Estimating Social Preferences and Gift Exchange at Work*

Stefano DellaVigna
UC Berkeley and NBER

John A. List
U Chicago and NBER

Ulrike Malmendier
UC Berkeley and NBER

Gautam Rao
Harvard University and NBER

This version: February 4, 2016

Abstract

We design a model-based field experiment to estimate the nature and magnitude of workers' social preferences towards their employers. We hire 446 workers for a one-time task. Within worker, we vary (i) piece rates; (ii) whether the work has payoffs only for the worker, or also for the employer; and (iii) the return to the employer. We then introduce a surprise increase or decrease in pay ('gifts') from the employer. We find that workers have substantial baseline social preferences towards their employers, even in the absence of repeated-game incentives. Consistent with models of warm glow or social norms, but not of pure altruism, workers exert substantially more effort when their work is consequential to their employer, but are insensitive to the precise return to the employer. Turning to reciprocity, we find little evidence of a response to unexpected positive (or negative) gifts from the employer. Our structural estimates of the social preferences suggest that, if anything, positive reciprocity in response to monetary 'gifts' may be larger than negative reciprocity. We revisit the results of previous field experiments on gift exchange using our model and derive a one-parameter expression for the implied reciprocity in these experiments.

*We thank James Andreoni, Supreet Kaur, Constanca Esteves-Sorenson, Johannes Hermle, Klaus Schmidt, Dmitry Taubinsky, and the audiences at Munich University, Zurich University, the ASSA 2016 Conference and the SITE Conference on Experimental Economics for useful comments. We also thank Debbie Blair, Alec Brandon, Natalie Cox, Johannes Hermle, Justin Holz, David Novgorodsky, Waldo Ojeda, Avner Shlain, Jessica Shui, Ferenc Szucs, Jonas Tungodden, and Brian Wheaton for excellent research assistance.

1 Introduction

What motivates workers? Incentives at the workplace, including career concerns, play a large role. Yet, the power of incentives is limited by the boundaries of what can be measured. Many jobs do not involve piece rates or other outcome-contingent pay.

Even when incentives are muted, workers may work hard because they care about their contribution to the firm. For this reason, Akerlof and Kranton (2005) and Besley and Ghatak (2005) make the case that organizations should select workers with pro-social preferences. But what is the strength and nature of this motivation? Understanding it better is important for workplace productivity and organizational design.

The literature points to examples of the role of workers' social preferences towards their employers. Kahneman, Knetsch and Thaler (1986) provide survey-based evidence on the importance of fairness in wage setting. Krueger and Mas (2004) chronicle a striking case of negative reciprocity at work: a break-down in trust between employers and employees leads to employee retaliation and production of defective goods leading to hundreds of lives lost.

Overall, however, there is little systematic field evidence about the nature of social preferences towards employers.[1] In particular, do workers' social preferences take into account the employer payoffs, as in pure altruism models à la Becker (1972)? In this case, workers work harder when their effort is of higher value to the employer. Or is the right model one akin to warm glow à la Andreoni (1989, 1990)? In that case, workers value contributing something to the employer, but are insensitive to the actual employer payoff. Also, does extra employer generosity matter, as in gift exchange models à la Akerlof (1982)? Conversely, is there a response to ungenerous employer behavior, as suggested by the Krueger and Mas (2004) findings?

Our paper showcases a model-based field experiment designed to estimate the nature and shape of social preferences at work. We build on gift exchange field experiments à la Gneezy and List (2006). In these experiments, employees are hired for a one-time task, to shut down repeated-game incentives. Employees in different treatments are then exposed to different employer actions, such as surprise pay raises (Gneezy and List, 2006), pay cuts (Kube, Marechal, and Puppe, 2013), and in-kind gifts (Kube, Marechal, and Puppe, 2012). The differences in productivity across the treatments provide evidence about gift exchange and reciprocity.

We build on the important lessons of these experiments, but closely link the design to a simple model of social preferences, such that the experimental results permit us to estimate the underlying social preference parameters. In addition, we evaluate the importance of social preferences even in the absence of any gifts from the employer.

[1] There is a large literature examining social preference in the laboratory where subjects assume the role of firms and workers. In addition, an extensive literature studies *horizontal,* as opposed to *vertical,* social preferences at the workplace, i.e. social preference between workers, e.g. Bandiera, Barankay, and Rasul (2005), Cohn et al. (2011), Hjort (2014).

How does the focus on parameter estimation affect the experimental design? We show that previous field experiments, while providing very valuable qualitative evidence on social preferences at work, did not allow for estimation of the social preferences.[2] In particular, two elements were missing from the design.

First, the experiments did not specify the value to the employer of the worker's effort: in tasks such as coding of library books, it is unclear how much the employer benefits. In a model of pure altruism towards the employer, the value of the work produced plays a critical role: holding all else constant, employee effort increases in the value of the effort. Consider an alternative model, which we label 'warm glow', in which the return to the employer does not matter. Note that we use the term 'warm glow' as a placeholder for any motive that increases a worker's utility from exerting effort on behalf of the employer, *independent* of the returns generated for the employer. This could be a positive feeling from doing meaningful work or adhering to a social norm of working hard. The important distinction from altruism is the non-dependence of effort on the employer's actual utility. This distinction could not be tested in prior experiments, since the value of effort was unobserved and held constant.

Second, a key unobservable is the cost of effort. Assume for example that an unexpected pay increase leads to 20 percent higher effort, as is the case initially in Gneezy and List (2006). The increased effort could reflect 20 percent higher altruism towards the firm, under a cost function with unit elasticity. Alternatively, it could reflect a 100 percent higher altruism under an inelastic cost function with elasticity 0.2. Without information on the curvature of the cost function, it is impossible to tell. Yet, the two estimates have very different implications for the quantitative importance of social preferences in workplace settings.

We design a field experiment to address both issues, while maintaining several of the advantages of previous experiments. We hire workers for a one-time 6-hour task to prepare mailing envelopes for multiple charities. To address the first issue, we explicitly inform the workers about the average per-envelope return to the charity (the employer) based on previous fundraising returns. Furthermore, we vary this return, informing workers (truthfully) that some envelopes raise money for a cause with a one-to-one fundraising match, and thus on average have twice the return. Finally, in some periods—paid training sessions—the workers fold practice envelopes which are not used by the charity, and thus have no return to the employer. We thus assume that this work does not generate any warm glow or altruism.

We address the second issue—the unknown cost of effort function—by varying the piece rate. Workers sometimes receive a lump-sum pay, while at other times they work for one of two piece rates. Observing the optimal effort for different piece rates allows us to back out the marginal cost and thus the curvature of the cost of effort function.

Our design goal is to keep the advantageous features of the gift-exchange experiments in the field, while allowing for variation in the return to the firm and in the worker's piece rate.

[2]Lab experiments on gift exchange, in contrast, did not suffer from these shortcomings, as we discuss below.

The ultimate design is a hybrid within-between experiment. Each worker goes through ten 20-minute rounds of folding envelopes for an employer. The rounds differ in the piece rate and in the return to the employer, hence the *within-subject* variation. To make the changes more plausible, the workers fold envelopes for three charities, each with a different compensation scheme. A first charity pays a lump-sum $7 for 20 minutes, a second charity pays a pure piece rate of 20c per envelope, and a third charity pays a combination with $3.50 lump-sum and a 10c piece rate. The compensation schemes are designed to result in equivalent payoffs for a worker of average productivity (35 envelopes in 20 minutes). To limit any gift response, the compensation for the different rounds is announced at the beginning of the experiment. In addition, we vary the return to the employer by having rounds with a donor match, as well as rounds in which the employer does not benefit from the effort (training rounds).

The between-subject part of the experiment takes place in the final 2 (out of 10) rounds. All participants work again for a charity that previously paid a $7 flat pay (for 20 minutes of work), but they now work under different conditions. In the control group, the charity pays again $7 for each of the last two rounds. In the *positive monetary gift* group (as in Gneezy and List, 2006), the charity that used to pay $7 now pays $14 per round. In the *in-kind gift* group (as in Kube, Marechal, and Puppe, 2012) the charity pays again $7, but in addition provides a gift-wrapped thermos of the value of $14. In the *negative monetary gift* group (as in Kube, Marechal, and Puppe, 2013), the charity pays only $3 per round. Importantly, all the gifts—and more generally all payments to workers—are paid for by the employers, and the returns to worker effort (the raised donations) also go directly to the employer.[3]

The combination of within- and between-subject structure increases statistical power. The within-subject variation in the first eight rounds identifies the sensitivity to private incentives and to the employer's return and thus pins down baseline social preferences. The between-subject variation in the final two rounds provides evidence on positive and negative reciprocity: whether the initial social preferences change in response to changes in the employers' generosity.

This gift exchange experiment is to our knowledge the largest run so far, with 446 workers. As such, the emphasis on model-based estimation does not come at the expense of precision in the reduced-form results. Having said that, our field experiment is not as natural as, say, an uninterrupted 6-hour job coding books under the same pay scheme. We acknowledge this design cost, but also emphasize that it yields as a benefit the identification of the social preferences.

The theoretical model, experimental design, and a pre-analysis plan including the structural model were registered at the AEA RCT Registry in November 2014, after a quarter of the data had been gathered and analyzed. To our knowledge, this may be the first social science experiment to pre-register a detailed structural model.

[3]As we tell the subjects (truthfully), the Becker Center at the University of Chicago is collaborating with the charities and facilitating the employment, while not paying for the work. The one exception is the two paid training sessions, which by design are paid by the Becker Center.

Turning to the results, we first characterize the evidence from the initial rounds. Productivity is clearly sensitive to incentives: moving from a 0c piece rate to a 20c piece rate increases output by 4 envelopes out of 35, a highly statistically significant 12 percent increase.

We have two key results on social preferences. First, worker effort increases by 3.5 envelopes (a 10 percent increase) when the envelopes are utilized by the charity compared to when they are not used, holding constant the piece rate. This suggests *some* form of social preferences: workers do not just care about earning money, but value that the work counts for the employer. Second, a doubling of the employer return in the form of a one-to-one donor match leads to a very modest productivity increase of 0.6 envelopes (1.6 percent), a difference that is not statistically significant. These findings are more consistent with a warm glow model than with pure altruism.

We estimate social preference models on the data from these first eight rounds. Workers choose optimal effort taking into account their private return (the piece rate), the cost of effort, and the return to the employer through the social preferences. We allow for pure altruism as well warm glow. The cost of effort function allows for individual fixed effects, learning by doing, and a stochastic component. Building on Shearer (2004) and Andreoni and Sprenger (2012), we estimate the model with non-linear least squares. Furthermore, we extend these and related papers by considering two classes of cost of effort functions, power and exponential.

The estimation results support the warm glow model. In a specification allowing for both warm glow and pure altruism, all the weight is on the warm glow parameters. In terms of magnitudes, the workers weight the firm's *average* return about 0.4 as much as their private payoffs. The cost of effort is estimated to be inelastic, with a (precisely estimated) elasticity of 0.1. The results are not sensitive to the assumptions about the cost of effort function.

We use these estimated social preferences to compute the optimal piece rate. In the absence of social preferences, piece rates are critical to incentivize workers. In contrast, the estimated warm glow in the experiment is strong enough that, holding constant the flat pay, the optimal piece rate is zero. This echoes results by Besley and Ghatak (2005), Bellemare and Shearer (2011), and Englmaier and Leider (2012) that incentives and social preferences are substitutes for motivating workers.

Next, we turn to the gift exchange treatments in the last two rounds. We find very limited effects of the gift treatments on productivity. In particular, we estimate no impact of the negative gift (a wage cut) and a slight *negative* impact of the in-kind gift. We find suggestive evidence that the positive monetary gift treatment leads to a small increase in output in the first round after the gift, but not in the final round. The gift effects are not larger when the return to the firm is higher, a prediction of the pure altruism model.

The estimated effects, which are smaller than in most previous papers, are precisely estimated thanks to the large sample size and the ability to control for individual productivity. We can reject that the negative gift decreases output by more than 5 percent, compared to a 20 percent decrease in Kube et al. (2013). Similarly, we can reject that the in-kind gift increases output by more than 2 percent, compared to a 25 percent increase in Kube et al. (2012).

We fit our model to the data and, not surprisingly, estimate relatively small reciprocity parameters. We measure reciprocity as a shift in the social preferences toward the employer due to kind (or unkind) employer behavior. The negative gift treatment is estimated to lower the warm glow coefficient (estimated at about 0.40) by only 0.01 to 0.07 (depending on specification). The estimated reciprocity effect is larger (though not significant) for the positive gift at a 0.05 to 0.20 increase in warm glow. Hence, our data points to, if anything, stronger positive than negative reciprocity in response to monetary 'gifts', though the results ought to be taken with caution given the small effect sizes.

What explains these small gift effects? We consider four possibilities and argue that neither appears to explain the results. First, our gift treatments may not have triggered the required surprise and mood response necessary to trigger reciprocity. We use a debriefing survey to show that the workers in fact do report the expected increases (or decreases) in happiness and surprise from receiving the gifts. Second, the worker effort towards the end of the experiment may have become habitual and thus unresponsive to inputs. Counter to this scenario, subjects are highly responsive to piece rate changes even in the later rounds. Third, the gifts may not be sufficiently large to induce gift exchange. However, the positive monetary and in-kind gifts are of similar value (in absolute terms) as used previously. The negative gift treatment, which more than halves pay for the last two rounds (from \$7 to \$3), implies a smaller pay reduction than in previous papers only as a share of total earnings (12 percent versus 33 percent). Fourth, a set of other explanations pointed out by Esteves-Sorenson (2015), such as peer effects, recruitment at above-average wages, and small sample size, do not apply to our setting.

We suggest two explanations for the results. The first is simply that previous estimates of gift exchange effects may have overestimated the strength of reciprocity at work. Consistent with this interpretation, Esteves-Sorenson (2015) also fail to find an effect in a well-powered gift exchange field experiment. A second explanation is that the employee social preferences are, to a first approximation, set at first contact. Gift treatments, under this explanation, are more effective when introduced initially, as in the previous experiments, as opposed to during an ongoing work relationship, as in our experiment. Under either interpretation, gift exchange in the workplace is likely to be of more limited application than initially conceived.[4]

We also discuss an application of our framework to previous gift exchange field experiments. Estimating social preferences in these experiments is infeasible, given missing design elements. However, one can derive a measure of reciprocity—the proportional change in social preferences

[4]We leave it to future research to separate out the two explanations. After having already completed the largest gift exchange field experiment to date, we attempted to recruit more workers for an additional field experiment with the traditional between-subject structure. However, we could not recruit enough subjects from the same population to guarantee adequate power for the test.

towards the employer due to the gift—with more limited information: the elasticity of effort with respect to incentives (assuming a cost of effort function with constant elasticity). We thus revisit some previous experiments and compute the implied reciprocity for calibrated values of the elasticity. For elasticity values as in our task, some previous papers imply very large reciprocity, such as a 400 percent increase in social preferences with a gift. The calibrated reciprocity falls to a 40-50 percent increase if the elasticity is five times larger than we estimate. Future experiments may want to incorporate the elasticity estimation in the design, as we do.[5]

This paper relates to the literatures on social preferences at work, providing evidence on workers' *vertical* social preferences towards their employers, complementing a larger literature about *horizontal* social preferences between co-workers (e.g. Bandiera, Barankay, and Rasul, 2005, Charness and Kuhn, 2007; Cohn et al., 2014, Hjort, 2014; Breza, Kaur, and Shamdasani, 2015). It also relates to the literature on gift exchange in the field, as outlined above. The main contribution of the paper is a novel design that allows for estimation not only of reciprocal response to a 'gift', but also of baseline social preferences.[6]

The field experimental literature we contribute to itself builds on a series of laboratory experiments on labor markets and gift exchange, starting from Fehr, Kirchsteiger and Reidl (1998). These experiments endow 'workers' with a 'cost of effort' function (a monetary transfer) and also inform subjects of how their 'effort' affects the payoffs of the other player (the 'firm'). Our field experiment methodologically builds a bridge towards this lab design by estimating the cost of effort and specifying the impact on the employer payoff. Our paper relates also to the literature on reciprocity (Fehr and Gächter 2000; Charness and Rabin 2002; Dufwenberg and Kirchsteiger 2004; Falk and Fischbacher 2006).

The paper also relates to a small number of papers estimating parameters in real-effort experiments, typically assuming a power cost of effort function (e.g., Augenblick, Niederle, and Sprenger, 2012; Gill and Prowse, 2012; Augenblick and Rabin, 2015). We show that using an alternative exponential cost of effort function does not affect the results in our case, but may well matter when the experimental variation in observed effort is larger.

Finally, the paper relates to a growing literature on *structural behavioral economics* (Laibson, Repetto, and Tobacman, 2007; Conlin, O'Donoghue, and Vogelsang, 2007; DellaVigna, Malmendier, and List, 2012; Barseghyan, Molinari, O'Donoghue, and Teitelbaum, 2013; DellaVigna, Malmendier, List, and Rao, 2015).

[5]Although not estimating the parameters themselves, Esteves-Sorenson and Macera (2015) similarly compare the effect of a (positive) monetary gift with the effect of a piece rate, and find significant response to a piece rate but insignificant gift effects. The experiment in Esteves-Sorenson and Macera (2015) is not designed to estimate parameters given that there is only one (convex) piece rate treatment, and the return to the employer is not made clear to the subjects.

[6]Bellemare and Shearer (2011) is the only previous paper on gift exchange in the field with parameter estimation which we are aware of.

2 A Simple Model

Worker i in each round t chooses optimal effort $e_{i,t}$ as a function of pay incentives and cost of effort. In addition, the worker has social preferences towards the employer and thus cares about the employer's payoff. For simplicity, we assume risk neutrality as well as additive separability between rounds t. This allows us to write the worker's problem separately for each round t:

$$\begin{aligned} \max_{e_{i,t}} u(e_{i,t}) &= L_t + p_{Wt} e_{i,t} - C_{i,t}(e_{i,t}) + A\left(Gift, p_{Et}, p_{Wt}\right) e_{i,t} \\ s.t. \quad e_{i,t} &\geq 0. \end{aligned} \tag{1}$$

The first component of the utility function captures the monetary payoff from exerting effort $e_{i,t}$: a lump-sum payment $L_t \geq 0$ and a piece rate $p_{Wt} \geq 0$ (where W stands for worker). The payment scheme (L_t, p_{Wt}) varies by round and thus is indexed by t.

The second component is the cost of effort $C_{i,t}(e_{i,t})$, which is allowed to differ across individuals i (to capture differences in individual ability) and over rounds t (to capture potential learning and tiredness).[7] For any i and t, we assume the regularity conditions $C'() > 0$, $C''() > 0$, and $\lim_{e\to\infty} C'(e) = \infty$, guaranteeing the existence of a unique solution.

The third component captures how the worker internalizes the return to the employer, p_E , of each unit of effort $e_{i,t}$. In the experiment, p_{Et} corresponds to the average donation raised per envelope mailed when the employer is a charity, and the average revenue raised per advertisement mailed when the employer is a grocery store. The per-unit return p_{Et} varies by round t, since in some rounds a match increases the return p_{Et} for the charity. The total return to the employer is $T + (p_{Et} - p_{Wt})e_{i,t}$, the sum of a lump-sum profit (or loss term) T from the fixed costs of employing the worker, and the variable return from effort $(p_{Et} - p_{Wt})e_{i,t}$. (Notice that the variable cost of labor needs to be subtracted from the return, and thus the social preference term A may depend on the piece rate p_W).

The worker cares about the payoff to the employer with a social preference coefficient A, which may depend upon unexpected gifts $Gift$ from the employer. We consider two special cases: (i) *altruism,* where the worker takes into account the employer's actual *marginal* payoffs, and thus $A = \alpha(p_{Et} - p_{Wt})$,[8] and (ii) *warm glow*, where the worker simply derives warm glow from doing his part by exerting effort for his employer, regardless of how the effort translates into payoffs for the employer; thus, $A = a$. This set-up is similar to the ones in Bellemare and Shearer (2011) and Englmaier and Leider (2012b), among others.

The maximization problem (1) yields the first-order condition

$$p_{Wt} + A\left(Gift, p_{Et}, p_{Wt}\right) - C'\left(e^*_{i,t}\right) = 0 \tag{2}$$

[7]In Section 4, we specify the functional forms we empirically estimate.

[8]The fixed profit term T drops out in the maximization.

or

$$e^*_{i,t}\left(Gift, p_{Et}, p_{Wt}\right) = C'^{-1}\left(p_{Wt} + A\left(Gift, p_{Et}, p_{Wt}\right)\right). \quad (3)$$

where $C'^{-1}()$ is the inverse function of $C'()$, which exists and is monotonically increasing by the assumptions above. The second order conditions are satisfied since $-C''(e^*) < 0$. In the following we will assume an interior solution.[9] The optimal effort e^* is increasing in the social preference parameter A and in the piece rate p_W (provided A does not decrease enough in p_W). The non-variable terms L (lump-sum pay) and T (flat profits) do not appear in the first order condition (2) since they do not affect the marginal effort.

Figure 1 presents an example for parameter values corresponding to the point estimates in Section 6. The marginal cost curves are initially flat and then steeply increasing; compared to the marginal cost curve near the beginning of the experiment (round 2), the curve for a later round (round 6) is shifted to the right to reflect increased productivity over time. The lower marginal benefit curve with no piece rate ($p_W = 0$) equals the social preference parameter A. The second, higher curve includes in addition a piece rate of 20 cents ($p_W = .2$). At the first marginal curve, the 20 cent piece rate increases output from 28 to 30 units.

