibrium in which two types pool by playing the same bonus with positive probability. So there are types $(\theta_1, \theta_2) \neq (\theta_1', \theta_2')$ and a bonus paid $(b, b') \in \mathcal{B}^2$ with $\sigma_{\theta_1 \theta_2}^{b,b'} > 0$ and $\sigma_{\theta_1' \theta_2'}^{b,b'} > 0$. The D1 criterion rules out any such equilibrium by considering a deviation to a slightly higher bonus $(b + \epsilon', b')$ (where $\epsilon' > \epsilon$) by the highest type in the pool. Any response from the agents that the lower type would prefer would be strictly preferred by the higher type. Thus in the first step of the D1 criterion, only the higher type would be considered. Next, the deviation is profitable if ϵ' is sufficiently small.[34] This follows similar reasoning to the no transparency (proof of Proposition 1). Note that parameters that satisfy Assumptions 1 and/or 2 mean that when type (H, L) and (L, L) pool at $b = (0, 0)$, type (H, L) finds it profitable to separate.

The argument above does not rule out a separating equilibrium with $\sigma_{HH}^{b',0} = \sigma_{HH}^{0,b'} = \frac{1}{2}$, $\sigma_{HL}^{b,b} = 1$ and $\sigma_{LL}^{0,0} = 1$, for some $b, b' > \epsilon$. Note in such an equilibrium, it must be that $b' > 2b$ to ensure that type (H, L) does not deviate to (H, H). Such an equilibrium can be ruled out by a deviation to a bonus pair $(b'', 0)$ by type (H, L) where $2b > b'' > \epsilon$. Note that the equilibrium must be supported by pessimistic beliefs for the off-path action $(b'', 0)$ to ensure that a deviation from (H, L) is not profitable. The first stage of the D1 criterion rules out type (L, L) from the deviation. Then to see that, given this, the deviation is profitable for type (H, L), note that the equilibrium payoff is less than the deviation payoff

$$V_{HL}^* = -2b + \tfrac{1}{2}(g_H^P + 1) < -b'' + g^P,$$

where the deviation payoff on the RHS uses the fact that both agents now learn their productivity perfectly.

Now considering separating equilibria as in the definition in the main text. The lowest

[33] Note given the symmetry assumption type (L, H) strategy is pinned down by (H, L).

[34] Note that since the limit $\epsilon \to 0$ is considered, then ϵ' can be made arbitrarily small.

cost separating equilibrium is the only one that satisfies the D1 criterion. The argument is exactly as in the case of no transparency (proof of Proposition 1).

Finally, note that the monotonicity assumption rules out a separating equilibrium, but σ_{HL}^{0,b_T^*}, so when only one agent has high productivity a bonus is paid only to the low productivity agent.

Assumptions 1 and 2 both are not satisfied (case 2 in text). Now the D1 selects the lowest cost partially separating equilibrium. The argument above applies, with the exception that it does not rule out pooling of types (L, L) and (H, L) at a bonus $b = (0, 0)$. The difference is now that a deviation to induce a belief concentrated on (H, L) for type (H, L) is not profitable for any bonus cost. Furthermore, given this, clearly no (fully) separating equilibrium exists.

A.4 Proof of Proposition 5

The state of the world is a triple (θ_1, θ_2, B). Denote a generic state of the world by α. In the game without commitment, the principal's strategy is given by a choice of transparency given the state,

$$a^P : \alpha \to \{T, N\},$$

and a choice of bonus given the level of transparency and the state,

$$b : \alpha \times \{T, N\} \to \Delta(\mathcal{B}^2).$$

The agents update their beliefs according to Bayes rule given the equilibrium strategy of the principal.

Given either choice of a^P, the principal pays bonuses as described in Propositions 1 and 2. So if the principal's strategy a^P does not depend on the state (i.e. $\Pr[a^P|\alpha] = k$ for a constant $k \in [0, 1]$), then the agents' beliefs are as described in the main text. However, if the principal makes different choices of transparency a^P for different states, then agents update their beliefs twice: first after the choice of a^P, then after b.

In the states that a bonus is not paid, i.e. either if $B = 0$, or if $\alpha = (L, L, \bar{B})$, then the principal's equilibrium strategy cannot depend on the state. Suppose, towards a contradiction, that it did. Then, depending on the choice of a^P, the agents would have different beliefs (since they do not update after b_i). The principal would always choose the a^P that induced the higher beliefs and higher payoff, thus this cannot be an equilibrium strategy.

