A Case for Pay Secrecy*

Tomer Blumkin†, David Lagziel‡, and Yoram Margalioth§

November 13, 2022

Abstract:

In this paper we study a labor market setup in which workers exhibit relative remuneration concerns with respect to their peers. We first characterize the optimal labor contract offered by the firm and provide necessary and sufficient conditions for the desirability of incorporating pay-secrecy clauses in such a contract. We then demonstrate that, in contrast to conventional wisdom viewing wage-secrecy arrangements as detrimental for workers, the latter may in fact gain from the lack of pay transparency.

Journal of Economic Literature classification numbers: E24, D82, J30, J31, J71.

Keywords: secrecy; wages; relative wage; labor contracts; wage compression; wage dispersion.

*The authors wish to thank the editor of the journal, two anonymous reviewers, Leif Danziger, Ada Gonzalez-Torres, Ori Heffetz, Nadav Levy, Omer Moav, Tal Regev, Moses Shayo, Ran Spiegler, Yoram Weiss, David Wettstein, Ro'i Zultan, and participants of the IDC Economics seminar, the Brown-Bag Economics seminar at BGU, the Behavioral and Experimental Economics seminar at the Hebrew University, the CESifo Area Conference on Employment and Social Protection in Munich (Germany), the Economic Theory seminar at Tel-Aviv University, the Economic Theory seminar at Bar-Ilan University, the Strategy seminar at Coller Business School at Tel-Aviv University for their valuable comments and suggestions. Blumkin and Lagziel acknowledge the support of the Israel Science Foundation, Grant #513/19.

†Ben-Gurion University of the Negev, Beer-Sheva 8410501, Israel. e-mail: Tomerblu@bgu.ac.il; CESifo, Poschingerstrasse 5, 81679 Munich, Germany; IZA, Schaumburg-Lippe Strasse 7/9, 53113 Bonn, Germany.

‡Ben-Gurion University of the Negev, Beer-Sheva 8410501, Israel. e-mail: Davidlag@bgu.ac.il.

§Tel Aviv University, Tel Aviv 6997801, Israel. e-mail:margalio@tauex.tau.ac.il

1 Introduction

1.1 "To disclose or not to disclose?": the debate on pay-transparency – some background

Many people are taught since childhood that it is rude to discuss money, and indeed, the sharing of salary information has long been taboo in American society. In addition, companies, especially private-sector, non-unionized organizations, foster cultures and policies that discourage the sharing of salary information. The National Labor Relations Board (NLRB), as well as federal courts, found non-disclosure agreements and similar pay secrecy practices, to be unlawful, as they violated the employees' rights to engage in concerted activity for collective bargaining (even if non-unionized) or other mutual aid or protection. But law enforcement has been weak, penalties were low, and employees were unaware of the law protecting their right to discuss their salaries [Bierman and Gely (2004)].

In recent years, a new motivation to prohibit pay secrecy became dominant. It is the desire to eliminate the gender pay gap [Lobel (2020)]. In 2014, President Obama declared that pay secrecy fosters discrimination and signed an executive order banning federal contractors from retaliating against employees for discussing their compensation. A dozen states and the District of Columbia passed their own legislation banning pay secrecy rules, joining six states that already had such laws in place. Finally, digital platforms provide crowdsourced salary information, and the social norm regarding privacy seems to be changing. As one author put it - "pay confidentiality has been eroding for years . . . Millennials share every thought it seems." [Risher (2014)].

Nevertheless, field studies and surveys tell us that only little has changed in practice. Cullen and Perez-Truglia (2022) conducted a natural field experiment with a sample of $2,060$ employees from a large commercial bank, typical of corporate business around the world in all relevant respects. They found no evidence of information diffusion: the information provided to an employee did not travel to their peers, not even to their closest peers.[1] Contrary to the predictions of information diffusion models, they found that employees with higher centrality in the email network and employees who talk more often with coworkers do not have lower mis-perceptions. The 2017-2018 national survey of the Institute for Women's Policy Research (IWPR) found that 60 percent of workers in the for-profit private-sector, work under a pay secrecy policy of some sort, just six percentage points lower than in 2010 (66 percent). The proportion of private-sector workers who reported that they are formally prohibited from discussing their pay fell from 25 percent in 2010 to 16 percent in $2017-18$, but at the same time, the share of private-sector workers reporting that they are discouraged from discussing their pay increased from 41 percent to 44 percent (IWPR 2021).

[1] The zero information diffusion was not due to lack of interest. Some employees were willing to pay days' and even weeks' worth of salary for information about the salaries of their managers or their peers. However, while employees value the salary information a lot, they may value their privacy even more, Cullen and Perez-Truglia (2022).

1.2 Review of related empirical evidence and theoretical research

Our discussion in the previous subsection raises a fundamental question - who benefits from pay secrecy? The latter has been the focus of a large number of empirical studies. According to Card et al. (2012), employers, as well as employees, do. At least in the short run. Taking advantage of a new website listing the pay of the University of California employees, Card et al. (2012) performed a randomized manipulation of access to information on co-workers' salaries. They found that the information caused a reduction in job satisfaction among workers with pay below the median for their department and occupation group, and an increase in their intention to look for a new job. By comparison, workers who were paid above the median reported no significant changes in job satisfaction or job search intentions.

The research concludes that employers have a strong incentive to impose pay secrecy rules to prevent decline in job and pay satisfaction of the lowest-earning workers. The research hypothesized, however, that in the longer run pay transparency may lead to endogenous changes in wage-setting policies and employee composition that ultimately affect the distribution of wages.[2]

In a lab experiment, Bracha et al. (2015) found that when participants were unaware of different pay rates given for the same task, labor supply was marginally lower under the lower pay rate, indicating that participants respond to incentives, namely, that higher pay results in greater labor supply. When pay transparency was introduced, lower-paid workers supplied significantly less work time relative to higher-paid workers and significantly less time than when they were unaware of the higher pay rates. The transparency effect disappeared when subjects were given a moderately trivial reason for the difference in relative pay, but this result did not hold when the reason given was patently arbitrary. It suggests that lower paid workers may not be adversely affected by pay transparency if they are given an explanation for the wage differences, namely, believe that the wage differences are fair.[3]

In a field experiment with 378 Indian workers in seasonal manufacturing jobs, Breza et al. (2018) tried to find out whether in a world with heterogeneous productivity, fairness norms are violated if pay levels correspond to workers' marginal products and are therefore unequal? They found that when workers can clearly perceive that their higher-paid peers are more productive than themselves, pay disparity has no discernible effect on output, attendance, or group cohesion. In such a setting there is no reason for pay

[2] In 2001, Norwegian tax records became easily accessible online to everyone in the country. This seems to provide a similar natural experiment setting. However, in contrast to the website studied in Card et al. (2012) the Norwegian website reveals the net income aggregated from all sources of income, making it useless for salary comparisons [see Perez-Truglia (2020)].

[3] However, in a post-study feedback question the researchers found that the moderately trivial reason was viewed by the participants to be less fair than the arbitrary reason. It is therefore unclear why was the moderately trivial reason sufficient to prevent the decrease in labor supply, whereas the arbitrary one was not. The moderately trivial reason was the assignment of pay rates on the basis of a deliberately arbitrary evaluation of an essay the participants were required to write. We think that this answer reflects disappointment that the participants, who were students, had with the way in which their essays were evaluated. Hence the peculiar result is likely to be the outcome of the specific choice of reason.

secrecy.

