

MANAGEMENT SCIENCE

Stand by Me. Experiments on Help and Commitment in Coordination Games

Journal:	<i>Management Science</i>
Manuscript ID:	MS-11-01483.R2
Manuscript Type:	Behavioral Economics
Keywords:	Economics: Behavior and Behavioral Decision Making, Economics: Microeconomic Behavior, Games-Group Decisions
Abstract:	<p>We present experiments studying how high ability individuals use help to foster efficient coordination. After an initial phase that traps groups in a low productivity equilibrium, incentives to coordinate are increased, making it possible to escape this performance trap. The design varies whether high ability individuals can offer help and, if so, whether they must commit to help for an extended period. If help is chosen on a round by round basis, the probability of escaping the performance trap is slightly reduced by allowing for help. The likelihood of success significantly improves if high ability individuals must commit to help for an extended time period. We develop and estimate a structural model of sophisticated learning that provides an explanation for why commitment is necessary. The key insight is that potential leaders who are overly optimistic about their ability to teach their followers are too fast to eliminate help in the absence of commitment.</p>

SCHOLARONE™
Manuscripts

1
2
3
4
5
6
7
8
9

Stand by Me

Experiments on Help and Commitment in Coordination Games

10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

May 10, 2014

Abstract: We present experiments studying how high ability individuals use help to foster efficient coordination. After an initial phase that traps groups in a low productivity equilibrium, incentives to coordinate are increased, making it possible to escape this performance trap. The design varies whether high ability individuals can offer help and, if so, whether they must commit to help for an extended period. If help is chosen on a round by round basis, the probability of escaping the performance trap is slightly reduced by allowing for help. The likelihood of success significantly improves if high ability individuals must commit to help for an extended time period. We develop and estimate a structural model of sophisticated learning that provides an explanation for why commitment is necessary. The key insight is that potential leaders who are overly optimistic about their ability to teach their followers are too fast to eliminate help in the absence of commitment.

Keywords: Incentives, Coordination, Experiments, Organizations, Heterogenous work teams

JEL Classification Codes: C92, D23, J31, L23, M52

1. Introduction

A group's performance is often constrained by its lowest performing individual. An assembly line moves no faster than the slowest person in the line, a report doesn't get finished until the last person completes their section, and a meeting can't start when a key attendee is late. Groups with such strong complementarities can easily become stuck in a performance trap: everyone understands that all would benefit if all group members put forth their best effort, but, since no individual can unilaterally change the outcome for the better, all group members shirk in the belief that any attempt to break out of the performance trap is wasted effort.

When a firm (or any organization) is stuck in a performance trap, spontaneous escape is unlikely given the need for coordinated change. Leadership is usually required – somebody has to take the initiative to start the process of change. Managerial leadership is an obvious possibility, but leadership by workers can also play an important role. Help, defined as a voluntary activity by high ability types that makes the job of low ability types easier at some cost for the high ability types, is a natural instrument that workers can use to provide leadership.¹ Intuitively, work teams often contain some workers that have higher ability than others. If workers are rewarded based on team production, then high ability workers can have an incentive to help their less able colleagues.² Help between workers has received little attention from researchers in economics and management.³ To the best of our knowledge, there exist no previous experimental papers on the topic of help.

The purpose of this paper is to present laboratory experiments exploring the use of help as a leadership tool. The provision and impact of help are complex phenomena involving interplay between the effects of material incentives, beliefs, and non-pecuniary motivations. Laboratory experiments are well suited for studying settings where such factors interact.⁴ The advantages of laboratory experiments for studying help are due to the control and observability available in experimental settings. We exogenously control whether help is possible and directly observe how much help is provided. We can

¹ Other natural instruments that can be used to lead the way out of a performance trap include leading by example (Brandts and Cooper, 2006; Hyndman, Terracol, and Vaksmann, 2009) and communication (Blume and Ortman, 2007, Brandts and Cooper, 2007).

² Hamilton et al. (2003, 2004) provide a real world example of this. They study the productivity of worker teams in a garment plant. Teams with heterogeneous abilities become more productive as the spread in ability increases (holding the average fixed). Hamilton et al. conjecture that this is due to mutual learning, a specific form of help where high ability types help low ability types understand how work should be done and pull them up to a higher level.

³ See Drago and Garvey (1998) for a questionnaire study on the effects of promotion incentives on help. See also Brown and Heywood (2009) for a field study on help in settings with individual performance pay and Antonetti and Rufini (2008) for a theoretical analysis of training of low ability types by high ability types.