Altruism. Under pure altruism, the worker's social preference utility is $\alpha(p_{E_t} - p_{W_t})e_{i,t}$. An altruistic worker values each dollar the employer makes (through their effort) the same as α dollars in their own pocket. Capturing reciprocity models, the altruism parameter α towards the employer may depend on the receipt of a gift from the employer. Thus, the social preference term is $Ae = (\alpha + 1_{Gift}\alpha_{Gift})(p_{E_t} - p_{Wt})e$ and the first-order condition becomes

$$p_{Wt} + (\alpha + \alpha_{Gift}1_{Gift})(p_{Et} - p_{Wt}) - C'\left(e^*_{i,t}\right) = 0.$$

Warm Glow. Under warm glow, workers do not take into account or care about the employer payoff, but instead care about the effort they put in: $A = ae$. This can be interpreted as individuals deriving utility from the process of helping rather than the actual utility of the receiver. This case is inspired by the idea of warm glow proposed by Andreoni (1989 and 1990), where donors derive utility from giving, but not necessarily from the public good itself.[10] This specification also captures, in reduced form, a social norm to put in effort for an employer, or a utility from exerting effort doing meaningful work (Ariely, Kamenica and Prelec 2008). As before, we allow for warm glow to change as a result of receiving an unanticipated gift. Thus, the first-order condition under warm-glow is

$$p_{Wt} + (a + a_{Gift}1_{Gift}) - C'\left(e^*_{i,t}\right) = 0.$$

[9] A sufficient condition to ensure an interior solution is $C'(0) = 0$ and $A > 0$. While one of the assumed cost of effort functions will not satisfy this assumption, in practice zero effort is not observed in our experiment.

[10] The warm glow could also depend on the return to the firm, in which case it would be indistinguishable from pure altruism in our setting.

Altruism and Warm Glow. In the general case with both altruism and warm glow the first order condition becomes:

$$p_{Wt} + (\alpha + \alpha_{Gift} 1_{Gift})(p_{Et} - p_{Wt}) + (a + a_{Gift} 1_{Gift}) - C'\left(e^*_{i,t}\right) = 0.$$

The altruism and warm glow can be distinguished only if the return to the charity p_E varies.

Estimation from Standard Gift Exchange Experiment. Consider a gift exchange experiment à la Gneezy and List (2006). For the combined altruism and warm glow case, taking into account that in these experiments there is no piece rate, the optimal efforts are:

$$\begin{aligned} e^*_{Contr} &= C'^{-1}(\alpha p_E + a) \quad \text{and} \\ e^*_{Gift} &= C'^{-1}((\alpha + \alpha_{Gift}) p_E + (a + a_{Gift})). \end{aligned} \tag{4}$$

Can one back out the social preference parameters from the observed effort e_{Contr} and e_{Gift}? As expression (4) clarifies, two crucial pieces of information are missing. First, we do not know what workers assume the return to the charity p_E to be, since they are not informed of this. Second, the econometrician does not know the cost of effort function $C(e)$. Hence, it is impossible to identify the social preference parameters.

It is helpful to consider the special case with pure altruism and a power cost function: $c(e) = ke^{1+s}/(1+s)$. This function is characterized by a constant elasticity $1/s$ with respect to the return to effort.[11] The two solutions then reduce to:

$$e^*_{Contr} = \left(\frac{\alpha p_E}{k}\right)^{1/s} \quad \text{and} \quad e^*_{Gift} = \left(\frac{(\alpha + \alpha_{Gift}) p_E}{k}\right)^{1/s}.$$

By dividing through and inverting, we obtain

$$\frac{\alpha + \alpha_{Gift}}{\alpha} = \left(\frac{e^*_{Gift}}{e^*_{Control}}\right)^{1/s}. \tag{5}$$

While we cannot back out the altruism parameters without knowledge of the return p_E, we can infer the proportional increase in altruism $(\alpha + \alpha_{Gift})/\alpha$, *provided* one knows the curvature s. In the simple quadratic case ($s = 1$), an observed x percent increase in effort due to a gift implies an x percent increase in altruism. But for higher curvature ($s > 1$), the underlying increase in altruism is higher than x percent. Thus the elasticity to the return to effort $1/s$ plays a key role in mapping from observed effort to the underlying preferences.

The power cost function has the special feature of constant elasticity. A plausible alternative is that the elasticity decreases as effort increases. A function with this feature is the exponential

[11] The first order condition is $k(e^*)^s = v$ (in our case equal to $p_W + \alpha(p_E - p_W)$ for the altruism case). Thus, $e^* = (v/k)^{1/s}$ where v is the return per unit of effort. Then $\partial e^*/\partial v = (1/ks) * (v/k)^{1/s-1}$ and the elasticity is $\eta_{e,v} = (1/ks) * (v/k)^{1/s-1} v (v/k)^{-1/s} = 1/s$.

cost function, $C(e) = k\exp(se)/s$.[12] In this case, the solutions are

$$e^*_{Contr} = \frac{1}{s}\log\left(\frac{\alpha p_E}{k}\right) \quad \text{and} \quad e^*_{Gift} = \frac{1}{s}\log\left(\frac{(\alpha+\alpha_{Gift})\, p_E}{k}\right).$$

We can transform the solution and divide through to obtain

$$\exp\left[s\left(e^*_{Gift} - e^*_{Contr}\right)\right] = \frac{(\alpha+\alpha_{Gift})}{\alpha}. \tag{6}$$

Expression (6) highlights another implication. Consider an experiment with a positive gift treatment, which increases output by x units, and a negative gift treatment, which decreases output by x units. Would these equal-sized impacts of the gifts on effort imply that positive reciprocity has the same magnitude as negative reciprocity? Expression (6) shows that it is not the case. Because of the steep curvature of the exponential function, the x unit increase for the positive gift would require a larger proportional change in altruism (positive reciprocity) compared to the corresponding change in altruism (negative reciprocity) for the negative gift. Intuitively, it is harder to increase effort at the margin than to reduce it.

Estimation from Generalized Gift Exchange Experiment. What design would then allow for estimation of social preferences? As outlined above, one needs to measure the return to the employer, p_E, and to identify the parameters of the cost of effort function, k and s.

The first part is easy to accomplish, as one can simply inform the subjects of the return p_E as part of the design, provided the task allows for it. Identifying the cost of effort parameters k and s, instead, requires additional treatments. Unlike the standard gift exchange experiments which do not involve a piece rate, it is useful to experimentally vary the piece rate p_W to identify the cost of effort function. From (3) notice that

$$\frac{\partial e^*}{\partial p_W} = \frac{\partial C'^{-1}(p_W + A)}{\partial p}\left(1 + \frac{\partial A}{\partial p_W}\right). \tag{7}$$

Expression (7) shows that variation in piece rate p_W helps pin down the curvature of the cost of effort function. Notice, however, that the cost of effort function will be identified jointly with the social preferences, given that A features in (7). For the altruism case, the parameter α appears also in the term in parenthesis since $\partial A/\partial p_W = -\alpha$ (an altruistic worker cares that the piece rate comes out of the employer's pocket and reduces their return). In the warm glow case, $\partial A/\partial p_W = 0$ but nonetheless A features in the argument of the function C'^{-1}.

Thus, it is useful to also observe the worker in a training period, in which the work does not benefit the firm and in which the incentive p_W is paid by a third party. In this case we can assume $A = 0$, and the effort of the worker is driven solely by piece rate incentives.[13]

[12] The first order condition is $k\exp(se^*) = v$ where v is the return per unit of effort (in our case equal to $p_W + \alpha(p_E - p_W)$ for the altruism case). Thus, $e^* = (1/s)\log(v/k)$. Then $\partial e^*/\partial v = (1/s) * (k/v)/k$ and the elasticity is $\eta_{e,v} = (1/sv) * v/((1/s)\log(v/k)) = 1/\log(v/k)$.

[13] An implicit assumption, typical in the experimental literature on social preferences, is that the social preferences do not extend to this third party (the experimenter).

The key test to distinguish warm glow from altruism is whether worker effort responds to changes in the return to the employer p_E. Consider the expression:

$$\frac{\partial e^*}{\partial p_E} = \frac{\partial C'^{-1}\left(p_W + A\right)}{\partial p} \frac{\partial A}{\partial p_E}. \tag{8}$$

Under warm glow, $\partial A/\partial p_E = 0$ and thus the optimal effort does not depend on the return to the charity: $\partial e^*/\partial p_E = 0$. In the case of altruism, instead, $\partial A/\partial p_E = \alpha$ and thus the response of worker effort to the return to the charity plays a key role in identifying altruism (see also Englmaier and Leider, 2012b). For the altruism case, combining (7) and (8) yields

$$\frac{\partial e^*}{\partial p_E} / \frac{\partial e^*}{\partial p_W} = \frac{\alpha}{1-\alpha}. \tag{9}$$

The ratio of the response to the employer's return p_E and the response to the piece rate p_W identifies the altruism α. Thus, variation in the return p_E plays an important role in the design.

3 Experimental Design

Motivation. We designed the experiment with two goals in mind. The first goal was to maintain the advantageous features of the field experiments on gift exchange: (i) workers work on a real-effort task (like cataloging library books); (ii) the task allows for a natural measure of effort (like the number of books coded); (iii) recruitment into a one-time task avoids confounds from repeated game incentives; (iv) workers are assigned randomly into a gift treatment or a control group to causally estimate the effect of a gift on productivity.

The second goal is to allow for identification of the social preferences, which the previous experiments do not. To achieve this, as Section 2 explains, we wanted to (i) (truthfully) inform workers of the value of their task to the employer, p_E; (ii) vary the piece rate p_W offered to the workers without triggering a gift effect; (iii) vary (truthfully) the return to the employer p_E; (iv) have treatments when the firm does not benefit from the worker's marginal productivity (the 'training' periods); (v) maximize the power of the study to identify the parameters.

To our mind, the design which best accomplishes these goals combines between-subject and within-subject variation. In keeping with previous experiments, we randomize *between* subjects the assignment into the gift or control treatments. We instead randomize *within* subject other aspects of the design to maximize power, allowing us to estimate individual-specific differences in the cost of effort. Subjects work through several rounds of a real-effort task. Across the rounds, the piece rate and the return to the charity are varied in a pre-announced fashion. The order of the rounds varies across individuals to separately estimate learning-by-doing (or tiredness) over the course of the day. Then, in the final rounds, there is an unexpected 'gift' from the employer (depending on the treatment). Thus, the within-person variation takes place before the gift treatments, so the response to the gift cannot confound the other treatments.

The next design choice involved finding a task for which we could plausibly convey, and vary, the value of effort to the employer, p_E. Coding of library books, for example, does not lend itself readily to such purpose. We decided to partner with multiple charities to assign subjects to prepare envelopes for fund-raising mail campaigns. Since similar campaigns have been done, we could convey the average employer return to the workers. Furthermore, the return could plausibly be higher for envelopes for which a donor (truthfully) pledged to match the raised funds. To check whether there was anything special about having a charity as employer, one of the rounds involved stuffing advertisement mailers for a firm (a grocery store).[14]

Finally, we wanted to generate plausible variation in piece rate p_W in the initial rounds, without triggering gift effects. For this purpose, subjects stuff envelopes for three different charities, each offering different piece rates. To minimize the gift effects from such piece rate variation, the associated lump-sum pay keeps the total earnings constant for a person of average productivity (about 35 envelopes per 20 minutes). Thus, the pay packages are ($7 fixed pay, no piece rate), ($3.5, 10c piece rate), and (no fixed pay, 20c piece rate). To further minimize the potential for triggering gift exchange, the pay scheme for the first eight rounds was announced at the beginning of the day's work, thus allowing expectations to settle. The pay scheme for the final two rounds, where the gift will take place, are indicated as 'TBD'.

We should be clear that the design features which allow for estimation of social preferences are not entirely without cost, even as we did our best to minimize the trade-offs. The most important issue is that the changing payment schemes over the workday surely must feel less natural to the workers than the simpler design of previous experiments. Nonetheless, we think that it is worthwhile to have at least some of the field experiments structured to allow for parameter estimation (Card, DellaVigna and Malmendier, 2013), to complement other experiments which have a more natural structure, but do not allow for estimation. The emphasis on parameter estimation led to new predictions (such as comparing the effect of returns to the worker and the employer) and allows us to measure the strength and nature of the underlying preferences, as well as implications for optimal pay schemes.

Design. The detailed design is as follows. We hire temporary workers for a single day's employment (about 6 hours of work) through ads on Craigslist.com. The ads make it clear that the job is a one-time opportunity, and we exclude anyone who attempts to sign up a second time.[15] The work takes place on Saturdays and Sundays on the University of Chicago campus. After showing up at work, the participants are taken to a classroom where a research assistant explains the nature of the work following a script. The Becker Center at the University of Chicago is presented as partnering with the employers – the charities – to facilitate the work.

[14] We did not make the grocery store our main employer because we could not find a compelling way to truthfully vary the return to the employer.

[15] A typical ad read: '*The Becker Friedman Institute is seeking individuals to help prepare letters for fundraising and advertising campaigns. No experience necessary. Employment is for six hours over a single day THIS weekend. [...] Employees can expect to earn around $60 for the day.*'

The participants also receive a sheet indicating a timeline with the pay conditions for the ten rounds of work, except for rounds 9 and 10 which are reported as TBD.

The workers prepare mailers by folding and placing materials in envelopes, working their way through a mailing list. The task is simple but requires attention to match the materials. The workers do the task for 20 minutes, take a 10-minute break, then move on to the next batch of letters for another 20 minutes, and so on for ten rounds. During the 10-minute break, the research assistants count the envelopes produced by each participant and check the accuracy of five of the envelopes per worker. The envelopes include fund-raising material for three charities (*Respond Now*, *Breakthrough Urban Ministries*, and the *Rehabilitation Institute of Chicago*) for 8 rounds and an advertising campaign for a local grocery store for the other 2 rounds.

Figure 2a shows Order A, the first of two orders which we randomize between. In the first four rounds of Order A, the participants fold envelopes at the 10-cent piece rate (and $3.50 flat pay), but with different treatments. The first round is a training period. We tell the participants that they '*will earn a fixed amount of $3.50 plus $0.10 per envelope completed during this training. [...] The training is paid for by the Becker Center. We will be discarding all of the envelopes prepared in this training session.'* (There is no deception in the experiment and the envelopes are discarded as announced). Thus, the employer – the charity – does not directly benefit from the productivity. The training rounds are presented as necessary to ensure that the actual mail solicitations be accurately prepared in the following rounds.[16]

In rounds 2 and 3, the workers stuff envelopes for charity 1 for a 10 cent piece rate: '*As mentioned before, [Charity Name] will be paying for your work. The pay is $3.50 plus $0.10 per envelope completed, as noted on your schedule.*' In round 2, but not in round 3, there is a higher return to the employer due to a donor match: '*Thanks to an anonymous donor, [Charity Name] has received a matching grant that will match every dollar raised by these letters 1 to 1 up to $2,000 total. A number of such matching grant campaigns have been run by charities similar to [Charity Name], and historically, charities like [Charity Name] have yielded roughly $0.60 per mailer with such campaigns, including the match. Given that [Charity Name] is offering a $0.10 per-envelope payment today, it expects to get roughly $0.50 for each additional envelope that you prepare during this session.*' Notice that we make plain to the workers the return net of the piece rate. In round 3, there is the same piece rate, but no match.[17] In round 4, the workers stuff envelopes for a grocery store at the same 10-cent piece rate and with a similar stated 30-cent return to the employer.

After a 50-minute lunch break, the workers restart with a new training period (round 5)

[16]We assume that the Becker Center, which pays the training round piece rate, does not enter the social preference utility of the workers.

[17]The script says: '*A number of such campaigns have been run by charities similar to [Charity Name], and historically, these charities have yielded roughly $0.30 per mailer with such campaigns. Taking account of [Charity Name]'s per-envelope payment for your help today, it expects to get roughly $0.20 for each additional envelope that you prepare during this session*'.

on the material for Charity 2, which pays 20 cents per envelopes.[18] Next, they engage in consequential work for Charity 2 in round 6 at the same 20-cent piece rate. In rounds 7 and 8, they then stuff envelopes for Charity 3 at the 0-cent piece rate, with a charity match in round 8. In rounds 9 and 10, the gift exchange randomization takes place, as we discuss below.

Figure 2b shows the treatments for Order B. Other than for the training rounds, which for logical reasons had to precede the other treatments, the two orders are mirror images of each other: round 8 in the order A becomes round 2 in the order B, round 7 becomes round 3, and so on. While the two training sessions remain in rounds 1 and 5, we do switch the pay schemes in the training period between the two orders. The randomization of the order allows us to observe each treatment in two different positions and thus disentangle the effect of the treatments from confounding effects due to learning and tiredness over the course of the day.[19]

The arrows in Figures 2a and 2b illustrate four planned comparisons between the rounds (which we pre-registered)[20] for the reduced-form results. First, comparing rounds 6 and 7 in order A (3 and 4 in order B) illustrates the impact of the piece rate change from 0 cent to 20 cents. By design, these treatments are contiguous to minimize the impact of learning by doing or tiredness. There is an additional comparison with the 10-cent piece rate in round 3 in order A (round 7 in order B), though the treatments are not contiguous. Second, comparing round 5 and round 6 provides an estimate of the effect of whether the effort counts for the employer, that is the difference between the envelopes being used for the charity and the envelopes being discarded (in the training round). This comparison provides evidence on the magnitude of the social preferences. Third, the comparison between rounds 7 and 8 provides evidence on the impact of the return to the employer, shedding light on the nature of the social preferences (altruism vs. warm glow). Fourth, the comparison between rounds 3 and 4 in order A (6 and 7 in order B) identify the difference between working for a grocery store versus for a charity.

In the final two rounds, round 9 and 10, we implement the between-subject gift treatments. At the beginning of the experiment, we indicated to subjects that in rounds 9 and 10 they would work once again for the charity that previously paid the ($7, $0-piece rate). After the eighth round, we split the workers into separate rooms, one for each gift treatment.[21] Once in

[18]This second training session is justified to the workers by the (slight) difference in materials for the different employer, as well as by restarting after a break.

[19]We considered a full randomization of the order of the treatments, as opposed to just two orders. We decided against it on two grounds. First, the implementation would have been difficult. Second, the choice of just two orders allowed us to maximize power by placing next to each other treatments we intended to compare, like match and no-match (for same piece rate) or training and "real" work (with same piece rate). This minimizes the impact of the counfounding factor of productivity changes over time.

[20]In the pre-registration we emphasize equally comparisons taking place in rounds 1-4 and in rounds 5-10. For example, the effect of training can also be estimated comparing rounds 1 and 2. However, the steep oberved learning by doing in rounds 1-4 confounds these comparisons. Thus, we focus on rounds 5-10 for the reduced form results. The structural estimates use all the variation in the data, including the early periods.

[21]We inform the workers that '*we will now have to split into a few rooms because our room reservation has*

the room, we read to them the script for their relevant gift treatment.

Control. In this treatment, we inform subjects that "*in this session and the next, [Charity name] will pay $7 just as it paid in a previous session.*"

Positive Monetary Gift. We inform subjects that "*in this session and the next, [Charity name] will pay $14 instead of the standard $7 that it paid in a previous session.*"

Negative Monetary Gift. We inform subjects that "*in this session and the next, [Charity name] will pay $3 instead of the standard $7 that it paid in a previous session.*" We follow Kube et al. (2013) in providing no explanation for the wage change.

Positive In-Kind Gift. We tell subjects that "*in this session and the next, [Charity name] will pay $7 as it paid in a previous session. As a token of appreciation, the charity is also giving you this thermos with a retail value of $14.*" We then offer them a gift-wrapped thermos with the name of the charity on it to make clear the gift is coming from the employer. This treatment is modelled on Kube et al. (2012), including the expression of appreciation.

After the announcement of the gift, the workers fold envelopes for rounds 9 and 10, allowing us to test for (quick) decay of the effect of the gift on productivity. One of the two rounds has a donor match raising the return to 60 cents. If a positive gift increases the altruism parameter α, the resulting increase in effort should be larger when the employer return is higher.

After the final treatment, we conduct a short debriefing survey, thank the subjects, pay them according to their accumulated earnings, and walk them to an exit. Participants assigned to different treatments are walked to different exists, minimizing the chance of a meeting.

Randomization. The experimental design involves three crossed between-session randomizations. The first randomization is into order A or B. The second is the assignment (into three orders) of the charities to the role of Charity 1, 2, and 3. The third is whether the charity match is in round 9 or in round 10. This produces 2x3x2=12 combinations. The order of the 12 types of sessions was randomly drawn at the beginning of the study, and then we looped through the 12 sessions six times. On each day that the experiment is run, we run either one or two sessions depending on the number of responses to the advertisement.

The final randomization, to the gift treatments, is made at the individual level within a session. We stratify on the pre-lunch performance to maximize statistical power.

Data. We ran 24 sessions between October 2013 and January 2014 for a total of 131 subjects.[22] We then stopped to estimate the model on the first round of data and ensure that the design that we had settled on based on simulations worked appropriately. On November 21, 2014, we registered the design, including the model, the treatments and randomization, the structural estimation, and the envisioned number of sessions (up to 72 in total, which is

---

expired.'

[22] In September and October 2013, we run 4 sessions with a pilot design. We used the data from the 17 subjects in this pilot to set the pay rate, since we aim to equate on average earnings across the three different piece rates. We are not otherwise using this pilot data in the paper.

what we ended up collecting). The only design change between the first and the second round of data collection was the addition of the in-kind gift treatment which we did not run initially. Between November 2014 and May 2015, we ran 49 sessions for a total of 319 subjects.[23]

After excluding 4 subjects who left the experiment early, the final sample includes 446 workers, the largest sample size that we are aware of among field experiments on gift exchange.

Summary Statistics. The sample (Column 1 of Table 1) is 52 percent female, covers a wide age range, and overrepresents unemployed individuals. Column 2 tests if these demographic variables and two additional self-reported variables predict effort, as measured by the average output over the ten rounds. Productivity is higher for employed individuals and females, as well as for 25-34 years olds relative to both younger and older participants. Using this specification, we form an index of predicted productivity based on demographics, which we then use to test for balance.