However, when coworkers' productivity is difficult to observe, pay inequality reduces output by 0.45 standard deviations and attendance by 18 percentage points. Most interestingly, the research found that not only the low wage workers reduced their attendance, but so did the high and median wage workers (by 10 and 13 percentage points, respectively) compared to their counterparts working in a setting of equal pay. It found no evidence of positive impacts on output for high and median wage workers. The intuitive explanation given by the researchers for the decrease in attendance of high-wage workers is the reduced desire to work alongside disgruntled coworkers in an unpleasant workplace atmosphere due to hostility, resentment, or social awkwardness. Finally, on the last day of the experiment, the workers played cooperative games. The workers who worked in a setting of unjustified unequal pay did worse when they played with coworkers from their own unit than with strangers. This, as well as their answers on an end-line survey, reveal that the unequal pay adversely affects social cohesion.

The implication of this study on the issue of pay secrecy is the following. When managers know which workers are more productive than others on average, but setting pay according to expected output is difficult to justify from coworkers' perspective, pay secrecy is warranted, assuming it can be enforced. [4]

The business management literature supports the economics literature reviewed above. In a quasi-experimental study, Vidal and Nossol (2011) studied a German warehouse with 65 employees, all performing the same task. Work was strictly individual and opportunities to observe each other‘s work were limited. Everyone received the same fixed base salary, which represented around 75 percent of the average worker total compensation. In addition, workers were paid according to the number of goods they have dispatched. The payment for each dispatched goods was the same.[5] In November 2001 the management started to communicate to workers their own rank-order position in the productivity/wage per hour distribution. The information was disclosed in a purely private manner to each employee, letting the worker know his/her relative position without revealing any names. This relative performance information exercise (henceforth, RPIE) had no effect on workers' quitting (separation) rates and led to an increase in individual productivity that is statistically and economically significant. The researchers found RPIE to be unambiguously beneficial to the firm's profitability.

The research supports the economics literature detailed above by providing another field experiment showing that when workers can clearly perceive that their higher-paid peers are more productive than themselves, pay disparity does not decrease productivity of low-pay workers and does not increase their quitting rates.

[4]It should be noted that in the field experiment of Breza et al. (2018), managers maintained pay secrecy, but were unable to enforce it. In fact, they have realistically assumed it would not be kept. At the end of the month long experiment, 86.6 percent of workers could accurately report the wages of the coworkers in their product units.

[5]There was also a negligible discount (less than 1 percent of total compensation) when a worker's monthly mistakes exceeded a certain threshold.

Moreover, Vidal and Nossol (2011) contribute to the literature the following insight. When workers are confident that performance is accurately measured, performance or wage transparency can motivate workers to increase their productivity in an effort to improve their rank for pure competitive reasons. That is, in addition to and beyond any pecuniary incentives.

The importance of employees' perception regarding the fairness of how their pay is affected by their effort and productivity was further developed in Belogolovsky and Bamberger (2014). They note that while there may be objectively valid and positive factors underlying the adoption of pay secrecy (most notably, to avoid employee jealousies), employees interpret pay secrecy as a signal of deception [Day (2007); Milkovich et al. (2013)].

Belogolovsky and Bamberger (2014) conducted a laboratory real-effort experiment consisting of 280 Israeli undergraduate students, in which information of peers' level of remuneration (but not performance) was provided to a subgroup of the participants. The participants were also subject to different methods of performance evaluations. The researchers found that pay secrecy adversely impacts individual task performance because it weakens the perception that an increase in performance will be accompanied by an increase in pay. Moreover, pay secrecy elicits detrimental sorting effects inducing workers (particularly the higher-performing ones) to look for a new job.

These detrimental effects are attenuated when performance assessment is objective (namely, does not involve employer's discretion) and when performance is evaluated against some predetermined standard and not determined relative to that of the employee's peers.

Cullen and Perez-Truglia (2022) highlight an important point that have mostly eluded previous research on pay secrecy. Vertical inequality makes up the vast majority of the pay inequality at firms. Less than 5 percent of the pay inequality is horizontal. This is important because pay inequality, hence, pay secrecy, have opposite horizontal and vertical effects. Revealing information regarding horizontal pay inequality has detrimental effects. It reduces job satisfaction, effort and productivity and increases separation rates. In contrast, revealing information regarding vertical pay inequality has positive effects. It improves the effort and performance of all employees who aspire to be promoted to that level. Indeed, this explains why most of the pay inequality is vertical. Firms are better off providing financial incentives vertically, in the form of promotions, rather than providing horizontal incentives such as pay-for-performance.

This, however, questions the wisdom of pay secrecy where there is opportunity for upward mobility in the organization. It seems that pay secrecy would be beneficial only where employees have little expectation of reaching higher echelons in their organizations, or where pay-for-performance has sufficiently high impact on productivity but the differences in productivity cannot be evaluated in a manner perceived by the employees as inherently fair.

Lastly, insights from the literature on gift exchange [Akerlof (1982)] and efficiency-wages suggest that paying employees more than they expect (e.g., based on their knowledge of co-workers salaries) may motivate them to exert higher efforts. DellaVigna et al. (2022) designed field experiments, hiring 446

workers to do clerical work, offering them varying piece-rates and surprising them with gifts in cash and in-kind,[6] right before and immediately after the completion of the task. There was very small elasticity of productivity to motivation (piece rate increases) hence the presence of gift exchange effects could be measured only on the willingness of workers to work extra hours, namely, labor supply effect. They found that gifts significantly increased labor supply and that the effect of the gift did not decay quickly over time. They found that workers were motivated by their feeling of 'warm glow' [Andreoni (1989, 1990)] namely, the positive feeling of being appreciated, adhering to social norm of working hard to justify the employer's kindness towards them, and signaling pro-sociality.

In addition to the empirical strand of the literature on pay secrecy there exists another (smaller) strand in the literature which explores the theoretical foundations of incorporating pay secrecy arrangements in labor contracts.

In an early study, Danziger and Katz (1997) demonstrate how, in equilibrium, a wage-secrecy convention facilitates risk shifting between firms and workers in response to productivity shocks. They argue that a spot market in labor is associated with a high degree of wage uncertainty, whereas binding workers to their firms results in production inefficiency. Wage secrecy reduces labor mobility by reducing the ability of workers to locate higher-paying firms. Hence, secrecy makes risk shifting feasible, but avoids the more extreme inefficiencies caused by the rigidity of a binding convention.

More recently, Long and Nasiry (2019) consider a principal-agent setup in which agents exhibit relative remuneration concerns. The firm jointly decides on the form and level of compensation and whether or not to disclose information about peers' level of remuneration (thereby controlling for the degree of social comparisons). When determining the optimal wage transparency policy, the firm is shown to trade-off a positive motivation effect that induces a worker to exert additional efforts to outperform his peer when exposed to social comparisons, and a negative participation effect that induces the firm to pay more (than in the absence of social comparisons) to ensure the worker's participation. The paper focuses on positive aspects and examines the role played by the work environment in shaping the conditions under which an individual firm would benefit from pay transparency. In particular, counter to conventional wisdom, pay transparency is shown to be desirable in work environments with strong inter-dependencies amongst workers.

Cullen and Pakzad-Hurson (2021) study the question of who benefits from pay transparency in equilibrium. They build a dynamic wage-bargaining model and test the model's theoretical predictions empirically in an event study framework, taking advantage of the roll-out of U.S. state laws between 2004 and 2016 aimed at facilitating communication about pay between coworkers. They also conduct a formal analysis of transparency policies investigated in other studies, and demonstrate how wage effects are predicted in these settings by the level of individual bargaining power. They find that under pay-transparency employers

[6]Kube et al. (2012) found the impact of in-kind gifts to be larger than the impact of monetary gifts. DellaVigna et al. (2022) found monetary and in-king gifts to have a similar impact.

credibly refuse to pay high wages to any one worker to avoid costly renegotiations with others. Employees are less concerned with securing a high initial wage because they know they will be able to renegotiate once someone else in their position receives a raise. The overall equilibrium effect is lower average wages compared to equilibrium under pay-secrecy, because the employer can commit to a maximum wage. It becomes monopsonistic. Each employee will either accept the transparent maximum wage, or will consume her outside option. When employees are unionized, namely, bargain collectively, wage transparency's effect is muted, because employees are less able to exploit differences in their outside options to secure heterogeneous wages.

