PAY INEQUALITY, PAY SECRECY, AND EFFORT:
THEORY AND EVIDENCE

Gary Charness
Peter Kuhn

Pay Inequality, Pay Secrecy, and Effort: Theory and Evidence
Gary Charness and Peter Kuhn
NBER Working Paper No. 11786
November 2005
JEL No. C91, J41, M52

ABSTRACT

We study worker and firm behavior in an efficiency-wage environment where co-workers' wages may potentially influence a worker's effort. Theoretically, we show that an increase in workers' responsiveness to co-workers' wages should lead profit-maximizing firms to compress wages under quite general conditions. Our laboratory experiments, on the other hand, show that --while workers' effort choices are highly sensitive to their own wages-- effort is not affected by co-workers' wages. As a consequence, even though firms in our experiment tended to compress wages when wages became public information, this did not raise their profits. Our experimental evidence therefore provides little support for the notion that inter-worker equity concerns can make wage compression, or wage secrecy, a profit-maximizing policy.

Gary Charness
Department of Economics
University of California, Santa Barbara
Santa Barbara, CA 93106
charness@econ.ucsb.edu

Peter Kuhn
Department of Economics
University of California, Santa Barbara
Santa Barbara, CA 93106
and NBER
pjkuhn@econ.ucsb.edu

1. Introduction

A central component of many efficiency wage models is the notion that workers will withhold effort when they perceive that they have not been paid a fair wage (e.g. Akerlof and Yellen 1990; Bewley 1999). An equally influential notion has been that workers' perceptions of fairness depend, at least in part, on the wages paid to their co-workers (e.g. Frank 1984). Together, these hypotheses have been invoked to explain two compensation practices in firms: wage compression (Akerlof and Yellen 1990, p.265), and wage secrecy (Lawler 1990, pp. 238-242).

Despite the intuitive appeal of this chain of thought, some potential gaps in the argument deserve scrutiny. First, to our knowledge, neither wage compression nor wage secrecy has been explicitly derived as an optimal firm policy in an efficiency wage model.[1] Second, aside from various ethnographic accounts (e.g. Bewley 1999), the only evidence that worker effort depends on perceptions of fairness of any sort appears to be experimental in nature. Beginning with Fehr, Kirchsteiger, and Riedl (1993) and continuing with studies such as Fehr, Kirchler, Weichbold and Gächter (1998), Charness (2004), and Fehr and Falk (1999), experimental gift-exchange labor markets usually show that 'workers' offered a 'gift' (or wage) by 'firms' tend to reciprocate with return gifts ('effort'). This occurs even in completely anonymous, one-shot interactions when the dominant strategy is for workers to provide no effort at all. While this evidence suggests that effort decisions might respond to perceptions of fairness, it does not speak

---

[1] In our reading, Akerlof and Yellen (1990) come closest to doing this, but only *suggest* wage compression as an optimal firm response to co-worker equity concerns. Harris and Holmstrom (1982) generate optimal wage compression, but from insurance rather than equity motives; Danziger and Katz (1997) extend this model to justify wage secrecy. Frank (1984) derives equilibrium wage compression from co-worker equity concerns using a compensating-differentials argument. Finally, Lazear (1989) argues that wage gaps in tournaments should be attenuated when worker co-operation is important; Lazear's argument does not involve worker equity concerns.

to the notion that wage comparisons among co-workers might influence compensation policy, thereby giving rise to such policies as wage compression or secrecy.

If workers' effort levels respond to co-worker wages, does it follow that a profit-maximizing firm will compress wages, relative to productivity or some other standard? *Do* workers' effort decisions in fact depend on their co-workers' wages? And if so, are relative-wage effects strong enough to justify either a substantial degree of wage compression, or wage secrecy as a profit-maximizing policy? In this paper we examine all of these questions, the first using a simple theoretical model; the latter two in a laboratory experiment. Because it is a "gift-exchange" experiment, we argue that our experiment is particularly well suited for detecting 'equity'-driven co-worker wage effects on effort: The very nature of the experiment (a) eliminates influences on worker behavior *other* than fairness considerations, and (b) makes it costless for workers to punish firms quite severely for wage differentials that are perceived to be unfair. Within the context of the experimental literature, our experiments thus extend (with a few modifications) gift-exchange labor-market models to the case where two (differently-productive) workers are employed by each principal.[2]

Our main theoretical results are as follows. We consider alternative cases in which an effort response to co-worker wages occurs equally when own wages exceed or fall short of co-workers' wages (*symmetric*) or only when one is paid less than a co-worker (*asymmetric*); in both cases, an increase in workers' responsiveness to co-worker wages should lead profit-maximizing firms to choose greater wage compression. In the symmetric case, however, we find a surprising degree of 'neutrality' of other outcomes (in particular effort, productivity, and

---

[2] While—as noted—there have been many experimental studies of labor markets, only a few allow for multiple workers per firm. Güth, Königstein, Kovács and Zala-Mezõ (2001) consider a two-part contract in a principal-agent relationship and find that making work contracts observable leads to a greater degree of pay compression. On the other hand, Cabrales and Charness (1999) find no evidence of one agent having much concern about the payoff of

profits) to workers' equity-related behavior. For example, in the case of a linear (own-wage) effort supply function, all these outcomes are unaffected by how much workers' effort depends on their co-worker's wage. Further, the intuition behind wage compression in the symmetric case is somewhat unexpected: Workers' responsiveness to relative wages give firms more 'leverage' in the sense that a smaller wage gap is needed to elicit the fixed, profit-maximizing effort differential between the workers. Finally in the symmetric case the amount of wage compression predicted by the model can be quantified under quite general conditions, and is relatively modest. For example, if workers respond *as much* to relative wages as they do to own wages —which seems to us a plausible upper bound to equity-motivated behavior— then in the linear effort supply case the profit-maximizing wage gap between differently-productive workers is reduced by exactly half. Under *no* circumstances (in the symmetric case) is it ever optimal to pay equal wages to differently-productive workers.

In the asymmetric case, in contrast, equilibrium levels of effort and output *are* affected by workers' equity-motivated behavior, and firms *are* harmed when workers respond to relative wages. Optimal wage compression is harder to quantify in general (for example the functional form of the production function now matters), but is greater than under symmetry in the following sense: Above a threshold level of workers' responses to co-workers' wages, strictly-egalitarian wage policies can now be profit-maximizing.

Empirically, while effort is highly sensitive to own wages in our data, we detect little or no response of effort to co-worker wages. This result surprised us, given the prevalence of the notion that worker 'jealousy' is so important. That said, we can think of at least four considerations that may help explain this unanticipated result. First, it is of course important to

---

the other agent, although the principal's payoff was important for agent behavior. Importantly, neither of these studies involves gift exchange, as principals choose contract menus with (essentially) take-it-or-leave-it options.

make a distinction between *saying* one cares about relative wages (perhaps for purely strategic reasons) and being willing to act differently because of them. Using a co-worker's wage as a reference point in verbal negotiations (Babcock *et al.*, 1996) is very different from withdrawing effort when one feels a co-worker is paid too much, and most of the anecdotal evidence on this issue refers to the former activity only. Our results thus do not rule out the possibility that wage compression or secrecy might allow managers to lead a "quiet life" (Bertrand and Mullainathan 2003), freer of equity-driven complaints than it otherwise would be.

A second possible factor explaining our surprising results might loosely be called the limited dimensionality of workers' 'signal space'. *If* a worker cared about a co-worker's wages, then we can imagine the worker as wanting to send signals to the firm about the two things the worker cares about: his own wage *and* his co-worker's wage. The problem, of course, is that the worker has access to only one signal—effort—with which to signal his happiness on these two dimensions. In this situation it does not seem implausible to us that when forced to choose, workers focus on their own pocketbooks rather than their neighbors'. We cannot think of any good reason why this inherent 'dimensionality' problem would be any less severe in the field than in the laboratory.[3]

Third, while the 'jealousy' hypothesis that motivated this paper has considerable intuitive appeal, we note that plausible arguments can also be made for *positive* concerns with co-worker wages: workers have been known to go on strike, or to refuse to cross picket lines, to raise the wages of their co-workers, even when those co-workers are paid more than they are. If plausible

---

[3] Of course, in some work environments, effort may have a second dimension (such helping co-workers) that *is* affected by relative pay. Clearly, this is an interesting question for further experimentation, but it seems obvious to us that the one-dimensional effort case is the place where experimental investigation of co-worker wage effects on effort should begin.

stories involving both positive and negative concerns with co-worker wages can be told, we should perhaps not be so surprised to find a zero effect.

Finally, we note that our results are in fact consistent with a considerable body of recent experimental evidence on three-person games. In fact, a behavioral regularity seems to be emerging that the payoffs of (third-party) agents without direct power in the decision-making process are essentially completely ignored when there is another party involved whose actions affect one's payoff. For example, Güth and van Damme (1998) report the results of a three-player ultimatum game with an inactive third party. A proposer suggests an allocation for the three players, and the designated responder decides whether to accept or reject the proposal; a proposal is implemented if it is accepted, but all players receive zero if it is rejected. The proposed share for the inactive player was always quite low; nevertheless, rejections were very infrequent.