Covariate Balance. In Columns 3-6 we examine the randomization with respect to each covariate separately (Panel A), as well as with respect to the index of predicted effort (Panel B). In Column 3, we regress the indicator for Order A on the various demographics. We find a statistically significant relationship with the female indicator, with the indicator for older workers and with self-reported donations. These relationships, which can emerge by chance, do not appear to reflect imbalance of higher-productivity workers into order A or B: Order A is somewhat overrepresented in males (who have lower productivity) and workers 55 years and older (who have somewhat higher productivity). Indeed, the regression using the index of productivity formed based on demographics (Panel B) indicates only a minor degree ($t=0.4$) of selection of higher predicted-effort individuals into Order A. The randomization into the various gift treatments (Columns 4 to 6) reveals no systematic patterns, not surprisingly since the assignment to the gift treatments was stratified on performance in rounds 1-4. The assignment into the different charity orders is similarly orthogonal to observables. Thus, there is no evidence overall of covariate imbalance, except to a limited extent for the order assignment. Particularly given the presence of individual fixed effects in the structural estimation, any small degree of imbalance is unlikely to have substantive implications.

4 Structural Estimation

To estimate the model, we build on Shearer (2004) and Andreoni and Sprenger (2012) and specify an error term that allows us to estimate the model by non-linear least squares.

[23] There actually are 73 sessions because in one of the sessions one of the letter materials was shown incorrectly, and the RA opted to repeat the session. In the spirit of intent-to-treat and transparency, we also retain this session, thus the 73 sessions.

Cost of Effort. We assume that workers maximize (1) with the following specification for the cost function $C_{i,t}(e_{i,t})$ associated with effort $e_{i,t}$ by worker i in round t:

$$C_{i,t}(e_{i,t}) = c\left(e_{i,t}\right) * exp(k_i) * \exp f\left(t\right) * exp(-s * \epsilon_{i,t}) \tag{10}$$

The first term in (10) is the cost of effort function $c\left(e\right)$, which we consider in two families, power, as used in some previous literature, and exponential. The power cost function is $c\left(e\right) = e^{1+s}/\left(1+s\right)$, with $s > 0$ denoting the inverse of the elasticity of effort to the return to effort. In the exponential specification, the cost function is $c\left(e\right) = \exp\left(se_{i,t}\right)/s$ with $s > 0$. Both cost functions satisfy the desired properties $C'\left(e\right) > 0$, $C''\left(e\right) > 0$, and $\lim_{e\to\infty} C'\left(e\right) = \infty$.[24]

The second term is the individual fixed effect: the higher is k_i, the lower the average productivity. We exponentiate the cost so that the individual term $\exp\left(k_i\right)$ is never negative.

The third term, $\exp f\left(t\right)$, captures the evolution of the cost of effort over time: for example, learning by doing entails a declining $f\left(t\right)$. As leading approach we use indicators for the different rounds, but for robustness we also use polynomials. For the indicator approach, we cannot dummy out every round, since that would take out the comparison to the training rounds which are always in rounds 1 and 5. We allow for indicators d_2, d_3, and d_4 for rounds 2, 3 and 4, d_{5-8} for rounds 5-8 and d_{9-10} for rounds 9-10. This specification is motivated by the overall flatness of the output function from round 5 on; by assuming a constant cost from round 5 to round 8, we can identify off, among others, the impact of training (rounds 5 versus 6). The indicator d_{9-10} ensures that the estimated gift effects are not biased by some change in the cost of effort in the last two rounds, and effectively captures the productivity in the Control group (no gift) in rounds 9 and 10.[25] In the alternative polynomial approach, we allow for a quadratic function $f\left(t\right) = \eta_1 t + \eta_2 t^2$ and similarly for a cubic.

The fourth term introduces the error term: we assume ϵ to be normally distributed with $\epsilon_{i,t} \sim N(0, \sigma_\epsilon^2)$, which implies that $\exp\left(-s\epsilon\right)$ has a lognormal distribution. A lognormal distribution avoids a negative error term that would imply negative marginal cost of effort.

Non-Linear Least Squares Derivation. Given these assumptions, we can return to the first-order condition (2). For the power cost of effort function, we obtain

$$p_{Wt} + A\left(Gift, p_{Et}, p_{Wt}\right) - \left(e_{i,t}\right)^s * \exp\left[k_i + f\left(t\right) - s * \epsilon_{i,t}\right] = 0.$$

[24]The exponential cost function does not satisfy the property $C'\left(0\right) = 0$, allowing for the possibility of optimal effort at the zero corner. In our case we can neglect this given the high average mean effort and the fact that the lowest effort ever observed in any round is 7.

[25]Our leading indicator function $f\left(t\right)$ in the pre-registration differs in two ways. First, we restricted the coefficient on d_2 to equal half of the coefficient on d_3, since our earlier estimates sugested that we could not estimate separately a coefficient on d_2. Since our results imply that we can estimate it (though we do not reject the one-half restriction), we allow for a more general specification. Second, we assume the same indicator for rounds 5 to 10. After the registration, we realized that allowing for a separate indicator for rounds 9-10 is more parallel to the experimental design, allowing for a control group for the gift treatments. We show that adopting the pre-registered specification leads to similar results.

Taking the second term to the right hand side and taking logs, we obtain

$$\log\left(p_{Wt} + A\left(Gift, p_{Et}, p_{Wt}\right)\right) = s\log\left(e_{i,t}\right) + k_i + f\left(t\right) - s * \epsilon_{i,t} = 0.$$

Solving out for $\log\left(e_{i,t}\right)$, we obtain the estimating equation

$$\log\left(e_{i,t}\right) = \frac{1}{s}\left[\log\left(p_{Wt} + A\left(Gift, p_{Et}, p_{Wt}\right)\right) - k_i - f\left(t\right)\right] + \epsilon_{i,t}. \tag{11}$$

We estimate equation (11) with a non-linear least squares regression. Equation (11) highlights the advantage of specifying costs as in (10): the fixed effects k_i can be interpreted as individual differences in (minus) log effort, and similarly $f\left(t\right)$ is interpretable as changes over the rounds in (minus) log effort. This equation also makes clear that $1/s$ is the elasticity of effort e with respect to the return of effort, captured by $p_W + A$.

Similarly, we derive the first order condition for the exponential cost of effort function:

$$e_{i,t} = \frac{1}{s}\left[\log\left(p_{Wt} + A\left(Gift, p_{Et}, p_{Wt}\right)\right) - k_i - f\left(t\right)\right] + \epsilon_{i,t}. \tag{12}$$

The exponential cost function leads to the same NLS specification, except with effort e as dependent variable. We thus consider the reduced form effect both on effort and on log effort.

Estimation. We use Stata's `nl` program for estimation, which employs the iterative Gauss-Newton method to converge to a solution, using starting values from a uniform distribution over a range of plausible parameter values. The convergence properties are generally very good. The standard errors for the parameter estimates are clustered by session.

5 Evidence on Baseline Social Preferences

We first analyze the evidence on baseline social preferences from rounds 1 to 8. In Section 6, we then analyze the gift exchange response in rounds 9 and 10.

5.1 Treatment Comparisons

Figure 3 plots average output by round for the two orders. The confidence intervals, as elsewhere in the paper, are clustered at the session level to account for correlation within a worker over time and across workers within a session.[26] In Appendix Figures we provide the corresponding evidence using log output as measure of effort.

There is substantial learning by doing: the average number of envelopes stuffed within 20 minutes increases from about 25 envelopes in round 1 to about 35 in round 4, a 40 percent

[26] Clustering at the session level produces very similar results to clustering at the individual level, indicating that the within-session correlation of errors is small.

increase. From round 5 on, after the lunch break, there are no more obvious gains (nor losses) in productivity, though the average output varies significantly in response to the treatments.

A second clear pattern is the response to piece rate. For example, the only two instances in which productivity decreases substantially from one round to the next are cases of piece rate decreases: rounds 6 to 7 in order A (20c to 0c) and rounds 8 to 9 in order B (10c to 0c).

Turning to the treatments indicative of social preferences, the figure shows a marked response to whether envelopes are used or discarded (holding constant the piece rate paid), comparing rounds 5 and 6. We observe instead a very small impact of changes in employer return due to a match: the productivity is similar in rounds 7 and 8.

The pattern across rounds and between orders in general lines up well with what one would expect. One exception is the comparison across order A and B in rounds 1 and 4: the productivity is higher in order A despite lower incentives than in order B (10c versus 20c). An imbalance in worker ability does not appear to account for much of this difference, since Table 1 provides very limited evidence of selection of more productive workers into order A. Furthermore, in the only two rounds in which the two orders have the same treatments and are thus comparable, rounds 9 and 10, the productivity is very similar across the orders. Instead, it appears that workers work harder for a firm (the grocery) than for a charity.

Effect of piece rate. Figure 4a presents the evidence on piece rate variation, comparing the rounds highlighted in Figure 2. The figure shows very strong response to increasing the piece rate from 0c to 20c – an increase of 4 envelopes, or 12 percent, a highly significant difference. Importantly, the piece rate effect is not confounded by an income effect, since the flat pay is proportionally lower for the higher piece rate. Figure 4a also provides a comparison to the 10c piece rate, though the rounds being compared are not contiguous in this case.

Effect of consequences for employer. The second comparison, in Figure 4b, highlights the impact of consequences to the employer: we compare round 4, when the envelopes are discarded (since it is a training round) to round 5, when the envelopes are sent; notice that the worker piece rate is held constant. This comparison aims to test whether workers display social preferences towards the employer, or more generally whether they care that their work is consequential and has meaning as in Ariely et al. (2008). Productivity is 3.5 envelopes (10 percent) higher when the letters are used, a difference nearly as large as the one induced by a 20c piece rate increase. The difference is even larger in the other training round, comparing rounds 1 and 2, but improvements in productivity over time bias that comparison upward. The structural model utilizes all the rounds and controls for learning across rounds.

Effect of return to employer. The third comparison, in Figure 4c, examines the impact of the precise employer return: do workers work more when the return from their effort is higher (due to a match on the resulting donations)? The answer is largely no: the higher match rate leads to a statistically insignificant increase in the number of envelopes stuffed of just 0.6 envelopes (1.7 percent), an effect size much smaller than the impact of piece rate or of consequential work. This evidence suggests that a warm glow model is likely to better capture the social preferences at work compared to a pure altruism model.

Effect of Charities and Firm. Appendix Figure 1a shows that the effort provided does not differ sizably between the three charities. The charity assignment is randomized, so this comparison holds constant the variation in other treatments. Thus, in most of the paper, we pool across the charities. Appendix Figure 1b compares the effort when the employer is a charity versus when it is a firm (a grocery store), holding constant the return to the employer and the piece rate. The effort is actually somewhat *higher* for the firm compared to the charities. This suggests that the substantial baseline social preferences we identify towards charities as employer may represent a lower bound for social preferences towards other employers.

Overall, we document three main results. First, worker effort is clearly sensitive to private incentives. Second, workers appear to have substantial social preferences towards employers, since they exert more effort when their work is directly useful to the employers (i.e. when the envelopes they produce are actually mailed out). Third, the social preferences are most likely of the warm-glow type, rather than pure altruism, given the very limited response to the higher return to the employer, conditional on the work being used at all.

5.2 Structural Estimates

To estimate the social preferences, and to utilize all the information, we estimate the model in Section 4 using data from the first eight rounds. (We include also data from the final two rounds when estimating gift exchange effects).

As we discussed in Section 4, the first-order conditions for optimal effort imply that we can estimate the model with non-linear least squares, allowing for individual fixed effects and for learning or tiredness over time. The specification is not an OLS regression just because the marginal benefit of effort appears in log format as $\log(p_W + A)$. Furthermore, the two cost functions—power and exponential—reduce to the same specification, with different dependent variables: log productivity for the power function and productivity for the exponential function.

Estimates. Column 1 in Table 2 shows the results for the pure altruism case. The social preference weight A equals $\alpha(p_E - p_W)$: the worker places value α on the net marginal return to the employer. The estimated altruism $\hat{\alpha}$ across the three charities is comparable, ranging from 0.20 to 0.28. The weight on the grocery store is significantly higher, capturing the higher productivity for the firm (Figure 3).[27] The curvature of the cost function is high at 11.3,

[27]The difference in altruism weight α between the charity employers and the firm may appear surprisingly large. One factor that biases upward the estimated altruism weight towards the grocery store is the fact that for the grocery store we only run a condition with normal return (30 cents) for the employer, while for the charities we also have conditions with higher return (60 cents). Since the workers do not respond to the higher return, the estimated altruism for the charities is compressed downward, while this is not the case for the firm. Consistent with this, the difference in social preference weight is much smaller in the warm glow case, which

implying an elasticity of effort to the value of work of $1/11.3 = .09$. The inelastic cost function is not surprising: in a fixed amount of time, it is hard to increase productivity much.

To highlight the identification, recall from (9) that in a pure altruism model to a first approximation the altruism coefficient is pinned down by the ratio of the response to the employer return and the response to the piece rate: $(\partial e^*/\partial p_E)/(\partial e^*/\partial p_W) = \alpha/(1-\alpha)$. In our setting, effort increases by 0.6 envelopes for a 30-cent increase in the employer return; thus, $\partial e^*/\partial p_E \simeq .02$. In contrast, effort increases by 4 envelopes in response to a 20-cent increase in the piece rate; thus, $\partial e^*/\partial p_W \simeq .2$. It follows that α is approximately 0.11. Why is the estimated α then larger? The issue is that for an α of about .1 the model badly undermatches the effect of going from training to a normal period. Moving from the 20-cent training (round 5 in order A) to a standard 20-cent period (round 6 in order A), equation (11) implies that log output should increase by

$$\log(e_{t=6}) - \log(e_{t=5}) = \frac{1}{s}\left[\log(.2 + \alpha(.3 - .2)) - \log(.2)\right],$$

since the return to the employer in round 6 equals 0.3 (the raw return) minus the piece rate, 0.2. For an $\alpha \simeq .1$, the right-hand side equals approximately $s^{-1}[\log(1.05)] \simeq .005$ log points, that is, a half percent increase in output. In the data, the output increase is instead 3.5 envelopes, or about 0.1 log points. Thus, the altruism model struggles to capture both the response to match and to training, as we further highlight below.

Column 2 in Table 2 shows the warm glow results. In this specification, the worker cares about the average return to the employer, which we set to 0.3, not the actual return $p_E - p_W$. Thus, $A = .3a$, where a is the warm glow weight. The workers exhibit significant warm glow, putting weight on the employer equal to about half the weight put on private payoffs. The estimated curvature s of the cost of effort function is similar to the one with altruism.

To assess the identification of the warm glow model, recall that this model does not predict any response to the charity match, consistent with the data. As far as the response to training, it predicts a response of

$$\log(e_{t=6}) - \log(e_{t=5}) = \frac{1}{s}\left[\log(.2 + .3a) - \log(.2)\right].$$

For the estimated $\hat{a} \simeq .4$, the right hand side equals $s^{-1}\log(1.6) = .05$, a predicted 5 percent increase in output from round 5 to 6. Thus, the warm glow model comes much closer to capturing the effect of the training.

How do we compare quantitatively the fit of the altruism and warm glow models? The explanatory power R^2 of the two models, which have the same number of parameters, is higher for the warm glow case at .8377 compared to .8350 for altruism. This difference may seem

assumes an average return to the employer. As a second factor, given how inelastic the production function is, it takes a large increase in social preference to match even a quite small change in output.

small, but one ought to keep in mind that (i) most of the variation is explained by the individual fixed effects and (ii) the fit of the two models effectively only differs in rounds with high returns to the employer and in the training rounds, and thus differences in fit cannot be large.

For a second comparison, we nest the two models and allow for both altruism and warm-glow in Column 3. The models are separately identified because of the variation in return to the employer and the presence of a training period. (As such, we cannot estimate two coefficients for the grocery store, since we do not vary the return there.) The results are striking: the data does not reject the null of no altruism towards any of the three charities, while the warm glow weights are very similar to the ones in Column 2.

The specifications so far follow the assumption of power cost of effort function. In Columns 4-6 we present the results assuming a cost of effort function with declining elasticity, the exponential function. The implied NLS specification is the same as in Columns 1-3 except that the dependent variable is number of envelopes prepared in the round, as opposed to the log. The results for the two specifications are nearly identical. Thus, the identification of the baseline social preferences is not sensitive to the exact assumptions about the cost of effort.

In Table 3, we present a set of robustness checks, adopting as benchmark the specification with warm glow and altruism for both power and exponential cost functions. First, we examine alternative modelling of the learning by doing: instead of set of indicators, we use a polynomial in the round number, quadratic in Columns 1 and 4 and cubic in Columns 2 and 5. The results are similar, supporting the warm glow model. Next, we use an alternative assumption for the training periods. Instead of assuming that during training periods the workers have no social preferences (since the letters are not used), we allow for warm glow, restricted to be half the size as in the periods where output is used. Under this alternative, the estimated warm glow is higher: it now takes a higher lever of warm glow to match the observed output increase from training to a real production period. There is also some evidence of pure altruism, although the altruism coefficients are smaller than the warm glow ones.

Quality of Fit. In Figures 5a-c we compare the observed patterns with the predictions of the altruism and warm glow models. Specifically, for the exponential cost specifications (Columns 4-5 of Table 2), we compute the predicted effort and average across the relevant observations. While both models fit well the response to incentives (Figure 5a), the altruism model has trouble fitting the combination of a large response to the training (Figure 5b) and a small response to the match rate (Figure 5c). The last piece of evidence would imply a small altruism parameter α, but then the model predicts a training effect that is too small relative to the data. The warm glow model instead predicts no response to the match rate by assumption and matches the training effect with a higher level of warm glow a.

Optimal Piece Rate. What are the implications of the estimated social preferences? As a step to address the question, we simulate the counterfactual productivity as we vary incentives. In particular, we increase the piece rate from 0 cents (per envelope) to a higher piece rate, holding constant the flat pay. We do the simulation for an individual with average cost k.

Consider first the no-warm-glow case ($a = 0$). At zero piece rate, the worker puts zero effort; as such, introducing a piece rate has dramatic impacts on output (blue continuous line in Figure 6a), and thus on profits (blue continuous line in Figure 6b). The optimal piece rate is 3 cents, as higher piece rates, while yielding higher output, do not pay for themselves.

Consider then an individual with our estimated warm glow ($\hat{a} = .4$). Effort is quite high even with no piece rate and it increases slowly with the piece rate increases (green dotted line in Figure 6a); thus, the profit rate is monotonically decreasing in the piece rate, as the productivity increases for a higher piece rate do not pay off enough.

Why the difference? Given the warm glow, the worker already puts in the productivity associated with small or moderate marginal costs, stopping only where the marginal cost of effort is elevated. A higher piece rate pushes the individual to tackle these extra envelopes as well, but given the high curvature, the productivity gains are small. More generally, the warm glow ensures that the employee takes care of the low-marginal-cost tasks, the ones that otherwise it would take incentives to address. Examples in typical workplaces include reporting a defect in a product or introducing briefly a new worker to a production line.

Does this result depend critically on the magnitude of the estimated warm glow? The red lines in Figures 6a-b show that this is not the case. For an employee with only one tenth the estimated warm glow as in our estimates ($a = .05$), the shape of the profit function is similar as for our estimated warm glow level (although in this case the optimal piece rate is interior). It does not take much social preference for the workers to take care of the low-marginal-cost tasks which, at least in our estimate, would do most of the work towards profit maximization.

This argument is consistent with the emphasis of Akerlof and Kranton (2008) and Besley and Ghatak (2005) that it is critical for firms to recruit workers that believe in the mission of the firm. It also reflects the theoretical results of Englmaier and Leider (2012) that incentives and social preferences are largely substitutes in motivating workers: one really needs one of the two, not necessarily both. Having said that, an obvious caveat is that the results may differ significantly in other settings where, for example, the marginal increment in productivity due to extra incentives may be high, even with social preferences.

6 Evidence on Gift Exchange and Reciprocity

After considering the evidence on baseline social preferences towards the employer, we now examine the evidence for reciprocity in response to unexpectedly generous, or ungenerous, behavior by the employer. In reciprocity models, the social preference weight that the person places on another person is affected by kind, or unkind, actions of the other person. In our setting, we do not model the underlying reciprocity model, but we capture its key property, which we measure as a shift in the warm glow weight in response to a 'gift'.

Following the literature, we randomize subjects into either a control group which receives the expected wage, or treatments groups which get more, or less, than expected. Unlike in previous papers, in our design the gift exchange treatments come after the other treatments discussed above. This allows us to benchmark gift-exchange effects to piece-rate incentives and thus obtain parameter estimates for both baseline social preferences and reciprocal response.

6.1 Treatment Comparisons

After round 8 subjects are randomized into four treatments and sent to separate rooms (Figure 2). All subjects once again work for a charity that previously paid $7 for 20 minutes.

In the *control* group, the charity once again pays the $7 flat pay in rounds 9 and 10, therefore plausibly matching expectations (as we document below with a debriefing survey).

The *positive monetary* gift group, modeled upon Gneezy and List (2006) and follow-up papers, is told that the charity which used to pay $7 will now pay $14 in rounds 9 and 10, a doubling of pay and a $14 monetary 'gift' relative to expectations. Since subjects earn an average of $70 overall, as share of total earnings this constitutes a 20 percent increase.

The *negative monetary* gift group, modeled upon Kube, Marechal, and Puppe (2013), is told that the charity which used to pay $7 will now only pay $3 per round, less than half than before, a $8 monetary cut, 12 percent as fraction of overall earnings. Following the literature, no specific reason or justification is provided for this pay cut.

The *positive in-kind* group, modeled upon Kube, Marechal, and Puppe (2012), is told that, on top of $7, they also receive a gift-wrapped thermos from the charity, for a value of $14. Thus, the cost of the gift to is the same in the positive in-kind and monetary gift treatments. Kube, Marechal, and Puppe (2012) find that the in-kind gift is more effective at triggering gift exchange, presumably because it is perceived to be a kinder gesture. The in-kind gift treatment was not run in the first 24 sessions and thus has a somewhat smaller sample size.

Figure 7a shows the results.[28] Compared to the control group, the positive gift treatment produces 0.45 additional envelopes, a difference that is not significant. The negative monetary gift treatment has an even smaller effect, keeping productivity essentially constant compared to the control group. Finally, the in-kind gift actually leads to a *decrease* of productivity of 1.15 envelopes, although the difference is not significant. Overall, these results suggest very limited gift-exchange reciprocity – both positive and negative – by workers towards the employer.