⁴ For examples, see Zwick and Chen (1999), Brandts and Charness (2003), Weber and Camerer (2003) and Harbring and Irlenbusch (2009).

1
2
3 therefore cleanly determine whether help is responsible for observed changes in performance as opposed
4 to other factors. We can also observe the details of the interactive process of giving help and reacting to
5 help, rather than only observing the final outcome, all of which makes it possible to pin down the
6 mechanism by which help affects performance. While it is risky to generalize conclusions from a lab
7 experiment to managerial choices in field settings, we believe that laboratory experiments are a valid tool
8 for studying the effects of help.
9

10
11
12
13 Our experimental results show that help is only effective if high ability types are forced to commit
14 to providing help for an extended period of time. To explain this result, we develop and estimate a
15 structural model of learning that includes sophisticated learners. This exercise indicates that the root
16 cause of poor performance with help, in the absence of forced commitment, is “over-optimism” by high
17 ability types: believing incorrectly that they have taught others to play an efficient equilibrium, high
18 ability types tend to prematurely abandon the provision of help. This finding has implications for
19 leadership in many settings and helps explain puzzling patterns from other studies of leadership.
20
21
22
23

24
25 Getting into the details, our experiments are based on the “corporate turnaround game.”⁵ This is
26 an experimental setting designed to represent a corporate environment in which a group has fallen into a
27 performance trap and needs to escape from it. The game involves repeated play between a “manager” and
28 four “employees” of a “firm”. We automate the role of the manager while employees are played by
29 experimental subjects. In each round, the manager first chooses a bonus rate which determines the
30 fraction of the firm's profits transferred to the employees. After seeing the bonus rate, the four employees
31 simultaneously choose how much costly effort to expend with firm output and profits determined by the
32 minimum effort. Critically, the four employees are not identical. Instead, each group has one high ability
33 type with relatively low effort costs and three identical low ability types with relatively high effort costs.
34
35
36
37
38

39
40 Employees initially face a low bonus rate that makes the minimum possible effort a dominant
41 strategy for low ability types, trapping groups in the worst outcome possible. The bonus rate is then
42 exogenously increased, turning the game into a weak link game (Van Huyck, Battalio and Beil, 1990)
43 with multiple Pareto ranked equilibria. Even with this change in incentives, the strong complementarities
44 between workers’ efforts make a spontaneous escape from the performance trap far from certain since
45 improved performance requires a unanimous switch to higher effort.
46
47
48

49
50 The key issue in our paper is whether high ability types can use help to assist the group’s escape
51 from the performance trap. We study two specific kinds of help, indirect and direct. We first consider
52 indirect help which involves high ability types improving the productivity of low types at the cost of
53

54
55 ⁵ The corporate turnaround game was introduced by Brandts and Cooper (2006). See also Hamman, Rick, and
56 Weber (2007). More broadly, there exists an extensive literature on coordination in weak-link games. Camerer
57 (2003) surveys the early literature. Prominent recent papers include Chaudhuri, Schotter and Sopher (2006), Weber
58 (2006), Blume and Ortmann (2007) and Kogan, Kwasnica and Weber (2011).
59
60

1
2
3 decreasing their own productivity. We refer to this as “indirect” because high ability workers do not
4 spend time on the low ability workers’ task, but instead make it possible for the low ability workers to
5 accomplish their tasks more efficiently. High ability types can gain by providing indirect help, even
6 though their productivity is lowered, if low ability types increase their effort enough that the group
7 escapes its performance trap. We defer introducing the direct help treatment until later, as the motivation
8 for this treatment grew out of the results for the indirect help treatment.
9

10
11
12
13 An initial set of treatments (Experiment 1) includes four treatments that vary the ability of high
14 ability types to provide indirect help. In the Endogenous Help treatment, the high ability worker can help
15 his low ability co-workers and the level of help can be changed on a round-to-round basis. Behavior in
16 this treatment is compared with outcomes from two control treatments, one where help is not possible (No
17 Help) and one where symmetric costs are imposed exogenously for the remainder of the interaction
18 (Symmetric Costs). More precisely, the latter treatment imposes the same cost structure as occurs when
19 the high ability type chooses a sufficiently high level of help as to equalize effort costs. We find that the
20 ability to provide symmetric help on a round-to-round basis has no positive effect on the likelihood of
21 efficient coordination as minimum effort levels are *lower* in the Endogenous Help treatment than in either
22 the No Help or Symmetric costs treatments, albeit weakly in the first case.
23

