Ceding control: an experimental analysis of participatory management

Philip Mellizo, Jeffrey Carpenter & Peter Hans Matthews

Journal of the Economic Science Association
A Companion Journal to Experimental Economics

ISSN 2199-6776
Volume 3
Number 1

J Econ Sci Assoc (2017) 3:62-74
DOI 10.1007/s40881-017-0034-1

Springer

J Econ Sci Assoc (2017) 3:62–74
DOI 10.1007/s40881-017-0034-1

ORIGINAL PAPER

Ceding control: an experimental analysis of participatory management

Philip Mellizo[1] · **Jeffrey Carpenter**[2] · **Peter Hans Matthews**[3]

Received: 1 September 2016 / Revised: 16 January 2017 / Accepted: 17 February 2017 /
Published online: 10 March 2017

Abstract We use an experiment to evaluate the effects of participatory management on firm performance. Participants are randomly assigned roles as managers or workers in firms that generate output via real effort. To identify the causal effect of participation on effort, workers are exogenously assigned to one of the two treatments: one in which the manager implements a compensation scheme unilaterally or another in which the manager cedes control over compensation to the workers who vote to implement a scheme. We find that output is between seven and twelve percentage points higher in participatory firms.

Keywords Voice · Control · Intrinsic motivation · Participatory management · Real effort · Experiment

JEL Classification C92 · J33 · J53 · J54 · M50

Electronic supplementary material The online version of this article (doi:10.1007/s40881-017-0034-1) contains supplementary material, which is available to authorized users.

✉ Jeffrey Carpenter
jpc@middlebury.edu

Philip Mellizo
pmellizo@wooster.edu

Peter Hans Matthews
pmatthew@middlebury.edu

[1] Department of Economics, College of Wooster, Wooster, OH, USA

[2] Departments of Economics, Middlebury College, IZA, Middlebury, VT, USA

[3] Department of Economics, Middlebury College, Middlebury, VT, USA

1 Introduction

Many "extraordinary claims have been made about [the effects of] employee involvement" (Lawler et al. 1995) on worker satisfaction and overall firm performance. Despite the strong intuition that employee participation in decision making is a "win-win" for employees and firms, the empirical record is mixed (e.g., Cappelli and Neumark 2001). Certainly, some of this owes to the list of well-known difficulties that accompany estimation: the inability to suitably control for unobservable heterogeneity, endogeneity, self-selection, or, more fundamentally, the availability of reliable and appropriate data. It may also be true that differences in sample characteristics, research designs, practices, and/or performance metrics have additionally contributed to mixed reported findings.

In this paper, we use an experiment to mitigate some of the hurdles that hinder the identification of any causal effect of participation on performance. In particular, unlike the field where important aspects of production and motivation (like participation) are often determined endogenously, experiments allow one to implement exogenous, ceteris paribus changes. In our case, we compare groups of workers that were exogenously allowed (or not allowed) by a manager to participate in the running of the firm by having their vote to determine the group's compensation policy. Because we collected compensation preferences before anyone learned the details of the experiment, they too are exogenous and allow us to control for worker selection. In this setting, we find that the average treatment effect of ceding authority to workers on effort in a real effort task is large (between seven and twelve percentage points), statistically significant and robust.

We allow managers to choose to either implement a compensation scheme unilaterally or cede the right to choose the workers in an environment in which nothing is known of the worker characteristics and managers (along with their preferences) are randomly assigned to firms. Interestingly, in this setting we find that managers are reluctant to cede decision-making authority despite the possibility that it might be beneficial to do so. This finding replicates a result that is now common in the related literature (e.g., Fehr et al. 2013; Bartling et al. 2014).

Considering our main contribution, we find a large and significant effect of participation, one that echoes the recent results on institutional choice and democracy found in Dal Bó et al. (2010). Like Dal Bó et al., we find that democratic participation affects motivation; however, while our experimental manipulation to account for possible selection effects is similar, it is also a bit simpler. Dal Bó et al. allow participants to vote on whether or not to modify the payoffs of a social dilemma game but a computer then decides whether or not to heed the results of the vote. When the computer ignores the vote, it decides to modify the game randomly. In our simplification, the context is also a bit more natural in that a human manager takes the role of the computer and decides to cede control to the workers before knowing the outcome of the vote.