Controls. Panel A of Table 4 shows the treatment effects in regression form in Column 1. In Columns 2 and 3, we add as a control the average previous productivity of a worker across rounds 1 through 8 (Column 2) and across rounds 5 through 8 (Column 3). Controlling for past performance raises the R^2 substantially. Since the specification in Column 3 has higher explanatory power, we use it for the subsequent analysis with controls.

[28] Appendix Figures 4a-d show that the results are parallel using log output.

Figure 7b plots the coefficients from Column 3. The addition of controls lowers the standard errors by about a quarter, without essentially changing the point estimates. We do not find any significant effect of the gift treatments, and in fact we can reject relatively small gift effects. We can reject that a negative gift lowers effort by more than 1.6 envelopes, a 4.4 percent decrease, a much smaller effect than the 20 percent decrease in Kube et al. (2013). We can also reject that the in-kind gift increases productivity by more than 0.7 envelopes, a 2 percent increase, again much smaller than the 25 percent increase in Kube et al. (2012).

Panel B of Table 4 display the results using log output as a measure, the measure implied by a power cost function. The pattern of results is very similar.

Decay. To examine the decay of gift effects over time, we pre-registered a comparison of rounds 9 and 10. To the extent that the gift exchange effects are very short-lived, the effects should be stronger in round 9 than in round 10. Figure 7c shows this comparison. While there is no discernible pattern for the negative gift and the in-kind gift treatments, for the positive monetary gift treatment there is an initial 1 envelope (3 percent) increase in round 9, with no effect instead in round 10. Columns 4 and 5 in Table 3 show the regression results. After controlling for average performance in rounds 5-8, the positive gift treatment is associated with an increase of 1.3 envelopes in round 9, a statistically significant increase, with no effect in round 10. These patterns are parallel using log output as dependent variable (Panel B).

This evidence of a gift effect with rapid decay could be related to the finding of Gneezy and List (2006) of positive gift exchange which decays over time. That being said, the effect in our setting is much smaller (3 percent versus 30 percent) and the rate of decay is faster (30 minutes versus 90 minutes). Overall, there is suggestive evidence on an initial positive effect in the positive gift treatment, with no evidence of gift effects for the other treatments.

Match. A second pre-registered comparison is the interaction with the return to the firm. The return to the charity was doubled due to the presence of a match in either round 9 or round 10 (depending on the randomization). We can thus test whether reciprocity is responsive to the return to the employer, although the power of the test is limited given the overall null effect of the gifts. Figure 7d and Columns 6 and 7 in Table 4 show that there is no effect of the match on the response to the gifts. This evidence is consistent with the evidence on baseline social preferences suggesting little scope for the precise return to the employer.

Heterogeneity. The estimates so far provide evidence only on the average impact of the gifts, which could mask substantial heterogeneity. Figure 8 plots the CDF of productivity in rounds 9 and 10 for the various treatments. To enhance power, we plot the residual of output, controlling for average output in rounds 5-8. While the differences between the treatments are small, there are hints of an output increase with the positive gift at lower quantiles.

6.2 Structural Estimates

We now estimate the social preferences on the full data set (rounds 1 to 10). Relative to the examination of baseline social preferences in Tables 2 and 3, the key addition is the estimation of how much the social preference parameters shift (if at all) in response to the gift treatments. These parameters capture reciprocity models which posit that the social preferences depend on how an individual is treated (Fehr and Gächter 2000; Charness and Rabin 2002; Dufwenberg and Kirchsteiger 2004; Falk and Fischbacher 2006). Whether reciprocity is intention-based or action-based, a reciprocal worker who receives a surprisingly generous treatment from the employer is likely to display more positive social preferences towards the employer.

We make the following additional assumptions. First, given the strong support for the warm glow model for baseline social preferences, we assume that social preferences are of the warm glow type. Second, since the results on baseline social preferences suggest similar social preferences across the three charities, we assume the same warm glow across the charities.

Table 5 displays the results for power cost (Columns 1-3) and exponential cost function (Columns 4-6). In addition to the specification with indicators for the various rounds (Columns 1 and 4), we also estimate a quadratic (Columns 2 and 5) and a cubic polynomial (Columns 3 and 6) in the number of rounds. The three specifications for the learning by doing have similar explanatory power. Appendix Figure 5 plots the implied productivity shift over time $-(1/s)\left[\bar{k}+f(t)\right]$ for an individual of average k for the three specifications, showing quite similar estimated time effects. Thus, it is not surprising that the results on baseline social preferences and curvature of the cost of effort are quite similar across the specifications, and are parallel to the estimates of baseline social preferences (Table 2).

Turning to the reciprocity parameters, we do not find statistically significant evidence of reciprocity for any of the gift treatments. The magnitudes however differ across the types of gift, in ways consistent with the reduced-form results. For the positive monetary treatments, there is an increase in warm glow ranging from .05 to .20. The latter estimate is sizeable, being a 50 percent increase over the baseline warm glow coefficient of .4.

How can a small increase of .6 envelopes (a 2 percent increase) in the reduced form translate into a 50 percent increase in warm glow? This is a case where the link to the model is important: given the inelastic cost of effort function, even small increases imply sizeable (if imprecisely estimated) shifts in social preferences.

The estimates for the negative gift treatment are smaller in absolute value, ranging from -.01 to -.08, consistent with the very small reduced-form estimates. Finally, the in-kind gift treatment is associated to a decrease in the warm glow parameter, ranging from -0.07 to -0.12.

In Table 6, we consider robustness checks to the results for both the power and exponential cost function. In Columns 1 and 4, we consider an alternative specification for the round indicators which we had pre-registered, yielding similar results.[29] Next, motivated by the observed decay of the gift effects in Gneezy and List (2009), we allow for the warm glow effect of a gift, a_{Gift}, to decay. Namely, the warm glow parameter for round 10 equals δa_{Gift}, with $\delta = 1$ indicating no decay and $\delta = 0$ indicating full decay.[30] As Columns 2 and 5 show, the estimates indicate substantial decay ($\hat{\delta} = .23$) and a larger and statistically significant effect of the positive monetary gift ($\hat{a}_{Gift} = .37$). This is consistent with the result in Table 4 and Figure 7c of a decay effect for the positive monetary gift treatment (though not for the other gift treatments). Finally, as a last robustness check, in Columns 3 and 6 we present estimates assuming pure altruism, rather than warm glow. The model has trouble converging, but at least for the exponential cost function reports qualitatively similar conclusions.

Finally, we report the fit of the model across all ten rounds. As for Figures 5a-c, we take the exponential cost specifications for either warm glow (Columns 4 of Table 5) or altruism (Column 6 of Table 6), compute the predicted effort and average across the relevant observations in a round-order. We in particular average across the different gift treatments in rounds 9-10. Figure 9 shows that the model overall does well in fitting the data.

7 Discussion and Conclusion

In this Section we take stock of the results and discuss them in the context of the literature.

Real-Effort Experiments. The experiment is geared towards the estimation of social preferences, but the use of varying incentives to back out the cost of effort function applies more broadly to experiments with real effort tasks. While most such papers focus on reduced-form estimates, some, like us, use variation in piece rate to pin down the cost of effort function.

The papers we locate in this literature all use variants of a power cost of effort function: Augenblick, Niederle, Sprenger (2015) and Augenblick and Rabin (2015) use a power function with a (fixed) intercept $c(e) = (\omega + e)^{\gamma}$. Gill and Prowse (2012) estimate a model of disappointment aversion with quadratic cost, a special case of power cost with unit elasticity.

We show that one can easily extend the estimation to an exponential cost function. Unlike the power function, the exponential function is not characterized by constant elasticity and thus provides a useful alternative. In our case, the two cost functions yield very similar results, which is not surprising given that most of our effect sizes are quite small, and thus the overall curvature does not matter as much. The different cost functions could matter much more, however, for larger effect sizes. We hope that future work will examine this further.

[29] Namely, the specification restricts the indicator for round 2 to equal 1/2 the indicator for round 3. Furthermore, the specification does not allow for a separate indicator for rounds 9-10, requiring it to equal the indicator for rounds 5-8.

[30] We pre-registered this specification with decay. In light of the small gift effects, we report this specification but in the benchmark specification we restrict $\delta = 1$.

Baseline Social Preferences. A key finding in the paper is the support for a 'warm glow'-type model, as opposed to a model of pure altruism towards the employer. The sizeable magnitude of the estimated warm glow implies that social preferences can play an important role as motivators at work, and have the potential to substitute for incentives, at least when incentives are hard to design or apply, for example due to multi-tasking or observability problems. The theme that social preferences towards employers may be of importance in the workplace is a classical one. However, we are aware of little previous field evidence pinning down the features of such social preferences, let alone obtaining structural estimates.

An important theme to return to is the interpretation in terms of warm glow. Clearly, our model of 'warm glow' is not a fully specified model. We see it as a simple alternative to the pure-altruism model that is typically used to model workplace social preferences (e.g. Bandiera, Barankay, and Rasul, 2006 and Hjort 2014). The warm glow model could be capturing norms in the workplace ('one needs to put effort') or value placed on 'meaningful' effort (Ariely, Kamenica and Prelec 2008). We see the model in the paper as a starting point to better understand social preferences in the workplace.

A related issue is the role of the training rounds, which play an important role for the estimates of social preferences. We assume that in the training rounds, in which the envelopes are discarded, the workers are working only because of the piece rate compensation. It is certainly possible that social preferences might still be active during training, since workers may reason that working harder in the training will prepare them to do good work for the employer in the following rounds. In Columns 3 and 6 of Table 3, we estimate a version of this alternative model, assuming that social preferences partially act also in the training rounds. The social preference estimates are altered, as expected, but the overall pattern is similar.

Another issue is the fact that the employers are charities. While a fraction of the population works for non-profits, employment under a profit-seeking employer is clearly more common. Would our results extrapolate to this situation? We chose charities as employers not out of a preference, but because it was the employer for which we could most plausibly convey, and vary, the return to effort. One concern in regard to this choice is that employees may work extra hard for a charity compared to a corporate employer. Thus, we included a round of working for a grocery store for comparison. To our surprise, the workers work at least as hard for the grocery store, which suggests that the estimates for warm glow are not necessarily overestimates of what we would find for a corporate jobs, though only follow-up work will tell.

Gift Exchange. Turning to Section 6 on gift exchange, a sizeable literature exists on gift exchange experiments in the field (partly surveyed in Esteves-Sorenson, 2015), even if the findings are not typically translated into the underlying social preferences. Our results differ from some previous papers which find larger impacts from some of the gift treatments.

What explains then the small gift effects for all three treatments in our setting? We consider four leading possibilities and argue that they do not appear to explain the results.

A first possibility is that our gift treatments did not trigger the required surprise and mood response to induce reciprocation. To address this important concern, we use a short debriefing survey that the workers complete after the task. For the last 65 workers[31], we ask them '*How did the pay in the last two periods make you feel? (Check all that apply) [ ] No particular reaction [ ] It made me happy [ ] Felt more motivated and energetic [ ] It was what I expected [ ] Surprised, it was more than I expected [ ] Surprised, it was less than I expected [ ] Felt unhappy [ ] Felt insulted [ ] It was unfair.*' We code the share that report being happy or unhappy, as well as the share reporting a positive surprise or a negative surprise. As Figures 10a and 10b show, in the positive gift treatments 70 to 80 percent of subjects report positive mood, compared to 20 percent in the control group and 5 percent in the negative gift group. The results are similar for positive surprise, and are the converse for unhappiness and negative surprise. Thus, the results are not due to lack of a 'first stage' on emotions and perceptions.

A second possibility is that the worker effort towards the end of the experiment has become habitual and unresponsive to inputs, thus explaining the lack of response to the gifts. We test for this hypothesis by examining the response to piece rate variation in the last rounds. As Figure 3 shows, in Order A there is a 3-envelope reduction in output from round 8 to round 9 as the piece rate decreases from 10c to 0c. There is no such change in Order B, where the piece rate stays at 0c. Thus, subjects are highly responsive to motivation even in the later rounds.

A third explanation is that the gifts were of smaller magnitude compared to the existing papers. However, the positive monetary gift and the in-kind gift have approximately the same value as those used previously. The negative gift treatment more than halves the pay for the last two rounds (from $7 to $3); the pay reduction is smaller than in previous papers only if considered as share of *total* earnings (12 percent versus 33 percent in Kube et al., 2013). Given our point estimate of essentially zero on the negative gift treatment, extrapolating linearly to the 33 percent cut used in previous experiments is unlikely to generate substantial effects.

A fourth set of explanations consists of various confounders, such as those pointed out by Esteves-Sorenson (2015). None of the confounders applies to our study. In particular, subjects are recruited at market wages (avoiding selection of workers hired at above-average wage), there are no peer effects by design as subjects are split into rooms, nor are there issues with sample size, as ours is the largest gift exchange experiment in the field run to our knowledge. Further, previous papers have pointed out that gift exchange may not occur if workers do not have a clear channel to reciprocate (Englmaier and Leider 2012) or do not know the return to firm. Neither applies in our case.

We suggest two explanations for the gift exchange results. The first is that gift exchange effects are simply of smaller magnitude than found in some of the earlier papers. Indeed, Esteves-Sorenson (2015) similarly fail to find an effect from a positive gift exchange treatment

[31] We are thankful to Uri Simonsohn and Eldar Shafir for suggesting the addition of the question and making suggestions on the wording.

in a large-sample field experiment. A second explanation is that the social preferences towards the employer are, to a first approximation, set at first contact. Gift treatments, under this explanation, are more effective when introduced initially, as in previous experiments, as opposed to during an ongoing work relationship, as in our experiment.[32] Under either interpretation, gift exchange in the workplace is likely to be of more limited impact than initially conceived. Still, our experiment also documents substantial baseline social preferences at work.

Calibration to Previous Gift Exchange Experiments. We present also an alternative approach to relate our estimates to previous experiments on gift exchange in the field. As we discussed, estimating social preferences in previous experiments is infeasible, given missing design elements. However, as the model in Section 2 shows, one can derive a measure of reciprocity, provided we know one key variable. Using (5), under a power cost of effort function, we can derive the proportional increase in warm glow $(a + a_{Gift})/a$ due to the gift. Such reciprocity measure equals $(e^*_{Gift}/e^*_{Control})^s$, where s is the measure of curvature of the cost of effort function, and thus the inverse of the elasticity of effort with respect to incentives. While we do not observe the elasticity in these tasks, a conjecture is that the elasticity may not be so different as the one in our task. After all, all the tasks used in the literature are measures of quantity produced in a fixed amount of time, making it hard to dramatically increase production in response to increased incentives or motivation. Provided a calibration for s, we can revisit some previous experiments and compute the implied reciprocity.

Table 7 shows the results for some classical papers in the literature. For elasticity values as in our task, some previous papers imply very large reciprocity effect, such as a 400 or 700 percent increase in social preferences with a positive gift, and an 88 percent decrease with a negative gift. The calibrated reciprocity falls to a 40-50 percent increase only the elasticity is five times larger than in our experiment.

We hope that future experiment will consider incorporating the estimation of the elasticity in the design, as we do.

[32] We leave it to future research to separate out the two explanations. After having already completed the largest gift exchange field experiment to date, we attempted to recruit more workers for an additional field experiment with the traditional between-subject structure. However, we could not recruit enough subjects from the same population to guarantee adequate power for the test.

References

[1] Akerlof, George A. 1982. "Labor Contracts as Partial Gift Exchange", *Quarterly Journal of Economics,* 97 (4): 543-569.

[2] Akerlof, George A. and Rachel E. Kranton. 2005. "Identity and the Economics of Organizations" *Journal of Economic Perspectives,* Vol. 19, No. 1, pp. 9-32.

[3] Andreoni, James. 1989. "Giving with Impure Altruism: Applications to Charity and Ricardian Equivalence." *Journal of Political Economy*, 97(6), 1447-1458.

[4] Andreoni, James. 1990. Impure Altruism and Donations to Public Goods: A Theory of Warm-Glow Giving. *The Economic Journal*, 100(401), 464-477.

[5] Andreoni, James and Charles Sprenger. 2012. "Estimating Time Preferences from Convex Budgets," *American Economic Review*, vol. 102(7), pages 3333-56.

[6] Ariely, Daniel, Emir Kamenica, Drazen Prelec. 2008. "Man's search for meaning: The case of Legos" *Journal of Economic Behavior & Organization*, Volume 67, Issues 3–4, Pages 671–677.

[7] Augenblick, Ned, Muriel Niederle and Charles Sprenger. 2015. "Working Over Time: Dynamic Inconsistency in Real Effort Tasks " *Quarterly Journal of Economics*, 130 (3): 1067-1115.

[8] Augenblick, Ned and Matthew Rabin. 2015. "An Experiment on Time Preference and Misprediction in Unpleasant Tasks" Working paper.

[9] Bandiera, Oriana, Iwan Barankay and Imran Rasul. 2005. "Social Preferences and the Response to Incentives: Evidence from Personnel Data", *Quarterly Journal of Economics,* Vol. 120, No. 3, pp. 917-962.

[10] Barseghyan, Levon, Francesca Molinari, Ted O'Donoghue, and Joshua C. Teitelbaum. 2013. "The nature of risk preferences: Evidence from insurance choices." *American Economic Review* 103, no. 6 (2013): 2499-2529.

[11] Becker, Gary S. 1974. "A Theory of Social Interactions" *Journal of Political Economy,* 82(6), 1063-1093.

[12] Bellemare, C. and B.S. Shearer (2011) "On the Relevance and Composition of Gifts within the Firm: Evidence from Field Experiments," *International Economic Review*, Vol. 52(3), pp. 855-882.

[13] Besley, Tim and Ghatak, Maitreesh (2005) "Competition and incentives with motivated agents". *American Economic Review*, 95 (3). pp. 616-636.

[14] Breza, Emily, Supreet Kaur, and Yogita Shamdasani. 2015. "The Morale Effects of Pay Inequality" Working paper.

[15] Card, David, Stefano DellaVigna, and Ulrike Malmendier. 2011. "The Role of Theory in Field Experiments". *Journal of Economic Perspectives,* 25(3), pp. 39-62.

[16] Charness, Gary and Peter Kuhn. 2007. "Does Pay Inequality Affect Worker Effort? Experimental Evidence", *Journal of Labor Economics*, Vol. 25, No. 4, pp. 693-723.

[17] Cohn, Alain, Ernst Fehr, Benedikt Herrmann and Frédéric Schneider. 2014. "Social Comparison And Effort Provision: Evidence From A Field Experiment", *Journal of the European Economic Association* ,pp. 877–898.

[18] Cohn, Alain, Ernst Fehr, and Lorenz Goette. Forthcoming. "Fair Wages and Effort: Evidence from a Field Experiment", *Management Science*.

[19] Conlin, Michael, Ted O'Donoghue, and Timothy J. Vogelsang. 2007. "Projection Bias in Catalog Orders." *American Economic Review,* 97(4), 1217-1249.

[20] DellaVigna, Stefano, List, John. A., and Malmendier, Ulrike. (2012). "Testing for Altruism and Social Pressure in Charitable Giving." *Quarterly Journal of Economics*, 127(1), 1-56.

[21] DellaVigna, Stefano, List, John. A., Malmendier, Ulrike, and Gautam Rao (2015). "Voting to Tell Others." working paper.

[22] Englmaier, Florian and Stephen Leider. 2012a."Contractual and Organizational Structure with Reciprocal Agents" *American Economic Journal: Microeconomics*, 2012, vol. 4, issue 2, pp. 146-83.

[23] Englmaier, Florian and Stephen Leider. 2012b."Managerial Payoff and Gift Exchange in the Field" Working paper.

[24] Charness, Gary, and Matthew Rabin. 2002. "Understanding social preferences with simple tests." *Quarterly journal of Economics*: 817-869.

[25] Esteves-Sorenson, Constanca. 2015. "Gift Exchange in the Workplace: Addressing the Conflicting Evidence with a Careful Test " Working paper.

[26] Dufwenberg, Martin, and Georg Kirchsteiger. 2004. "A theory of sequential reciprocity." *Games and economic behavior* 47.2: 268-298.

[27] Falk, Armin, and Urs Fischbacher. 2006. "A theory of reciprocity." *Games and economic behavior* 54.2: 293-315.

[28] Fehr, Ernst, Georg Kirchsteiger, and Arno Riedl. 1998. "Gift exchange and reciprocity in competitive experimental markets." *European Economic Review* 42.1: 1-34.

[29] Fehr, Ernst and Simon Gächter, 2000. "Fairness and Retaliation - The Economics of Reciprocity", *Journal of Economic Perspectives*, 14, 159-181.

[30] Gill, David and Victoria Prowse. 2012. "A structural analysis of disappointment aversion in a real effort competition" *American Economic Review*, 102(1), 469-503.

[31] Gilchrist, Duncan S., Michael Luca, and Deepak Malhotra. Forthcoming. "When 3+1>4: Gift Structure and Reciprocity in the Field." *Management Science*.

[32] Uri Gneezy and John A List. 2006. "Putting Behavioral Economics to Work: Testing for Gift Exchange in Labor Markets Using Field Experiments" *Econometrica*, Volume 74, Issue 5, pages 1365–1384.

[33] Hjort, Jonas. 2014. "Ethnic Divisions and Production in Firms", *Quarterly Journal of Economics*, vol.129, no. 4, pp. 1899-1946.

[34] Kahneman, Daniel, Jack L. Knetsch and Richard Thaler. 1986. "Fairness as a Constraint on Profit Seeking: Entitlements in the Market" *American Economic Review*, Vol. 76, No. 4, pp. 728-741.

[35] Krueger, Alan B. and Alexandre Mas. 2004. "Strikes, Scabs, And Tread Separations: Labor Strife And The Production Of Defective Bridgestone/Firestone Tires," *Journal of Political Economy*, v112(2), 253-289.

[36] Kube, Sebastian, Michel André Maréchal and Clemens Puppe. 2012 "The Currency of Reciprocity: Gift Exchange in the Workplace" *American Economic Review,* Vol. 102, No. 4, pp. 1644-1662.