Now I show that in any equilibrium the principal must choose transparency ($a^P = T$) in every state. I begin by considering what happens in the state $(H, H, \bar{B})$. Here, since the beliefs of the agents will both be that $\theta_i = H$ with probability 1 following the bonuses, the principal strictly prefers between $a^P = T$ to $a^P = N$ since this results in lower bonuses. Now consider the other sets of states, first if $B = 0$, or if $\alpha = (L, L, \bar{B})$ (call this set of states A_1), and second if $B = \bar{B}$, and $\theta_1 \neq \theta_2$ (call this set of states A_2). Assume that in an equilibrium the principal plays $a^P = N$ in some state. I show by contradiction that such an equilibrium cannot exist.

First, consider a possible equilibrium in which the principal plays $a^P = N$ only in states A_1. Here the posterior belief of agent i following $a^P = N$ and seeing $b_i = 0$ is $\Pr[\theta_i = H|a^P = N, b_i = 0] = p_T$ and following $a^P = T$ and seeing $b_i = 0$ is $\Pr[\theta_i = H|a^P = T, b_i = 0] = 0$. So the principal can gain from a deviation in which she chooses $a^P = N$ when the state is in A_2. The posterior belief of the agent with $\theta_i = L$ is increased from 0 to $p_T > 0$ and the posterior belief of agent $j \neq i$ remains unchanged.

Second, consider a possible equilibrium in which the principal plays $a^P = N$ only in the states A_2. Here the posterior belief of the agents will remain unchanged if $a^P = T$ is played in place of $a^P = N$ (since when agent i sees $a^P = N$ and $b = 0$ he has a posterior of 0). This means that by assumption the principal will play $a^P = T$ instead of $a^P = N$.

Finally, consider a possible equilibrium in which the principal plays $a^P = N$ in states in A_1 and A_2. Here the posterior belief of agent i will be $\Pr[\theta_i = H|a^P = N, b_i = 0] = p_N$ if they get no bonus and 1 if they get a bonus. When the state is in A_1 the principal can deviate to $a^P = T$ and increase the posterior belief of agent i from p_N to p'_T and so this cannot be an equilibrium.

This means that in any equilibrium it must be the case that $a^P = T$.

Appendix B Definition of D1 criterion with multiple receivers

As mentioned in the main text, the D1 criterion is not defined for the class of games with multiple receivers (agents in my terminology). In general, this potentially creates complications. For example, following the action of the sender (and beliefs induced), if the receivers actions influence each other's payoffs (i.e. there is a game), it is not clear what set of potential outcomes should be considered for each possible belief induced. In the setting with a single receiver, the D1 criterion compares the payoff for the principal when the receiver best responds to the sender's equilibrium payoff. In my setting, since the receivers (agents) have payoffs that do not depend of the belief of the other agent—and so a game is not induced following the sender (principal's) action—the best responses can still be used as before, thus allowing the D1 criterion to be extended in a natural way.

The procedure I propose is as follows. As in the standard D1 criterion, at the stage where some types are excluded from a particular deviation, I consider whether each type can get a higher payoff than their (expected) equilibrium payoff from that particular deviation given some set of beliefs of *all* agents (receivers).[35] Then having excluded these types, as in the single receiver definition, an equilibrium fails the D1 criterion if there is a type such that a deviation gives a strictly higher payoff compared to the equilibrium payoff for any best response of the agents (receivers) given that the type cannot be any of those excluded.

More formally, I adapt the Definition 11.4 and 11.6 in Fudenberg and Tirole (1991). Before going into the definition, I simplify the latter part of the game tree so that the agents' action is a choice of cutoff for which outside option they will accept, rather than having a realisation of outside option then a binary stay/quit decision—this is strategically equivalent to the game described in the text. Denote the vector of cutoffs that the agents choose by $\bar{u}^A = \left(\bar{u}_1^A, \bar{u}_2^A\right) \in [0,1]^2$, so that $a_i^A = Q$ if and only if $u_i < \bar{u}_i^A$. I also introduce notation for best responses given an action from the principal and belief that it induces. Let $T \subseteq \Theta$ where Θ is the set of all possible $\theta = (\theta_1, \theta_2)$ 'types' of the principal (productivity of the agents she employs). Let T_i denote the i^{th} projection map of T, $T_i = proj_i(T)$. Let $\mu_i \equiv \mu(\theta_i|b)$ be the beliefs of agent i following bonuses $b = (b_1, b_2)$.[36] Denote the (expected) payoff of agent i in terms of the cutoff strategy $\bar{u}_i^A$ by $U_i(b, \bar{u}_i^A, \theta_i)$—note that $\bar{u}_i^A$ and θ_i do not affect the payoff of agent $j \neq i$. Denote the principal's (expected) payoff by $V(b, \bar{u}_1^A, \bar{u}_2^A, \theta)$. The set of best response vectors for agent i when $\theta \in T$ and a bonus b is paid is given by

$$\mathrm{BR_i}(T, b) = \bigcup_{\mu_i : \mu(T_i|b)=1} \mathrm{BR}(\mu_i, b),$$

where

$$\mathrm{BR}(\mu_i, b) = \underset{\bar{u}_i^A \in [0,1]}{\arg\max} \sum_{\theta_i \in T_i} \mu(\theta_i|b) U_i(b, \bar{u}_1^A, \theta_i).$$