Kagel and Wolfe (2001) study a three-person ultimatum game in which one person suggests a three-way allocation. Each other person individually chooses whether to reject or to accept the proposal if he or she is (later) chosen to be the active responder, with one response selected at random for implementation; the 'consolation prize' received by the inactive player in the event of a rejection is varied across sessions. Despite the predictions of all of the distributional models, the data show essentially no effect of the consolation prize on rejection rates. Fréchette, Kagel, and Lehrer (2003), Fréchette, Kagel, and Morelli (2004a), and Fréchette, Kagel, and Morelli (2004b) perform experiments involving voting and coalitions. In all of these papers, minimal winning coalitions (three voters out of five, for example) are the norm, with voters almost invariably showing little or no concern for the payoffs of the voters who are not part of this coalition.

Thus, our result is not really an isolated one in the experimental literature, but is quite consistent with an emerging behavioral regularity: when responding to a party who has influence on one's payoffs, one is primarily concerned with the relationship with the influential party and any consideration of the powerless party or parties is largely lip service. Further, in combination with the above results, the lack of regard for third parties identified in this paper suggests that models such as Fehr and Schmidt (1999), Bolton and Ockenfels (2000), and Charness and Rabin (2002) do not effectively capture preferences in asymmetric multi-player games, despite their considerable success in explaining data from two-player experimental games.

Interestingly (and perhaps related to the above discussion), despite the lack of worker response to co-worker wages in our experiment, it appears that the subjects representing firms in our sample anticipated some sort of adverse reaction from an unequal-wage policy, as they were more likely to compress wages when wages were public information than when wages were private. In fact, the most commonly-selected wage policy when wages were public was an egalitarian one, even though a number of other (unequal) wage policies consistently yielded higher profits. Based on the results of an exit survey, we believe this represents an overestimate of the importance of workers' responses to co-worker wages by profit-motivated firms, rather than a deliberate sacrifice of profits by firms in the interests of inter-worker wage equity.

Finally, in our experiment we find that requiring firms to share wage information with all their workers does not significantly reduce their profits. This is true in two senses. First, under the actual (non-profit-maximizing) wage policies adopted in the public regime, firms' profits are lower than in the private wage regime, though the difference is quantitatively small and statistically insignificant. Second, if we hold firms' wage offer mix fixed at the private-wage-regime level, the point estimate of the profit differential goes to zero. Thus, at least in our

experiment --which was expressly designed to identify equity-related effort responses--, we find little support for the notion that inter-worker equity concerns can make wage compression, or wage secrecy, a profit-maximizing policy.

2. A Model

In this section we model the expected implications of co-worker wage effects on worker's effort choices for profit-maximizing wage levels, wage compression, worker effort and profits. We consider two main cases: In the case of *symmetric* worker responses to co-workers' wages, workers' effort decisions respond equally to an extra dollar of 'overpayment' relative to their co-workers as they respond to a dollar less of 'underpayment'. In the *asymmetric* case, workers' effort decisions respond to underpayment only. In both cases, we consider the effects of an increase in the strength of co-worker wage-sensitive behavior (henceforth the parameter "*b*") on a variety of outcomes, assuming firms set wages to maximize profits.

To be clear, our main motivation in constructing these models (*and* in designing our experiments) is not to provide a complete model of pay setting when co-workers' wages affect utility, but to critically examine the conventional wisdom behind informal discussions of why pay compression exists. In our view, a reasonable summary of these arguments, common in Human Resource Management (HRM) textbooks and some economics writings as well, runs as follows: Employers have workers who are differentially productive, which—in a gift-exchange environment—clearly gives employers a profit-based incentive to pay workers differently. But employers are also concerned that large pay gaps between differently-abled workers may reduce profits by creating 'morale' problems. Capturing this fundamental tension between employers' desires to reward abler workers and workers' possible aversion to pay inequality is at the heart of our modeling strategy.

In view of the above goals, we have made two key modeling decisions: first, we keep the model as simple as possible, for example focusing on the case of linear effort supply functions (generalizations are discussed at various points). Second, on the worker side, we begin with an assumption about *behavior* (i.e. an effort-supply function) rather than utility. This allows us to focus on the implications of co-worker wage-sensitive *behavior* for firms' profit-maximizing wage policies, while remaining as agnostic as possible about the precise nature of social preferences underlying worker behavior (an area that, as noted earlier, remains in considerable dispute). The cost of this is that we cannot, and do not, derive any conclusions about worker welfare, which almost certainly would be highly sensitive to the precise structure of social preferences assumed.

a. The Symmetric Case

Imagine two workers in a firm, each of whose effort is given by:

(1) $E = aw + b(w - w_c)$

where w is the worker's own wage, w_c is his/her co-worker's wage, $a > 0$, and $b \geq 0$. Workers respond to co-worker wages when $b > 0$; we refer to $b = 0$ as *difference-neutral* behavior and $b > 0$ as *difference-sensitive* behavior. The formulation in (1) has three noteworthy features.

First, by assumption, when the two workers receive the same wage $(w = w_c)$, a worker who responds to co-worker wages $(b > 0)$ exerts the same effort as a worker who does not $(b = 0)$. Second, when workers respond to co-worker wages, they reduce their effort below the difference-neutral level if paid less than a co-worker, and raise effort above the difference-neutral level when paid more than a co-worker. These effects are equal in magnitude, their

strength given by a single parameter *b*. This second assumption is our 'symmetry' restriction and is relaxed later in this section. Third, our effort-supply functions are assumed for simplicity to be linear; the objective of our model is not to provide a general characterization of optimal wage-setting when co-workers' wages affect effort, but to illustrate as simply as possible why the most common intuitive arguments for wage compression are not correct.[4] Finally, since we allow output to be a general, nonlinear function of each worker's effort (see below), linearity of the effort-supply function is not as restrictive as it might seem.

Let total revenues produced by a type-1 (low-productivity) worker be given by $R(E)$, revenues from a type-2 worker by $\theta R(E)$, where $R' > 0, R'' < 0,$ and $\theta > 1$. Total profits earned by a firm employing one worker of each type are then:

(2) $\Pi = R\big(aw_1 + b(w_1 - w_2)\big) + \theta R\big(aw_2 + b(w_2 - w_1)\big) - w_1 - w_2 .$

In this base case, first-order conditions for a maximum of profits with respect to w_1 and w_2 can be written respectively as[5]:

(3) $aR'(E_1) + b\big[R'(E_1) - \theta R'(E_2)\big] = 1$

(4) $\theta aR'(E_2) + b\big[\theta R'(E_2) - R'(E_1)\big] = 1$

---

[4] We have also examined the more general case, where $E = F(w) + bG(w - w_c)$, where $F(0) = G(0) = 0$, $F' > 0$, and $G' > 0$. In this case, symmetry is defined by the condition $G(x) = -G(-x)$, and difference-neutral behavior by $b = 0$. In this framework, Result 1 (below) is unchanged. Further, the "neutrality" results –that effort levels and firms' maximized profits are invariant to *b* continues to hold in the special case where $F(w)$ –but not necessarily $G(w)$—is linear. In the linear *F* case we can also show that wage compression is always optimal for *b*>0. Finally, (and once again in sharp contrast to widespread intuition) when *F* is concave ($F'' < 0$), increases in *b* now have the effect of *raising* the firm's maximized level of profits. Proofs are available at:
http://www.econ.ucsb.edu/~pjkuhn/Data/WageCompression/WageCompMaterialsIndex.html.

[5] One can imagine solutions to this problem where at least one worker's effort level is zero. For simplicity, we assume that $R''(0)$ is large enough to rule this out. Also, for consistency with the standard formulation and for ease of presentation, we model firms' decisions as not subject to any constraints on the total wage bill, $w_1 + w_2$. Such a constraint is imposed for practical reasons (discussed below) in our experiment. In our base-case model, the profit-maximizing wage bill is in fact invariant to *b*; as a consequence the model's predictions are *identical* to the case of a fixed budget for wages.

where $E_1 \equiv aw_1 + b(w_1 - w_2)$ and $E_2 \equiv aw_2 + b(w_2 - w_1)$ are worker 1's and worker 2's effort respectively.

Result 1. When workers' behavior is difference-neutral ($b = 0$), profit-maximizing firms will pay higher wages to their more-productive workers ($w_2 > w_1$), who in turn supply greater effort than the less-productive workers ($E_2 > E_1$).

Proof. When $b = 0$, (3) and (4) simplify respectively to $aR'(E_1) = 1$ and $\theta aR'(E_2) = 1$. It follows directly (from $R'' < 0$) that $E_2 > E_1$. Since, in the difference-neutral case, each worker's effort is proportional to his/her own wage only ($E_1 = aw_1$; $E_2 = aw_2$) it also follows that $w_2 > w_1$.

Result 2. When workers respond to relative wages ($b > 0$), profit-maximizing effort levels for each worker (E_1 and E_2) are identical to the difference-neutral levels identified in Result 1, *regardless* of the value of b. Wages, however, are not identical to the difference-neutral case: profit-maximizing firms compress wages relative to the difference-neutral equilibrium; i.e. $w_1 > w_1^n$ and $w_2 < w_2^n$, where w_1^n, w_2^n denote wages under difference neutrality. As b rises, w_1 rises and w_2 falls, but workers' wage rankings are never reversed, i.e. $w_1 < w_2$ regardless of b. Further, regardless of the level of b, the total wage bill ($w_1 + w_2$), and total profits are *identical* to the difference-neutral case.

Proof. Suppose that $R'(E_1) < \theta R'(E_2)$ in the difference-sensitive equilibrium. It then follows from (3) that $aR'(E_1) > 1$, and from (4) that $\theta\, aR'(E_2) < 1$. Together these contradict the supposition. A parallel argument rules out the possibility that $R'(E_1) > \theta R'(E_2)$. The only remaining possibility that satisfies both (3) and (4) equates effort levels to those in the difference-neutral equilibrium, i.e., $E_1 = E_1^n$ and $E_2 = E_2^n$.