24
25
26
27
28
29 This surprisingly poor performance reflects neither a failure to use help nor a lack of
30 responsiveness by minimum effort to changes in help. Instead, the problem seems to stem from allowing
31 high ability types to change their level of help on a round-by-round basis. Many high ability types
32 frequently shift the amount of help they provide up and down, causing the minimum effort to ratchet
33 downwards for the group since the negative response to cutting help is stronger than the positive effect of
34 increasing help. This suggests that high ability types would do better if they committed to a level of help
35 for an extended time period, but the Endogenous Help treatment is insufficient to reach this conclusion
36 since high ability types might be changing help frequently because their group is doing badly. In other
37 words, there is a possibility of reverse causality.
38

39
40
41
42
43
44 The final treatment of Experiment 1 (Forced Commitment) addresses this by forcing high ability
45 types to commit to their chosen level of help (including zero help) for an extended period of time. Forced
46 Commitment yields significantly higher minimum effort levels than either the Endogenous Help or the
47 No Help treatments. Exogenously imposing commitment makes it easier to escape a performance trap
48 because it gives groups time to equilibrate before help can be changed. Having stabilized at an efficient
49 outcome, subsequent reductions in help have little effect on the effort of low ability types.
50

51
52
53
54 We develop and estimate a structural model that provides an explanation for why high ability
55 types reduce help too rapidly in Endogenous Help, and, by extension, captures the main treatment effects
56 observed in Experiment 1. This is a simplified version of Camerer, Ho, and Chong’s (2002) sophisticated
57
58
59
60

1
2
3 EWA model, combining the spirit of level-k reasoning with a learning model. Learning models are
4 inherently dynamic and hence well suited to studying the dynamic phenomena observed in our data.⁶
5 Including sophisticated learners in the model allows for individuals who anticipate learning and
6 optimization by others. Without this it becomes difficult to explain why help is used by high ability
7 types, let alone why it is switched frequently.
8
9

10
11 Subjects are assumed to be one of two types, unsophisticated or sophisticated learners. Note that
12 a subject being a high (low) ability type does not imply they are a sophisticated (unsophisticated) learner,
13 as the former is randomly and exogenously assigned by the experimenters while the latter is a behavioral
14 characteristic of the individual. As will be seen, the case of greatest interest occurs when a sophisticated
15 learner is assigned to the high ability role as sophistication affects their use of help. Unsophisticated
16 learners follow a simple rule closely akin to Cournot learning. Sophisticated learners noisily optimize
17 versus unsophisticated learners and are forward looking. Sophisticated learners anticipate the effect of
18 changing incentives on others' behavior and understand the benefits of strategic teaching.⁷ For high
19 ability types who are also sophisticated learners, help is an attractive tool. By giving short-term
20 incentives for higher effort, help not only changes the behavior of low ability types in the short run, but
21 also affects the experiences on which their learning is based. This pushes the learning dynamic towards
22 equilibria with coordination at higher levels of effort. Two additional features were added to the model
23 based on specific features of our data. (1) Over-optimism: Sophisticated learners overestimate their ability
24 to affect the beliefs of unsophisticated types and (2) Reciprocity: Low ability workers reward kindness
25 (increased help) by choosing higher effort levels and punish unkind behavior (decreased help) by moving
26 to lower effort levels.
27
28
29
30
31
32
33
34
35
36

37 To generate predictions from the model, we fit its parameters using data from the three treatments
38 we are trying to track, No Help, Endogenous Help, and Forced Commitment. The estimation is done via
39 the simulated method of moments, minimizing the squared difference between the observed and
40 simulated minimum effort averaging across periods and treatments. The resulting parameters give the
41 best fit averaging across treatments, but it does not follow that the model tracks differences between
42 treatments especially well and it is easy to generate examples where models fit to the data make
43 inaccurate predictions about differences between treatments. Indeed, if we fit a version of the learning
44 model with over-optimism and reciprocity removed and then simulate data for all three treatments using
45 the estimated parameters, the constrained model fails to predict that Forced Commitment increases
46
47
48
49
50
51

52
53 ⁶ Crawford and Broseta (1998) provide a notable example of using a learning model to track data from coordination
54 game experiments.