Finally, Gürtler and Struth (2021) provide a formal analysis of wage transparency rules. They assume that the firm manipulates the information the workers receive to lower the workers' wage demands. The firm's strategic behavior depends on whether the hazard rate of productivity is increasing or decreasing. The hazard rate determines how likely it is that a marginal increase in a worker's wage demand leads to a rejection of that demand. If the hazard rate is relatively large, workers act cautiously by demanding relatively low wages. As a result, the firm wishes the workers to believe that productivity is in the region where the hazard rate is large. It does so by reducing the hiring threshold below the productivity level, repelling some workers who were profitable hires. On the contrary, if the hazard rate is decreasing, the firm wishes to signal a relatively high productivity, and it therefore decides to hire some workers whose wage demands exceed their productivity, thereby suffering a loss. Wage transparency rules have another effect, which is the workers' learning of the firm's productivity and wage-setting behavior. This allows them to make better decisions than if they stayed uninformed, leading to relatively greater payoffs. When productivity has an increasing hazard rate, the strategic and the learning effect oppose each other, and the change in the workers' payoffs depends on the dominating effect. It could lead to lower average wages compared to wages under secrecy rules. In contrast, when the hazard rate is decreasing, the workers' payoffs surely increase.

1.3 Overview of our setup and the key insights

In this study, we aim to explore theoretically the role played by pay secrecy arrangements in labor contracts. We consider a setup in which workers engage in social comparisons and exhibit what we refer to as 'relative ambition concerns' with respect to members of their production team. In particular we assume that an increase in the worker's teammates' level of remuneration adversely affects his incentives. This 'negative externality' notably applies to both underpaid and overpaid workers. The firm can jointly determine the level of compensation offered to the worker (a piece-rate remuneration per unit of effort) and the degree of pay transparency – namely whether or not to inform the worker about his peers' level of compensation. Thus, under a pay-secrecy regime, the worker is only informed about the average wage level of his teammate, whereas with pay transparency in place, the worker knows exactly how much his

teammate earns.

Our model provides two key novel insights. On the positive side, we characterize necessary and sufficient conditions for the desirability of incorporating pay secrecy clauses in the optimal labor contract offered by the firm. Pay secrecy is shown to be desirable when the degree of relative ambition exhibited by the workers is sufficiently large, namely they care a lot about their 'rank' in the team, remuneration-wise. With a high level of relative ambition, the firm may gain from inducing social comparisons to elicit efforts via offering an asymmetric pay policy to ex-ante homogeneous workers. However, implementing a transparent discriminating pay policy that introduces asymmetric pay policy across all production teams (which indeed triggers social comparisons) is too costly due to the complementarity in production, which requires high effort levels from all team members. A better strategy is to resort to a confidential discriminating pay policy, in which the wage rate is reduced only in fraction of the production teams and workers are only informed about their own level of compensation. By doing so one can maintain production teams in which workers are remunerated with a high wage rate (as dictated by complementarity), but perceive that their teammates earn (on average) a low wage rate.

On the normative side, we defy the conventional wisdom reflected in the popular debate on pay transparency, which suggests that secrecy arrangements serve the firms' interest and are unequivocally harmful for workers. We demonstrate that workers may in fact gain from pay secrecy, which serves to mitigate the negative externality associated with social comparisons, when the degree of relative ambition is sufficiently high.

1.4 Structure of the paper

The paper is organized as follows. In Section 2 we present the basic set-up and discuss the key assumptions. Section 3 provides the main results of the paper divided as follows: in Section 3.1 we study the single-firm problem and explore the positive foundations of pay secrecy, in Section 3.2 we test the robustness of our key positive insights, and in Section 3.3 we focus on the market equilibrium and the normative implications of the model. Section 4 provides concluding remarks.

2 The model

Consider a group of employees working for a single representative firm. The goal of the firm is to minimize costs, subject to a fixed level of production.[7] The firm rewards employees through a fixed piece-rate per unit of effort (which is assumed to be observable by the firm). In turn, the (ex-ante) homogenous

[7] Our initial focus is set on the positive foundations of pay secrecy in a partial equilibrium setup; namely, exploring the potential role of pay secrecy as a means to economize on the firm's cost of production. We later on extend the analysis to a general equilibrium setup, which allows for free entry of profit-maximizing firms, and examine the normative implications of pay-transparency regulation.

employees choose their efforts levels, taking into account their own pay structure, as well as that of their counterparts.

Acknowledging that employees are concerned with their relative compensation, firms may desire to employ some form of a non-disclosure policy regarding wages, that prevents employees from sharing wage related information with their peers, so as to mitigate the associated dis-incentivizing effects. In its quest for an optimal outcome, the firm, in addition to determining the level of remuneration via setting the piece-rate, can exercise a *confidential-pay policy*. We capture this latter secrecy aspect by assuming that employees are matched into productive couples such that individual partners are used as *benchmarks* where, under a confidential pay policy, a worker is only informed about the average level of remuneration of his/her teammate.

More formally, the process begins when all employees are matched into couples. For simplicity, we assume that there is a continuum of employees with a mass of 2, such that the total mass of couples is normalized to unity. Next, the firm commits to a *feasible policy* $F \in \Delta\mathbb{R}^2_+$ which dictates the distribution of wages across all matched couples. A feasible policy must sustain two conditions. First, the marginals must be identical to reflect the true distribution of wages among employees. Second, F is either a product distribution (i.e., featuring independent marginals) or satisfies a property that both coordinates are deterministically dependent. A product distribution is considered a *confidential-pay policy* (CP policy) since partners' wages are independent, whereas the case of deterministically dependent wages reflects an *observable-pay policy* (OP policy) among matched employees.[8] Denote the set of feasible policies by $\mathcal{F}$. To abstract from technical issues (which are beyond the scope of this paper), we assume that all feasible policies have finite support.[9]

Once wages are distributed according to F, the employees are privately informed of their realized pay. All employees possess the same utility function

$$U(e, w; w_p) = ew - \frac{e^2}{2}\left[1 - \gamma(w - w_p)\right],$$

where $e \geqslant 0$ is the employee's effort level, $w \geqslant 0$ is the employee's wage (per unit of effort), $w_p \geqslant 0$ is the wage of the employee's partner (also, per unit of effort) as perceived by the employee based on the firm's policy, and $\gamma \geqslant 0$ is a measure of the employees' sensitivity towards their partners' income. Note that in case $\gamma = 0$ we revert to the neo-classical model without the social-comparison component of $w - w_p$. For concreteness and to simplify the exposition, we assume the existence of a single consumption good

[8] An alternative way to consider the firm's strategy is through a two-stage decision process. First, the firm chooses a distribution over $\mathbb{R}_+$, which dictates the wages of all employees before they are matched. Next, employees are either randomly matched to generate independent wages within couples (a CP policy), or employees are deterministically matched to produce an OP policy.

[9] We simplify by considering only two polar cases: perfect observability (OP) and no obserbility (CP). In general, one could allow for partial observability so that the firm can determine the optimal extent of pay transparency. Clearly, this would only strengthen the case for pay secrecy.

produced by the representative firm, the price of which is normalized to unity. Note that under a CP policy, the employee only knows that distribution over w_p, and not the realized wage of his partner (as in the case of an OP policy), and hence maximizes his expected utility accordingly. By the linearity of the utility function w.r.t. w_p, however, the latter is equivalent to maximizing the utility w.r.t. the effort level while replacing w_p with the expected wage.