Next, recall that wages are related to efforts via the system of linear equations

$E_1 = aw_1 + b(w_1 - w_2)$ and $E_2 = aw_2 + b(w_2 - w_1)$. Solving these for w_1 and w_2 yields:

(5) $w_1 = \dfrac{bE_2 + (a+b)E_1}{a^2 + 2ab}$

(6) $w_2 = \dfrac{bE_1 + (a+b)E_2}{a^2 + 2ab}$

From previous results we know that $E_1 = E_1^n < E_2 = E_2^n$, and that E_1 and E_2 in (5) and (6) are invariant to b. Since both w_1 and w_2 are weighted averages of E_1 and E_2, but the latter assigns a higher weight to E_2, it follows that $w_1 < w_2$. Summing (5) and (6) yields $w_1 + w_2 = (E_1 + E_2)/a$, which is independent of b. Independence of profits from b follows from result in conjunction with the independence of effort levels from b. Finally, differentiating (5) and (6) with respect to b (and using $dE_1 = dE_2 = 0$) yields:

(7) $dw_1 / db = \dfrac{(w_2 - w_1)a}{a^2 + 2ab} > 0$

(8) $dw_2 / db = \dfrac{(w_1 - w_2)a}{a^2 + 2ab} < 0.$

A key implication of Result 2 is that firms are not hurt by worker behavior that is sensitive to co-workers' wages. The generality of this result is also noteworthy: it holds

regardless of the *strength* of workers' effort responses to co-worker wages (b), regardless of the form of the production function (R), and regardless of the slope of the own labor-supply function (a).[6] To see the intuition for this result recall that in equilibrium total costs are given by $w_1 + w_2$ = $(E_1 + E_2)/a$; thus *because of symmetry the marginal cost to the firm of inducing an extra unit of effort from either worker is independent of b*.

The other key implication of Result 2 is that, despite the invariance of effort to b, some wage compression is in the interests of a profit-maximizing firm. To see why, consider the effects of a small increase in b, beginning at the difference-neutral ($b = 0$) equilibrium. If after this increase we kept wages at their difference-neutral levels, low-productivity (type-1) workers will work less than before; their effort is reduced by the fact they are underpaid ($w_1 < w_2$). By the same argument, type-2 workers will work harder. But we have just shown that —because the marginal cost of both E_1 and E_2 is independent of b—firms maximize profits by keeping effort levels unchanged. To achieve this, firms must raise w_1 and cut w_2; i.e., compress wages. In sum, wages are compressed because, when workers respond symmetrically to each other's wages, a smaller wage gap is needed to elicit the fixed, profit-maximizing effort levels from both workers. It is perhaps noteworthy that this intuition differs somewhat from that in popular discussions, perhaps because the latter discussions seem to focus only on the effects of wage discounts on low-ability workers, ignoring the potential for wage premia to motivate high-ability workers.

Result 3. When workers respond to relative wages, the optimal amount of wage compression relative to the difference-neutral case is given by:

---

[6] It does, however, require our assumption that the two workers respond equally to each others' wages, i.e. have the same b. The asymmetric case (considered below) provides a clear illustration of this, since in that case only the worker receiving the lower wage has a positive b (at the margin).

$$(9) \quad \frac{w_2 - w_1}{w_2^n - w_1^n} = \frac{a}{a + 2b} = \frac{1}{1 + 2(b/a)}$$

Proof. Subtracting (5) from (6) for the cases $b>0$ and $b=0$ respectively, noting (from Result 2) that effort levels are identical in the two cases, then taking the ratio of the two cases yields the result shown.

Equation (9) implies a very specific relation between the strength of workers' responses to co-worker wages and the profit-maximizing degree of wage compression. This relation is true for *any* production function R. For example, if workers' effort is one-tenth as sensitive to relative wages as to their own wage ($b/a = .1$), the wage gap should be reduced by 16.7 percent, to 83.3% of its value under difference neutrality. If workers respond equally to own and relative wages ($b=a$), the wage gap should be 50% of its difference-neutral level. Finally, note that (9) approaches zero from above as b grows without bound. Even in the most extreme case imaginable, it is therefore never optimal to pay equal wages to workers of differing ability.

Because of its relative simplicity, extensions to the symmetric case are relatively easy to explore. A key question allows for an arbitrary number (N) of workers instead of two, and considers the effects of changing the relative numbers of high- and low-productivity workers. These questions are explored in background materials available from the authors, which show that the invariance of effort and profits to b extends to the N-worker case as well.[7] Interestingly, we can also show that profit-maximizing wages offered to workers of *both* types of workers must rise when a high-productivity worker replaces a low-productivity worker. This (profit-

---

[7] These and other supporting materials referred to in this paper are accessible at:
http://www.econ.ucsb.edu/~pjkuhn/Data/WageCompression/WageCompMaterialsIndex.html

maximizing) rent-sharing among workers within firms may help explain why firms seem to prefer to hire more-productive workers.[8]

b. The Asymmetric Case

Suppose now that workers only respond to relative wages when they are paid *less* than their co-workers; this corresponds to the Bolton (1991) model of social utility. Specifically, let $E = aw + b\delta(w - w_c)$ where the function $\delta(x) = x$ for $x < 0$, $\delta(x) = 0$ for $x \geq 0$.

Result 4. When workers' responses to relative wages are asymmetric in the sense described above, profit-maximizing effort levels deviate from the difference-neutral levels. In particular, $E_1 > E_1^n$ and $E_2 < E_2^n$, i.e. low-ability workers provide *more* effort than in the difference-neutral equilibrium, while high-ability workers provide less. E_1 is monotonically increasing in b; E_2 monotonically decreasing.

Proof. Returning for simplicity to the two-worker case, the first-order conditions for a profit maximum ((3) and (4)) now become:

(10) $(a + b)R'(E_1) = 1$

(11) $\theta a R'(E_2) - bR'(E_1) = 1$

The result for E_1 follows directly from monotonicity of R' in (10). For E_2, solve (10) for $R'(E_1)$ and substitute into (11), yielding:

(12) $\theta a R'(E_2) = \dfrac{a + 2b}{a + b}$

Since the RHS of (12) is increasing in b, E_2 must be decreasing in b.

---

[8] In a strict neoclassical model where wages equal marginal products, firms (and workers) would of course be

Result 5. When workers' responses to relative wages are asymmetric, profit-maximizing firms pay high-ability workers less than their difference-neutral wage, i.e. $w_2 < w_2^n$. Also, for positive b, the wage ratio, w_2 / w_1, must be below its difference-neutral level.

Proof. Solving for wages as a function of effort, equations (5) and (6) now become:

$$(13) \quad w_1 = \frac{aE_1 + bE_2}{a(a+b)}$$

$$(14) \quad w_2 = \frac{E_2}{a}$$

as long as $w_1 < w_2$. Since E_2 declines with b, w_2 must do so as well. Taking the ratio of (14) to (13) yields:

$$(15) \quad \frac{w_2}{w_1} = \frac{E_2 + (b/a)E_2}{E_1 + (b/a)E_2}.$$

Recall that under difference neutrality ($b=0$), the wage ratio is given by $w_2^n / w_1^n = E_2^n / E_1^n > 1$. According to Result 3, $E_1 > E_1^n$ and $E_2 < E_2^n$ for $b > 0$, which in turn implies that $w_2 / w_1 < w_2^n / w_1^n$.

Result 6. In contrast to the symmetric case, for high-enough values of b a firm's profit-maximizing wage policy in the asymmetric case could involve equal wages for differently-productive workers ($w_1 = w_2$).

---

indifferent to the quality of a marginal worker hired.

Proof. By example: Let $R(E) = E^{.5}$, choose units of effort so that $a = 1$ and let worker 2 be twice as productive as worker 1, i.e. $\theta = 2$. Computing profit-maximizing effort and wage levels using (10) and (12)-(14) generates values of $w_2 < w_1$ for any b in excess of about .355. Since firms will never wish to reduce w_2 *below* w_1 (note that (10)-(14) no longer apply when $w_2 < w_1$ since they are predicated on the low-ability worker receiving the lower wage), we conclude that there is a critical value of b above which a strict egalitarian wage policy maximizes profits.

Result 7. When workers' responses to co-worker wages are asymmetric, maximized profits are strictly declining in the strength of those responses (b).

Proof. Applying the envelope theorem to the expression for profits yields:

$$(16) \qquad \frac{d\Pi}{db} = \frac{\partial\Pi}{\partial b} = R'(E_1)(w_1 - w_2),$$

which is strictly negative in the relevant region ($w_2 > w_1$).

In contrast to the symmetric case, where the firm's cost function $C(E_1, E_2)$ was unaffected by workers' equity-related behavior (b), a higher b does raise total costs in the asymmetric case. Compared to symmetry, asymmetry also raises the 'marginal productivity' of w_1 relative to w_2, in the sense that an increase in worker 1's wage no longer generates an adverse effort response from the other worker, whereas an increase in worker 2's wage continues to do so. Together these two factors explain the lower profits and smaller wage gaps under asymmetry than symmetry. The greater familiarity of this intuition compared to that for symmetric case (described earlier) suggests that the informal, implicit theorizing behind previous discussions may have taken for granted the asymmetric scenario. In the empirical work that follows, we test for the presence of both symmetric *and* asymmetric worker responses to co-worker wages.