55 ⁷ Strategic teaching refers to attempts to alter others' future choices by manipulating their learning processes. For
56 experimental evidence of strategic teaching see Terracol and Vaksmann (2009), Hyndman et al. (2009), Fehr et al.
57 (2012) and Hyndman et al. (2012).
58
59
60

1
2
3 minimum effort relative to No Help and Endogenous Help. Performing the same fitting and simulation
4 exercise with the full unconstrained model (i.e. including over-optimism and reciprocity), the simulations
5 successfully tracks the differences between these three treatments. The estimated parameter for over-
6 optimism is significant and positive, but, surprisingly, the estimated reciprocity parameter is tiny and not
7 significant. Over-optimism helps the model track the data because it increases the use of help and, more
8 importantly, leads to premature reductions of help as observed in the data. A high ability type who is
9 sophisticated and over-optimistic will abandon help too rapidly, believing he has led his group to the
10 efficient equilibrium before this has actually occurred. Reciprocity is a powerful force in many
11 experimental settings (see Cooper and Kagel, forthcoming), but adding a simple, natural version of
12 reciprocity to our model harms its ability to track the data. In our model, sophisticated learners believe
13 the other subjects are unsophisticated learners who noisily optimize subject to beliefs. If reciprocity is
14 added to the model, the objective function of low ability types changes. A high ability type who is
15 sophisticated anticipates that reciprocity affects the choices of low ability types, expecting a more positive
16 effect when help is increased and a more negative effect when help is reduced. Like over-optimism,
17 adding reciprocity increases usage of help, but, unlike over-optimism, reciprocity makes reducing help
18 less attractive. This limits the model's ability to capture the overly rapid reductions of help that play such
19 an important role in our data.
20
21
22
23
24
25
26
27
28
29
30

31 To recapitulate, we learn three important things from the structural model: (1) sophisticated
32 learners must be included to generate use of help, (2) adding over-optimism allows the model to track the
33 differences between treatments by causing high ability types to abandon help too quickly, and (3) adding
34 reciprocity does not help the model track the data because this feature discourages high ability types from
35 abandoning help too quickly.
36
37
38

39 Experiment 1 studies indirect help, but this is just one of many possible ways in which help could
40 be provided. To study whether the importance of commitment generalizes to other natural versions of
41 help, Experiment 2 examines "direct" help. This involves high ability types spending time on the low
42 ability types' task, allowing the low ability types to do less work to complete their tasks. Help is direct
43 because high ability workers do a portion of the low ability types' work rather than making it easier for
44 low ability types to do their own work as with indirect help. Experiment 2 studies direct help with and
45 without forced commitment. In the first treatment high ability types choose help on a round-by-round
46 basis while in the second they commit to their chosen level of help for an extended period of time.
47 Without commitment the effect of allowing direct help is negative, with minimum effort levels greatly
48 reduced relative to the relevant control, the No Help treatment. Forced commitment again improves
49 performance. The reasons for this improvement parallel those from the indirect help treatments. The
50
51
52
53
54
55
56
57
58
59
60

1
2
3 results of Experiment 2 reinforce our main conclusion: the efficacy of help as an instrument for leading
4 groups out of performance traps depends critically on the commitment of leaders to providing help.
5

6
7 More broadly, our theoretical and experimental results suggest that stable behavior is a
8 characteristic of successful leadership with a broad variety of tools employed for escaping performance
9 traps. For example, in Brandts and Cooper (2007) leadership could be exercised by managers, played by
10 subjects rather than the computer, who controlled the bonus rate. The bonus rate was changed by
11 managers in almost 2/3 of the periods. Managers who were below the median frequency of switching
12 earned about 50% more than managers at or above the median (398 vs 260 ECU/round).⁸ This parallels
13 the negative effect due to excessive switching of help observed in the Endogenous Help and Direct Help
14 treatments. The insight is the same in all cases. Re-coordinating on a good equilibrium is a gradual
15 process. Once this process has finished, the good outcome can persist even if incentives are changed, but
16 if incentives change before equilibration has a chance to occur then coordination on a good equilibrium
17 may likely not occur. Institutions that tie leaders' hands can have a positive effect by preventing
18 premature abandonment of costly yet effective instruments for overcoming coordination failure.
19

26 **2. The Turnaround Game and Indirect Help**

27
28 The "turnaround game" is played by a fixed group ("firm") consisting of a manager and four
29 employees who interact repeatedly over a number of rounds. The firm's productivity in each round is
30 determined by employees' effort choices for the round, with employees' incentives to exert effort
31 depending on an ex ante profit sharing decision made by the manager. In the experiments reported below
32 the experimenter plays the role of manager while experimental subjects fill the roles of the employees.
33 Since our focus is on interactions between high and low ability employees, making the manager
34 exogenous generates a more controlled environment to study their relations. Even though the manager's
35 decisions are exogenous, it is useful for expositional purposes to treat her as a player in the game.
36