The bracketed expression in the utility formulation denotes the cost of effort and can be decomposed into two components. The first, given by $e^2/2$, is a standard cost function which satisfies monotonicity and convexity properties. The second, given by $\frac{e^2}{2}\gamma(w - w_p)$, captures our notion of relative ambition. The parameter γ measures the degree of relative ambition exhibited by the worker and reflects the extent to which social comparison considerations affect the worker's effort choice. In particular, for a given level of w, an over (under)- paid worker is willing to exert a higher (lower) level of effort; respectively. We henceforth refer to the relative ambition effect on the effort choices associated with an over (under-) paid worker, correspondingly, as ahead-seeking and behind-aversion motive. The wage rate w_p associated with the team-mate serves as a benchmark, relative to which the worker measures his ranking within the team. An increase in the benchmark exerts a negative externality on the worker and induces him to exert a lower effort level. Note that the second component is plausibly dropped once $w = w_p$, which means that a symmetric pay structure reverts our formulation to the neo-classical one with no social comparison considerations in place.

The output per matched couple, exerting efforts e and e_p, is determined by a symmetric Cobb–Douglas production function $Q(e, e_p) = e^{1/2} \cdot e_p^{1/2}$. Assuming that the required production is exogenously fixed to X, the firm is confronted with the following cost-minimization problem,

$$\begin{aligned} \min_F C &= \min_{F\in\mathcal{F}} \mathbf{E}[ew + e_p w_p], \\ \text{s.t.} \quad \mathbf{E}[Q(e, e_p)] &= e^{1/2} e_p^{1/2} \geqslant X > 0, \\ e &= \operatorname{argmax}_{\tilde{e}} \mathbf{E}[U(\tilde{e}, w; w_p)], \\ e_p &= \operatorname{argmax}_{\tilde{e}} \mathbf{E}[U(\tilde{e}, w_p; w)], \end{aligned}$$

where the expectation operator $\mathbf{E}[\cdot]$ represents the aggregation across all matched couples at the firm's level, and taken w.r.t. F. A direct optimization shows that $e = \max\{\frac{w}{1-\gamma(w-w_p)}, 0\}$ and a similar equality holds for e_p. To avoid the unrealistic possibility of infinite effort levels, we cap all effort levels by a sufficiently large $\overline{e} > X$. Notice the crucial difference between OP and CP policies in the eyes of the employees. Under an OP policy, every worker knows his partner's realized wage w_p, whereas a CP policy suggests that employees only know their partners' expected wage based on F. To ensure the feasibility of a CP policy, we assume that workers are unable to observe the effort levels of their teammates (otherwise, they could infer wages based on observed level of effort, which will render these policies irrelevant). Notice that wage policies depend on the employees' effort levels, rather than on the couples' output. This

eliminates the moral hazard problem of evaluating workers according to their aggregate production.[10]

We simplify our exposition by setting the workers' reservation utility to zero. Thus, a scenario in which a worker is paid a zero wage-rate and exerts no effort should be interpreted as one in which the worker is paid his reservation wage and exerts a minimal effort level. Furthermore, both behind-aversion and ahead-seeking are captured by the same parameter, γ, in the given utility function, which implies a symmetric effect, upwards and downwards. One should note that all these assumptions are invoked for tractability purposes. In Subsection 3.2, we extend the baseline set-up by incorporating *asymmetry between the behind-aversion and the ahead-seeking effects*. In addition, we let the reservation utility of the workers (and hence their effort levels in equilibrium) to be bounded away from zero.

3 Main results

3.1 Micro foundations and positive implications

We first turn to explore the positive foundations of pay secrecy. Our first main result establishes the case for pay secrecy, by providing a sufficient condition for a confidential-pay policy to dominate any observable-pay policy. Formally,

Theorem 1. *If $\gamma X > 3$, then an observable-pay policy is suboptimal.*

Proof. Fix $\gamma X > 3$. Take an OP policy where every matched couple is paid w and w_p w.p. 1. Evidently, there is no need for additional wage levels due to the structure of the production function. In case $w = w_p$, then $e = e_p = w$, $\mathbf{E}[Q(e, e_p)] = e = w = X$ and $C = 2ew = 2X^2$. In case $w \neq w_p$, we shall prove that $C > 2X^2$. The cost C in this case is given by

$$C = ew + e_p w_p = \frac{w^2}{1 - \gamma(w - w_p)} + \frac{w_p^2}{1 + \gamma(w - w_p)},$$

[10] Our setup bears some resemblance to the model in Charness and Kuhn (2007) who employ a reduced form gift-exchange framework in which workers' effort levels are assumed to be a linear function of their own- as well as their co-workers' wage rates, capturing both the behind aversion and ahead-seeking features discussed before. Charness and Kuhn (2007) consider the case of heterogeneous workers (differing in productivity) and examine the role of wage compression used by the firms to elicit work efforts. Notice that in our case labor contracts are complete, so there are no moral hazard issues to overcome. Still, the fact that some of the ex-ante identical workers are "overpaid" can by interpreted as a gift.

and the production constraint implies that $e \cdot e_p = \frac{w \cdot w_p}{[1-\gamma(w-w_p)][1+\gamma(w-w_p)]} = X^2$. Then,

$$\begin{aligned} C - 2X^2 &= \frac{w^2}{1-\gamma(w-w_p)} + \frac{w_p^2}{1+\gamma(w-w_p)} - 2\frac{w \cdot w_p}{[1-\gamma(w-w_p)][1+\gamma(w-w_p)]} \\ &= \frac{w^2 + w^2\gamma(w-w_p) + w_p^2 - w_p^2\gamma(w-w_p) - 2w \cdot w_p}{[1-\gamma(w-w_p)][1+\gamma(w-w_p)]} \\ &= \frac{(w-w_p)^2 + \gamma(w^2-w_p^2)(w-w_p)}{[1-\gamma(w-w_p)][1+\gamma(w-w_p)]} \\ &= \frac{(w-w_p)^2 + \gamma(w+w_p)(w-w_p)^2}{[1-\gamma(w-w_p)][1+\gamma(w-w_p)]} > 0, \end{aligned}$$

and the strict inequality follows from the fact that the denominator is positive and $w \neq w_p$. Thus, we conclude that an optimal OP policy can achieve a minimal cost of $2X^2$.

Now, fix a CP policy so that every employee receives either a wage of w w.p. $p \in (0,1)$, or nothing, otherwise.[11] There is a probability of p^2 for a couple to be productive, hence

$$\mathbf{E}[Q(e,e_p)] = p^2 e = p^2 \frac{w}{1-\gamma(w-wp)} = \frac{p^2 w}{1-\gamma w(1-p)} = X,$$

which translates to $w = \frac{X}{p^2+\gamma X(1-p)}$ and $e = \frac{X}{p^2}$. The expected cost under this CP policy is

$$C = 2pew = 2p \cdot \frac{X}{p^2} \cdot \frac{X}{p^2+\gamma X(1-p)} = \frac{2X^2}{p^3+\gamma Xp(1-p)}.$$

We conclude that the given CP policy dominates any OP policy if there exists a probability $p \in (0,1)$ so that $p^3+\gamma Xp(1-p) > 1$. Notice that $p = 1$ yields an equality as expected, since the symmetric OP policy is embedded in the CP policy given $p = 1$. If indeed $\gamma X > 3$, the LHS of the inequality is decreasing when p tends to 1, so that the result follows for $\gamma X > 3$ and p sufficiently close to 1. ■