[37] Kube, Sebastian, Michel André Maréchal and Clemens Puppe. 2013. "Do Wage Cuts Damage Work Morale? Evidence From A Natural Field Experiment", *Journal of the European Economic Association* Volume 11, Issue 4, pages 853–870.

[38] Laibson, David, Andrea Repetto, and Jeremy Tobacman. 2007. "Estimating Discount Functions with Consumption Choices over the Lifecycle," Working paper.

[39] Shearer, Bruce. 2004. "Piece Rates, Fixed Wages and Incentives: Evidence from a Field Experiment" *Review of Economic Studies*, 71 (2): 513-534.

Figure 1. Example of Equilibrium Effort Determination – Marginal Benefit and Marginal Cost

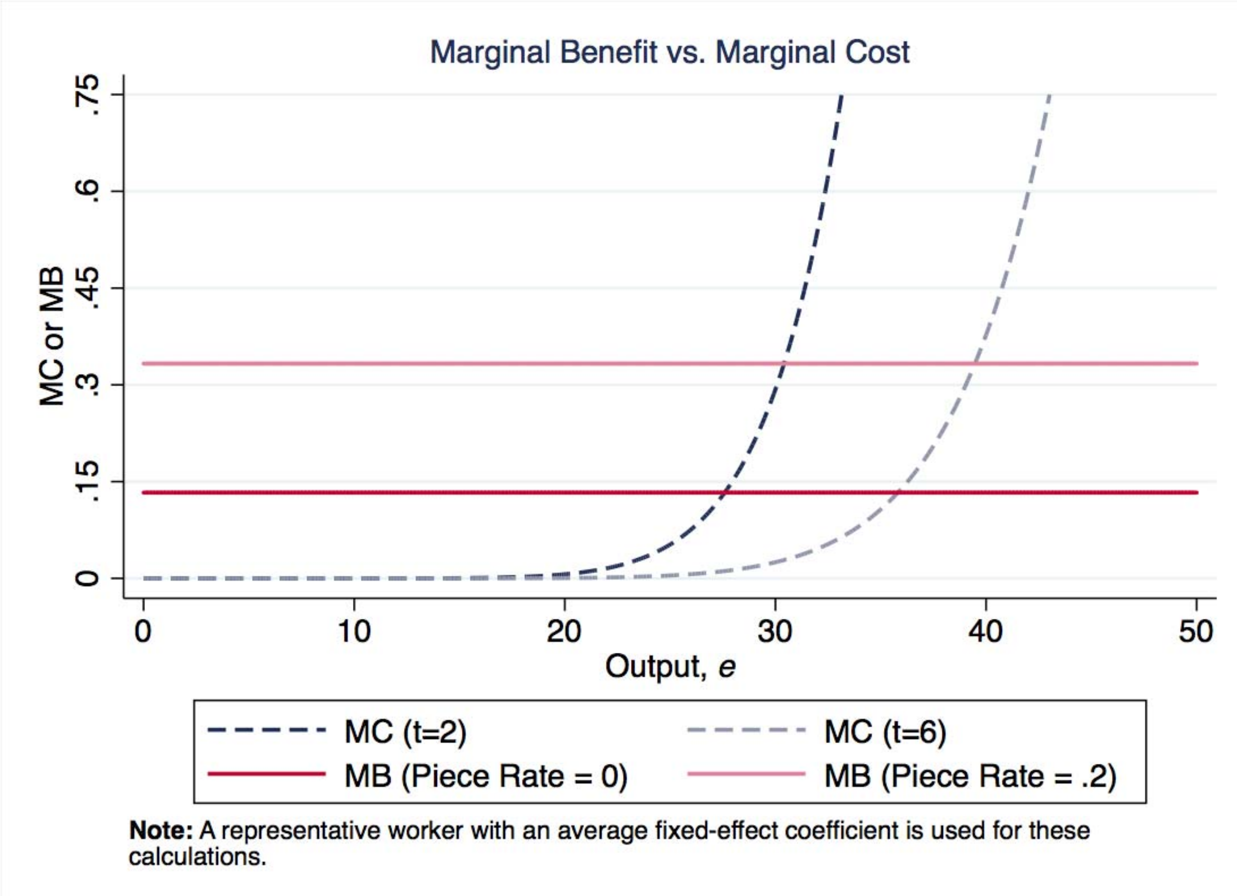

Notes: Figure 1 illustrates the equilibrium effort determined as the intersection of the relevant MC and MB curves. The plot is made for the parameters corresponding to the benchmark structural estimates in Table 2. We take the parameters from Column 2 in Table 2, taking the average warm glow across the three charities, assuming an individual with an average fixed effect k, and assuming the learning by doing as in periods 5-8. The MC curves plot the marginal cost of effort for different effort levels, in the specific case using the power cost of effort function, for a subject of average productivity and for an average realization of the error term. (Both subject fixed effects and the realization of the error term shift the marginal cost curve) The marginal cost of effort shifts out for later rounds (round 6 versus round 2 in the figure) to capture learning by doing (see Figure 3). The MB curve captures the marginal benefit for the subject which equals the sum of the private benefit (the piece rate) and the social preferences towards the employer. The figure plots two cases, piece rate of 0c (in which case the marginal benefit equals just the warm glow) and piece rate of 20c (in which case the marginal benefit equals the warm glow plus 20 cents).

Figure 2. Design: 10 Rounds of Experiment, Order A and B

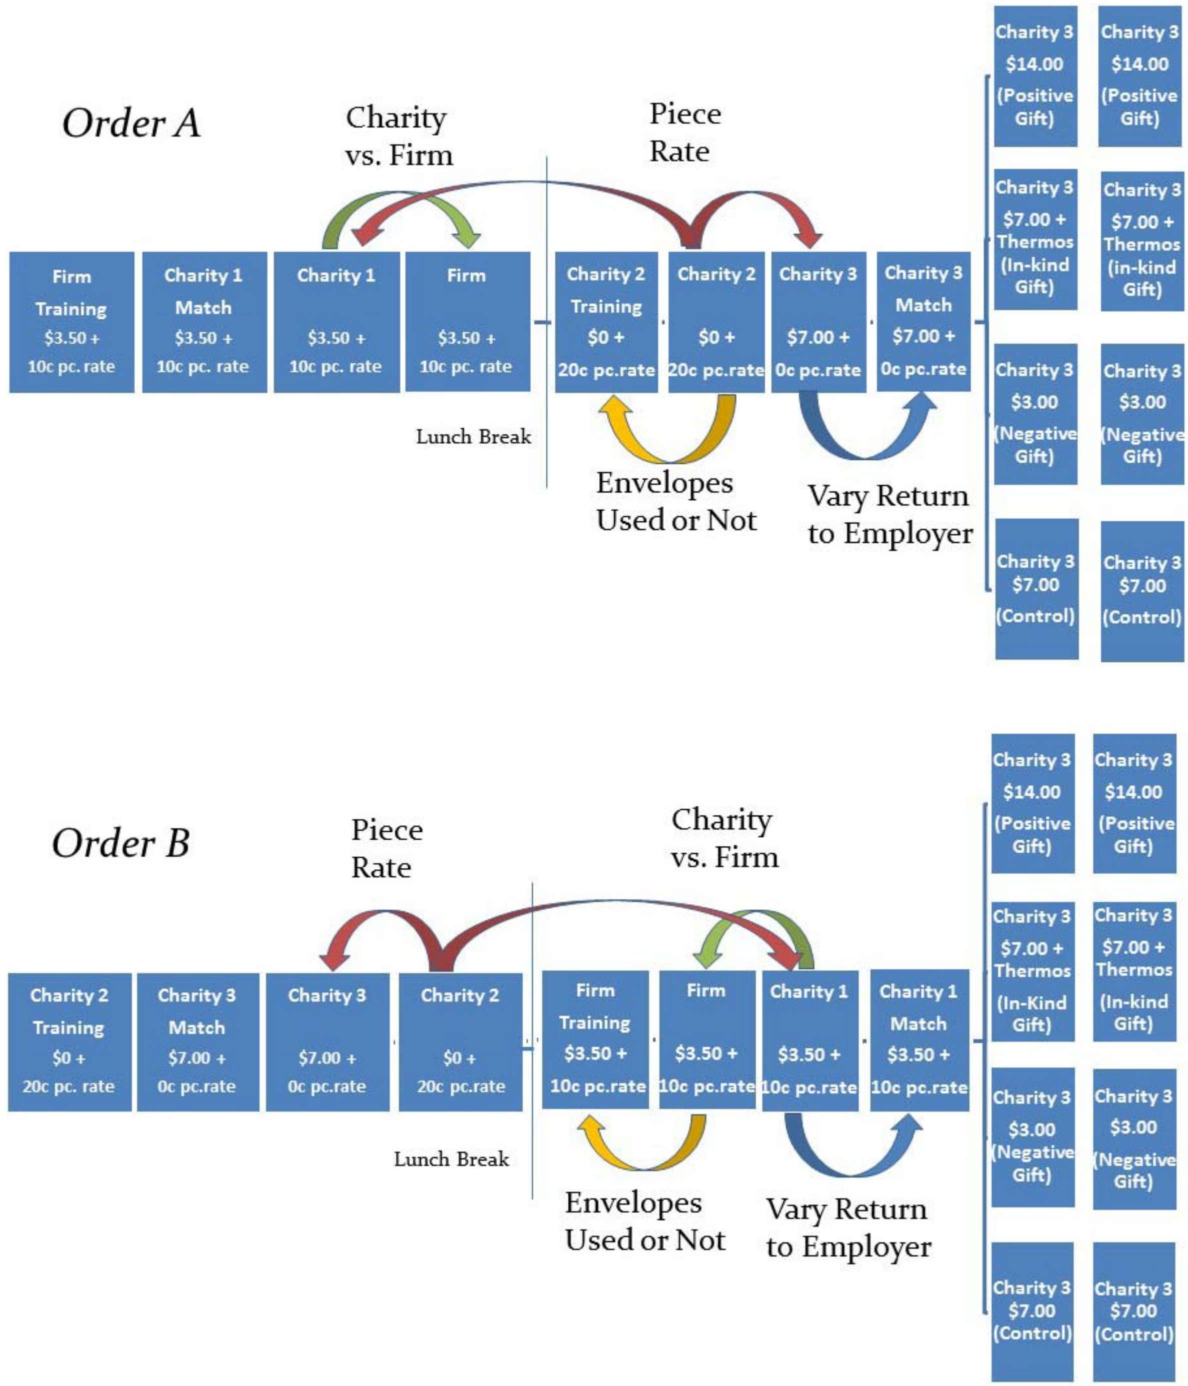

Notes: Figure 2 displays the sequence of the 10 experimental rounds of envelope preparation, each of which lasting 20 minutes. Between each round there is a 10 minute break, except between rounds 4 and 5 when there is a longer break for lunch. Subjects are randomized across sessions into Order A or Order B, as well as into three assignments of charities to be Charity 1, 2, and 3. In rounds 9 and 10, subjects are split within session into four gift exchange treatments (in the first 24 experimental sessions we did not run the in-kind gift treatment). Depending on randomized session assignment, either session 9 or session 10 involves a charity match (high return for the employer). The arrows indicate the main experimental comparisons evaluated in Figure 4a-c and in Appendix Figure 1a-b.

Figure 3. Average Effort over the 10 Rounds, by Order

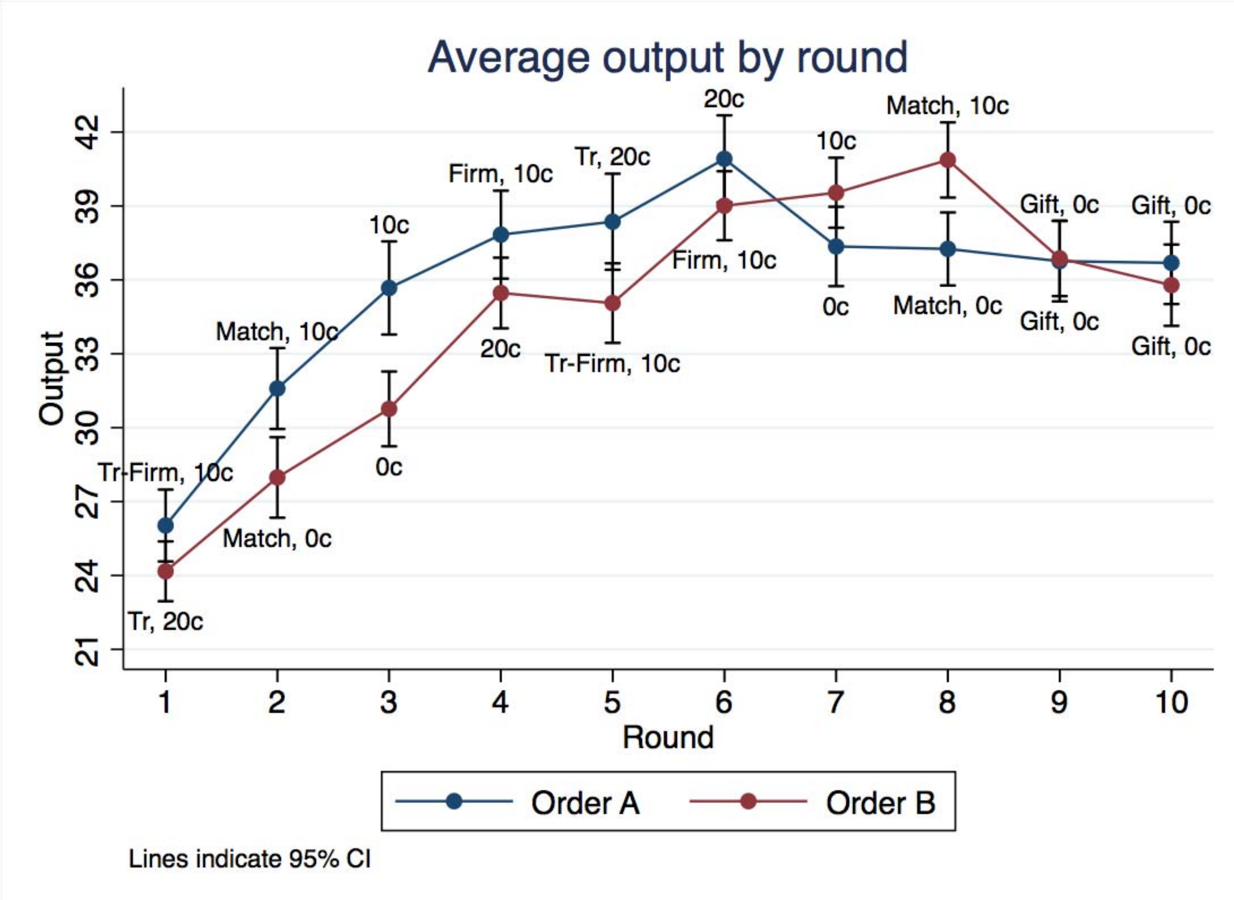

Notes: Figure 3 displays the average output (number of envelopes folded within a 20-minute round) in a round. The figure indicates 95% confidence intervals computed clustering by session, thus allowing for correlation of errors among subjects in a session. Subjects are randomized into Order A or Order B. See Figure 2 for more detailed labeling of the 10 rounds in each order. The output for rounds 9 and 10 averages across the gift treatments displayed in Figure 2. The figure indicates clear learning by doing in rounds 1-4, with average output about stable since, but highly responsive to variation in piece rate in and training (envelopes not used by the employer).

Figure 4. Key Experimental Results on Social Preferences (pre-gift)

Figure 4a. Variation in Piece Rate

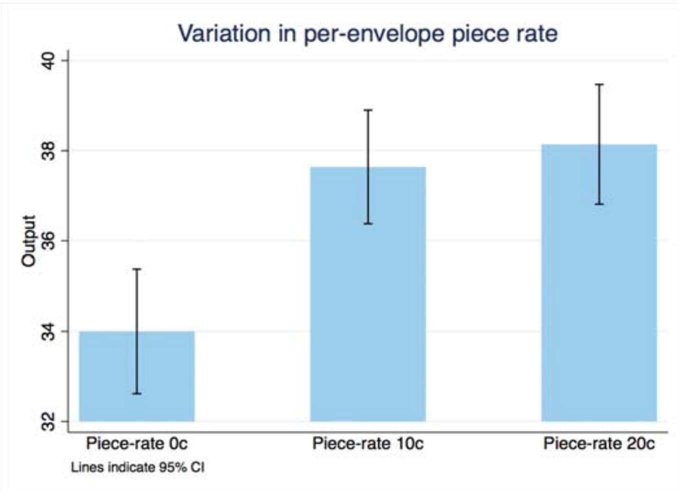

Figure 4b. Consequences to the Employer

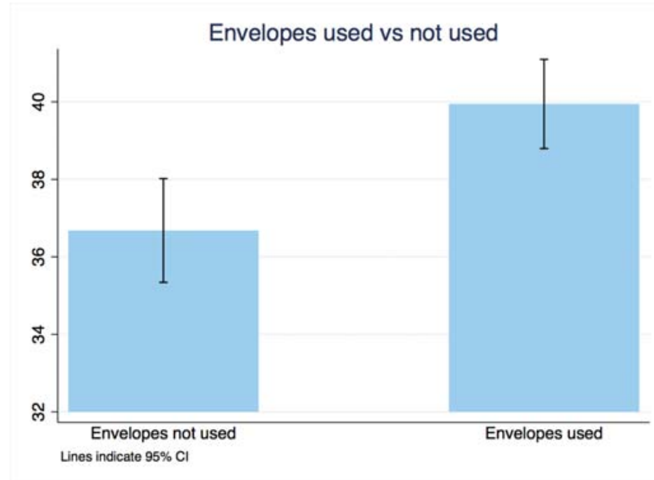

Figure 4c. Variation in Return to Employer (Match)

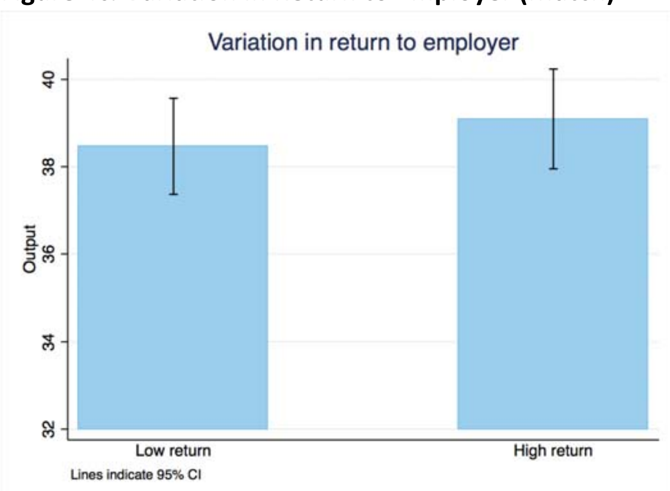

Notes: Figures 4a-c display key comparisons of average output (number of envelopes folded within a 20-minute round) across rounds, as outlined by the arrows in Figure 2. The comparisons average across order A and B. Figure 4a compares the piece rates for 0c, 10c, and 20c (respectively, rounds 7, 3, and 6 in Order A and rounds 3, 7, and 4 in order B). Figure 4b compares the impact of envelopes being used (rounds 5 and 6). Figure 4c compares the impact of high return to the employer (charity match) (rounds 7 and 8). The figures indicate 95% confidence intervals computed clustering by session.

Figure 5. Fit of warm Glow versus Altruism Model, Key Experimental Comparisons

Figure 5a. Variation in Piece Rate, Fit

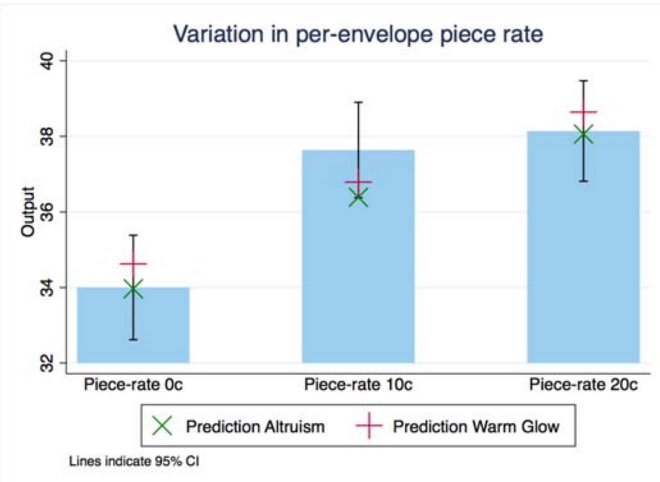

Figure 5b. Consequences to the Employer, Fit

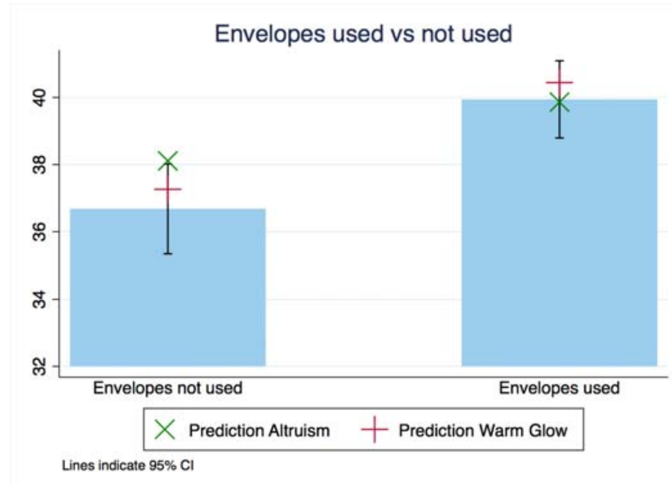

Figure 5c. Variation in Return to Employer (Match), Fit

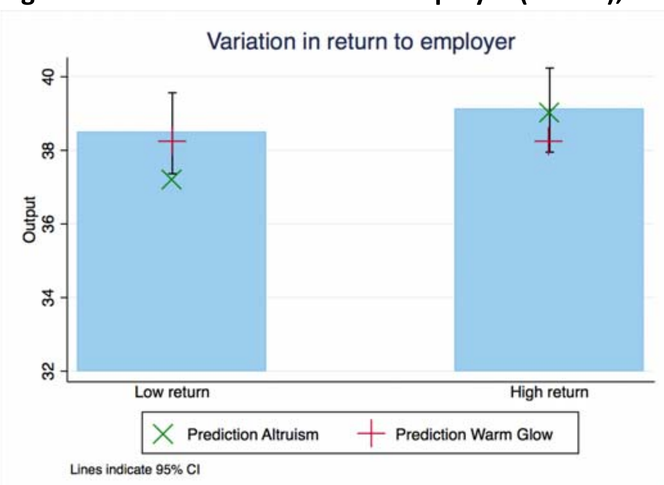

Notes: Figures 5a-c display key comparisons of average output (number of envelopes folded within a 20-minute round) across rounds outlined by the arrows in Figure 2 and summarized in Figure 4a-c. In addition to the evidence produced in Figures 4a-c, the figures also indicate the average prediction for the model estimated with altruism (Table 2, Column 4) or with warm glow (Table 2, Column 5). We use the specification with the exponential cost function since the output variable here is number of envelopes, as opposed to log number of envelopes.