37
38 The turnaround game embodies three basic design choices. First, the firm's technology has a
39 weak-link structure, with production (as well as profits) depending in every round on the minimum effort
40 chosen by an employee. No employee can unilaterally increase the firm's production. Second, the
41 manager observes the output produced and, hence, observes the minimum effort but does not observe any
42 of the individual effort levels. This implies that any incentives the manager gives employees cannot
43 depend on individual effort. Finally, the firm manager can only reward employees with bonuses based on
44 the minimum effort observed. She can change this bonus rate but cannot otherwise influence the
45 employees' choices. In what follows we present the main features of the turnaround game in more detail.
46
47
48
49
50
51
52
53
54
55

56 ⁸ This figure pools data from all three treatments with subjects playing as managers and controlling the bonus rate.
57 The three treatments varied the ability of managers and workers to communicate.
58
59
60

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

An experimental firm in the turnaround game consists of a fixed grouping of four subjects (employees) who interact for thirty consecutive rounds, broken into three ten-round blocks. Each block starts with the announcement of a common bonus rate (B) for the ten rounds of the block that determines how much additional pay each employee receives for each unit increase in the minimum effort of the four employees. While playing in a block with a particular bonus rate, subjects did not know what the bonus rate would be in subsequent ten-round blocks. The bonus transfers part of the firm's profits from the manager to the employees. All four employees observe B and then simultaneously choose effort levels, where E_i is the effort level chosen by the i^{th} employee. We restrict an employee's effort to be in ten hour increments: $E_i \in \{0, 10, 20, 30, 40\}$. Intuitively, employees spend forty hours per week on the job, and effort measures the number of these hours that they actually work hard rather than loafing. Employees' payoffs are determined by Equation 1 below. Note that effort is costly, with C_i denoting the cost of a unit of effort for the i^{th} employee. All payoffs are denominated in "experimental currency units" (ECUs). These were converted to monetary payoffs at a rate of 1 euro equals 500 ECUs:

$$\text{Employee } i: \pi_e^i = 200 - C_i E_i + \left(B \times \min_{j \in \{1,2,3,4\}} (E_j) \right)$$

In all treatments the average cost of effort equals 7, (i.e. $\frac{1}{4}(C_1 + C_2 + C_3 + C_4) = 7$). Employee 1 is a high ability type with a low initial effort cost $C_1 = 1$ while the other three employees are low ability types with high initial effort costs $C_2 = C_3 = C_4 = 9$. In treatments with help, as described in Section 3, the final effort costs will differ from these as a function of how much help is provided.

In all treatments the first ten rounds were played with $B = 8$ and without the possibility of help. Table 1 shows the resulting payoffs for the two ability types for the case of $B=8$. With $B=8$ the game is *not* a weakest-link game. For a high ability (low cost) type the best response depends on the minimum effort of the other workers, but for low ability (high cost) types it is a dominant strategy to choose zero effort. In the unique Nash equilibrium all employees exert zero effort. We use the game with $B=8$ in the first block of ten rounds to get play stuck in a performance trap. This cannot be classified as coordination failure since there is a unique equilibrium, but sets a strong precedent of low effort for games where coordination at high effort is possible in equilibrium.

[Table 1 around here]

The final twenty rounds in all treatments are played with $B = 14$. Table 2 shows the new payoff tables for the two types of employees if effort costs are held fixed at their original levels. The game is now a weakest-link game. Coordination by all four employees on any of the five available effort levels is an equilibrium, but consider the incentives faced by employees if they try to move upwards from the lowest minimum effort level. For simplicity assume that all employees choose 0 or 40. For a high ability

1
2
3 (low cost) type, the incentives are fairly good. Increasing effort to 40 incurs a sunk effort cost of 40
4 ECUs in exchange for a potential gain of 520 ECUs. For this increase to have positive expected value,
5 the probability of the three low ability (high cost) types *all* increasing to effort level 40 must only be
6 greater than 1/14. The odds are more foreboding for low ability (high cost) types. They must sink an
7 effort cost of 360 ECUs and can only potentially gain 200 ECUs. For a positive expected payoff, the
8 probability of the high ability type *and* both low ability types increasing their effort to 40 must be greater
9 than 9/14. Low ability types must be far more optimistic than high ability types to be willing to take the
10 risk of increasing their effort levels.