Define the $D(\theta, T, b)$ to be the set of best responses to action b and beliefs concentrated on T that make type θ strictly prefer b to his equilibrium strategy,[37]

$$D(\theta, T, b) = \bigcup_{\mu : \mu(T|b)=1} \left\{ (\bar{u}_1^A, \bar{u}_2^A) \in BR(\mu, b) \text{ s.t. } V^*(\theta) < V(b, \bar{u}_1^A, \bar{u}_2^A, \theta) \right\},$$

and let $D^0(\theta, T, b)$ be the set of best responses that make type θ exactly indifferent.

Definition 6 (Multi-receiver D1 criterion)**.** *A type θ is deleted for strategy b if there is*

[35] The standard definition considers only the belief and action of the single receiver.

[36] Note that since θ_1 and θ_2 are independent, $\mu(\theta_i|b)$ does not impose any restrictions on the values that $\mu(\theta_j|b)$ can take.

[37] Fudenberg and Tirole (1991) use mixed-strategy best responses in their definition, but this is not relevant for my setting so I omit this detail.

a θ' such that

$$\left\{D(\theta,\Theta,b)\cup D^0(\theta,\Theta,b)\right\}\subset D(\theta',\Theta,b).$$

Let $J(b)$ be the set of all θ that are not deleted. If there exists θ' and b such that

$$V^*(\theta') < \min_{(\bar{u}_1^A,\bar{u}_2^A)\in(\mathrm{BR}_1(\Theta\setminus J(b),b),\mathrm{BR}_2(\Theta\setminus J(b),b))} V(b,\bar{u}_1^A,\bar{u}_2^A,\theta'),$$

then the equilibrium fails the multi-receiver D1 criterion.

The key features of the setting that allow the D1 criterion to be extended in this way are:

1. Agent i's payoff does not depend on agent j's action. If this were not the case then $\mathrm{BR}(\mu_i,b)$ would not be defined in the way it has been above. Instead the principal's choice of b and the corresponding belief would induce a game played between the two agents which means that the definition would have to be adapted further;

2. Agents' productivities are independent and so $\mu(\theta_i|b)$ does not impose any restrictions on the values that $\mu(\theta_j|b)$ can take;

3. Agent i's payoff does not depend on θ_j. Combined with independence assumption above, this means that for a given b, $\mathrm{BR}_i(T,b)$ does not restrict $\mathrm{BR}_j(T,b)$ in any way. This means that the 'max' and 'min' over the principal's expected utility within the definition can be done dimension by dimension.

Appendix C Correlation in agents' productivity and outside option and agents receiving informative signals

In this section I provide formal analysis of the discussion in Section 4.2. I assume that high and low productivity agents get outside offers drawn from different distributions. Formally, I assume that they are drawn from the distributions with c.d.f.$F_H(u)$ and $F_L(u)$ (and respective p.d.f. f_H and f_L). They both have full support on $u\in[0,1]$. I assume that $F_H(u)\leq F_L(u)$ for all u, so the distribution for high productivity agents first order stochastically dominates.

I show that the incentive constraints of the principal can be obtained in a similar way as before, but to ensure existence of a separating equilibrium there need to be restrictions on the parameters. First, I show how the agents update their prior following a bonus from the principal's perspective—i.e. the principal's expected belief for an agent after paying him a bonus b_i given his talent. Under no transparency (in a separating equilibrium) after a bonus $b_i > 0$ agent i updates to a belief of 1 regardless of θ_i;[38] while after a bonus b_i agent i's expected belief is $\hat{p}_N(\theta_i)$ where

$$\hat{p}_N(H) \equiv \int_0^1 \frac{p_N f_H(u)}{p_N f_H(u) + (1-p_N) f_L(u)} du,$$

and $\hat{p}_N(L)$ is similarly defined.

The equivalent constraint to inequalities 3.3 is

$$1 - F_L(\hat{p}_N(L)) \leq b_N \leq \big(1 - F_H(\hat{p}_N(H))\big) g_H^P. \tag{C.1}$$

For a b^N to exist that satisfies this it must be that $1 - F_L(\hat{p}_N(L)) < \big(1 - F_H(\hat{p}_N(H))\big) g_H^P$. Note that $\hat{p}_N(H) > \hat{p}_N(L)$—which intuitively is the case since a high productivity agent is more likely to get a higher outside option and become less pessimistic compared to a low productivity agent. This means that condition above is not necessarily satisfied, but for g_H^P above a given cutoff it will be.