3. The Experiment

Following Fehr, Kirchsteiger, and Riedl (1993), we model the labor market as a simple 'gift exchange'. The firm moves first, by offering the worker a salary, S, which can depend on the worker's type. The worker then selects an effort level, E. Payoffs are then:

Principal's payoff ("Profits"): $\Pi = Q - S = Q(E) - S$

Agent's payoff ("Utility"): $U = S - V(E)$.

Clearly, the perfect equilibrium to this game is not efficient. According to standard reasoning, agents should expend no effort and (anticipating this) the principal will pay no salary. In practice, however, it is well known that much more cooperation than this occurs.

As noted, our question in this paper is whether the amount of cooperation (and the surplus generated) is influenced by pay comparisons made between two types of workers (high-productivity and low-productivity) employed by the same firm. In each period, each firm is therefore matched with one worker of each type. We vary whether wages are public (both workers know both wages) or private (each worker knows only his or her own wage).

The experimental instructions are provided in the Appendix.[9] Each firm is endowed with $4 (lab dollars) in each period, and can pay total wages (in integer amounts) in each period that do not exceed the $4 endowment.[10] If a firm chooses not to spend the entire endowment, it keeps the unspent money, but cannot use any such savings to pay higher wages in later periods. The wages chosen are subtracted from the $4 endowment, and the firm receives the benefits of any revenues produced by the workers. Earnings accumulate over the course of the session, and are

---

[9] Since the ordering of treatments varied across sessions, instructions differed slightly to reflect this. The instructions shown are the exact ones for sessions 1 through 4.

then converted from lab dollars to real dollars. Both types of workers had the same conversion rate, while the firm's conversion rate differed. Each worker saw only his or her own (productivity) schedule, while the firm saw both. All of the above was common information to the participants.

Also common information in the experiment was the fact that the two workers had different productivity schedules (though the direction and magnitude of these differences was known only to firms). The rationale for this was to approximate real-world labor markets in which workers can be fairly certain their productivity is not identical to their co-workers', but do not have any good way to determine their true relative value to the employer. The focus of our experiment, therefore, is on a case where workers differ in productivity, but do not have good information on the direction or magnitude of those productivity differences.[11]

Firms were thus allowed to choose from among five salary levels (zero, one, two, three or four lab dollars), and workers could respond with one of four effort levels (zero, low, medium or high). The $Q(E)$ and $V(E)$ functions for both worker types are shown in Table 1.

[Table 1 about here]

As can be seen, increasing effort is increasingly costly for the workers. Workers receive no direct benefit from providing costly effort, while the firms' profits depend critically on the effort levels chosen.

---

[10] We considered the idea of allowing firms to use unspent endowments from previous periods to make wage offers, but rejected it because of the non-stationarity it would add to the firm's decision problem.

[11] A number of seminar participants have argued that we might detect greater behavioral responses to co-worker wages if workers who knew they were equally productive were paid different wages. This may very well be true, but strikes us as posing a very different question than the one posed in this paper: arbitrary differentials paid by firms to identical workers for hard-to-understand reasons may very well produce a different worker response than wage differentials driven at least in part by the fact that some workers have a greater capacity to reward the firm. But worker response to these arbitrary differentials strikes us as much less interesting than the responses studied here.

We conducted seven sessions at the University of California at Santa Barbara, with 18 students in four of the sessions, 15 students in one of the sessions, and 12 in two sessions (the differences are due to the variance in show-ups at the laboratory). Participants were recruited using an e-mail message to the general student population. Since no person participated in more than one session, there were thus 111 different participants. Average earnings were about $16 for the one-hour sessions.

At the beginning of each session the students were randomly divided into three groups of equal size: firms, low-productivity (type-1) workers, and high-productivity (type-2) workers. Each person stayed in his or her assigned role for the duration of the 30 periods in the session. After each period, the firms and workers were randomly re-matched (with no 3-person group remaining the same from one period to the next). All of this was common information.

Within each period, each firm first makes a salary payment to both of his or her agents. After these payments are entered in the workers' accounts, all workers decide on how much of a transfer ("effort") to make to the firm, given the costs shown in Table 1. The calibration in Table 1 was chosen (a) to generate non-zero effort levels from the majority of workers in a one-on-one gift-exchange game, based on past results with those games; (b) to embody large productivity differences between the worker types (thus giving firms an incentive to differentiate wages), but (c) to also allow workers to impose high costs on firms by choosing zero effort levels if (for example) the worker were to feel unfairly treated.[12]

---

[12] Table 1 also has the following properties (when all payoffs are measured in lab dollars): the effort level that maximizes total surplus (defined as profits + utility, or $Q(E) - V(E)$) is 'medium' for the low-productivity worker and 'high' for the high-productivity worker. A low effort level by type 1 workers in the presence of a $1.00 wage results in equal sharing of the $1.80 surplus from production between the worker (Net Receipts = $1.00 - $0.10 = $0.90) and the firm (Profit = $1.90 - $1.00 = $0.90). Similarly, a high effort level by type 2 workers when assigned a $3.00 wage results in equal sharing of the $4.80 surplus that occurs in that event.

Different information structures were implemented over the course of the various sessions. Holding all else constant, with *private wages* each worker is told only his or her own wage in the period, while with *public wages*, each worker learns both wages chosen by the firm. Participants knew that there would be regime changes during the session, but were not told in advance the nature of these changes. In our first four sessions, we had public wages during periods 11-25 and private wages in all other periods. In the remaining three sessions, we had private wages during periods 11-25 and public wages in all other periods.[13]

4. Worker Behavior

a) Means

In this section we analyze how workers' effort decisions responded to the wage they were offered, and to the wage that was offered to the other worker employed by their firm. We do this separately for two information regimes: 'wage secrecy' (where workers were informed only of their own wage); and 'public wages' (where they are informed of their own wage and that received by their co-worker). As we do not expect co-workers' wages to affect effort in the wage secrecy regime, this serves as a useful specification check for our experimental design and econometric procedures.

Unadjusted counts of all possible wage-offer combinations and the mean effort levels of both worker types for each combination are presented in Table 2:

[Table 2 about here]

---

[13] A referee has asked whether, for example, the private wage treatments that occur in periods 26-30 should be combined with earlier private treatments, since workers will acquire some knowledge about the *typical* amounts that firms pay the 'other' worker may during (public wage) periods 11-26. To address this and related possibilities, we replicated our main analysis of worker behavior (Table 3 below) two ways: excluding the final five periods in all our sessions, and excluding the final 20 periods. In the former case, the results are essentially identical. In the second, they are similar but considerably noisier due to the much smaller sample size. That said, since our main

The table shows, for example, that the most common wage-offer pair chosen by firms when wages were secret was a wage of 1 to their low-productivity worker and 2 to their high-productivity worker. This combination was chosen 137 times; when it was chosen, the average effort level chosen by type-1 workers (with possible effort choices ranging from zero to 3) was 0.715; the average effort chosen by type-2 workers was 1.226. In the second-most-common wage-offer pair—(1,3), chosen 79 times—, the average effort level chosen by type-1 workers was 0.937; the average effort chosen by type-2 workers was 1.747.

Part A of Table 2 shows two results very clearly. First, workers' effort decisions respond very strongly to their own wage: reading down the columns for type-1 workers, or across the rows for type-2 workers, mean effort levels rise essentially monotonically, and precipitously, with own wages. There are only two exceptions to these monotonic increases, involving cells with relatively few observations and 'extreme' wages. Second, no such pattern is visible for co-workers' wages (going across rows for type-1 workers or down columns for type-2 workers). Since workers were not informed of their co-worker's wage in this wage-setting regime, this is exactly what we should expect.

Part B of Table 2 presents results in exactly the same format for all the experimental rounds in which workers were informed of the wage the co-worker was offered before choosing their own effort levels. In our view, the most striking aspect of Part B is its similarity to Part A: Own wages matter (a lot), but no strong or consistent pattern emerges for the effect of co-worker wages. Holding type 1's wage fixed at $1, we do see a small but monotonic decline in type 1's effort with increasing type 2 wages. At the same time, other comparisons—for example, holding type 1's wage fixed at zero, or holding type 2's wage fixed at 1— go the other way. While we

---

interest is in worker responsiveness to co-worker wages within the *public* wage regime, we are not overly concerned that workers may know more about co-worker wages in some private-regime periods than in others.

combine all this information formally in the following Table, our main impression from Table 2

is of a small and inconsistent effect of co-workers' wages on effort.

b) Regressions

Table 3 summarizes the information in Table 2 in a regression context. This serves

several purposes, one of which is to control for period effects that could bias the Table 2 results.

Examination of the data reveals significant (though not dramatic: even in the final period of each

session, 70 percent of firms offered positive wages to their type-2 workers, and 59 percent

offered positive wages to their type-1 workers) 'unraveling' in the sense of declining effort levels

across periods even within treatments. The regression context also lets us parameterize the effect

of other workers' wages in simple ways and conduct significance tests for co-worker wage

effects that (a) treat every round as a separate observation, but (b) allow for correlated error

terms within subjects by 'clustering' on individuals.[14] Finally, the regression context allows us

to ask whether estimated effects of offered wages are different when we look only 'within

subjects' (i.e. allowing each subject his/her own effort intercept and examining the effects of

different wage offers to the same person).