11
12
13 [Table 2 around here]
14
15
16

17
18
19 The advantage of high ability (low cost) types only matters if they can get low ability (high cost)
20 types to increase their effort. The flip side of high ability types having good incentives to try to
21 coordinate is that they also have strong incentives to try to influence the turnaround process. They can do
22 this by helping low ability types.
23
24
25

26 27 **3. Experiment 1: The Effect of Indirect Help**

28
29 There are many ways in which workers can help each other, and no one experiment can capture
30 all of the possibilities. We sought kinds of help that were easily implemented in the context of the
31 turnaround game, did not fundamentally change the nature of the game being played, and had natural
32 analogs in relevant field settings. When implementing help in our experiment, we made choices along
33 several dimensions. We restricted ourselves to implementations where help is “efficiency neutral,”
34 meaning the total surplus at the efficient equilibrium is unaffected by the provision of help. Our goal was
35 to emphasize the use of help as a leadership tool rather than as a means of directly improving
36 productivity. We also restricted ourselves to types of help that have a temporary effect on costs rather
37 than a permanent one. Using help with a temporary effect goes hand in hand with making help efficiency
38 neutral. It also simplifies the dynamic optimization problem facing subjects. We want subjects focused
39 on the tradeoff between the current costs of help and the long term gains from coordinating on an efficient
40 equilibrium without also needing to consider the long term implications for effort costs. Finally, both
41 indirect and direct help have natural interpretations in the context of escaping performance traps. We
42 examine indirect help first.
43
44
45
46
47
48
49
50

51
52 We model indirect help as a voluntary activity by which the high ability (low effort cost) worker
53 reduces the effort cost of a low ability (high effort cost) type at the expense of increasing his own effort
54 cost. To make this more concrete, imagine an experienced worker who volunteers to keep an eye on a
55 junior colleague. She stops what she’s doing occasionally to check on her colleague. Possibly she
56
57
58
59
60

1
2
3 catches a mistake before it becomes difficult (and time-consuming) to fix. The junior colleague does his
4 own work, but the advice from his senior colleague makes it easier for him to do it. Of course, this has a
5 cost for the person giving advice since she must take time and focus away from her own work to monitor
6 her colleague.
7
8
9

10 To make indirect help efficiency neutral, we model the cost of help as being 1 to 1 – lowering a
11 low ability type's effort cost by one unit raises the high ability type's effort cost by one unit. For
12 example, if the high ability worker decides to reduce the effort cost of all three low ability workers by
13 one, the effort cost of these three workers decreases from 9 to 8, while the effort cost of the high ability
14 worker increases from 1 to $1 + (3 \times 1) = 4$. To simplify the experiment, low ability types are not allowed
15 to help others. Only allowing one player to provide help eliminates additional coordination problems
16 such as who is supposed to be providing leadership if more than one individual provides it.
17
18
19
20

21 Experiment 1 comprises a total of four treatments which are summarized in Table 3. The results
22 of the first three treatments led us to the conjecture that commitment is crucial for performance
23 improvement. The fourth treatment, Forced Commitment, is designed to demonstrate that there is a causal
24 relationship between commitment and escaping a performance trap. We defer describing this treatment
25 until after the results of the first three treatments have been presented.
26
27
28
29

30 [Table 3 around here]
31

32 All treatments in Experiment 1 are identical in rounds 1-10 but vary with respect to how the effort
33 costs are determined in rounds 11 – 30. Recall that during the first ten rounds effort costs were
34 exogenously fixed in all treatments at $C_1 = 1$, $C_2 = C_3 = C_4 = 9$. In the No Help treatment, help is not
35 possible and effort costs remained exogenously fixed at the initial levels throughout rounds 11 – 30. In
36 the Endogenous Help treatment, the high ability employee (Employee 1) has, in rounds 11 – 30, the
37 option of providing each low ability employee (Employees 2 – 4) with 0, 1, 2, 3 or 4 units of help ($H \in \{0, 1, 2, 3, 4\}$) before effort decisions for the current round are made. In other words, the high ability
38 type worker commits to one of the possible levels of help for the current round *but not for future rounds*.
39 To keep matters simple, each low ability type must be given the same amount of help. After each round
40 all employees are informed about effort choices in the round.
41
42
43
44
45
46
47
48

49 For the first three levels of help, used in 94% of all observations, offering help corresponds to:

- 50 a. ($H = 0$) Leaving the effort cost distribution at: 1, 9, 9, 9.
 - 51 b. ($H = 1$) Reducing each of the three co-worker's effort cost by 1 and increasing own effort cost
52 by 3, leading to a cost distribution of 4, 8, 8, 8.
 - 53 c. ($H = 2$) Reducing each of the three co-worker's effort cost by 2 and increasing own effort cost
54 by 6, leading to a distribution of 7, 7, 7, 7.
- 55
56
57
58
59
60