Reflecting on our previous work, the difference in output between workers in participatory firms and those in no-voice, traditional firms is larger, but in the same "ballpark" as a related estimate described in Mellizo et al. (2014). This original

study estimated an effect of voting on effort to be between 7 and 9% points, at most while our current upper bound estimate is close to 12% points. However, there are substantial differences between the current experiment and the previous one, including one that might account for the larger effect. In addition to new design elements that allow us to more convincingly estimate a causal relationship, the elicitation of *ex ante* preferences to control for selection effects, changes in the compensation schemes available (a piece rate instead of revenue sharing) and a larger sample, we added an aspect of relational contracting (Macneil 1985) to the current experiment, which might partially explain why the current effect is more pronounced. In the current experiment, managers could either trust workers to pick a compensation scheme that would benefit everyone or not trust them and pick the scheme themselves. The manager's confidence in the workers' vote could very well interact positively with standard intrinsic motivational effects resulting from just allowing the workers more autonomy (Falk and Kosfeld 2006; Charness et al. 2012). In our previous experiment, things were simple but less realistic: there were no bosses and compensation schemes were either imposed randomly or via the worker's vote—an environment in which only the intrinsic motivation channel was likely to affect effort.

A more extensive review of the literature on participatory management and worker voice can be found in the online appendix that accompanies this paper, along with the experimental instructions and various robustness tests. What follows is a description of our experiment and a detailed analysis of our main results.

2 Experimental design and procedures

To provide a fresh estimate of the causal effect of participation on motivation, we conducted five one hour sessions with a total of 320 experimental subjects who earned an average of 15.25 € (standard deviation of 7.2 €) including a 5 € show-up fee. After arriving, participants were seated randomly at computer terminals where they found scratch paper, a pencil, a sheet of paper with a number place (Sudoku) logic puzzle, and instructions for an initial 2-min practice period during which all subjects familiarized themselves with the work task that would be used in the experiment.

In the first stage of the experiment, a 2-min practice period, subjects added different sets of five two-digit numbers that appeared on their computer screen. We did not allow the use of calculators but subjects could use the scratch paper and pencil we provided them. The numbers to be added in each problem were randomly generated. We used this work task since it yields low intrinsic reward, requires little skill for a college student, and because previous work has found that it does not result in biased performance in any systematic manner (Niederle and Vesterlund 2007). When the 2-min practice period ended, subjects saw a summary screen that indicated to them how many problems they correctly solved before proceeding to the next stage in the experiment. To prevent strategic voting over compensation schemes as much as possible, participants only saw their own output in the practice period, not the distribution of output for their group or the session.

At the beginning of the second stage of the experiment, subjects were informed that they had been randomly and anonymously grouped with three other subjects to form a firm connected through the computer network. Subjects then learned that they would be assigned to one of the two roles with different responsibilities—either a manager or a worker—though at this point in the study they were not told their roles. In the next phase of the instructions, subjects learned that all the firm's earnings were tied directly to the number of correct answers by workers to simple math problems presented in exactly the same (random) way as in the practice period. Specifically, all subjects learned that each correct answer provided by any worker would generate 0.75 € of revenue for the firm, 0.25 € of which went towards compensating the manager, and the remaining 0.50 € going towards the compensation of the workers. Participants then learned that workers would have 14 min to produce correct answers.

The determination of the worker's compensation was ultimately the responsibility of the manager. That is, the manager was given decision rights over implementing one of the two possible compensation schemes for workers. Firms would either operate under a piece-rate scheme equal to 0.50 € per unit of output, or a rank-order tournament where the number of correct answers from all group-members would be first summed, multiplied by 0.50 €, and then distributed to workers based on their ranks. The highest performer would receive 60% of the compensation proceeds, the second highest performer 30%, and the lowest 10%. We use these two compensation schemes since, from the perspective of rational agents, both stimulate the effort (Lazear and Rosen 1981).