[Table 3 about here]

Table 3 includes data only from the public-wage regime, i.e., from those rounds in which

workers were told (and were thus able to respond to) the co-worker's wage. Part A of the table

focuses on type-1 (low-productivity) workers, presenting estimated coefficients from regressions

---

[14] An alternative that allows for even more general heteroscedasticity would be clustering within sessions. When we do this, the standard errors on the own-wage coefficients in Table 3 rise by 20 to 25 percent (for example the column 1 standard error rises from .081 to .098 for Type 1 workers, and from .071 to .089 for Type 2 workers). Thus, while the own-wage coefficients remain highly significant, the co-worker wage coefficients become even less significant.

in which effort is the dependent variable. In column 1 we simply allow for a linear effect of the

worker's own wage on effort. The effects of own wages on effort are very strong, and highly

statistically significant (with a *t*-ratio in excess of 7).[15] Columns 2 and 3 add two alternative

measures of the co-worker's wage to the column 1 regression; the first of these corresponds to

the "symmetric" model in Section 2.[16] The second, a dummy variable for whether the worker's

own wage is less than his/her co-workers, captures the asymmetric model.[17] Coefficients on both

of these measures are small and statistically insignificant. Columns 4 and 5 replicate columns 2

and 3 with a more flexible measure of the worker's own wage; while own wage effects remain

strong and monotonic, there is essentially no change in the estimated co-worker wage effect.

Finally, columns 6 and 7 add worker and period fixed effects in turn.

Overall, in all specifications but one the estimated effect on a low-productivity worker's

effort of being paid less than his/her co-worker is negative—as one might expect from a

'jealousy' hypothesis—, but is statistically insignificant and economically very small. To see this

last point, consider the effect in column 7 of being 'underpaid' by *at least* one lab dollar on effort

of -.095. This effect is about *one seventh* the effect of having one's own wage raised from one to

two lab dollars. Another way to assess the strength and magnitude of our results is to assess their

implications for the relative sensitivity of worker effort to co-worker versus own wages, i.e. for

the parameter *b/a* in equation 9. To do so, we re-estimated columns 2 and 3 of Table 3 as

---

In fact, since our main result is that co-worker wages have no significant effect on effort, clustering on persons rather than sessions can be seen as the more conservative approach.

[15] When we do not cluster the standard errors at all, the *t*-ratios on own wages rise to around 12 and 13 (depending on the exact specification) while co-worker wages remain highly insignificant.

[16] Because it seems a more straightforward way to summarize the data, column 2 of Table 3 simply enters the co-worker's wage (rather than the wage gap between the workers) as an additional regressor. Thus, referring to equation (1), the coefficient on the co-worker's wage estimates the parameter *b*, while the coefficient on the own wage estimates $(a + b)$.

[17] Strictly speaking, the asymmetric model in Section 1 would require us to enter the difference between the own and co-worker wage $(w - w_c)$ interacted with this dummy variable. Results with this specification, plus a variety of others (including effects of being paid equally, of being paid more than one's co-worker, etc.) were all very similar in showing no robust effects of co-worker wages.

follows: first, we replaced the co-worker wage in column 2 by $(w - w_c)$, and the indicator for

own wage less than co-worker's wage in column 3 by an interaction between that indicator and

$(w - w_c)$; thus the estimated coefficients now refer exactly to the parameters a and b in equation

1. Second, we clustered the standard errors on sessions to ensure that our standard errors were

estimated as conservatively as as possible. We then conducted hypothesis tests on the ratio

between the two coefficients, a and b. In the symmetric model (column 2) we can reject $b=a$,

with a p-value of .0002; p-values for b/a = .5, .2, and .1 are 0.0005, 0.0188 and 0.1960,

respectively. Thus, taking the most conservative approach to our standard errors, our data are

still rich enough to decisively reject the hypothesis that workers value relative wages by any

more than one-fifth the value placed on their own wage. For the asymmetric model in column 3

our results are even stronger: the p-value for $b/a = .1$ is 0.0334.

Thus we can be more than 95 percent confident that workers respond to their relative

wages at no more than a *tenth* the rate they respond to the level of their own wage. Clearly,

effects of this size do not suggest there will be strong beneficial effects on profits of wage

policies that accommodate workers' preferences for inter-worker equity. We shall address the

effects of wage compression on profits more directly in Section 4 of the paper.

Part B of Table 3 focuses on type-2 (high productivity) workers. Overall the results are

very similar, with strong and monotonic own-wage effects and insignificant co-worker wage

effects. Interestingly, there is no statistically-significant effect of being paid less than one's co-

worker on effort in any specification, with one coefficient actually being positive. Note also that

the standard errors on receiving a wage less than one's co-worker are much higher than in part A

of the Table. This is because the abler workers (type 2's) were only rarely paid less than their

type 1 co-workers.

In our model (as in most principal-agent models), disutility of effort is a convex monotonic function of effort, while output is a concave monotonic function of effort. Perhaps effects of other workers' wages on a worker's 'performance' would be more apparent if we were to focus on one of these other metrics. To address this question we replicated Table 2B and selected regressions in Table 3 for two alternative dependent variables: effort *costs* and revenues *produced* by the worker. The results – available on line—are very similar.[18]

We also administered an exit survey to participants. In the survey, workers were asked, "In the periods where you saw the wage the firm offered its 'other' worker, to what extent did you consider the other worker's wage when deciding how much effort to supply?" Responses to this question were coded on a five-point scale with 1 indicating "not at all" and 5 indicating that the other worker's wage was the respondent's "primary consideration in choosing [his/her] effort". In all, 40 of 73 workers (55%) chose a value of 3 or greater, indicating a "moderate influence" or greater of co-worker wages on their choices.[19] In order to reconcile these survey results with our subjects' behavior, we asked whether estimated responses to co-worker wages were larger for the subsample of workers who said they considered them. To that end, columns 2 and 4 of Table 4 replicate the regressions in column 6 of Table 3 for the subsample of workers who indicated that their co-worker's wage had at least a moderate effect on their effort decision (columns 1 and 3 consider an identical specification for the symmetric model).

[Table 4 about here]

In three of the four cases, estimated effects of co-worker wages are negative and considerably stronger than in Table 3 (in the remaining case they are essentially zero). Despite being

---

[18] See http://www.econ.ucsb.edu/~pjkuhn/Data/WageCompression/WageCompMaterialsIndex.html

[19] One worker failed to complete this information. We also collected information on the subjects' college major and gender. In fairly detailed exploratory analysis, we detected no robust correlations between major and either worker or firm behavior, and we found modest gender differences, at most. Details are available upon request.

stronger, however, these effects remain very small in magnitude compared to the own-wage effects, and remain statistically insignificant in all cases. In sum, just as casual evidence from real workplaces suggests, workers' *claims* that they care about their co-workers' wages are not at all rare in our data. When we examine worker behavior however, evidence of such behavior is stronger among workers who say they care, but surprisingly weak even in that subgroup.

5. Firm Behavior

a) Comparisons Across Regimes

Even if our examination of worker behavior suggests that relative wages did not affect worker behavior, it does not necessarily follow that firms in our experiment behaved as though this were the case. If the undergraduate 'firms' in our experiment had priors that were similar to ours, they would have entered the experiment with a belief that within-firm wage equity *does* matter and behaved accordingly. Did they?

Descriptive statistics concerning firms' wage offer behavior across the two wage-setting regimes in our experiment are provided in Table 5. The detailed counts of wage-offer pairs underlying Table 5 are reported and discussed later, in Table 7.

[Table 5 about here]

According to Table 5, low-productivity (type 1) workers were offered an average wage more than 20 percent higher with public wages than with wage secrecy, while the average wage for type 2 workers is just slightly lower with public wages. Thus it appears that firms 'wanted' to offer lower wages to the less-productive workers when unconstrained by perceived co-worker equity considerations. But firms voluntarily increase this wage ratio when they know that workers will observe both wages.

Comparing columns 1 and 2, a particularly striking result is the difference in the share of cases in which equal wages were offered to the two workers, from 30 percent when wages were secret to 45 percent when wages were made public. While most of this difference is due to a reduction in cases where less-productive workers were paid less, Table 5 indicates that a reduction in the number of cases where more-productive workers were paid less also took place. The total wage bill is largely unaffected by the regime, being only five percent higher with public wages.

The remainder of Table 5 presents means of other outcome variables by regime. Comparing the secret and public regimes, essentially all these outcomes are very similar. The only exception is that overall profits are fifteen percent higher when wages are secret, as firms earn net profits from their less-productive workers when wages are secret but not when wages are public. While this might suggest a subtle way in which co-worker wages could affect effort,[20] the regressions reported in the following Table 6 show that—unlike the wage-compression results discussed above—these differences are not statistically significant. Further, pooled regressions of effort on *regime* and wages (not reported) show no significant effect of regime on effort, holding wages fixed.

Overall, Table 5 thus suggests that, in the public-wage regime, firms anticipated an adverse response from paying different wages to their differently-productive workers. This resulted in a narrowing of within-firm wage differentials when wage information became public, with the total wage bill largely unaffected by this reduction in wage differentials.

[Table 6 about here]

---

[20] The argument goes like this: In the private-wage regime, type-1 workers are not aware that they are paid, on average, less than type-2 workers. In the public-wage regime, they are. Even though workers do not respond to co-worker wages offered by any *particular* firm within the public-wage regime–Tables 2B and 3 show this very

As noted, Table 6 examines the same effects as Table 5 but in a regression context that controls for period and firm effects. The regression without covariates in column 1 provides significance tests for differences between the means reported in Table 5; by construction these coefficients are exactly equal to the differences between columns in Table 5. Columns 2 and 3 of Table 6 control for period and firm effects in turn.