The rationale behind the sufficient condition stated in theorem 1 is twofold. First, recall that a higher value of γ reflects a higher sensitivity towards others' wages in the payoff function of a worker. The latter has a two-sided effect when the firm resorts to an asymmetric pay structure - inducing underpaid workers to reduce their effort levels and, correspondingly, inducing over-paid ones to increase their efforts. With an OP policy in place such asymmetry turns out to be suboptimal, due to the complementarity and symmetry of the production function and the ex-ante homogeneity of workers, that push in the direction of wage compression. However, with a CP policy in place, the firm does not have to implement the asymmetric pay-structure across the board (for each and every couple). The firm may maintain symmetry amongst a large fraction of its productive couples (in which both workers are being remunerated generously), but can still reduce the benchmark (average pay), by reducing the level of remuneration for a fraction of the workers. Confidentiality serves to mitigate the tension between the desirable ahead-seeking effect amongst

[11] Throughout the proof, we assume that fixed wages and probabilities meet the production requirement.

over-paid workers and the undesirable behind-aversion effect amongst under-paid ones. When γ is large enough the ahead-seeking effect prevails and switching to a CP policy economises on the firm's costs.[12]

Second, cost minimization implies that the optimal effort level is given by $e = \frac{w}{1-\gamma(w-w_p)}$. Thus, under a symmetric OP regime, effort grows linearly with the wage rate. With an asymmetric pay structure, however, a proportional increase in the wage rates induces a larger-than-proportional increase in effort, due to the enhanced bench-marking effect captured by the term on the denominator $\gamma(w - w_p)$, which reflects the worker's other regarding preferences. An increase in output, which requires an upward adjustment in wages, serves hence to amplify the desirable ahead-seeking effect associated with a CP regime. The latter makes the case for a switch from a symmetric OP policy to an asymmetric CP regime yet stronger. This also explains the dependency on X (namely, the condition $\gamma X > 3$), since a higher production level dictates higher wages, thus boosting the advantages of CP policies.

Our next two results show that CP policies are not uniformly optimal, i.e., across all values of γ and X. We first show that for every CP policy one can find a sufficiently small γ for which the CP policy is dominated by the symmetric OP policy. We then provide an additional mild condition under which one can find a sufficiently small γ for which all CP policies are dominated by the symmetric OP policy, that is, a CP policy is suboptimal.

Proposition 1. *For every confidential-pay policy F, there exists $\gamma^F > 0$, such that for every $\gamma < \gamma^F$, the symmetric observable-pay policy dominates F.*

Proof. Fix a CP policy F and $\gamma > 0$. Recall that F has finite support, so denote the wage levels and respective probabilities by (w_i, p_i) for every $1 \leqslant i \leqslant n$. Let $\mu = \sum_i p_i w_i$ denote the expected wage. The individual effort level is $e_i = \frac{w_i}{1-\gamma(w_i-\mu)}$. So the cost and the production constraint are given by

$$\mathbf{E}[Q(e, e_p)] = \sum_i \sum_j p_i p_j \sqrt{e_i e_j} = \sum_i p_i \sqrt{e_i} \sum_j p_j \sqrt{e_j} = \left(\sum_i p_i \sqrt{e_j} \right)^2 = X,$$

and

$$C = 2 \sum_i p_i w_i e_i.$$

[12]To further illustrate the desirable role of a confidential pay policy, notice that in the CP policy invoked in the proof of Theorem 1, bench-markers are offered their reservation wage rate that induces a zero effort level. By virtue of the complementarity of the production function, this implies that no production takes place in couples to which bench-markers are assigned. Thus, the output loss due to the behind-aversion effect is not sensitive to the degree of relative ambition. An increase in γ, however, contributes to a larger output gain due to the ahead-seeking effect in productive couples to which no bench-markers are assigned. For a sufficiently large increase in γ the enhanced output gain more than offsets the fixed output loss, thereby establishing the case for a CP policy. In contrast, an OP policy in which bench-markers are remunerated according to their reservation wage induces no production across the board, hence yielding no output gain from the ahead-seeking effect.

In case γ tends to 0, we obtain that $e_i = w_i$, and the equations above translate to $\sum_i p_i y_i = \sqrt{X}$ and $C = 2\sum_i p_i y_i^4$, where $y_i = \sqrt{w_i}$ for every i. This yields a constrained-minimization problem with a convex objective function and a linear constraint. Using Jensen's Inequality, the strict and unique minimal value is achieved by contracting all wages $\{w_i\}_{i=1}^n$ to μ. Namely, the strict minima is achieved when all employees receive the same wage of μ, which coincides with a symmetric OP policy. Since this is a strict minima and all functions are continuous, we can deduce that there exists $\gamma^F > 0$, such that for every $\gamma \in [0, \gamma^F]$, the symmetric OP policy dominates F. ■

If one plausibly assumes that enforcing a CP policy entails some costs to the firm, an immediate implication of Proposition 1 is that for a CP policy to dominate an OP policy the degree of relative ambition, captured by γ, must be sufficiently high. Formally,

Observation 1. *Let $K > 0$ denote the fixed cost of enforcing any* CP *policy. Then, there exists $\gamma^K > 0$, such that for every $\gamma < \gamma^K$, a confidential pay policy is suboptimal.*

Observation 1 states that γ has to be bounded away from zero to support a substantial gain from a shift from a symmetric OP policy to a CP one to cover the entailed enforcement costs.

3.2 Robustness

In this subsection, we extend our basic set-up to fit the empirical data which alludes to asymmetry between the upwards and downwards effects concerning relative ambition. In addition, we allow for the employees' reservation utility to be bounded away from zero, thus bounding their wages and effort levels from zero. Formally, we consider the following, more general, utility specification:

$$U(e, w; w_p) = ew - \frac{e^2}{2}\left[1 - (w - w_p) \cdot \begin{cases} \gamma_+, & \text{for } w \geqslant w_p, \\ \gamma_-, & \text{for } w < w_p, \end{cases}\right],$$

where $\gamma_+ > 0$ and $\gamma_- > 0$ denote the ahead-seeking and the behind-aversion effects, respectively. In addition, the employees' reservation utility is given by $\underline{u} > 0$. To be clear, any worker with any wage rate must secure an expected utility exceeding $\underline{u}$.

In-line with previous empirical studies showing that the behind-aversion effect is stronger than the ahead-seeking one, we assume that $\gamma_- > \gamma_+ > 0$. In simple terms, this assumption states that the under-paid worker's incentives to reduce his level of effort, due to social comparisons, are stronger than the corresponding incentives of the over-paid worker to increase his level of effort. Theorem 2 below shows that for a sufficiently high production level, the CP policy prevails over any OP policy, thus extending the result stated in Theorem 1 for the main utility specification.

Theorem 2. *For every $\gamma_\pm$ and for every reservation utility $\underline{u}$, there exists a sufficiently large production level X, such that an observable-pay policy is suboptimal.*

Proof. A straightforward verification, similar to the computation given in the proof of Theorem 1, shows that an optimal OP policy dictates a uniform pay $w = w_p$ across all workers. Again, similar to previous proofs, this generates a cost of $C = 2X^2$, where $e = w = X$, and the employees' utility is given by $u = e^2 - \frac{e^2}{2} = \frac{X^2}{2} \geqslant \underline{u}$. So, a necessary condition for a feasible symmetric OP policy is $X^2 \geqslant 2\underline{u}$.