The payoff going to the manager was significant because we wanted managers to carefully consider their choices and anticipate how their choices might affect worker effort. At the same time, we picked the parameters of the compensation schemes so that workers consistently earned more than the manager did. We chose to do this to minimize any invidious comparisons between workers and the manager. Under the piece rate, for each unit produced by worker *i* the boss earns 25 cents and the worker earns 50. In the tournament, winners and runners up always do better than the manager (earning 60 and 30 cents of their own value added, respectively). In addition, even the tournament loser will earn more than what the boss receives from her effort in many instances. In the end, fewer than 10% of workers earned less than the manager.

We then collected the compensation scheme preferences of the participants. These preferences were conditioned on the yet-unknown role assignment. Participants could register a preference for the piece rate scheme or the tournament and condition their choices on ending up as a worker or a manager. Once preferences were recorded, subjects were randomly assigned to one of the two roles by the computer program and informed of their job on their screen. While our preference elicitation was not incentivized, it was not clear why incentives would be necessary to get participants to reveal this information truthfully (an intuition corroborated by the strong associations found between stated preferences and choices in Tables 3, 5 of the next section).

Once participant roles were revealed, we explained that the workers would vote for a compensation scheme while the managers decided whether to cede authority

and let the worker vote determine (majority rule) the pay structure or implement a scheme unilaterally. Whether the manager ceded control or not determined the treatment into which workers were sorted. In the "No Participation" treatment, the manager decided on the compensation scheme unilaterally and in the "Participation" treatment, the manager allowed the worker vote to determine the compensation scheme.

Despite ceding seeming like an unusual choice to face for our managers, it is externally valid. For example, as discussed at some length in Semler (1989, 1993), the Brazilian manufacturing firm Semco has experimented with letting its workers set their own wages. Further, ceding control over other related parameters of the job (work hours and working from outside the office) is now standard practice.

Importantly, the worker and managerial choices were made simultaneously so managers could not condition their choices on the outcomes of the vote and workers could not vote based on the votes of the other workers or the manager's decision to cede control or not. After the votes and managerial decisions were recorded, participants learned the decision taken by the manager in their firm, the outcome of the vote among workers regardless of whether the votes actually determined compensation or not, and the compensation scheme to be implemented for the work task of solving addition problems.

During the 14-min production stage, workers added while managers waited. Recall that all subjects were given puzzles with the instructions and were never instructed against using them. The primary reason for the puzzle was to allow managers to have something to do while they waited for production to end. If they worked on the puzzle, they could preserve role anonymity in the experiment. We also provided the puzzles to reduce the opportunity cost of not adding for the workers. If they did not want to add, workers could work on the puzzle and appear as busy as the others.

At the completion of the work period, all subjects received a summary of their performance (i.e., the number of correct answers), the number of correct answers solved by the firm's highest performer, the total number of correct answers produced by the firm, their relative rank, their payout, the manager's earnings, and the total firm revenue generated during the work period. At this point, the experiment was completed and subjects were asked to fill out a brief post-study questionnaire. While subjects filled out the questionnaire final payments were gathered and distributed privately to each subject, one-by-one, as they left the lab. All sessions were conducted at the LINEEX laboratory at the University of Valencia and the experiment was programmed and conducted with the software z-Tree (Fischbacher 2007).

3 Results

Although our primary interest is the relationship between participation and worker productivity, we first report the results of "balance tests" for our experiment and examine both manager choices and worker votes. As part of the post-experiment questionnaire, participants reported their gender, their competitiveness (on a 1–5

Likert scale) and their "enjoyment of math" (also on a 1–5 scale). The means of these variables, along with the means of output in the practice period (our measure of ability), are collected in Table 1. The wording of the competitiveness question was put in the context of sports, a context in which competitiveness is seen as a virtue: "concerning just sports and leisure activities, how competitive do you think that you are?" and the math enjoyment question was phrased simply: "I enjoy solving math problems".