As expected, essentially all the differences observed in the raw data are replicated in these regressions, though not all are statistically significant. By several measures, the increase in voluntary wage compression when wages are made public is substantial and statistically significant. Profit effects of making wages public are small and statistically insignificant.

b) Effects of Wage Policies on Firms' Outcomes Within Regimes.

The second question we address in this section is the impact of specific pay structures on profits *within* each compensation regime. The goal is to match firms' perceptions with reality: Within each regime, which exact wage pair maximized firms' profits, and did firms disproportionately choose wage pairs that yielded higher profits? Did firms' decisions to compress wages in the public-wage regime, documented above, reduce their profits given the apparent lack of equity-motivated worker behavior?

Table 7 presents mean profits earned by each possible wage-offer combination, separately for the wage-secrecy (Part A) and public wage regime (Part B).

[Table 7 about here]

According to Part A, the wage structure that maximized firms' profits under wage secrecy was a wage of 0 for the low-productivity worker and a wage of 3 for the high-productivity worker,

---

clearly—for any given own wage they might work less *on average* in the public-wage regime because of the knowledge that, on average, they are paid less than their co-workers.

yielding a total net profit of $1.73 (in lab dollars) per period; however, this wage package was only chosen 18 times. On the other hand, the second-highest level of profits was earned by the package $(w_1, w_2) = (1,2)$, which also involved higher pay for the more-productive worker but a smaller pay differential. Interestingly, this was by far the most prevalent wage package, suggesting at least a rough correspondence between firms' perceptions and reality. In contrast, the 'egalitarian' wage package that exhausts the firm's budget (2,2) yielded negative profits in the private-wage regime. Significance tests on the above differences show that profits under both (0,3) and (1,2) were significantly greater than the negative net profit from choosing (2,2), with t-ratios of 2.39 and 2.33 respectively.[21]

Within the public-wage regime, if we exclude the idiosyncratic case (only 11 observations) of the (1,0) wage package, we see that the highest profits were made by wage packages involving paying the type-2 worker a wage of 2, with profits highest when the type-1 worker received 0. The next-best profits were generated by wage packages paying the type-2 worker a wage of 3. Importantly, all of these strategies give higher wages to the more-productive worker. However, the most commonly-chosen pay package was now the egalitarian (2,2); using the same criterion as above this package yielded significantly lower profits than (0,2) ($t = 2.65$), and (marginally) (1,2) ($t = 1.76$). Taken together, these results reinforce our suspicion that firms' *incorrect* anticipation of an adverse worker response explains the greater prevalence of wage compression under the public-wage regime.

A final piece of evidence concerning the optimality of firms' wage policies under the public wage regime arises from some simple counterfactual calculations. In particular, recall

---

[21] Significance tests were conducted by regressing profits on a set of dummies for all 15 wage combinations in Table 7, with (2,2) as the omitted category. Adding period effects changes these t-ratios to 2.47 and 2.37 respectively; adding both period and firm effects changes them to 3.88 and 2.12. All *t*-ratios adjust for correlation of disturbances within firms.

from Table 5 that average profits per period were .485 and .557 under the public versus secret wage regimes respectively. Now, suppose that in the public wage regime firms offered exactly the same mix of wage pairs that they offered in the *secret* wage regime (shown by the counts (N) in Table 2A) instead of the mix of wage pairs they actually offered (in Table 2B). Suppose further that, for each wage pair, workers continued to choose exactly the same effort mix that they did under the public wage regime.[22] What would firms' average profits be? Straightforward calculations reveal they would be .550, which is essentially identical to profits under wage secrecy. Thus, *all* of the (statistically insignificant) shortfall in profits associated with the public wage regime in our experiment would be eliminated if firms simply continued to maintain the (larger) ability-related wage differentials they paid when wages were secret.[23]

Is it possible that firms chose to equalize worker wages because firms themselves had preferences for equity among their workers? While concerns like this might motivate firms' behavior in general, they cannot explain why firms equalized wages *more* when wages were public than private.[24] Also, in the exit survey, 70 percent of 'firms' agreed with the statement "I wanted to offer higher wages to the high-productivity workers in order to maximize my profits". Of these, 73 percent also agreed that "when the workers knew each others' wages, I felt constrained in my ability to offer different wages by the reaction I expected from my low-productivity worker". In contrast, only 22 percent of firms agreed with the statement "I wanted to offer equal wages to the two workers, independently of their ability, because this is the fair thing to do". In sum, firms' wage-equalizing behavior seems mostly driven by a (mistaken) attempt to raise profits, rather than by a concern for worker equality *per se*.

---

[22] As a result, firms' cell-specific profits would thus remain the same as in Table 7B.
[23] It is worth noting that the same reasoning does not quite apply to the wage-secrecy regime: imposing the public-wage distribution on the wage-secrecy regime causes mean profits to fall slightly, from .557 to .530.

6. Conclusion

Anecdotal evidence suggests that it is common, in both individual and collective salary negotiations, for workers to point to wages received by a 'comparison group' as a justification for a wage increase. In the case of individual negotiations, these comparisons often involve an individual's co-workers within the same firm. Recently, some economic theorists have speculated that co-worker comparisons of this nature might lead profit-maximizing firms to compress wage differentials, relative to productivity differentials (e.g. Frank 1984; Akerlof and Yellen 1990).

In this paper, we show that wage compression is in fact profit maximizing in an efficiency-wage context when workers' effort responds to co-workers' wages. If workers' equity-related behavior is symmetric (effort responds equally to reductions in underpayment as to increases in overpayment), this wage compression occurs even in model specifications where other key dimensions of the firm's optimal policy (including effort and profit levels) are unaffected by the presence of such behavior. Relatedly, the intuition behind this wage compression is somewhat unexpected: It occurs because worker's responsiveness to their co-workers' wages provide the firm with extra 'leverage', in the sense of reducing the wage gap that is required to generate an given effort gap between the workers. In the asymmetric case, these neutrality results disappear, and workers' responses to co-workers' wages will in general reduce

---

[24] Of course, one might argue further that firms wanted (for reasons *distinct* from profit maximization) to be *seen* as caring about worker equity. Since the only agents who ever see a firm's wages in our experiment are its two workers, however, it is not clear this hypothesis can be distinguished from an instrumental concern with equity.

profits.[25] Wage compression is now predicted to be greater than in the symmetric case, in the

sense that fully egalitarian wages *can* be optimal for high enough values of the equity parameter,

b.

Consistent with previous research on gift-exchange labor markets, our experimental

results exhibit a strong and robust degree of apparent reciprocity between workers and firms:

The larger the wage 'gift' received from the firm, the larger the effort gift each worker provides

in return. Surprisingly (to us) however, we can detect little or no overall response of worker

effort to co-worker wages in our data. This lack of response occurs despite the fact that, in our

experiment, workers who might feel unfairly treated can impose substantial, anonymous

'punishments' (in the form of zero effort) on firms at zero costs to themselves. (In fact, by

choosing zero effort after the wage has been paid, workers actually reap an immediate *gain*).

What might reconcile our experimental results with the widespread intuition that

'jealousy' plays a key role in workers' perceptions of fairness in compensation (and with our

priors before we ran the experiment)? One factor that might play a key role is a distinction

between words and actions: It is one thing to express a desire to be paid a (strategically chosen)

co-worker's wage, and quite another to take an action such as withholding effort to 'protest'

within-firm wage differentials. Second, suppose we think of the worker's effort decision –both

here and in the real world—as a summary 'signal' of the perceived fairness of the firm's total

compensation package. Then workers who 'care' (either positively or negatively) about co-

worker wages face the problem of signaling two things –the appropriateness of their own *and*

their co-worker's wage—with a one-dimensional effort level. In such a situation, the desire to

---

[25] We assume it will be obvious to the reader that our results for profits should be interpreted in a partial-equilibrium
context. If, for example, long-run profits are fixed by a free-entry constraint, it is consumers and not firms who will
bear any extra production costs generated by workers' concerns with each others' wages. Relatedly, since we do
not explicitly model the effects of equity concerns on worker *utility*, our model cannot address the interface between

reciprocate 'gifts' of high wages made by the employer may simply dominate any desire to 'protest' high co-worker wages.

Third, while the 'jealousy' hypothesis that motivated this paper has considerable intuitive appeal, we again note that plausible arguments can also be made for *positive* concerns with co-worker wages. If plausible stories involving both positive and negative concerns with co-worker wages can be told, we should perhaps not be so surprised to find a zero effect. And finally, the fact that agents in our labor-supply experiment respond very little to payments made to a third party does not of course imply that concerns (in either direction) with the welfare of such parties might not be important in other contexts. That said, however, our findings regarding 'third parties' appear to be consistent with a number of recent experimental results framed in a variety of ways.

In sum, our main empirical result is that workers do not protest 'underpayment' relative to a co-worker by withdrawing effort in a simple gift-exchange labor market. Thus, at least in our experiment --which was expressly designed to identify equity-related effort responses--, this behavioral response –which is essential to the argument that wage secrecy, or wage compression might be a profit-maximizing policy— does not appear to be present.

---

within-firm equity concerns and external labor markets. Thus our model identifies only 'first-round' effects that occur within firms, conditional on the composition of those firms' workforces.