Now, fix a CP policy so that every employee receives either a wage of w_1 w.p. $p \in (0,1)$, or a wage of $w_2 < w_1$, otherwise. Note that by taking $p = 1$ we revert back to the symmetric OP policy. The employees' effort levels are given by $e_1 = \frac{w_1}{1-\gamma_+(1-p)(w_1-w_2)}$ and $e_2 = \frac{w_2}{1+\gamma_- p(w_1-w_2)}$, respectively. The production condition is given by

$$Q = p^2 e_1 + (1-p)^2 e_2 + 2p(1-p)\sqrt{e_1 \cdot e_2} \geqslant X.$$

Since we follow a cost-minimization problem, the last inequality could be weakened to $p^2 e_1 = X$. Thus, $\frac{X}{p^2} = e_1 = \frac{w_1}{1-\gamma_+(1-p)(w_1-w_2)}$, which yields the following linear relation $w_1 = \frac{X[w_2\gamma_+(1-p)+1]}{p^2+X\gamma_+(1-p)}$ between the two wage levels. We can also compute e_2, as a function of w_2, and get

$$e_2 = \frac{w_2[p^2 + X\gamma_+(1-p)]}{p^2 + X\gamma_+(1-p) + \gamma_- pX - \gamma_- w_2 p^3}.$$

Using the given effort levels and wages, we can explicitly write the cost function

$$\begin{aligned} C &= 2pe_1w_1 + 2(1-p)e_2w_2 \\ &= 2p \cdot \frac{X}{p^2} \cdot \frac{X[w_2\gamma_+(1-p)+1]}{p^2 + X\gamma_+(1-p)} + 2(1-p)w_2 \frac{w_2[p^2 + X\gamma_+(1-p)]}{p^2 + X\gamma_+(1-p) + \gamma_- pX - \gamma_- w_2 p^3} \\ &= 2X^2 \frac{[w_2\gamma_+(1-p)+1]}{p^3 + X\gamma_+ p(1-p)} + 2(1-p)w_2^2 \frac{[p^2 + X\gamma_+(1-p)]}{p^2 + X\gamma_+(1-p) + \gamma_- pX - \gamma_- w_2 p^3}. \end{aligned}$$

Holding w_2 fixed, we can differentiate C w.r.t. p, and plug in $p = 1$ to get

$$\left.\frac{\partial C}{\partial p}\right|_{p=1} = 2X^2(-w_2\gamma_+ - 3 + X\gamma_+) - 2w_2^2 \frac{1}{1 + \gamma_- X - \gamma_- w_2}.$$

Thus, assuming that the reservation-utility condition holds, a sufficient condition for the CP policy to dominate the OP one is $\left.\frac{\partial C}{\partial p}\right|_{p=1} > 0$, implying that there exists an interior solution where $p < 1$ which yields a lower cost (since the symmetric OP policy is embedded in the given CP policy by taking $p = 1$).

The reservation-utility condition for the low-pay worker is given by

$$\begin{aligned} u &= e_2w_2 - \frac{e_2^2}{2}[1 - \gamma_-(w_2 - pw_1 - (1-p)w_2)] \\ &= e_2w_2 - \frac{e_2^2}{2}[1 + p\gamma_-(w_1 - w_2)] \\ &= \frac{w_2^2}{2(1 + \gamma_- p(w_1 - w_2))} \geqslant \underline{u}. \end{aligned}$$

Taking the limit as p tends to 1, the last inequality reverts to $\frac{w_2^2}{2(1+\gamma_-(X-w_2))} \geqslant \underline{u}$, and we shall use it as a strict inequality to ensure that the condition is also satisfied for smaller values of p (close to 1).

To conclude, the following inequalities establish sufficient conditions for a feasible CP policy to dominate the optimal OP policy

$$X^2 \geqslant 2\underline{u},$$
$$\frac{w_2^2}{2(1+\gamma_-(X-w_2))} > \underline{u},$$
$$X^2[\gamma_+(X-w_2)-3] - \frac{w_2^2}{1+\gamma_-(X-w_2)} > 0.$$

Notice that in the special case in which $\underline{u}=0$, substituting $w_2=0$ into the last inequality and re-arranging reverts back to $\gamma_+ X > 3$, as given in Theorem 1. It can further be verified that for $w_2=0$, both the symmetric OP feasibility condition (as a strict inequality) and the reservation utility condition (as an equality) are satisfied. In the general case where the reservation is bounded away from zero, one can fix $w_2 = X/2$, for example, and verify that all three conditions are satisfied (as strict inequalities) for a sufficiently large X. Moreover, in the limit when $p \to 1$, we know that $w_1 = e_1 = X$, so that $w_1 > w_2$, as needed, for p sufficiently close to 1, by continuity considerations. ■

Theorem 2 provides a natural extension to Theorem 1 and suggests that for a sufficiently high degree of relative ambition (bounded from below by γ_+) and a sufficiently high level of output X, provided that the feasibility of a symmetric OP policy and the reservation utility are guaranteed, a CP policy would dominate any OP policy. The mechanism at work is the same as in Theorem 1. A CP policy enables the firm to implement an asymmetric remuneration scheme in a probabilistic manner rather than across the board, and thereby exploiting the ahead-seeking effect to elicit additional efforts from overpaid workers, without the need to sacrifice significantly through the behind-aversion effect inducing lower efforts from underpaid workers.

Before concluding this section, two remarks are in order. First, notice that a reservation utility bounded away from zero sets a lower bound for the level of remuneration offered to the bench-markers (the underpaid workers who exert a relatively low effort). This in turn requires, for any given output level, a higher degree of relative ambition to establish the case for pay secrecy. Second, notice that the wage offered to underpaid workers induces a positive effort level in this case.

3.3 General equilibrium analysis and normative implications

In this section, we extend our analysis from the individual-firm problem to a market equilibrium. Specifically, we extend our analysis by allowing for free entry of firms, which, in turn, enables workers to optimally choose a firm. We use this extension to supplement the positive analysis with an examination of the normative implications of our model. In particular, we address the question of the social desirability of pay-transparency regulation; namely, whether ruling out the use of CP policies by the firms would be socially desirable from the employees' perspective. Before we proceed a brief remark is in order.

Whereas a general equilibrium setup is the natural environment to examine the normative implications of our model, a simple observation about a partial equilibrium setting (in which the number of firms is fixed and workers are on the long side of the market) defies the conventional wisdom about the merits of pay transparency. [13] In equilibrium, the firm would set its output at the level which maximizes profits. Based on our previous derivations, the optimal OP policy would yield an output level of $X = 1/4$, generating a profit of 1/8 for the firm, and an expected utility of 1/32 for workers. It is straightforward to find CP policies that produce a higher profit and a higher expected utility, for the firm and workers, respectively. Thus, the common perception that pay-secrecy policies benefit firms at the expense of their employees, is far from being a forgone conclusion.

We turn next to study the general-equilibrium setting. By virtue of the homogeneity of both the firms and the workers, we will construct a symmetric Nash equilibrium. Denote by U the (ex-ante) expected utility derived by a typical worker in equilibrium. A typical firm is solving the following constrained maximization problem:

$$\begin{aligned} \max_{F\in\mathcal{F}} \quad & \mathbf{E}[Q(e,e_p)] - \mathbf{E}[ew + e_p w_p], \\ \text{subject to} \quad & \mathbf{E}\left[ew - \frac{e^2}{2}\left[1-\gamma(w-w_P)\right]\right] \geqslant U, \end{aligned}$$

where the expectation operator $\mathbf{E}[\cdot]$ is taken with respect to F, the wage distribution chosen by the firm.

In equilibrium, the constraint is clearly binding, hence all firms offer their workers the same expected utility. Moreover, in the face of free entry, incumbent firms are threatened by potential entrants, hence firms' rents are fully dissipated in equilibrium, so each operating firm derives zero expected profits. The value of U is calibrated to satisfy the expected zero-profit condition.