Balance, according to *t* tests, has been achieved because none of the participant characteristics differs significantly between the two treatments. What differences do exist, however, should to work *against* finding a participation effect. Fewer men were assigned to the Participation treatment in which their votes counted and the extensive previous literature (e.g., Niederle and Vesterlund 2011; Carpenter et al. 2010) suggests that (1) men prefer tournaments (relative to women) and (2) tournaments tend to be more productive. Those in the No participation treatment were also slightly more competitive, on average, and, most importantly, the ability is a little higher in the No participation treatment. Significant or not, we shall control for these differences in our analysis.

The preferences and choices of our managers are summarized in Table 2 (raw preferences, not dependent on roles, are reported in the online appendix Table A1). Starting with the bottom row, we see that there were 80 firms (composed of 240 workers and 80 managers), and the managers were hesitant to cede control to the workers. In just one-quarter (20/80) of the firms, did the manager let the vote of the workers determine the compensation scheme. In the other sixty cases, managers were more likely to pick the tournament (37/80) than the piece rate (23/80). This unavoidable imbalance will make it more difficult to estimate a participation effect precisely.

Another important feature of the manager data is that compensation scheme preferences elicited before any roles were revealed seem to strongly determine choices. Considering the rest of Table 2, one can see the clustering of observations in which manager preferences are in accord with choices. The exception is among the "ceding" managers. Here, manager preferences were more evenly split between the pay schemes. Despite the dilution of the link between preferences and choices coming from the last column, Cramer's V, a measure of association for categorical variables, is quite large (0.67) in Table 2, indicating that the non-ceding managers followed their preferences closely. A more controlled demonstration of this result comes from Table 3 in which we report multinomial logit results showing that

Table 1 Mean observables by treatment

	Male	Ability	Enjoy math	Competitiveness
No participation	0.50 (0.50)	7.31 (4.18)	3.51 (1.11)	4.11 (0.89)
Participation	0.43 (0.50)	7.13 (5.02)	3.45 (1.20)	3.98 (1.02)

(s.d.); no participation refers to the treatment in which the manager decided the compensation policy unilaterally and participation indicates the worker vote determined the policy

Table 2 Manager preferences and choices

	Piece rate	Tournament	Cede	Total
Prefer piece rate	21	5	12	38
Prefer tournament	2	32	8	42
Total	23	37	20	80

Table 3 Manager preferences and choices

	Chooses piece rate		Chooses tournament	
	(1a)	(2a)	(1b)	(2b)
Prefers tournament	−1.946** (0.875)	−1.965* (1.032)	2.262*** (0.667)	3.624*** (0.833)
Ability		−0.138 (0.502)		1.515*** (0.586)
Male		0.860 (0.707)		0.205 (0.725)
Enjoy math		0.238 (0.346)		−0.275 (0.303)
Competitiveness		−0.414 (0.387)		−0.420 (0.393)
Observations	80	80	80	80

Dependent variable is whether a manager decided to use the piece rate or the tournament and ceding control to the workers is the baseline; manager compensation preferences elicited *ex ante*; multinomial logit coefficients; (robust standard errors)

* $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$

Table 4 Worker preferences and votes

	Vote piece rate	Vote tournament	Total
Prefer piece rate	127	21	148
Prefer tournament	15	77	92
Total	142	98	240

managers who prefer tournaments, *ex ante*, are less likely to pick piece rates and more likely to pick the tournament when ceding control is the point of reference. Table 3 also suggests that managers of higher ability were more likely to pick the tournament, while, interestingly, none of the other participant characteristics seemed to matter (i.e., although men might be more likely to opt into a competition, we find they are no more likely than women to impose one on someone else).

Like the managers, worker votes are very consistent with their stated preferences. In Table 4, we illustrate the association between worker preferences, elicited before the experiment began, and the votes the workers cast during the experiment. Overall, the workers seem to shy away from competition: 62% (148/240) registered a preference for the piece rate over the tournament and the vote ended with a similar proportion, 59% (142/240), casting a ballot for the piece rate. Looking at the cross tabulation, one sees the same clustering in the worker data as was evident in Table 2 . Here, the value of the association statistic is also quite high ($V = 0.69$).