Appendix: Experimental Instructions

Instructions for all Participants:

There are equal numbers of three types of participants in this experiment: firms, type-I workers, and type-II workers. **Once you have been randomly assigned a type, you will have that same type for the whole experiment.** The experiment consists of 30 periods. In each period, each firm will be grouped with two workers (one of each type). Firms and workers are randomly re-matched every period, subject to the restriction that no two workers will ever be paired with the same firm in consecutive periods.

Here is some relevant information:

1. Each firm is paired with two workers, one of each type. **Pairings change every period.**
2. Firms have a fixed endowment each period, and pay total wages each period not to exceed the level of this endowment; wages are restricted to be in whole lab dollars, e.g. $0, $1, $2., etc. **However, firms are not required to offer anyone a wage greater than 0 at any time;** this is a choice for each individual firm in each period.
3. After firms pay wages, **workers observe the wage assigned and choose effort** from one of 4 feasible levels: zero, low, medium, and high. **Firms are informed about each worker's choice of effort.**
4. Each firm receives the endowment plus the revenue generated by each worker's effort level, less the total wages paid.
5. Each worker receives the wage assigned, less the cost of the effort level chosen.
6. Zero effort costs the worker nothing, and yields zero revenue for the firm. The cost of effort increases with the effort level, as does the revenue produced for the firm.
7. Each worker only sees his or her own productivity schedule, while the firm sees the productivity schedules for both worker types.
8. Earnings accumulate for firms and workers over the course of the session.
9. For a given participant type, each lab dollar is worth a fixed number of real dollars. **The conversion rate is the same for both types of workers, but differs for firms.** Your conversion rate will be given on your other instructional materials.

After answering questions regarding these procedures, we will randomly divide participants into the three types of agents.

Do you have any questions?

Instructions for Firms:

At the beginning of each period your account will be credited with 4 lab dollars (**12 lab dollars = $1**), and you will be matched with a team of two workers. Your team will consist of one type-I worker and one type-II worker. You then decide what wages to offer to each individual worker in that team. Wages must be in whole dollars (e.g. $0, $1, $2. etc.) and total wages cannot exceed your $4. endowment. Each of your two workers will be informed of his or her wage and will then choose how hard to work for you. The more "effort" supplied, the more revenues you earn. But effort is costly to the workers.

Each team of workers has one type-I member and one type-II member. Each type faces a different cost for each level of effort. Table 1 shows these, as well as the revenue produced for each effort level. Workers see the same table as you do, except that the column reflecting costs and revenues for the other type of worker is deleted.

Table 1

Effort Level	Cost to Worker	Revenue produced by Type 1 Worker	Revenue produced by Type II Worker
Zero	0	0	0
Low	.10	1.90	2.80
Medium	.30	2.50	4.20
High	.60	2.70	5.40

1) For either worker, zero effort has no cost and produces no revenue for the firm.

2) Low effort by a Type I worker generates revenues of 1.90 lab dollars, while low effort by a Type II worker generates 2.80 lab dollars.

3) Medium effort by a Type I worker generates revenues of 2.50 lab dollars, while medium effort by a Type II worker generates 4.20 lab dollars.

4) High effort by a Type I worker generates revenues of 2.70 lab dollars, while high effort by a Type II worker generates 5.40 lab dollars.

From period 1 through 10, workers will NOT know the wage you pay the other worker in your team; they will only see their own wage. (Workers know that additional information becomes available starting in period 11, but won't know the nature of this information till then.)

From period 11 through period 25, each worker will be told the wage paid to the other worker.

From period 26 through 30, the rules revert back to what they were in periods 1-10: workers see their own wage only.

At the end of round 30, you will be paid $5 for participating in this experiment, plus **$1 for every 12 lab dollars** in your account at that time.

If you have any questions, please raise your hand.

Instructions for Type-I workers:

In each period you and another worker are paired with a firm. This firm will choose wages for you and for the other worker. After being informed of your wage, you then choose how hard to work for your firm. The more "effort" you supply, the more revenues your firm will earn. But effort is costly to you. Table 1 shows these costs and the revenues produced for each effort level. For you, **3 lab dollars = $1**.

Table 1

Effort Level	Cost to **You**	Revenue produced for the Firm
Zero	0	0
Low	.10	1.90
Medium	.30	2.50
High	.60	2.70

1) Zero effort has no cost and produces no revenue for the firm.

2) Low effort generates revenues of 1.90 lab dollars.

3) Medium effort generates revenues of 2.50 lab dollars.

4) High effort generates revenues of 2.70 lab dollars.

From period 11 through period 26, the experimenter will provide you with some extra information about the actions of "your" firm for that period before you decide how much effort to supply.

Your choice of effort will be conveyed to the firm at the end of the period.

At the end of period 30, you will be paid $5 for participating in this experiment, plus $1 for every 3 lab dollars in your account at that time.

If you have any questions, please raise your hand.

Instructions for Type-II workers:

In each period you and another worker are paired with a firm. This firm will choose wages for you and for the other worker. After being informed of your wage, you then choose how hard to work for your firm. The more "effort" you supply, the more revenues your firm will earn. But effort is costly to you. Table 1 shows these costs and the revenues produced for each effort level. For you, **3 lab dollars = $1**.

Table 1

Effort Level	Cost to **You**	Revenue produced for the Firm
Zero	0	0
Low	.10	2.80
Medium	.30	4.20
High	.60	5.40

1) Zero effort has no cost and produces no revenue for the firm.

2) Low effort generates revenues of 2.80 lab dollars.

3) Medium effort generates revenues of 4.20 lab dollars.

4) High effort generates revenues of 5.40 lab dollars.

From period 11 through period 25, the experimenter will provide you with some extra information about the actions of "your" firm for that period before you decide how much effort to supply.

At the end of period 30, you will be paid $5 for participating in this experiment, plus $1 for every 3 lab dollars in your account at that time.

If you have any questions, please raise your hand.

Table 1: Effort Costs and Revenues

Effort Level	Cost to Worker	Revenue produced by Type 1 Worker	Revenue produced by Type 2 Worker
Zero (0)	0	0	0
Low (1)	.10	1.90	2.80
Medium (2)	.30	2.50	4.20
High (3)	.60	2.70	5.40

Table 2: Mean Effort Levels as a Function of Offered Wages

A. Wage-secrecy Regime

		Type-2 Worker's Wage (lab $)				
Type-1 Worker's Wage (lab $)		**0**	**1**	**2**	**3**	**4**
0	Worker 1 effort	0.229	0.000	0.080	0.333	0.333
	Worker 2 effort	0.114	0.727	1.220	2.167	1.590
	N	70	22	50	18	39
1	Worker 1 effort	0.692	0.648	0.715	0.937	
	Worker 2 effort	0.077	0.722	1.226	1.747	
	N	13	54	137	79	
2	Worker 1 effort	1.923	0.800	1.341		
	Worker 2 effort	0.000	0.400	1.098		
	N	13	10	41		
3	Worker 1 effort	2.000	1.000			
	Worker 2 effort	0.000	0.667			
	N	3	3			
4	Worker 1 effort	1.000				
	Worker 2 effort	0.000				
	N	3				

B. Public-wage Regime

		Type-2 Worker's Wage (lab $)				
Type-1 Worker's Wage (lab $)		**0**	**1**	**2**	**3**	**4**
0	Worker 1 effort	0.020	0.167	0.160	0.259	0.192
	Worker 2 effort	0.040	0.533	1.400	1.815	1.654
	N	50	30	25	27	26
1	Worker 1 effort	0.818	0.703	0.648	0.559	
	Worker 2 effort	0.364	0.560	1.276	1.898	
	N	11	91	105	59	
2	Worker 1 effort	1.700	1.429	1.321		
	Worker 2 effort	0.000	0.714	1.179		
	N	10	7	106		
3	Worker 1 effort	3.000	1.429			
	Worker 2 effort	0.000	0.000			
	N	1	7			
4	Worker 1 effort	-				
	Worker 2 effort	-				
	N	0				

Note: mean effort levels were calculated by assigning values of 0, 1,2, or 3 to 'zero', 'low', 'medium' and 'high' respectively. Subjects also used these numbers to enter their effort decisions into the computer.

Table 3: Effects of Wages on Workers' Effort, Public-wage Regime

A. Type 1 Workers

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Own wage	.589 (.081)	.589 (.081)	.576 (.075)				
Own wage =1				.516 (.079)	.515 (.077)	.480 (.084)	.423 (.092)
Own wage =2				1.219 (.166)	1.216 (.164)	1.237 (.163)	1.077 (.164)
Own wage =3				1.478 (.464)	1.483 (.460)	1.504 (.485)	1.279 (.431)
Own wage =4*							
Co-worker's wage		-.007 (.035)		-.009 (.036)			
Own wage less than Co-Worker's			-.039 (.068)		-.003 (.072)	.001 (.077)	-.095 (.082)
Period effects?	No	No	No	No	No	Yes	Yes
Worker effects?	No	No	No	No	No	No	Yes
R squared	.252	.252	.253	.256	.256	.294	.605

B. Type 2 Workers

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Own wage	.536 (.071)	.536 (.071)	.530 (.073)				
Own wage =1				.479 (.125)	.448 (.111)	.360 (.116)	.271 (.151)
Own wage =2				1.217 (.161)	1.159 (.132)	1.065 (.132)	.960 (.165)
Own wage =3				1.802 (.199)	1.785 (.194)	1.739 (.205)	1.519 (.216)
Own wage =4				1.541 (.361)	1.567 (.353)	1.508 (.338)	1.593 (.300)
Co-worker's wage		.013 (.055)		-.063 (.059)			
Own wage less than Co-Worker's			-.074 (.156)		-.011 (.132)	-.020 (.136)	.043 (.171)
Period effects?	No	No	No	No	No	Yes	Yes
Worker effects?	No	No	No	No	No	No	Yes
R squared	.253	.254	.254	.276	.275	.318	.654

Robust standard errors, adjusted for clustering on 37 individual workers, in parentheses. Sample Size for all Regressions is 554.   * A wage of 4 was never offered to type-1 workers in any session.