An alternative way to derive the value of U in equilibrium, is to solve for the wage policy that maximizes the expected utility of the typical worker subject to a constraint that the firm derives non-negative expected profits. It is straightforward to verify that the two formulations are equivalent. We will henceforth stick to the alternative formulation. Formally, a wage policy is given by the solution to the following constrained-maximization problem:

$$\begin{aligned} \max_{F\in\mathcal{F}} \quad & \mathbf{E}\left[ew - \frac{e^2}{2}\left[1-\gamma(w-w_P)\right]\right], \\ \text{subject to} \quad & \mathbf{E}[Q(e,e_p)] - \mathbf{E}[ew + e_p w_p] \geqslant 0. \end{aligned}$$

The formulation of the constrained maximization program implies that a firm is unable to attract new workers. Offering its prospective employees a higher level of expected utility than that obtained on equilibrium renders the firm unprofitable. In a symmetric equilibrium, all firms would offer the same wage distribution to workers and produce the same level of expected utility, $U = \mathbf{E}\left[ew - \frac{e^2}{2}\left[1-\gamma(w-w_P)\right]\right]$.

[13] Notice that a partial equilibrium setting in which workers are on the short side of the market is tantamount to a general equilibrium with free entry of firms, which obviates the need to analyze this scenario separately

Before we proceed to state and prove the main result of this section, two remarks are in order. First, notice that the existence of an equilibrium is premised on our assumption that all policies are supported on a finite number of (bounded) wage levels. The latter assumption implies that the above constrained-maximization problem is embedded in a compact set (finite dimension, closed, and bounded). Continuity of the objective function ensures, hence, that a solution for the constrained maximization problem exists. Second, notice that the maximization problem reflects an ex-ante notion of stability. Workers make their employment choices based on ex-ante expected remuneration considerations, which in turn induce firms, in equilibrium, to offer a labor contract which maximizes the expected utility of the representative worker. There are two possible ways to address ex-post stability concerns. First, since the production is in couple-formations, which embeds a production technology which exhibits complementarity between the two team-mates, the only ex-post feasible unilateral deviation of a worker is to claim his reservation wage (presumably the output of self-employed workers, or the imputed value of the worker's leisure time, if unemployed). Alternatively, one could assume that switching/re-negotiation costs are prohibitively costly, which would effectively allow each worker 'a single draw' from the wage distribution.

We turn next to prove the main result of this section, which serves as a counterpart to Theorem 1. The following theorem states that when the degree of relative ambition is sufficiently high, the only possible equilibrium is one in which all firms resort to CP policies. Formally,

Theorem 3. *Assume $\overline{e} > \frac{1}{2}$. When $\gamma > 6$, an observable-pay policy is suboptimal.*

Notice that the lower bound on $\overline{e}$ is needed to maintain the optimal OP policy feasible. Otherwise, under the symmetric OP regime, there is no full rent-dissipation by the employees, which allows firms to maintain strictly positive profits.

Proof. Fix $\gamma > 6$. We begin by showing that an asymmetric OP policy is suboptimal. Consider an OP policy with two distinct wages levels w_1 and w_2, so that after plugging-in the employees' effort levels and the production function, the maximization problem translates to

$$\begin{aligned} \max_{w_1 > w_2} \quad & U = \frac{1}{2} \cdot \frac{w_1^2}{2[1-\gamma(w_1-w_2)]} + \frac{1}{2} \cdot \frac{w_2^2}{2[1+\gamma(w_1-w_2)]}, \\ \text{subject to} \quad & \sqrt{\frac{w_1 w_2}{[1-\gamma(w_1-w_2)][1-\gamma(w_2-w_1)]}} \geqslant \frac{w_1^2}{1-\gamma(w_1-w_2)} + \frac{w_2^2}{1+\gamma(w_1-w_2)}, \\ \text{and} \quad & \frac{w_1}{1-\gamma(w_1-w_2)} \leqslant \overline{e}. \end{aligned}$$

Firms can increase wages until either the maximal-effort condition, or the profit condition, become binding. If the maximal-effort condition is non-binding, then a binding profit condition dictates that $U = \sqrt{\frac{w_1 w_2}{16[1-\gamma^2(w_1-w_2)^2]}}$. Thus, for every $d > 0$ such that the product $w_1 w_2 = d$, the expected utility is maximized whenever the difference $w_1 - w_2$ is maximized. In other words, to maximize U under an asymmetric OP regime, one should increase w_1 while holding $w_1 w_2$ fixed, until the maximal effort $e_1 = \overline{e}$

is reached. Therefore, we can now assume that $e_1 = \frac{w_1}{1-\gamma(w_1-w_2)} = \overline{e}$, and $w_1 = \frac{\overline{e}(1+\gamma w_2)}{1+\gamma\overline{e}}$. We can now insert $e_1 = \overline{e}$ and w_1 to the optimization problem and obtain

$$\max_{w_2 \geqslant 0} \frac{1}{4}\left[\frac{\overline{e}^2(1+\gamma w_2)}{1+\gamma\overline{e}} + \frac{w_2^2}{2-\frac{1+\gamma w_2}{1+\gamma\overline{e}}}\right].$$

Again, this follows from the fact that the maximal effort condition must be binding. Notice that the given function is increasing in w_2, so that the expected utility increases as w_2 tends to w_1. We thus conclude that an asymmetric OP policy is dominated by the symmetric one, as long as the latter is indeed feasible.

Now we can revert to a specific CP policy with two wage levels, 0 and w, w.p. $1-p$ and p, respectively. Again, this policy embeds the symmetric OP policy once $p = 1$, so we need to prove that $p = 1$ is suboptimal. The high-pay employee's effort level is $e = \frac{w}{1-\gamma w(1-p)}$, so that the production level is $p^2 e$ and the cost is $2pwe$, and the profit condition translates to $p^2 e - 2pwe \geqslant 0$, or equivalently, $p \geqslant 2w$. Thus, under the given CP regime, the maximization problem is

$$\begin{aligned} \max_{w,p} \quad & p\frac{w^2}{2[1-\gamma w(1-p)]}, \\ \text{subject to} \quad & p \geqslant 2w, \\ \text{and} \quad & \frac{w}{1-\gamma w(1-p)} \leqslant \overline{e}. \end{aligned}$$

In case the profit constraint is binding, one can plug-in $p = 2w$, and differentiate w.r.t. p to show that, for every $\gamma > 6$, the expected utility is strictly decreasing at $p = 1$. Note that the effort level is given by $e = \frac{w}{1-\gamma w(1-p)} = \frac{p}{2-\gamma p(1-p)}$, so that $e \to \frac{1}{2}$ as $p \to 1$. Since $\overline{e} > 0.5$, it follows by continuity that for p close to 1 the maximal-effort condition is satisfied. Thus, we conclude that there exists a CP policy with $p < 1$ and sufficiently close to 1, which dominates the symmetric OP policy (i.e., yields a higher expected utility than that evaluated at $p = 1$), as needed. ■

The rationale underlying the sufficient condition stated in Theorem 3 resembles the key insight emerging from our previous analysis. With a high degree of relative ambition, asymmetric-wage policies are desirable as means to raise the workers' expected utility, via an enhanced ahead-seeking effect through benchmark reduction. The latter can either be achieved through an asymmetric OP policy or via a CP policy. The complementarity and symmetry exhibited by the production function renders benchmark reduction across the board (that is, across all production teams) too costly, and hence an asymmetric OP policy becomes suboptimal. Resorting to a CP policy serves to mitigate the cost entailed by benchmark reduction.