To corroborate the conclusions drawn from Table 4, we report marginal effects after probit regressions in Table 5. First, replicating the vast literature on gender

Table 5 Worker preferences and votes

	(1)	(2)
Worker prefers tournament	0.695*** (0.048)	0.703*** (0.049)
Ability		−0.033 (0.063)
Male		0.207*** (0.077)
Enjoy math		0.014 (0.034)
Competitiveness		0.009 (0.042)
Observations	240	240

Dependent variable is one when the worker voted for a tournament; tournament preferences elicited *ex ante*; probit marginal effects reported; (robust standard errors)

* $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$

differences mentioned above, we find that men are, indeed, more likely (21pp, $p < 0.01$) to vote for the tournament. However, no other characteristic predicts the vote. The fact that ability does not correlate with voting for the tournament suggests that not showing workers the distribution of ability prior to the start of the experiment did help attenuate strategic voting. Nevertheless, the most important result from Table 5 is that workers with a preference for the tournament were 70pp more likely to vote for it ($p < 0.01$).

In terms of the results of the worker votes, when the vote was determinative (that is, in the Participation treatment), 60% of the firms adopted piece rates and 40% adopted the tournament. As one may recall from Table 2, when the manager chose the compensation scheme, the percentages are almost exactly reversed: 38% of managers picked the piece rate and 62% picked the tournament. Using a proportions test, the difference in the incidence of tournaments between the two types of firms is only marginally significant ($p = 0.09$). Overall, 35 of the 80 firms used the piece rate and 45 used the tournament.

We now switch the focus of our analysis to worker output. As one can see in Fig. 1, there is a nice, symmetric, distribution of the number of correct answers that our workers produced. The mean number of sums created was 27 (s.d. 10.14) and output varies from zero to sixty sums.

The main purpose of the experiment is to identify the effect of participation, manifested in the right to vote, on worker effort and output. To identify any effect of participation as causal, we must exploit exogenous variation in the decisiveness of the vote. This variation is assured by our design. While the assignment of workers to the voting treatment is not completely random because managers decide to cede control or not, it is clearly exogenous for the workers—they have no say in this choice. Further, though managers may act on their preferences and, therefore, not decide whether to cede control to the workers randomly, they are randomly assigned to the firms. And, given the restrictions we placed on information about the characteristics of the workers (ability, in particular), managers could never condition their choices to cede or not on these attributes.

With this in mind, the natural place to start our analysis is a simple comparison of mean outputs between the participation treatments. This comparison is illustrated in

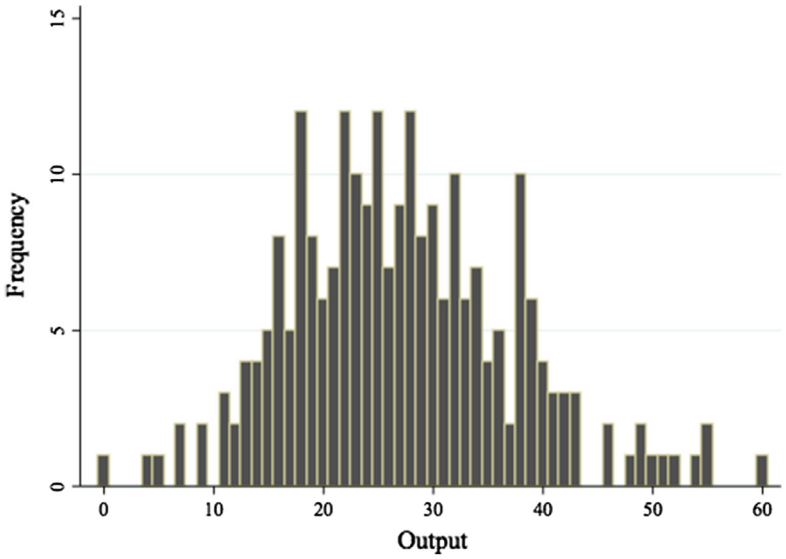

Fig. 1 The distribution of worker output

Fig. 2 where we compare the natural log of output in the two treatments (one was added before taking the log of the one zero observation seen in Fig. 1), allowing us to speak more conveniently in terms of percentage point differences. Indeed, the mean output is almost 7pp greater in the Participation treatment as hypothesized; however, because managers were hesitant to yield control and only 60 of the 240 workers are in participatory firms, the confidence intervals are larger than one might have wished. As a result, the difference in means is not significant using a simple t test ($p = 0.34$).