Table 4: Effects of Wages on Workers' Effort, Public-Wage Regime: Regressions for the Subsample of Workers who Indicated they took Others' Wages into Account

	Type 1 workers		Type 2 workers	
	(1)	(2)	(3)	(4)
Own wage =1	.517	.503	.345	.231
	(.097)	(.098)	(.217)	(.178)
Own wage =2	1.206	1.110	1.219	1.015
	(.195)	(.226)	(.290)	(.219)
Own wage =3	2.145	2.045	1.710	1.626
	(.485)	(.474)	(.433)	(.419)
Own wage =4*			2.196	2.209
			(.348)	(.344)
Co-worker's wage	.007		-.148	
	(.052)		(.102)	
Own wage less than Co-Worker's		-.114		-.126
		(.120)		(.152)
N	330	330	270	270
R squared	.326	.328	.478	.472

Robust standard errors, adjusted for clustering on workers, in parentheses. All regressions include a full set of period effects.

* A wage of 4 was never offered to type-1 workers in any session.

Table 5: Descriptive Statistics by Regime

	Regime	
	Public Wages	Secret Wages
Wages		
w_1 (mean)	0.966	0.795
w_2 (mean)	1.746	1.787
$w_1 < w_2$ (share of cases)	0.490	0.622
$w_1 = w_2$ (share of cases)	0.445	0.297
$w_1 > w_2$ (share of cases)	0.065	0.081
$\mid w_1 - w_2 \mid$ (mean)	0.978	1.267
$W = w_1 + w_2$ (mean)	2.712	2.582
Effort and Costs		
E_1	0.677	0.640
E_2	1.038	1.050
C_1	0.095	0.091
C_2	0.164	0.167
Revenues		
r_1	0.970	0.901
r_2	2.226	2.237
$R = r_1 + r_2$ (total revenues)	3.197	3.139
Profits		
Earned from worker 1 (r_1 - w_1)	0.005	0.107
Earned from worker 2 ($r_2 - w_2$)	0.480	0.450
Total (R-W)	0.485	0.557
N	555	555

Table 6: Regression coefficients for the difference in selected outcomes between public and private wage regimes.

	Regression Specification		
Outcome Variable	No controls	Period effects	Period and Firm Effects
Wages/Costs:	(1)	(2)	(3)
w_1	.171 (.056)	.201 (.051)	.201 (.051)
w_2	-.041 (.089)	-.038 (.102)	-.038 (.103)
$w_1 < w_2$	-.132 (.049)	-.133 (.051)	-.133 (.052)
$w_1 = w_2$	.148 (.049)	.149 (.050)	.149 (.052)
$w_1 > w_2$	-.016 (.017)	-.016 (.017)	-.016 (.017)
$\lvert w_1 - w_2 \rvert$	-.288 (.114)	-.324 (.125)	-.324 (.127)
$W = w_1 + w_2$ (total costs)	.130 (.100)	.163 (.100)	.163 (.102)
Revenues:			
R_1	.069 (.076)	.076 (.075)	.076 (.076)
R_2	-.011 (.155)	.029 (.158)	.029 (.160)
$R = r_1 + r_2$ (total revenues)	.058 (.177)	.105 (.181)	.105 (.184)
Profits:			
$r_1 - w_1$	-.102 (.070)	-.124 (.070)	-.124 (.071)
$r_2 - w_2$	.030 (.130)	.067 (.130)	.067 (.132)
$R - W$ (total profits)	-.072 (.140)	-.057 (.145)	-.057 (.147)
N	1110	1110	1110

Robust standard errors, adjusted for clustering on firm ID's, in parentheses.

Note: All regressions, including those with dichotomous dependent variables (($w_1 < w_2$), ($w_1 = w_2$), ($w_1 > w_2$)), estimated via OLS.

Table 7: Mean Profit Levels as a Function of Offered Wages

A. Wage-secrecy Regime

		Type-2 Worker's Wage (lab $)				
Type-1 Worker's Wage (lab $)		**0**	**1**	**2**	**3**	**4**
0	Profit	0.513	0.709	0.790	1.733	-0.428
	N	70	22	50	18	39
1	Profit	0.231	0.730	0.872	0.805	
	N	13	54	137	79	
2	Profit	0.131	-0.760	-0.024		
	N	13	10	41		
3	Profit	-1.200	-1.233			
	N	3	3			
4	Profit	-2.533				
	N	3				

B. Public-wage Regime

		Type-2 Worker's Wage (lab $)				
Type-1 Worker's Wage (lab $)		**0**	**1**	**2**	**3**	**4**
0	Profit	0.150	0.577	1.352	0.719	-0.592
	N	50	30	25	27	26
1	Profit	1.054	0.523	0.945	0.641	
	N	11	91	105	59	
2	Profit	-0.230	0.329	0.266		
	N	10	7	106		
3	Profit	-0.300	-2.514			
	N	1	7			
4	Profit*					
	N					

* A wage of 4 was never offered to type-1 workers in any session.

References

Akerlof, G. and J. Yellen (1990), "The Fair-Wage Effort Hypothesis and Unemployment," *Quarterly Journal of Economics*, **105(2)**: 255-284.

Babcock, Linda, Xianghong Wang and George Loewenstein (1996), "Choosing the Wrong Pond: Social Comparisons in Negotiations that Reflect a Self-Serving Bias" *Quarterly Journal of Economics*, **111(1)**: 1-19.

Bertrand, Marianne and Sendhil Mullainathan (2003), "Enjoying the Quiet Life? Corporate Governance and Managerial Preferences" *Journal of Political Economy* **111(5)**: 1043-75

Bewley, T. (1999), *Why Wages Don't Fall During a Recession*, Cambridge, MA: Harvard University Press.

Bolton, G. (1991), "A Comparative Model of Bargaining: Theory and Evidence," *American Economic Review*, **81**, 1096-1136.

Bolton, G. and A. Ockenfels (2000), "ERC: A Theory of Equity, Reciprocity, and Competition," *American Economic Review*, **90**, 166-193.

Cabrales, A. and G. Charness (1999), "Optimal Contracts, Adverse Selection, and Social Preferences: An Experiment," mimeo

Charness, G. (2004), "Attribution and Reciprocity in an Experimental Labor Market," *Journal of Labor Economics*, **22**, 665-688.

Charness, G. and M. Rabin (2002), "Understanding Social Preferences with Simple Tests," *Quarterly Journal of Economics*, **117**, 817-869.

Danziger, L., and E. Katz (1997), "Wage Secrecy as a Social Convention," *Economic Inquiry*, **35(1)**: 59-69.

Fehr, E. and A. Falk (1999), "Wage Rigidity in a Competitive Incomplete Market," *Journal of Political Economy*, **107(1)**: 106-134.

Fehr, E., G. Kirchsteiger, and A. Riedl (1993), "Does Fairness Prevent Market Clearing? An Experimental Investigation," *Quarterly Journal of Economics*, **108(2)**: 437-459.

Fehr, E., E. Kirchler, A. Weichbold and S. Gächter (1998), "When Social Norms Overpower Competition: Gift Exchange in Experimental Labor Markets," *Journal of Labor Economics*, **16(2)**: 324-351.

Fehr, E. and K. Schmidt (1999), "A Theory of Fairness, Competition, and Cooperation," *Quarterly Journal of Economics*, **114**, 817-868.

Frank, R. (1984), "Are Workers Paid Their Marginal Products?," *American Economic Review*, **74(4)**: 549-571.

Fréchette, G., J. Kagel, and S. Lehrer (2003), "Bargaining in Legislatures: An Experimental Investigation of Open versus Closed Amendment Rules," *American Political Science Review*, **97**, 221-232.

Fréchette, G., J. Kagel, and M. Morelli (2004a), "Behavioral Identification in Coalitional Bargaining: Am Experimental Analysis of Demand Bargaining and Alternating Offers," forthcoming in *Econometrica*.

Fréchette, G., J. Kagel, and M. Morelli (2004b), "Nominal Bargaining Power, Selection Protocol, and Discounting in Legislative Bargaining," forthcoming in *Journal of Public Economics*.

Güth, W., M. Königstein, J. Kovács and E. Zala-Mezõ (2001), "Fairness within Firms: The Case of One Principal and Many Agents," *Schmalenbach Business Review*, **53(2)**: 82-101.

Güth, W. and E. van Damme (1998), "Information, strategic behavior and fairness in ultimatum bargaining - An experimental study," *Journal of. Mathematical Psychology*, **42**, 227-247.

Harris, M. and B. Holmstrom (1982), "A Theory of Wage Dynamics," *Review of Economic Studies*, **49(3)**: 315-333.

Kagel, J. and K. Wolfe (2001), "Tests of Fairness Models Based on Equity Considerations in a Three-Person Ultimatum Game," *Experimental Economics*, **4**, 203-220.

Lawler, E. (1990), *Strategic Pay: Aligning Organizational Strategies and Pay Systems*, San Francisco: Jossey-Bass.

Lazear, E. (1989), "Pay Equality and Industrial Politics," *Journal of Political Economy*, **97(3)**, 561-580.