3.3.1 Policy implications

We turn next to examine welfare implications of pay transparency regulation under the general equilibrium set-up. In Theorem 3 we have established that when the degree of relative ambition is sufficiently high, an OP policy equilibrium does not exist. It follows that a regulatory restriction which prohibits the use of pay secrecy clauses in wage contracts has a detrimental effect on workers' well-being. Formally,

Observation 2. *For sufficiently high levels of relative ambition, the exclusion of* CP *policies lowers workers' expected utility.*

To see this, consider the setting analyzed in the proof of Theorem 3 ($\gamma > 6$), with a regulatory restriction of wage secrecy, such that firms are obliged to maintain wage transparency. Once CP polices become infeasible, the only supported equilibrium is the (previously unattainable) suboptimal OP policy profile, which imposes an expected welfare loss on workers.

Before we conclude, a final remark is called for. One should notice that wage secrecy entails ex-post wage differences amongst ex-ante homogeneous workers. Thus, taking a broader welfare perspective, in order to assess the social desirability of pay secrecy, one should in principle account for the trade-off between ex-ante efficiency considerations and ex-post equity concerns. The popular normative debate on pay transparency (see the concluding section below) has been primarily confined to the equity dimension, arguing that pay transparency can serve to reduce pay gaps in the labor market. We, in contrast, emphasize the efficiency dimension, by alluding to the potential efficiency gains associated with wage secrecy arrangements.

4 Concluding remarks

Our paper provides a positive explanation for a pay secrecy convention to arise in equilibrium. The key insight of our analysis hinges on the combination of the relative ambition exhibited by workers and the complementarity of the production function. A worker cares about the level of remuneration of a subset of her peers, which serves as her reference group and are complementary to her in the production process. The firm is employing a confidential pay policy to strike a balance between the desire to pay those in the reference group a high wage in order to elicit high productive efforts (by virtue of the complementarity) and the need to pay those workers a low wage in order to mitigate the dis-incentivizing effect of relative ambition. The idea underlying the wage secrecy policy is to use the workers outside the reference group as bench-markers, serving to reduce the expected wage of the workers in the reference group without actually doing so. In our set-up, the worker's reference group was comprised of the members of her production team.

Clearly, a worker's reference group may be defined more broadly. Considering the employees' social network, one can assume that a worker's reference group is comprised of a weighted average of adjacent peers, where closer peers are weighted more heavily than others. Our set-up is, in essence, a specific version of this concept since only team members are positively weighted. Though the general set-up is left for future research, the key feature that would maintain our qualitative insights is that, as a whole, the reference group would exhibit complementarity with respect to the worker's productive effort.

Somewhat surprisingly, pay secrecy, which is often described as a strategic tool used by employers to improve their bargaining position in wage negotiations and as a means to mitigate the potentially

demoralizing effect of pay gaps on employees, may actually improve employees' welfare. In particular, in case relative ambition considerations are sufficiently manifest, we demonstrate that in a general equilibrium setting with free entry (which implies full dissipation of firms' rents), a confidential pay policy would maximize the ex-ante utility of workers.

Our focus in the paper was on the efficiency enhancing features of pay secrecy. Most of the popular debate on the desirability of pay transparency, however, revolves around equity aspects. A notable example is the ongoing public discourse on executive excessive pay (that hogged the limelight in the early 90's), which was the trigger for legislation mandating the disclosure of this information in financial statement of publicly traded firms and setting salary caps on executive levels of remuneration. More recently, the issue of pay transparency has resurfaced, in the context of gender pay gaps, where transparency has been suggested as a means to address persistent gender inequities in the labor market.[14]

To the extent that executives' high compensation schemes reflect economic rents (potentially driven by poor corporate governance) and gender gaps in the film industry are a byproduct of gender-based discrimination, pay transparency should be promoted as a means to reduce inequities without entailing efficiency costs (or better, mitigating those). Our analysis demonstrates, however, that pay secrecy may be desirable on efficiency grounds. Thus, determining the optimal extent of pay transparency involves resolving an equity-efficiency trade-off, which in our case is captured by the choice between an OP and a CP policy.

References

Akerlof, G. A. (1982): "Labor Contracts as Partial Gift Exchange," *Quarterly Journal of Economics*, 97, 543, publisher: Oxford University Press.

Andreoni, J. (1989): "Giving with Impure Altruism: Applications to Charity and Ricardian Equivalence," *Journal of Political Economy*, 97, 1447–1458, publisher: University of Chicago Press.

——— (1990): "Impure Altruism and Donations to Public Goods: A Theory of Warm-Glow Giving," *Economic Journal*, 100, 464, publisher: Oxford University Press (OUP).

Belogolovsky, E. and P. A. Bamberger (2014): "Signaling in Secret: Pay for Performance and the Incentive and Sorting Effects of Pay Secrecy," *Academy of Management Journal*, 57, 1706–1733, publisher: Academy of Management Briarcliff Manor, NY.

Bierman, L. and R. Gely (2004): "Love, Sex and Politics - Sure - Salary - No Way: Workplace Social Norms and the Law," *Berkeley Journal of Employment and Labor Law*, 25, 167.

[14] A recent exposure of e-mails by executives in Sony Pictures revealed substantial gender wage variation in wage contracts signed with top stars in the Film industry in Hollywood.

Bracha, A., U. Gneezy, and G. Loewenstein (2015): "Relative Pay and Labor Supply," *Journal of Labor Economics*, 33, 297–315, publisher: University of Chicago PressChicago, IL.

Breza, E., S. Kaur, and Y. Shamdasani (2018): "The Morale Effects of Pay Inequality," *Quarterly Journal of Economics*, 133, 611–663, publisher: Narnia.

Card, D., A. Mas, E. Moretti, and E. Saez (2012): "Inequality at Work: The Effect of Peer Salaries on Job Satisfaction," *American Economic Review*, 102, 2981–3003.

Charness, G. and P. Kuhn (2007): "Does Pay Inequality Affect Worker Effort? Experimental Evidence," *Journal of Labor Economics*, 25, 693–723, publisher: The University of Chicago Press.

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

Cullen, Z. B. and B. Pakzad-Hurson (2021): "Equilibrium Effects of Pay Transparency," *NBER Working Paper Series, No. 28903, August 2022*, [Revise and resubmit at Econometrica].

Danziger, L. and E. Katz (1997): "Wage Secrecy as a Social Convention," *Economic Inquiry*, 35, 59–69, publisher: Western Economic Association International.

Day, N. E. (2007): "An investigation into pay communication: Is ignorance bliss?" *Personnel Review*, 36, 739–762.

DellaVigna, S., J. A. List, U. Malmendier, and G. Rao (2022): "Estimating Social Preferences and Gift Exchange at Work," *American Economic Review*, 112, 1038–1074.

Gürtler, O. and L. Struth (2021): "Do workers benefit from wage transparency rules?" *ECONtribute Discussion Papers Series 105, University of Bonn and University of Cologne, Germany.*

Kube, S., M. A. Maréchal, and C. Puppe (2012): "The Currency of Reciprocity: Gift Exchange in the Workplace," *American Economic Review*, 102, 1644–1662.

Lobel, O. (2020): "Knowledge Pays: Reversing Information Flows and The Future of Pay Equity," *Columbia Law Review*, 120, 547–612, publisher: Columbia Law Review Association, Inc.

Long, X. and J. Nasiry (2019): "Wage Transparency and Social Comparison in Salesforce Compensation," *Management Science*.

Milkovich, G. T., J. M. Newman, and B. A. Gerhart (2013): *Compensation*, McGraw-Hill/Irwin.

Perez-Truglia, R. (2020): "The Effects of Income Transparency on Well-Being: Evidence from a Natural Experiment," *American Economic Review*, 110, 1019–54, publisher: American Economic Association.

Risher, H. (2014): "Pay Transparency Is Coming," *Compensation \& Benefits Review*, 46, 3–5.

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