1
2
3 Help levels 3 and 4 modify the distribution of effort costs in an analogous way. At these help
4 levels Employee 1 assists the others so much that her costs of effort are higher than those of low ability
5 types. It is therefore not surprising that in our results the use of help levels 3 and 4 is rare.
6
7

8 In the Symmetric Costs treatment, the distribution of effort costs is exogenously switched for
9 Blocks 2 and 3 to the symmetric distribution with a common effort cost of 7. This is equivalent to
10 exogenously imposing $H = 2$, although without any of the intentionality (and hence the scope for
11 reciprocity) that may be attributed to the active provision of help. Effort costs in the Symmetric Costs
12 treatment were held constant for ten-round blocks, a fact which was common knowledge among
13 employees. In this treatment there is effectively a commitment to $H = 2$ for the ten-round block.
14
15
16
17

18 [Table 4 around here]
19

20 Section 6 contains a formal model designed to explore why Forced Commitment increases
21 minimum effort levels while Endogenous Help does not, but our original hypothesis about differences
22 between the No Help and Endogenous Help treatments were more intuitive in nature. Help, whether
23 provided endogenously or imposed exogenously, does not change the basic structure of the resulting
24 game. Any subgame that occurs after choosing a level of help is still a weakest-link game with five Pareto
25 ranked equilibria, but the costs and benefits of trying to coordinate on an efficient equilibrium are
26 changed. For example, set $H = 2$. This is the level of help imposed in the Symmetric Costs treatment and
27 most common level of help provided in the Endogenous Help treatment. The resulting payoff table is
28 shown in Table 4 (there is only one payoff table since all players have an effort cost of 7). Consider the
29 incentives to move from 0 to 40. This requires a (sunk) effort cost of 280 ECUs versus a potential gain of
30 280 ECUs. To have a positive expected payoff, the probability of *all* three other players increasing their
31 effort to 40 must be greater than $1/2$. This is an improvement over the equivalent figure of $9/14$ for low
32 ability types in the absence of help. Incentives are worsened for high ability types compared to $H = 0$, but
33 we conjecture that they provide help in the expectation of changing the behavior of low ability types and
34 hence should be willing to provide high effort even with lowered incentives.
35
36
37
38
39
40
41
42
43
44
45

46 ***Hypothesis 1:*** *Minimum effort will be higher in the Endogenous Help treatment than in the No Help*
47 *treatment.*
48
49

50 4. Procedures

51 Subjects were students of the University of Valencia recruited through an electronic recruitment
52 system. All sessions were run at the LINEEX computer lab of the University of Valencia. Subjects were
53 only allowed to participate in a single session and had no previous experience in similar experiments.
54
55
56
57
58
59
60

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

At the beginning of each session subjects were randomly seated. Printed instructions were distributed and read aloud by the experimenter. The instructions stress that there are two types of employees with differing payoff tables and that the groups were fixed for the duration of the experiment. Before beginning play, all subjects were asked to complete a short quiz about the payoffs and the rules of the experiment. The instructions for the endogenous help treatment are shown in Appendix A.

At the beginning of each ten-round block employees were informed of the bonus rate for that block. Employees were not told what bonus rates would be in subsequent blocks. Treatments differed in what happened before employees made their effort decision. In each round of the No Help and Symmetric costs treatments the four employees of a firm simultaneously chose their effort levels for the round. While choosing, the employees were shown a payoff table, similar Tables 1 and 2 above, showing their payoff as a function of their own effort level and the minimum effort level chosen by the other employees. This payoff table was automatically adjusted to reflect the current bonus rate. Subjects also had a printed copy of the payoff table for employees with different effort costs.

At the end of each round, all employees saw a feedback screen showing them their effort level, the minimum effort for their firm, their payoff for the round, and their running total payoff for the experiment. Separate windows on the feedback screen showed a summary of results from earlier rounds and the individual effort levels selected for all four employees in their firm. These effort levels were sorted from highest to lowest and did not include any identifying information about which employee was responsible for which effort level.⁹

Subjects in treatments involving help received additional printed instruction before the start of the second block explaining how help worked. These instructions explained the way in which help could take place and provided all possible payoff tables that could occur for either type of player with some feasible level of help. Subjects were not told about the possibility of help prior to this point in time, so all treatments are parallel until the beginning of the second block. Subjects in the No Help and Symmetric costs treatments did not receive new instructions prior to the second block, but did have a pause in play where they were told about changes in the bonus rate and (when relevant) cost structure.