Figure 6. Optimal Piece Rate for Estimated Social Preferences

Figure 6a. Effort as Function of Piece Rate

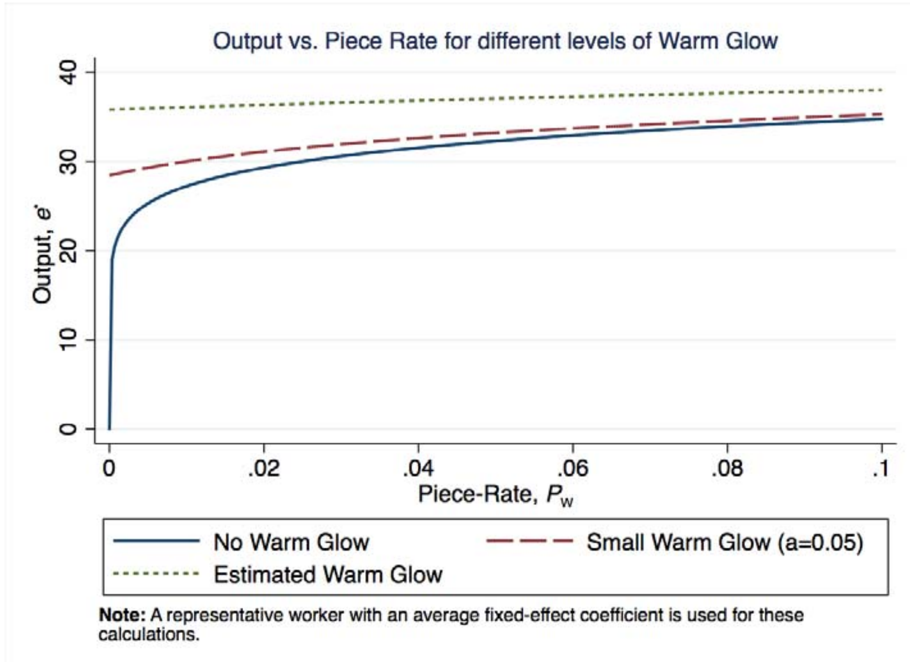

Figure 6b. Profit Rate as Function of Piece Rate

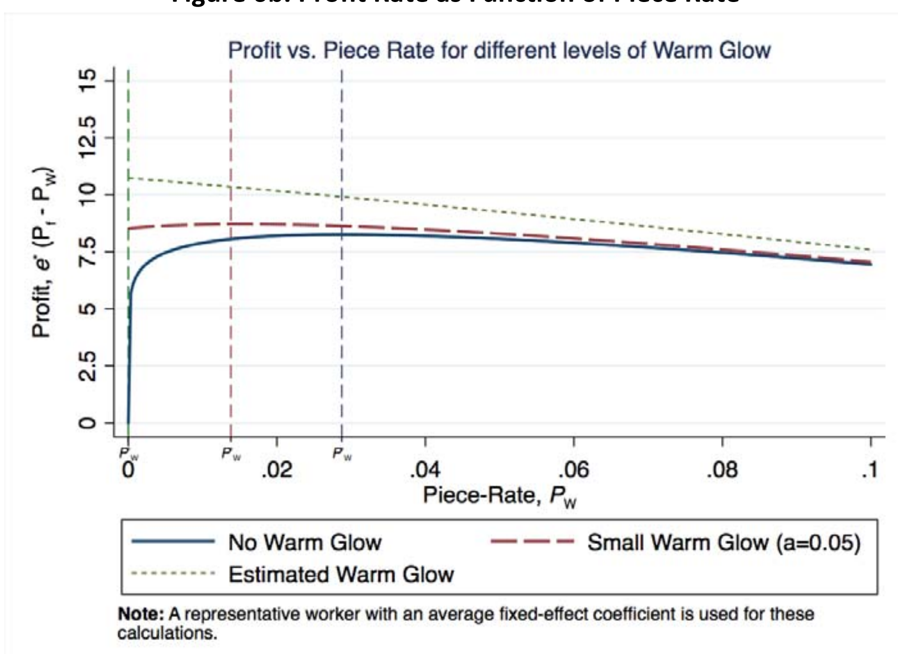

Notes: Figure 6 takes the estimated parameters in the warm-glow specification and predicts the implied effort *e** (Figure 6a) and profit rate e*(Pf-Pw) (Figure 6b), for different levels of the piece rate Pw. Specifically, the plots examine the impact on profits of increasing the piece rate holding constant all else (including the lump-sum pay). We take the parameters from Column 2 in Table 2, taking the average warm glow across the three charities, assuming an individual with an average fixed effect k, and assuming the learning by doing as in periods 5-8. The continuous blue line indicates the counterfactual for the case with no social preferences. In this case, effort and thus profits steeply increase initially with a higher piece rate as output steeply responds to incentives, then profits start declining as the extra productivity effect of piece rates is not worth any more the extra wage payment. The dotted green line indicates the curves for the estimated warm glow and shows that with social preferences the optimal piece rate is zero, as social preferences already motivate workers enough. The dashed red line shows that even if the warm glow were only about one tenth of the estimated one, holding all other parameters the same, the optimal piece rate would still be barely above zero.

Figure 7. Gift Exchange Effects

Figure 7a. Effect of Gift Treatments (No Controls)

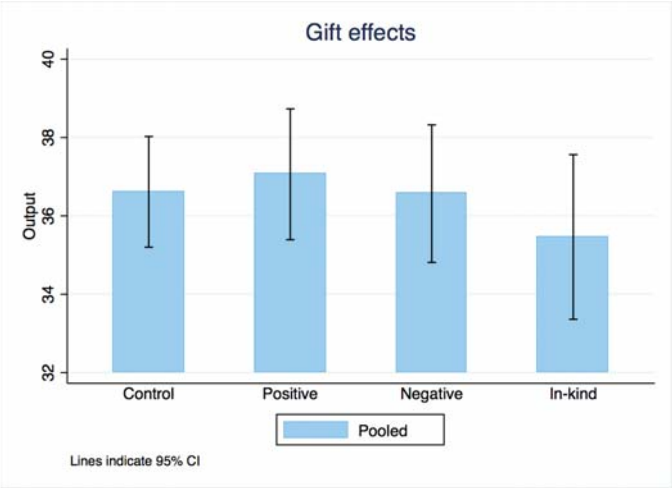

Figure 7b. Effect of Gift Treatments (With Controls)

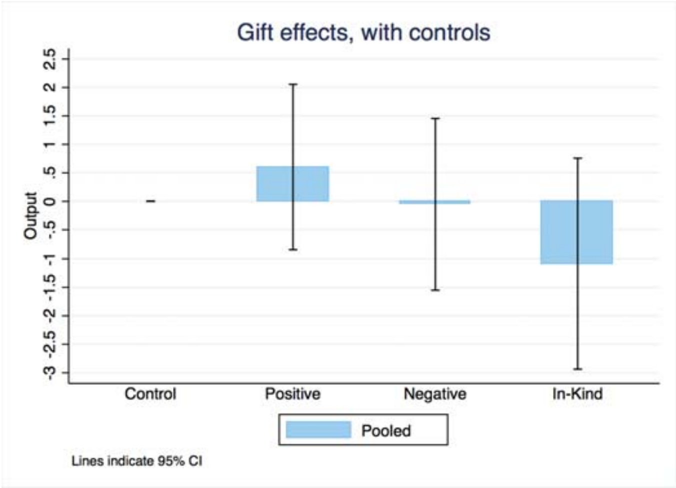

Figure 7c. Evidence on Decay of Gift Effects

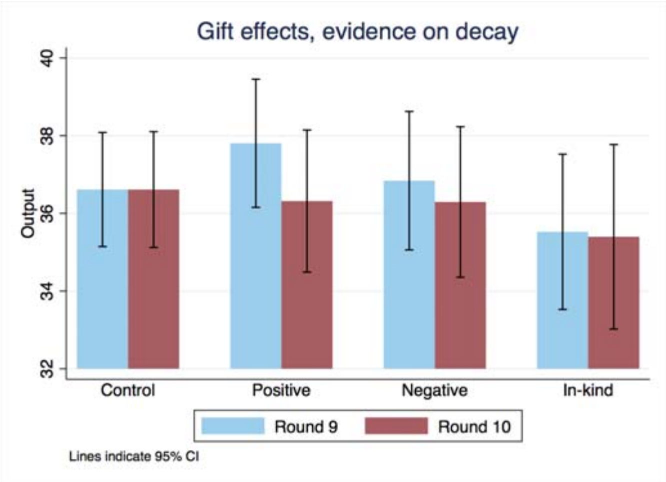

Figure 7d. Interaction with Return to Employer

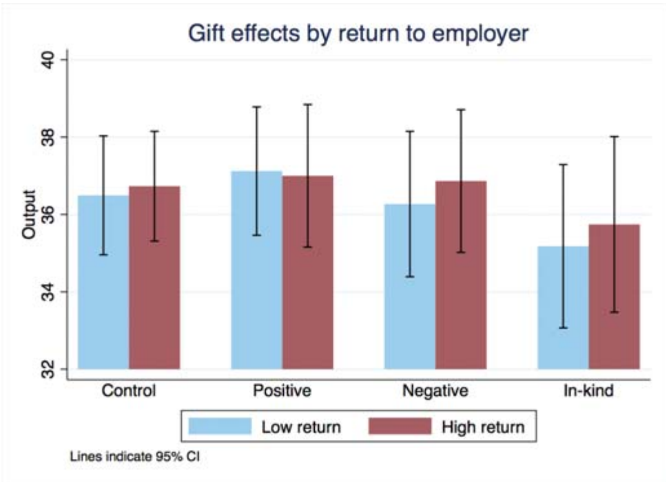

Notes: Figures 7a-d presents the key results for the gift exchange treatments in rounds 9 and 10 (see Figure 2). The Figures include 95% confidence intervals obtained after clustering for session. Figure 7a plots the average output (number of envelopes stuffed in 20 minutes) for the four treatments. Figure 7b presents the regression-adjusted coefficients after controlling for average productivity in rounds 5-8 (Table 4, Column 3). Figure 7c splits the treatment comparison of Figure 7a into round 9 and 10 to examine the impact of possible decay of gift effects. Figure 7d splits the results by return to the firm: in either round 9 or round 10 (depending on the randomization) the employer earns a higher return due to a charity match.

Figure 8. Heterogeneity in Gift Exchange Effects

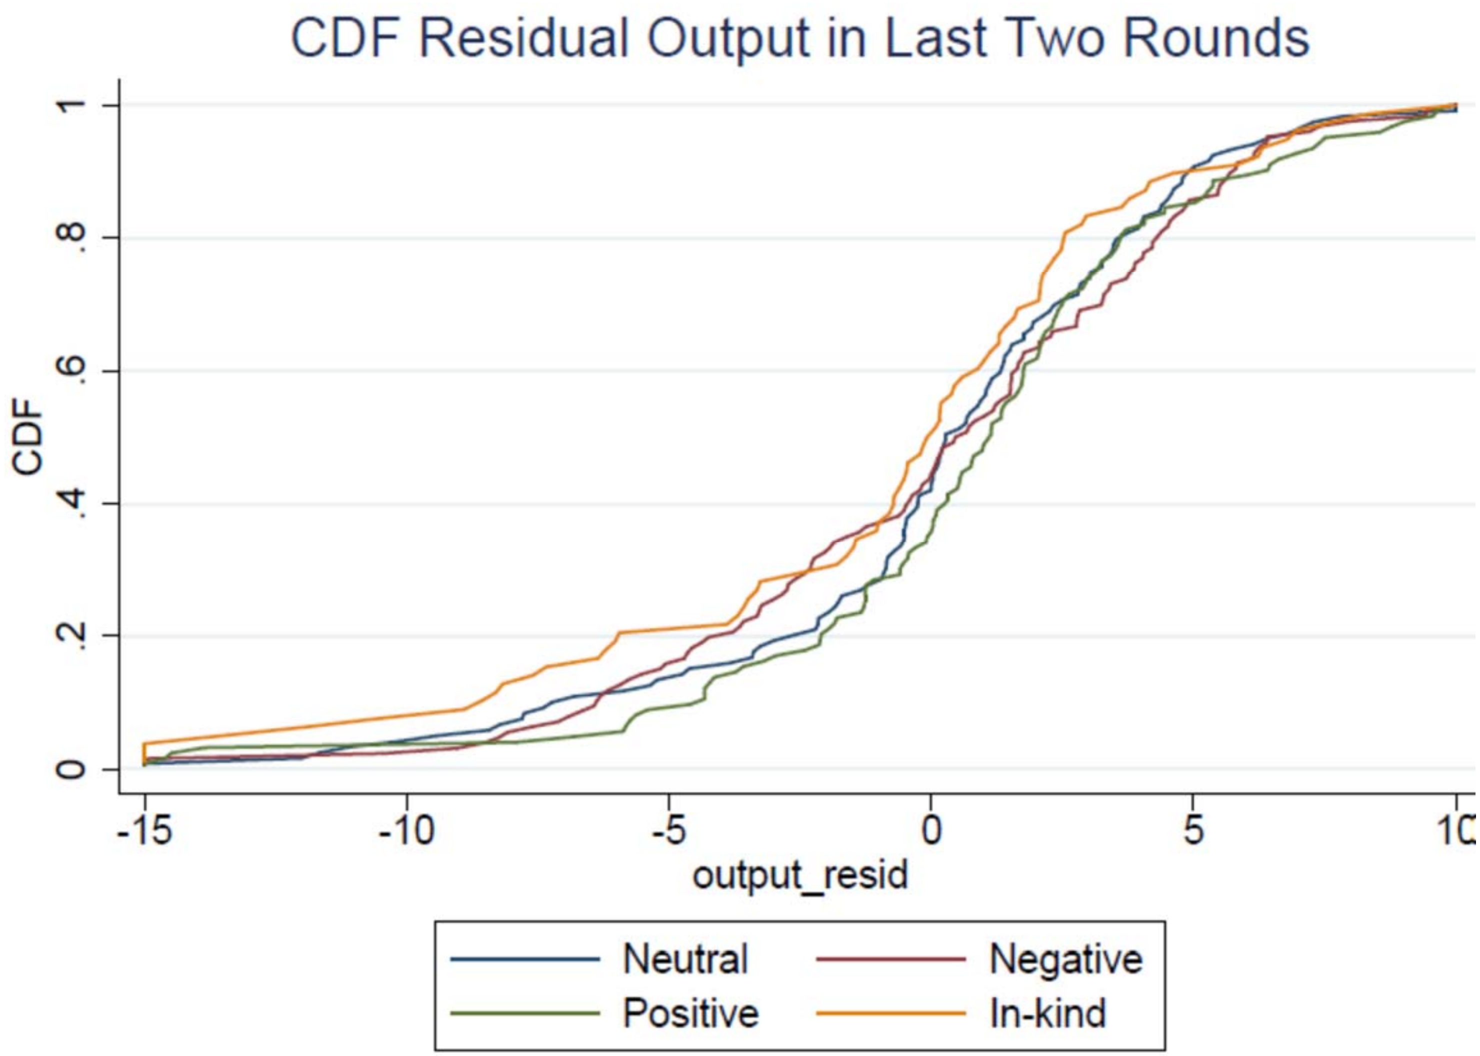

Notes: Figure 8 presents the cumulative distribution function of the effort measure for rounds 9 and 10, controlling for average productivity in rounds 5-8. More precisely, we regress productivity in rounds 9 and 10 on average productivity in rounds 5-8, take the residuals and average the two residuals for each worker, and plot them.

Figure 9. Fit of warm Glow versus Altruism Model, All 10 Rounds, Order A and B

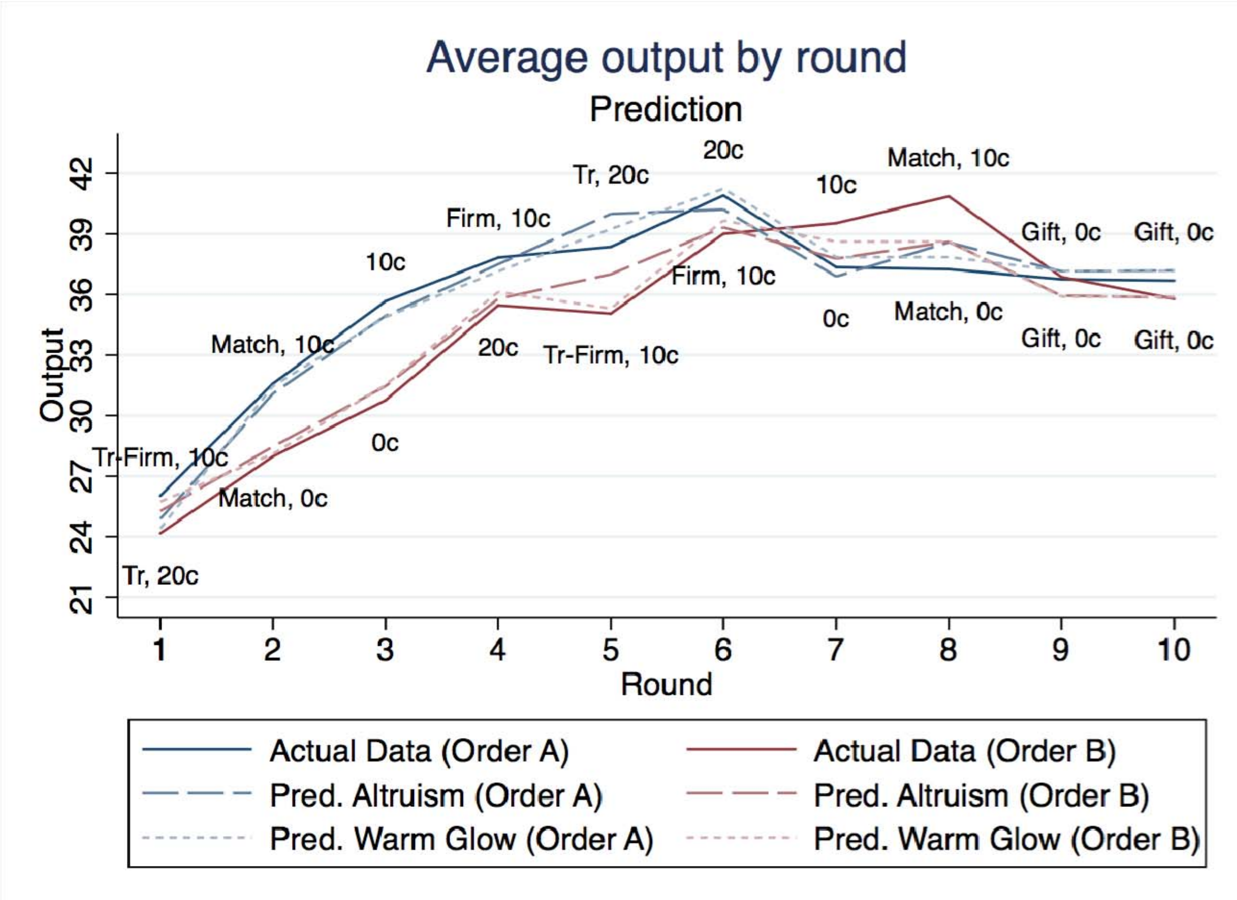

Notes: Figure 9 displays the average output (number of envelopes folded within a 20-minute round) in a round for Order A and Order B, together with the predicted output according to the warm glow model (Column 4 in Table 5) and according to the altruism model (Column 6 in Table 6). See Figure 2 for more detailed labeling of the 10 rounds in each order. The output for rounds 9 and 10 averages across the gift treatments displayed in Figure 2.