Recall, however, that the differences in observables seen in Table 1 are likely to mask the effect of voting in the summary statistics. To examine this possibility, we report regressions in Table 6 that account for these differences. In the first column, we replicate our t test and find a 6.6pp ($p = 0.23$) difference when no controls are added. In the second column, we add the controls for ability, gender, math enjoyment and competitiveness and see that the point estimate on the voice treatment does increase substantially. Here, the estimate is that workers with voice produce 11.3pp more than those without ($p = 0.02$), a substantial effect, especially given the subtlety of the treatment differences (i.e., the only difference is that the vote counted in one treatment and it didn' t in the other). From the rest of column (2), we also see that our measure of ability is, indeed, a strong predictor of output as is one's math enjoyment. However, we do not find significant effects of gender or competitiveness. Though the differences in ability and math enjoyment are not significant between the treatments (Table 1), they exist and because the effects of these characteristics on output are so large, the omission of even small differences from the analysis biases our estimate of the participation treatment effect.

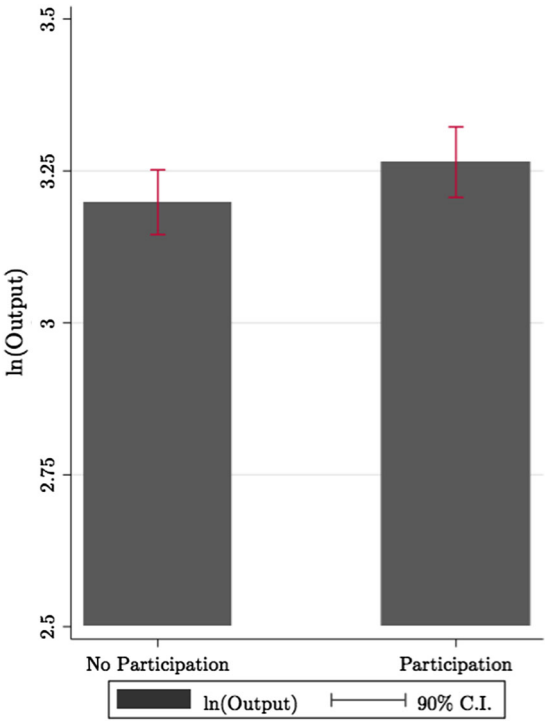

Fig. 2 Mean output by participation

Table 6 Does participation increase output?

	(1)	(2)	(3)
Participation	0.066 (0.055)	0.113** (0.048)	0.116** (0.047)
Ability		0.310*** (0.039)	0.307*** (0.039)
Male		0.023 (0.049)	0.018 (0.049)
Enjoy math		0.109*** (0.023)	0.107*** (0.022)
Competitiveness		0.038 (0.029)	0.037 (0.029)
Worker prefers tournament			0.088* (0.047)
Constant	3.198*** (0.037)	2.082*** (0.187)	2.067*** (0.189)
Observations	240	240	240
R^2	0.004	0.298	0.307

Dependent variable is the natural log of worker output; OLS; (robust standard errors)

* $p < 0.10$; ** $p < 0.05$, *** $p < 0.01$

Given assignment to a participatory firm is exogenous, we can interpret the point estimate on the Participation treatment causally. Yet one's intuition might suggest that the mixture of firm types (piece rate or tournament) between the conditions will influence this (average) treatment effect. We know, for example, that tournaments are a bit less common among our participatory firms and so would the point estimate be even larger if we accounted for this? The problem is that one cannot address this

concern by simply adding an indicator for the compensation scheme to the regressions in Table 6. To do so would be an example of a "bad control" (Angrist and Pischke 2009). Specifically, the compensation scheme implemented was the result of choice and not fixed at the start of the experiment. Nevertheless, our experiment was designed in anticipation of this issue. Instead of conditioning on the compensation scheme, we can add worker compensation preferences to the regression. These preferences are exogenous (fixed before any knowledge of the experiment was revealed) and, as seen in Table 5, worker preferences correlate highly with votes and, therefore, with which scheme is implemented.