There are equivalent conditions for the relevant inequalities under transparency. Again these will be satisfied if g_H^P is above a given cutoff. Assuming all these conditions are satisfied, then the equilibrium selected will be as in Propositions 1 and 2 (with b_N^*, $\underline{b}_T^*$ and $\bar{b}_T^*$ adjusted accordingly). Then equivalent results to Propositions 3 and 4 can be obtained.

[38] This follows from the full support assumption.

References

Baker, M., Halberstam, Y., Kroft, K., Mas, A., and Messacar, D. (2019). Pay Transparency and the Gender Gap. *NBER Working Paper Series*.

Battaglini, M., Bénabou, R., and Tirole, J. (2005). Self-control in peer groups. *Journal of Economic Theory*, 123(2):105–134.

Bewley, T. (1999). *Why Wages Don't Fall during a Recession*. Harvard University Press.

Blanes i Vidal, J. and Nossol, M. (2011). Tournaments Without Prizes: Evidence from Personnel Records. *Management Science*, 57(10):1721–1736.

Bénabou, R. and Tirole, J. (2003). Intrinsic and Extrinsic Motivation. *The Review of Economic Studies*, 70(3):489–520.

Bénabou, R. and Tirole, J. (2006). Incentives and Prosocial Behavior. *American Economic Review*, 96(5):1652–1678.

Cabrales, A., Calvó-Armengol, A., and Pavoni, N. (2008). Social Preferences, Skill Segregation, and Wage Dynamics. *The Review of Economic Studies*, 75(1):65–98.

Cho, I.-K. and Kreps, D. M. (1987). Signaling Games and Stable Equilibria. *The Quarterly Journal of Economics*, 102(2):179–221.

CIPD (2017). Reward Management: Focus on Pay. *Chartered Institute of Personnel and Development*.

Cullen, Z. and Pakzad-Hurson, B. (2021). Equilibrium Effects of Pay Transparency in a Simple Labor Market. *working paper*.

Cullen, Z. and Perez-Truglia, R. (2022). How Much Does Your Boss Make? The Effects of Salary Comparisons. *Journal of Political Economy*, 130(3):766–822.

Ederer, F. (2010). Feedback and Motivation in Dynamic Tournaments. *Journal of Economics & Management Strategy*, 19(3):733–769.

Frankel, A. and Kartik, N. (2019). Muddled Information. *Journal of Political Economy*, 127(4):1739–1776.

Fuchs, W. (2015). Subjective Evaluations: Discretionary Bonuses and Feedback Credibility. *American Economic Journal: Microeconomics*, 7(1):99–108.

Fudenberg, D. and Tirole, J. (1991). *Game Theory*. MIT Press.

Habibi, A. (2020). Pay Transparency in Organizations. *working paper*. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3248704.

Halac, M., Kartik, N., and Liu, Q. (2017). Contests for Experimentation. *Journal of Political Economy*, 125(5):1523–1569.

Halac, M., Lipnowski, E., and Rappoport, D. (2021). Rank Uncertainty in Organizations. *American Economic Review*, 111(3):757–786.

Hastings, D. F. (1999). Lincoln Electric's Harsh Lessons from International Expansion. *Harvard Business Review*, (May–June 1999).

Jehiel, P. (2015). On Transparency in Organizations. *The Review of Economic Studies*, 82(2):736–761.

Lam, L., Cheng, B., Bamberger, P., and Wong, M.-N. (2022). Research: The Unintended Consequences of Pay Transparency. *Harvard Business Review*.

Lazear, E. P. (2000). Performance Pay and Productivity. *The American Economic Review*, 90(5):1346–1361.

Li, J. and Matouschek, N. (2013). Managing Conflicts in Relational Contracts. *The American Economic Review*, 103(6):2328–2351.

Long, X. and Nasiry, J. (2020). Wage Transparency and Social Comparison in Sales Force Compensation. *Management Science*, 66(11):5290–5315.

MacLeod, W. B. (2003). Optimal Contracting with Subjective Evaluation. *The American Economic Review*, 93(1):216–240.

Mas, A. (2017). Does Transparency Lead to Pay Compression? *Journal of Political Economy*, 125(5):1683–1721.

Milgrom, P. R. (1981). Good News and Bad News: Representation Theorems and Applications. *The Bell Journal of Economics*, 12(2):380–391.

Obloj, T. and Zenger, T. (2022). The influence of pay transparency on (gender) inequity, inequality and the performance basis of pay. *Nature Human Behaviour*, 6(5):646–655.

Prendergast, C. (1999). The Provision of Incentives in Firms. *Journal of Economic Literature*, 37(1):7–63.

Spence, M. (1973). Job Market Signaling. *The Quarterly Journal of Economics*, 87(3):355–374.