Figure 10. Effect of Gifts on Worker Happiness and Surprise

Figure 10a. Fraction Stating a Happy or Unhappy Reaction

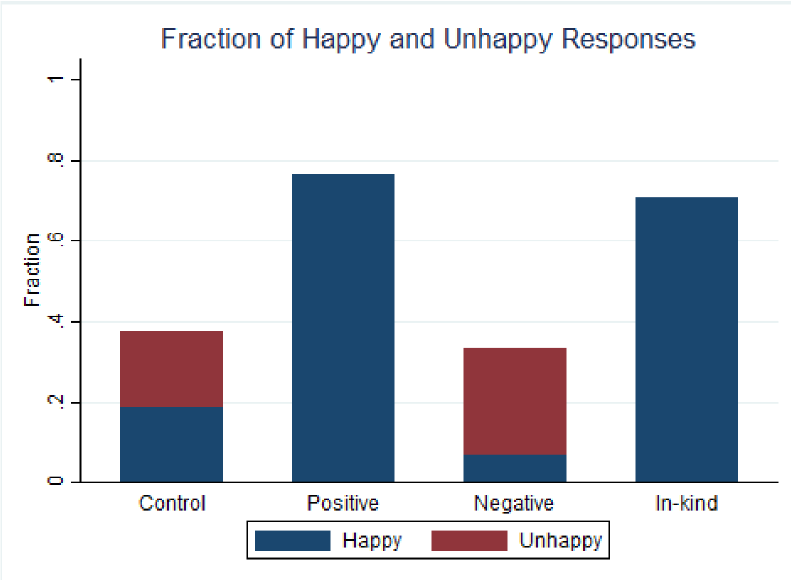

Figure10. Fraction Stating Positive or Negative Surprise

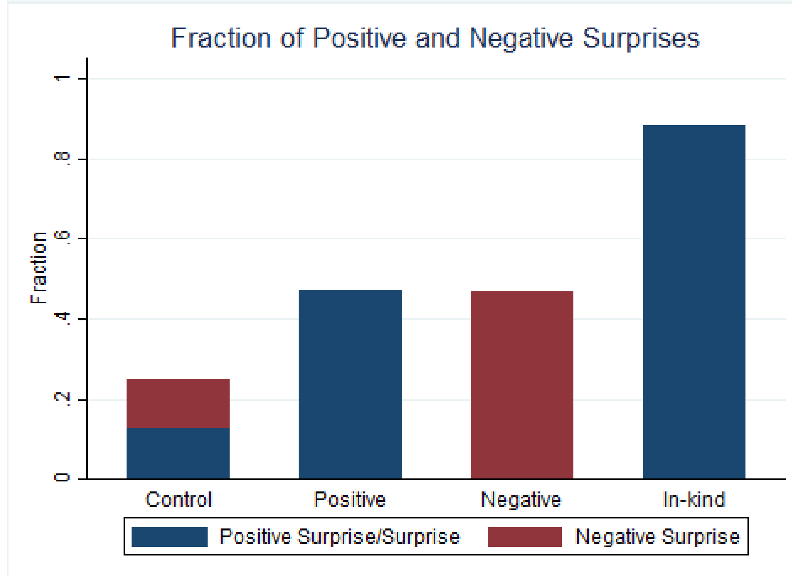

Notes: Figures 10a-b present the average response to a short debriefing questionnaire administered after the end of the experiment. The sample size includes 65 subjects, since the questions were only asked for the last 65 subjects in the experiments. Figure 10a presents the fraction that indicates being happy and the fraction that indicates being unhappy for each of the various treatments. Figure 10b indicates the fraction stating a positive surprise versus negative surprise (with the other categories being “as expected” or “none”). For the in-kind treatment, the bar shows the fraction that reported being surprised (we did not ask for the share with negative surprise).

Table 1. Summary Statistics and Covariate Balance

Specification:		OLS Regressions				
	Summary Statistics	Output Predictors	Checks of Randomization			
Dep. Var.:		Average Output	Indicator for Order A	Indicator for Positive Gift	Indicator for Negative Gift	Indicator for In-Kind Gift
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. Individual Demographics						
Is employed (self-reported)	0.397	2.022**	0.070	0.025	0.012	-0.045
	(0.490)	(0.799)	(0.050)	(0.047)	(0.044)	(0.039)
Female	0.522	2.535***	-0.099**	0.040	-0.010	0.011
	(0.500)	(0.691)	(0.049)	(0.043)	(0.044)	(0.036)
Age 25-34	0.361	2.959***	0.062	-0.110*	-0.005	0.063
	(0.481)	(0.875)	(0.058)	(0.060)	(0.052)	(0.048)
Age 35-44	0.191	0.989	0.039	-0.122*	-0.063	0.052
	(0.393)	(1.184)	(0.068)	(0.067)	(0.064)	(0.054)
Age 45-54	0.128	-2.122*	-0.035	-0.082	-0.160**	0.083
	(0.334)	(1.256)	(0.093)	(0.082)	(0.066)	(0.057)
Age 55+	0.058	1.305	0.243**	-0.002	-0.012	-0.021
	(0.235)	(1.753)	(0.100)	(0.111)	(0.120)	(0.069)
Has donated to charity	0.691	0.183	-0.131**	0.004	-0.000	0.026
(self-reported)	(0.463)	(0.946)	(0.059)	(0.057)	(0.048)	(0.043)
Has volunteered before	0.843	1.159	0.096	-0.043	0.042	0.056
(self-reported)	(0.364)	(1.051)	(0.065)	(0.060)	(0.070)	(0.056)
Mean of Dependent Variable		35.19	0.491	0.276	0.283	0.175
R squared		0.097	0.038	0.017	0.015	0.013
N	N = 446	N = 446	N = 446	N = 446	N = 446	N = 446
Panel B. Index of Demographics						
Predicted Effort Based on			0.004	-0.001	0.013	0.001
Demographics (Col. 2)			(0.010)	(0.008)	(0.008)	(0.007)
R squared			0.000	0.000	0.005	0.000
N			N = 446	N = 446	N = 446	N = 446

Notes: Column 1 in Panel A reports summary statistics on the sample of 446 participants in the experiment. Column 2 in Panel A reports the estimates of an OLS regression of average output (over the 10 rounds) on subject characteristics. Based on the estimate in Column 2 we form an index of predicted productivity based on demographics which we use in Panel B. In Columns 3-6 of Panels A and B we regress the assignment to different conditions (order A/B and assignment to the different gift treatments) on the subject characteristics (Panel A) and on the index of characteristics (Panel B). The standard errors are clustered at the session level.

* significant at 10%; ** significant at 5%; *** significant at 1%

Table 2. Estimation of Baseline Social Preferences, Non-Linear Least Squares

Dependent Variable:	Log (Number of Envelopes in a Round)			Number of Envelopes in a Round		
	(1)	(2)	(3)	(4)	(5)	(6)
Baseline Social Preferences:						
Altruism towards Charity B	0.195***		-0.046	0.218***		-0.010
	(0.051)		(0.049)	(0.046)		(0.044)
Altruism towards Charity RN	0.214***		-0.035	0.240***		-0.003
	(0.050)		(0.055)	(0.050)		(0.048)
Altruism towards Charity RIC	0.282***		0.068	0.302***		0.115
	(0.057)		(0.082)	(0.049)		(0.072)
Altruism towards Grocery Store	0.761***			0.736***		
	(0.091)			(0.079)		
Warm Glow towards Charity B		0.400***	0.463***		0.418***	0.427***
		(0.077)	(0.100)		(0.074)	(0.094)
Warm Glow towards Charity RN		0.427***	0.474***		0.449***	0.447***
		(0.073)	(0.085)		(0.079)	(0.082)
Warm Glow towards Charity RIC		0.505***	0.422***		0.521***	0.369***
		(0.080)	(0.108)		(0.074)	(0.096)
Warm Glow towards Grocery Store		0.720***	0.727***		0.715***	0.690***
		(0.073)	(0.081)		(0.074)	(0.075)
Incidental Parameters:						
Cost Function Curvature (s)	11.268***	9.476***	9.535***	0.296***	0.264***	0.258***
	(1.481)	(0.734)	(0.829)	(0.031)	(0.018)	(0.019)
Cost of Effort Function:	Power			Exponential		
Type of timetrend	Indicators for Rounds 2, 3, 4, 5-8					
Std. Deviation of Error Term	0.131	0.130	0.130	3.987	3.947	3.944
Std. Dev. of Individual f.e.s *(1/s)	0.248	0.248	0.248	8.134	8.148	8.160
R Squared	0.8350	0.8377	0.8379	0.8505	0.8535	0.8537
N	3568	3568	3568	3568	3568	3568

Notes: Specifications are from non-linear least squares regressions as in Section 4, with each observation being a worker-round combination. The sample is restricted to the first 8 rounds. The dependent variable is the log of the number of envelopes produced in that round in Columns 1-3 and is the number of envelopes produced in Columns 4-6. The specifications in Columns 1 and 4 allow for pure altruism towards the firm, in which the worker puts weight alpha on the return to the employer. The specifications in Columns 2 and 5 allow for a form of warm glow, that is, the worker puts a weight on the employer, but on the *average* return (30 cents per envelope), not the actual return (which varies by round). The specifications in Columns 3 and 6 include both altruism and warm glow coefficients, except for the grocery store for which there is no variation in return and thus one cannot separate altruism from warm glow. All specifications include fixed effects for worker i as well as indicators for rounds 2, 3, 4, and 5-8. The standard deviations listed are the standard deviation of the error term and the standard deviation of the individual fixed effects divided by the curvature s. The latter ratio indicates the variation in the individual productivity. The standard errors are clustered at the session level.
* significant at 10%; ** significant at 5%; *** significant at

Table 3. Estimation of Baseline Social Preferences, Robustness

Dependent Variable:	Log (Number of Envelopes)			Number of Envelopes in a Round		
	(1)	(2)	(3)	(4)	(5)	(6)
Baseline Social Preferences:						
Altruism towards Charity B	-0.029	-0.091**	0.082*	-0.027	-0.071**	0.090**
	(0.052)	(0.036)	(0.047)	(0.038)	(0.030)	(0.042)
Altruism towards Charity RN	-0.018	-0.107***	0.138**	-0.030	-0.089***	0.139***
	(0.057)	(0.032)	(0.055)	(0.038)	(0.030)	(0.051)
Altruism towards Charity RIC	0.100	-0.066	0.223***	0.086	-0.023	0.232***
	(0.103)	(0.073)	(0.060)	(0.082)	(0.068)	(0.053)
Warm Glow towards Charity B	0.397***	0.274***	0.883***	0.334***	0.259***	0.854***
	(0.087)	(0.083)	(0.108)	(0.071)	(0.067)	(0.106)
Warm Glow towards Charity RN	0.401***	0.299***	0.840***	0.355***	0.292***	0.821***
	(0.080)	(0.077)	(0.109)	(0.069)	(0.070)	(0.113)
Warm Glow towards Charity RIC	0.354***	0.314***	0.786***	0.293***	0.270***	0.752***
	(0.118)	(0.108)	(0.113)	(0.097)	(0.092)	(0.110)
Warm Glow towards Grocery Store	0.586***	0.650***	1.232***	0.541***	0.579***	1.176***
	(0.072)	(0.117)	(0.100)	(0.069)	(0.096)	(0.103)
Incidental Parameters:						
Cost Function Curvature (s)	10.866***	15.689***	3.654***	0.323***	0.413***	0.105***
	(0.896)	(1.984)	(0.245)	(0.025)	(0.044)	(0.006)
Cost of Effort Function:	Power			Exponential		
Type of timetrend	Quadratic in Rounds	Cubic in Rounds	Indicators for 2, 3, 4, 5-8	Quadratic in Rounds	Cubic in Rounds	Indicators for 2, 3, 4, 5-8
Specification	Benchmark		Partial Warm Glow During Training	Benchmark		Partial Warm Glow During Training
Std. Deviation of Error Term	0.130	0.129	0.129	3.930	3.900	3.908
Std. Dev. of Individual f.e.s *(1/s)	0.248	0.248	0.248	8.137	8.157	8.158
R Squared	0.8374	0.8411	0.8406	0.8548	0.8570	0.8564
N	3568	3568	3568	3568	3568	3568

Notes: Specifications are from non-linear least squares regressions as in specification in Section 4, with each observation being a worker-round combination. The sample is restricted to the first 8 rounds. The dependent variable is the log of the number of envelopes produced in that round in Columns 1-3 and is the number of envelopes produced in Columns 4-6. The specifications in Columns 1 and 4 allow for a quadratic function in the round number, while the specifications in Columns 2 and 5 allow for a cubic function in the round. The specifications in Columns 3 and 6 include indicators for rounds 2, 3, 4, and 5-8 and also assume that there is warm glow (but not altruism) even in the training rounds, assumed to be half the size as in the periods in which the envelopes are used. All specifications allow for both pure altruism towards the firm and a form of warm glow, that is, the worker puts a weight on the employer, but on the average return (30 cents per envelope), not the actual return (which varies by round). All specifications include fixed effects for worker i. The standard deviations listed are the standard deviation of the error term and the standard deviation of the individual fixed effects divided by the curvature s. The latter ratio indicates the variation in the individual productivity. The standard errors are clustered at the session level.

* significant at 10%; ** significant at 5%; *** significant at 1%

Table 4. Reduced-Form Effects of Gift Treatments

Specification:	OLS Regressions						
Dependent Variable:	Output in Rounds 9 and 10						
Panel A. Measure of Output:	Number of Envelopes Stuffed in 20 Minutes						
Sample:	Rounds 9 and 10			Round 9	Round 10	Match	No Match
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Positive (monetary) gift Treatment	0.448	0.903	0.603	1.350**	-0.145	0.428	0.778
	(0.966)	(0.737)	(0.729)	(0.636)	(0.904)	(0.801)	(0.771)
Negative (monetary) gift Treatment	-0.046	-0.014	-0.047	0.226	-0.321	0.133	-0.227
	(0.953)	(0.745)	(0.754)	(0.738)	(0.949)	(0.840)	(0.859)
Positive In-kind (Thermos) gift Treatment	-1.152	-1.011	-1.090	-1.024	-1.155	-0.924	-1.256
	(1.242)	(0.973)	(0.927)	(0.907)	(1.080)	(1.013)	(0.977)
Average Output Measure In Rounds 1-8		0.867***					
		(0.028)					
Average Output Measure In Rounds 5-8			0.815***	0.833***	0.797***	0.834***	0.796***
			(0.027)	(0.024)	(0.035)	(0.032)	(0.028)
Constant	36.613***	6.223***	5.149***	4.446***	5.852***	4.537***	5.761***
	(0.709)	(1.045)	(1.118)	(1.022)	(1.459)	(1.313)	(1.192)
R squared	0.003	0.585	0.608	0.668	0.556	0.622	0.595
N	N = 892	N = 892	N = 892	N = 446	N = 446	N = 446	N = 446
Panel B. Measure of Output:	Log of Number of Envelopes Stuffed in 20 Minutes						
Positive (monetary) gift Treatment	0.006	0.026	0.015	0.039**	-0.008	0.008	0.023
	(0.028)	(0.022)	(0.021)	(0.018)	(0.027)	(0.023)	(0.023)
Negative (monetary) gift Treatment	-0.017	-0.012	-0.018	-0.009	-0.027	-0.010	-0.026
	(0.035)	(0.031)	(0.031)	(0.031)	(0.035)	(0.032)	(0.035)
Positive In-kind (Thermos) gift Treatment	-0.040	-0.029	-0.033	-0.027	-0.039	-0.030	-0.036
	(0.037)	(0.029)	(0.028)	(0.026)	(0.034)	(0.031)	(0.029)
Average Output Measure In Rounds 1-8		0.782***					
		(0.028)					
Average Output Measure In Rounds 5-8			0.831***	0.851***	0.812***	0.843***	0.820***
			(0.030)	(0.029)	(0.036)	(0.031)	(0.034)
Constant	3.572***	0.826***	0.561***	0.490***	0.631***	0.523***	0.599***
	(0.020)	(0.099)	(0.110)	(0.108)	(0.130)	(0.112)	(0.125)
R squared	0.003	0.483	0.519	0.574	0.473	0.535	0.505
N	N = 892	N = 892	N = 892	N = 446	N = 446	N = 446	N = 446

Notes: Estimates from an OLS regression of output (Panel A) and log output (Panel B) in the final two rounds (Rounds 9 and 10) on the gift treatments. The omitted category is a Control treatment with no "gift" (pay is the same as previously experienced with the same charity). The standard errors are clustered at the session level.
* significant at 10%; ** significant at 5%; *** significant at 1%

Table 5. Estimation of Social Preferences with Gift Treatments, Non-Linear Least Squares

Dependent Variable:	Log (No. Envelopes in a Round)			Number of Envelopes in a Round		
	(1)	(2)	(3)	(4)	(5)	(6)
Baseline Social Preferences:						
Warm Glow towards Charity	0.443***	0.405***	0.343***	0.462***	0.337***	0.307***
	(0.063)	(0.043)	(0.043)	(0.065)	(0.032)	(0.035)
Warm Glow towards Grocery Store	0.720***	0.632***	0.539***	0.716***	0.551***	0.506***
	(0.072)	(0.064)	(0.062)	(0.073)	(0.058)	(0.062)
Reciprocal Social Preferences:						
Warm Glow Change -- Positive Monetary Gift	0.151	0.200*	0.086	0.135	0.098	0.053
	(0.128)	(0.114)	(0.089)	(0.121)	(0.085)	(0.075)
Warm Glow Change -- Negative Gift	-0.042	-0.016	-0.076	-0.001	-0.018	-0.047
	(0.123)	(0.125)	(0.093)	(0.095)	(0.072)	(0.061)
Warm Glow Change -- In-Kind Gift	-0.095	-0.074	-0.118	-0.106	-0.103	-0.118*
	(0.104)	(0.097)	(0.072)	(0.099)	(0.072)	(0.060)
Incidental Parameters:						
Cost Function Curvature (s)	9.440***	10.637***	11.366***	0.263***	0.316***	0.329***
	(0.737)	(0.835)	(0.894)	(0.018)	(0.024)	(0.025)
Cost of Effort Function:	Power			Exponential		
Type of timetrend	Indicators for Rounds	Quadratic in Rounds	Cubic in Rounds	Indicators for Rounds	Quadratic in Rounds	Cubic in Rounds
Std. Deviation of Error Term	0.144	0.144	0.144	4.318	4.308	4.302
Std. Dev. of Individual f.e.s *(1/s)	0.241	0.241	0.241	8.014	8.015	8.008
R Squared	0.7915	0.7908	0.7923	0.8184	0.8192	0.8197
N	4460	4460	4460	4460	4460	4460

Notes: Specifications are from non-linear least squares regressions as in specification in Section 4, with each observation being a worker-round combination. The sample includes all 10 rounds. The dependent variable is the log of the number of envelopes produced in that round in Columns 1-3 and is the number of envelopes produced in Columns 4-6. All specifications include fixed effects for worker i. Columns 1 and 4 also include indicators for rounds 2, 3, 4, 5-8, and 9-10. Columns 2 and 5 include a quadratic polynomial in the round. Columns 3 and 6 include a cubic polynomial in the round. The standard deviations listed are the standard deviation of the error term and the standard deviation of the individual fixed effects divided by the curvature s. The latter ratio indicates the variation in the individual productivity. The standard errors are clustered at the session level.

* significant at 10%; ** significant at 5%; *** significant at 1%

Table 6. Estimation of Social Preferences with Gift Treatments, Robustness

Dependent Variable:	Log (No. Envelopes in a Round)			Number of Envelopes in a Round		
	(1)	(2)	(3)	(4)	(5)	(6)
Baseline Social Preferences:						
Social Preferences towards Charity	0.457***	0.444***	na	0.447***	0.463***	0.187***
	(0.057)	(0.063)		(0.055)	(0.065)	(0.043)
Social Preferences towards Grocery Store	0.732***	0.720***	na	0.704***	0.716***	0.797***
	(0.068)	(0.072)		(0.065)	(0.073)	(0.108)
Reciprocal Social Preferences:						
Social Pref. Change -- Positive Monetary Gift	0.065	0.374**	na	0.041	0.314**	0.092
	(0.082)	(0.149)		(0.071)	(0.137)	(0.087)
Social Pref. Change -- Negative Gift	-0.099	0.032	na	-0.068	0.067	-0.001
	(0.096)	(0.135)		(0.067)	(0.100)	(0.060)
Social Pref. Change -- In-Kind Gift	-0.144*	-0.044	na	-0.152**	-0.079	-0.062
	(0.080)	(0.099)		(0.074)	(0.089)	(0.056)
Estimated Persistence of Social Preferences From Round 9 to 10		0.233	na		0.246	
		(0.251)			(0.248)	
Incidental Parameters:						
Cost Function Curvature (s)	9.039***	9.439***	na	0.257***	0.263***	0.410***
	(0.648)	(0.738)		(0.017)	(0.018)	(0.052)
Cost of Effort Function:	Power			Exponential		
Type of timetrend	Alternative Round Indicators	Standard Round Indicators (rounds 2, 3, 4, 5-8, 9-10)		Alternative Round Indicators	Standard Round Indicators (rounds 2, 3, 4, 5-8, 9-10)	
Specification	Benchmark (Warm Glow)	Estimated Decay of Gift Effect	Altruism (instead of warm glow)	Benchmark (Warm Glow)	Estimated Decay of Gift Effect	Altruism (instead of warm glow)
Std. Deviation of Error Term	0.144	0.144		4.321	4.315	4.365
Std. Dev. of Individual f.e.s *(1/s)	0.241	0.241		7.995	8.012	8.013
R Squared	0.7912	0.7918		0.8182	0.8187	0.8144
N	4460	4460		4460	4460	4460

Notes: Specifications are from non-linear least squares regressions as in specification in Section 4, with each observation being a worker-round combination. The sample includes all 10 rounds. The dependent variable is the log of the number of envelopes produced in that round in Columns 1-3 and is the number of envelopes produced in Columns 4-6. All specifications include fixed effects for worker i. Columns 1 and 4 also include indicators for rounds 2, 3, 4, 5-10. The estimated coefficient on round 2 is restricted to equal one half of the estimated coefficient in round 3. Columns 2 and 5 allow for a decay of the warm glow gift parameter in round 10, to equal delta*a_gift. Thus, delta=1 indicates no decay, delta=0 indicates full decay. The delta does not apply to round 9. Columns 3 and 6 estimate a model with pure altruism instead of warm glow. The model in Column 3 did not converge. The standard deviations listed are the standard deviation of the error term and the standard deviation of the individual fixed effects divided by the curvature s. The latter ratio indicates the variation in the individual productivity. The standard errors are clustered at the session level.