As seen in column (3) of Table 6, worker preferences do appear to matter: those who prefer tournaments work harder but what is really being captured is that workers who prefer tournaments are more likely to vote for and end up in one and work harder as a result. Most importantly, the inclusion of these preferences does increase the participation point estimate, however, not by much. The estimate is now closer to 12pp than 11 and we conclude that the marginally significant difference in the number of tournaments between the two sectors is not large enough to matter much.

To be thorough, we also examined a specification in which we interacted participation with all the observables (including Participation $\times$ Worker Prefers Tournament) but the resulting F test indicated that the interactions added little or nothing to the analysis ($p = 0.19$). We also added the manager compensation scheme preferences which were never significant, nor did they have any impact on the Participation point estimate. The results of other, more standard, robustness checks appear in appendix Table A2. In a first regression, we use the specification in column (3) of Table 6 and cluster the standard errors at the level of the work group. This has little effect and the point estimate remains significant at better than the 5% level. In a second regression, we revert back to using robust standard errors but include session fixed effects. In this case, the point estimate increases to 0.119 and is significant at better than the 1% level. In two additional regressions, we examined using the raw output amount as the dependent variable, instead of its natural log and, again, the results were very similar, though the point estimate looses significance in one specification. Lastly, we considered the hypothesis that workers may have been disappointed when their preferred compensation scheme was not used and that this might have demotivated them; however, we found no evidence of this effect, either overall or within the voting treatment. Given our various point estimates all fall within a tightly bound interval (between 11 and 12 percentage points) and remain significant in all but one specification, we conclude that our measure of the causal effect of participation on effort is robust.

Turning to the managers, the question is whether there is an opportunity cost for ceding managers—do managers who empower workers to participate in the running of the firm sacrifice profits? Given the productivity bump accruing to participatory firms, one might expect that ceding is good for profit; however, as mentioned above this will also depend on the frequency of productive tournaments in the two "sectors". Looking just at mean firm profit, participatory firms actually appear to do a bit better (i.e., 3.5pp). Digging a bit deeper in Table 7, in which we use the same empirical strategy as we used for worker output (i.e., controlling for mean worker

Table 7 Do ceding managers earn less?

	(1)	(2)	(3)
Participation	0.035 (0.056)	0.079 (0.051)	0.087* (0.050)
Ability (firm mean)		0.184*** (0.0347)	0.192*** (0.049)
Male (firm mean)		0.153 (0.093)	0.139 (0.099)
Enjoy math (firm mean)		0.119*** (0.056)	0.127** (0.056)
Competitiveness (firm mean)		0.054 (0.040)	0.056 (0.039)
Manager prefers tournament			0.042 (0.050)
Constant	4.071*** (0.036)	3.020*** (0.1307)	2.949*** (0.307)
Observations	80	80	80
R^2	0.003	0.337	0.343

Dependent variable is the natural log of firm revenue; OLS; (robust standard errors)

* $p<0.10$; ** $p<0.05$; *** $p<0.01$

characteristics and the compensation preferences of the manager) the difference rises to almost 9 percentage points and is significant at better than the 10% level. Ceding control in our experiment also seems to have increase the profitability of the firm.

4 Discussion

In his seminal 1937 essay, Ronald Coase used Dennis Robertson's imagery describing firms as "*islands of conscious power in [an] ocean of unconscious cooperation*" to underscore his view that firms are planned economies based on the managed coordination of resources and, further, that the ability for firms to skip the market by making decisions over how to organize production makes them more efficient than if production was coordinated by the market.

For most economists, the primary lesson taken away from Coase is that all coordinating mechanisms have positive transactions costs, which, once considered, give rise to the large class of "make, buy or integrate" problems at the heart of the post-war economic theory of the firm. And although much of this theory accepts the role of management coordination in economizing production, both the particulars of managerial procedures and their potential for affecting the efficacy of a given policy are generally not considered. For example, Alchian and Demsetz (1972) theorize over the role of a monitor–manager of team production, and Williamson (1971) suggests that "*fiat is frequently a more efficient way to settle minor conflicts over [appropriable quasi-rents]*," yet neither express an interest in understanding how a coordinator would go about doing this, or why we should expect compliance. Demsetz (1997) goes so far as to making a normative claim that "*[the] objective is to understand price-guided, not management guided, resource allocation. The firm ...is that well known black-box into which resources go and out of which goods come, with little attention paid to how this transformation is accomplished.*" Our

study suggests that to the extent there are economic benefits to managerial practices, it may also mean that they can be systematically exploited to increase the efficiency of the firm at relatively little cost.

Acknowledgements Thanks to the Beyster Fellow Programs for funding and The editors, two thoughtful referees, Michal Goldberg, Avner Ben-Ner, Alex Bryson, Joseph Blasi, Richard Freeman, Trevor Hyman Young, Penelope Hernandez, James Warner, Peter Thompson, the participants at the 2013 Midyear Fellows Workshop in honor or Louis O. Kelso and the Analytical Political Economy Workshop for valuable comments on earlier versions of this paper.

References

Angrist, J., & Pischke, J.-S. (2009). *Mostly harmless econometrics*. Princeton: Princeton University Press.

Alchian, A. A., & Demsetz, H. (1972). Production, information costs, and economic organization. *The American Economic Review*, *62*(5), 777–795.

Bartling, B., Fehr, E., & Herz, H. (2014). The intrinsic value of decision rights. *Econometrica*, *82*(5), 2005–2039.

Cappelli, P., & Neumark, D. (2001). Do high-performance work practices improve establishment-level outcomes? *Industrial and Labor Relations Review*, *54*(4), 737–775.

Carpenter, J., Matthews, P., & Schirm, J. (2010). Tournaments and office politics: Evidence from a real effort experiment. *American Economic Review*, *100*(1), 504–517.

Charness, G., Cobo-Reyes, R., Jimenez, N., Lacomba, J., & Lagos, F. (2012). The hidden advantage of delegation: Pareto improvements in a gift exchange game. *American Economic Review*, *102*(5), 2358–2379.

Coase, R. H. (1937). The nature of the firm. *Economica*, *4*(16), 386–405.

Dal Bó, P., Foster, A., & Putterman, L. (2010). Institutions and behavior: experimental evidence on the effects of democracy. *American Economic Review*, *100*(5), 2205–2229

Demsetz, H. (1997). The firm in economic theory: A quiet revolution. *The American Economic Review*, *87*(2), 426–429.

Falk, A., & Kosfeld, M. (2006). The hidden cost of control. *American Economic Review*, *96*(5), 1611–1630.

Fehr, E., Herz, H., & Wilkening, T. (2013). The lure of authority: Motivation and incentive effects of power. *American Economic Review*, *103*(4), 1325–1359.

Fischbacher, U. (2007). z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, *10*(2), 171–178.

Lawler, E. E., Mohrman, S. A., & Ledford, G. E. (1995). *Creating high performance organizations: Practices and results of employee involvement and total quality management in Fortune 1000 companies*. San Francisco, CA: Jossey-Bass.

Lazear, E., & Rosen, S. (1981). Rank-order tournaments as optimum labor contracts. *Journal of Political Economy*, *89*(5), 841–864.

Macneil, I. (1985). Relational contract: What we do and do not know. *Wisconsin Law Review,* (3), 483–526.

Mellizo, P., Carpenter, J., & Matthews, P. H. (2014). Workplace democracy in the lab. *Industrial Relations Journal*, *45*(4), 313–328.

Niederle, M., & Vesterlund, L. (2007). Do women shy away from competition? Do men compete too much? *Quarterly Journal of Economics*, *122*(3), 1067–1101.

Niederle, M., & Vesterlund, L. (2011). Gender and competition. *Annual Review of Economics*, *3*, 601–630.

Semler, R. (1989). *Managing without managers*. Cambridge, MA: Harvard Business Review.

Semler, R. (1993). *Maverick: The success story behind the world's most unusual workplace*. New York: Warner Books.

Williamson, O. E. (1971). The vertical integration of production: Market failure considerations. *The American Economic Review*, *61*(2), 112–123.