* significant at 10%; ** significant at 5%; *** significant at 1%

Table 7. Calibration of Implied Reciprocity Effects in Select Gift Exchange Papers

	Gift in Treament Condition	Task Assigned	% Effort Change With Gift	Implied Percent Altruism Change (Reciprocity) Due to Gift		
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. Findings from this paper:						
DellaVigna, List, Malmendier, Rao (2016)	Pay Increase from $7 to $14	Folding Charity Envelopes	2%	34%	.	.
	Pay Decrease from $7 to $3		-2%	-9%	.	.
	Gift of Thermos		-3%	-21%	.	.
Assumption about Cost Function:				Power Cost Function		
Estimated Curvature s				9.4(0.9)***		
Implied Elasticity				0.11		
Panel B. Selected Previous Findings on Gift Exchange in Field:						
Gneezy and List (2006) *Study 1*	Pay Increase from $12 to $20	Library Book Coding	27% (first 90 min)	846%	230%	61%
Gneezy and List (2006) *Study 2*	Pay Increase from $10 to $20	Door-to-door Fundraising	72% (first 3 hours)	16267%	1405%	196%
Kube, Marechal, and Puppe (2012) *Non-monetary gift condition*	Gift of Thermos	Library Book Coding	25%	715%	205%	56%
Kube, Marechal, and Puppe (2012) *Monetary gift condition*	7 Euro raise (from 36 euro pay)	Library Book Coding	5%	58%	28%	10%
Kube, Marechal, and Puppe (2013)	Pay cut from 15 to 10 euro/hr	Library Book Coding	-20%	-88%	-67%	-36%
Gilchrist, Luca, and Malhotra (forthcoming)	Pay increase from $3 to $4	Entering CAPTCHAs	18%	374%	129%	39%
Cohn, Fehr, and Goette (2014)	Pay increase from 22 to 27 ChF	Newspaper Distribution	3%	32%	16%	6%
Esteves-Sorenson (2015)	Pay Increase from $12 to $20	Enter data	2%	20%	10%	4%
Assumptions about Cost Function:				Power Cost Function		
Assumed Curvature s				9.4	5.0	2.0
Implied Elasticity				0.11	0.20	0.50

Notes: Table 7 revisits some of the findings in the previous gift exchange experiments in the field, with summary of the key gift treatments and findings in Columns 1-3. Panel A summarizes the effects from this paper: Column 2 reports the findings from Table 4, Column 3, Panel B (on log output). Column 3 reports the results from Table 5, Column 1, taking the ratio of the estimated warm glow change to baseline warm glow. For example, for the positive monetary gift .151/.443=34%. In Pnel B we revisit some classic experiments on gift exchange in the field. In Columns 4-6 we compute the implied percent increase in altruism or warm glow implied by the effort increase (or decrease), for a calibrated value of the elasticity of effort. The calibration holds for a power cost of effort function, which is characterized by constant elasticity. Column 4 uses the elasticity estimated for our task (Table 5, Column 1). Columns 5 and 6 report the results assuming higher elasticities.

Appendix Figures 1a-b. Additional Experimental Findings

Appendix Figure 1a. Effort Provided For Three Different Charities

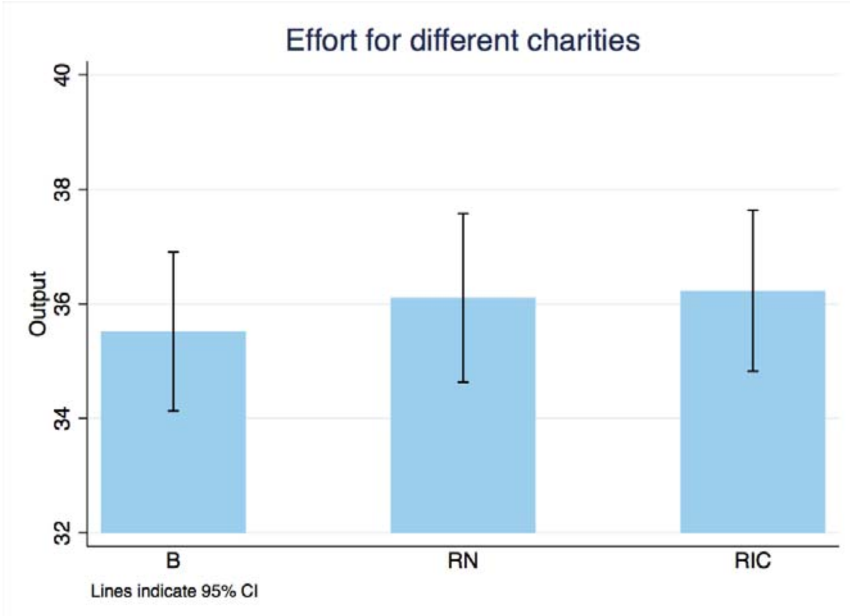

Appendix Figure 1b. Charity Employer versus Grocery Store Employer

Notes: Appendix Figure 1a-b displays additional experimental results on average output (number of envelopes folded within a 20-minute round). Appendix Figure 1a compares productivity across the three different charities used in the experiment. The charities are randomized in a rotating way to take the role of Charity 1, 2, and 3. The comparison uses output in all rounds except for the training rounds. Appendix Figure 1b compares output when producing for a charity versus for a firm (a grocery store) holding constant the piece rate at 10 cents and holding constant the perceived return to the employer at 30 cents per envelope. The rounds compared are outlined in Figure 1. The figures indicate 95% confidence intervals computed clustering by session.

Appendix Figure 2. Average Effort over the 10 Rounds, by Order, Log Output

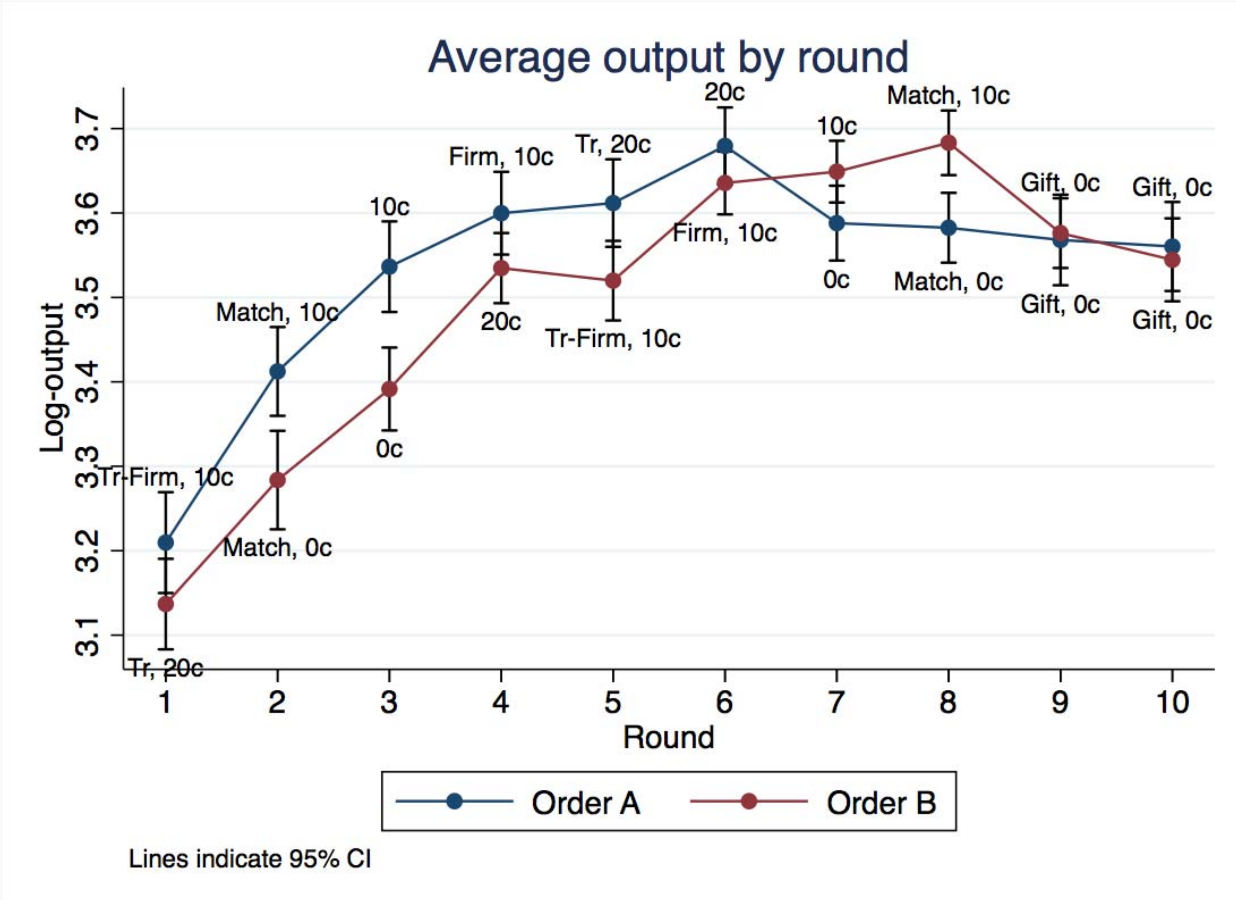

Notes: Appendix Figure 2 displays the mean of the output measure across subjects, where the output measure is the log of the number of envelopes folded within a 20-minute round. The figure indicates 95% confidence intervals computed clustering by session, thus allowing for correlation of errors among subjects in a session. Subjects are randomized into Order A or Order B. See Figure 2 for more detailed labeling of the 10 rounds in each order. The output for rounds 9 and 10 averages across the gift treatments displayed in Figure 2. The figure indicates clear learning by doing in rounds 1-4, with average output about stable since, but highly responsive to variation in piece rate in and training (envelopes not used by the employer).

Appendix Figure 3. Key Experimental Results on Social Preferences (pre-gift), Log Output

Appendix Figure 3a. Variation in Piece Rate

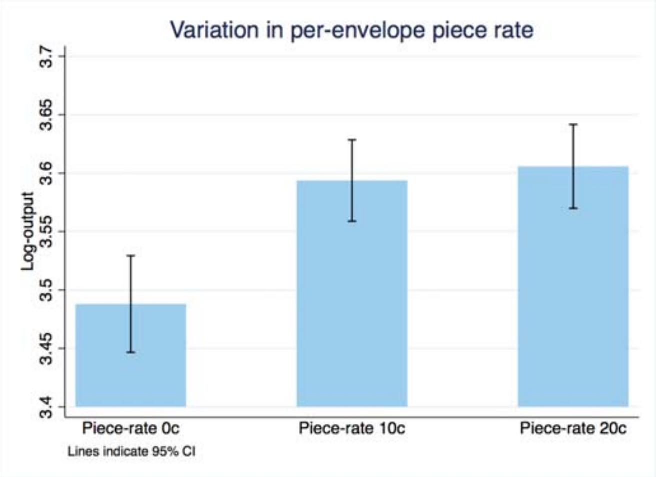

Appendix Figure 3b. Consequences to the Employer

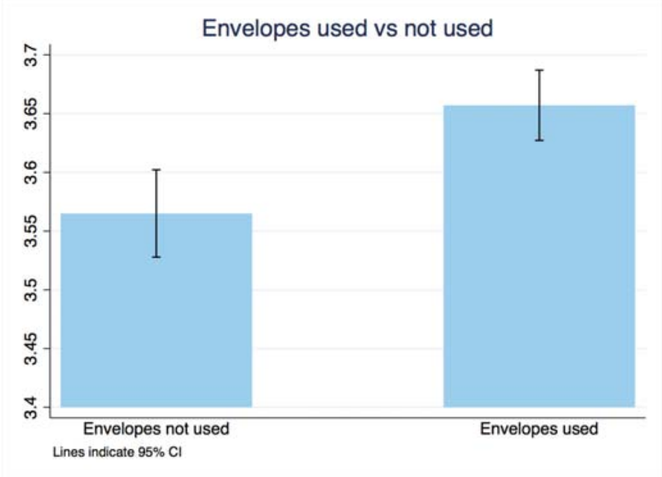

Appendix Figure 3c. Variation in Return to Employer (Match)

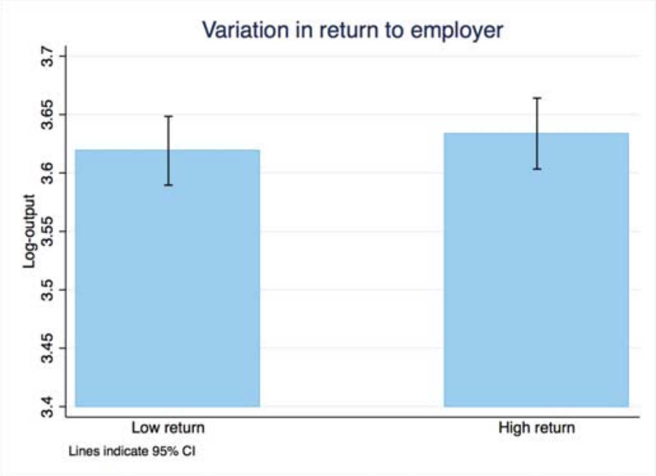

Notes: Appendix Figures 3a-c display key comparisons of average low output (log of number of envelopes folded within a 20-minute round) across rounds outlined by the arrows in Figure 2. The comparisons average across order A and B. Appendix Figure 3a compares the piece rates for 0c, 10c, and 20c. Appendix Figure 3b compares the impact of envelopes being used (rounds 5 and 6). Appendix Figure 3c compares the impact of high return to the employer (charity match) (rounds 7 and 8). The figures indicate 95% confidence intervals computed clustering by session.

Appendix Figure 4. Gift Exchange Effects, Log Output

Appendix Figure 4a. Effect of Gift Treatments (No Controls)

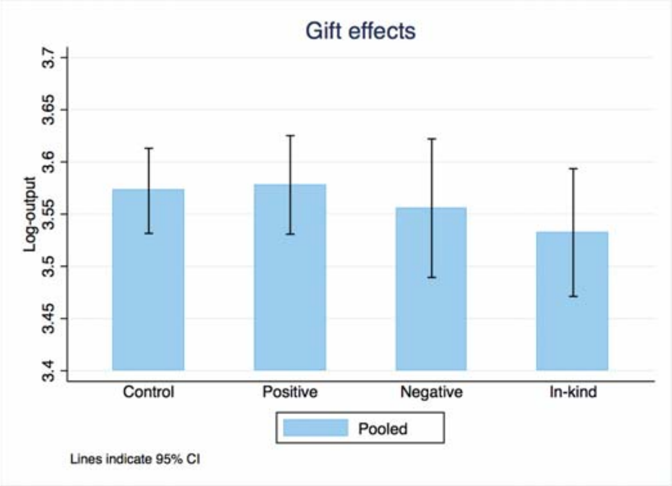

Appendix Figure 4b. Effect of Gift Treatments (With Controls)

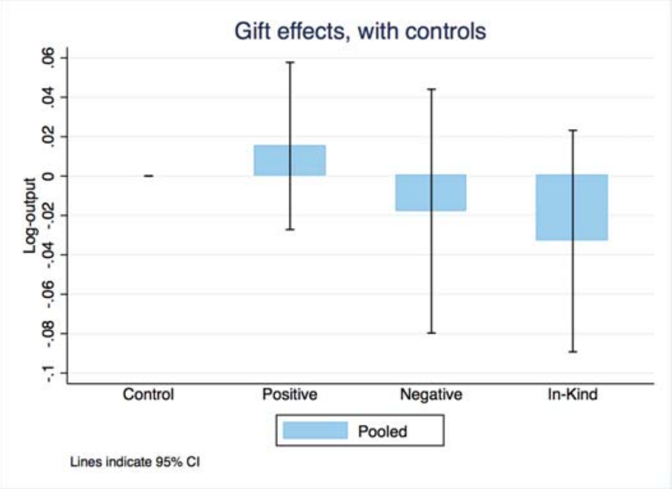

Appendix Figure 4c. Evidence on Decay of Gift Effects

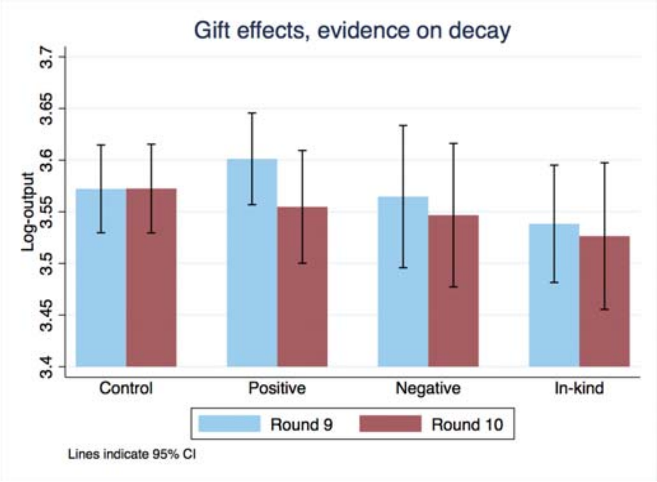

Appendix Figure 4d. Interaction with Return to Employer

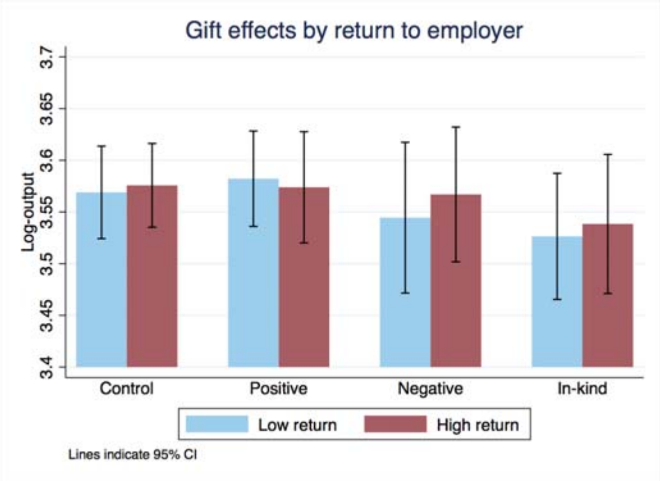

Notes: Appendix Figures 4a-d presents the key results for the gift exchange treatments in rounds 9 and 10 (see Figure 1) using log output in each round as the key variable. The Figures include 95% confidence intervals obtained after clustering for session. Appendix Figure 4a plots the average output (number of envelopes stuffed in 20 minutes) for the four treatments. Appendix Figure 4b presents the regression-adjusted coefficients after controlling for average productivity in rounds 5-8 (Table 4, Column 3). Appendix Figure 4c splits the treatment comparison of Appendix Figure 4a into round 9 and 10 to examine the impact of possible decay of gift effects. Appendix Figure 4d splits the results by return to the firm: in either round 9 or round 10 (depending on the randomization) the employer earns a higher return due to a charity match.

Appendix Figure 5. Estimated Productivity Effects, Different Models

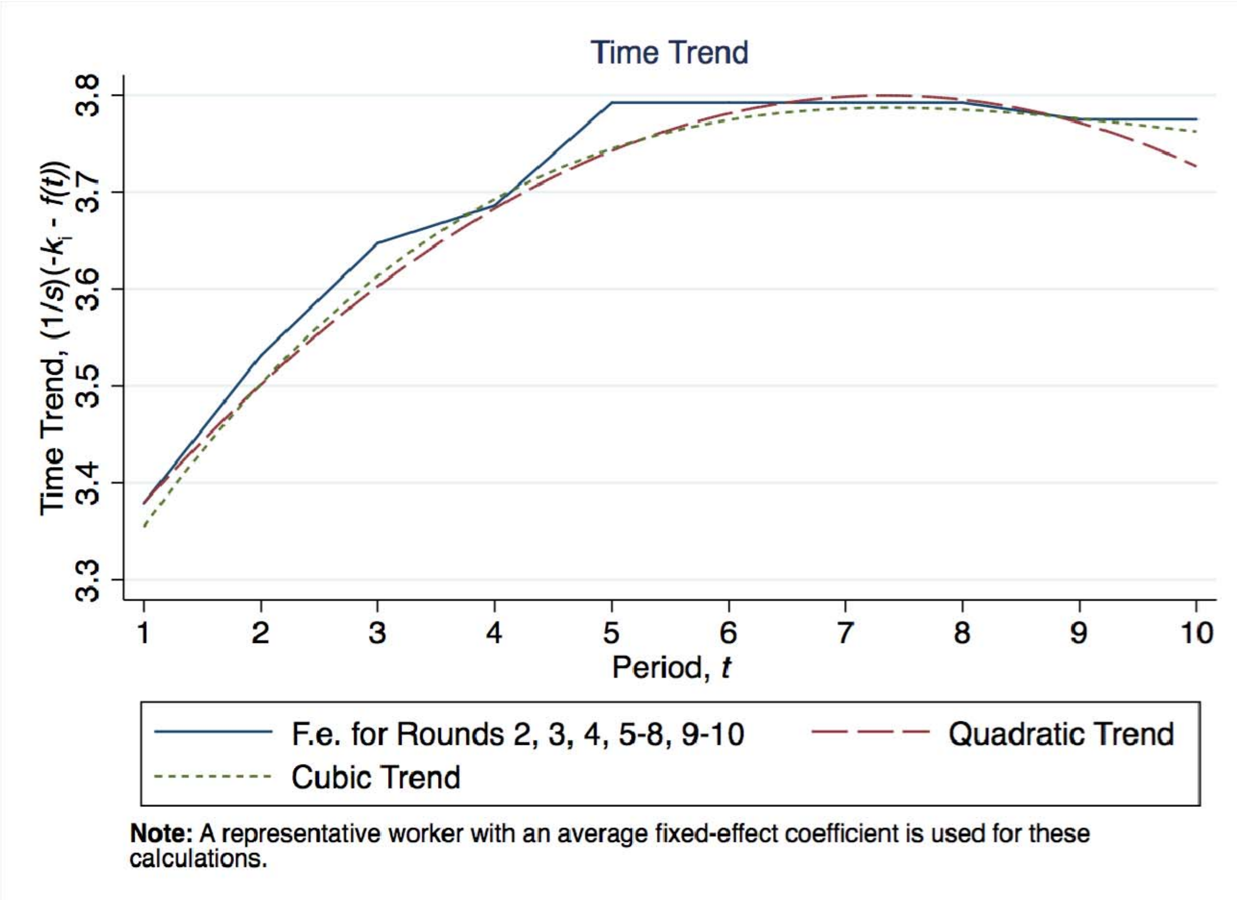

Notes: Appendix Figure 5 plots the estimated *(1/s)(-k-f(t))* function, that is, how the cost of effort function is estimated to change over time for an individual with representative *k*. The estimated coefficients are from specifications in Table 5, respectively Column 1 (indicators for rounds), Column 2 (quadratic polynomial), and Column 3 (cubic polynomial).