The timing of rounds 11 – 30 in the Endogenous Help treatment (as well as subsequent treatments with help) differs from that in the other two treatments of Experiment 1. After seeing the feedback from the previous round, the high ability type selected a level of help. All employees saw the level of help selected, as well as the resulting costs for the two types, and then simultaneously made effort decisions. The payoff table shown while making this decision adjusted to reflect the amount of help provided by the high ability type. Printed payoff tables were available to see the payoffs of the other type.

⁹ Given the symmetry of help the absence of identifying information seemed to us to be the simplest choice.

1
2
3 At the end of the session, each subject was privately paid in cash for all rounds played plus a
4 show-up fee. The average total payoff was €21.60, including a five euro show-up fee.
5
6
7

8 **5. Initial Results and a New Treatment**

9

10 Figure 1 shows average minimum effort in the four treatments of Experiment 1, including the
11 Forced Commitment treatment that we introduce below. In round 10, the last round prior to the bonus
12 rate increase, all firms had a minimum effort of zero. This is important since a history of low effort is a
13 precondition for our environment to be of interest.
14
15

16 [Figure 1 around here]
17
18

19 Figure 1 shows that minimum effort is slightly lower in the Endogenous Help treatment than in
20 the No Help sessions and higher in the Symmetric Costs treatment than in either the Endogenous Help or
21 No Help treatments. Comparing average effort at the employee level (rather than minimum effort at the
22 firm level) yields a similar picture.
23
24

25 Regression analysis contained in Appendix B (Table B.1) provides formal statistical backing for
26 the preceding observations. The ordered probit regressions reported in Appendix B correct for clustering
27 at the group level and include controls for behavior in the initial phase, Rounds 1 – 10. The regression
28 results reinforce our conclusions from the raw data: Endogenous Help does not significantly increase
29 firm production (as measured by minimum effort) relative to the baseline given by the No Help treatment.
30 Although Symmetric Costs somewhat improves matters, the difference between this treatment and No
31 Help is also not significant. The only significant difference we find between the three initial treatments is
32 between the Endogenous Help and Symmetric Costs treatments in Rounds 21 – 30, and this is very weak.
33 The data rejects Hypothesis 1.
34
35
36
37
38
39

40 Comparisons of payoffs across the three initial treatments reveal surprising patterns. Payoffs for
41 high ability (low cost) types are identical in the Symmetric Costs and Endogenous Help treatments,
42 averaging 350 ECUs over Rounds 11 – 30, but substantially lower in both cases than the average payoff
43 of 455 ECUs with No Help.¹⁰ Payoffs for low ability (high cost) types are roughly the same in the
44 Endogenous Help and No Help treatments, averaging 278 and 264 ECUs respectively across Rounds 11 –
45 30. Both of these figures are substantially lower than the average payoff for the Symmetric Costs
46
47
48
49
50
51
52
53

54 ¹⁰ Running OLS regressions with high ability type profits as the dependent variable and using equivalent
55 specifications to those described in Appendix B, the difference in high ability type payoffs between the No Help and
56 the Endogenous Help (Symmetric Costs) Help treatment is statistically significant at the 5% (1%) level in Rounds
57 11 – 20 and at the 10% (10%) level in Rounds 21 – 30.
58
59
60

1
2
3 treatment across this time frame, 348 ECUs.¹¹ It isn't a shock that Endogenous Help makes high ability
4 types worse off – after all, help is costly – but it is surprising that Endogenous Help does not make low
5 ability types better off even though they benefit from receiving help.
6
7

8
9
10 *Conclusion 1: Endogenous Help and Symmetric Costs do not improve firm productivity, as measured by*
11 *minimum effort. The data does not support Hypothesis 1.*
12

13
14 One possible explanation for why the Endogenous Help treatment does not increase minimum
15 effort levels relative to the No Help treatment could be that help is used infrequently or at very low levels.
16 However, (strictly) positive help is provided in 70% of observations, 37 of 39 high ability types (95%)
17 provide positive help at least once, and 36 of 39 high ability types (92%) provide positive help in at least
18 25% of the rounds. An average of 1.21 units of help per low ability type are provided, a level which
19 changes little over time. The lack of a positive effect from the Endogenous Help treatment does not reflect
20 a failure by high ability types to use their ability to provide help.
21
22

23
24 Looking at changes in help rather than levels provides a better idea for why minimum effort is
25 unexpectedly low in Endogenous Help. High ability types frequently changed how much help they
26 provided. The median number of changes in 20 rounds was 6, or roughly one change every three rounds.
27 About a quarter of the high ability types, 10 of 39, changed the level of help in at least half of the rounds.
28 There are only three help levels that get frequent use, so many high ability types oscillated back and forth
29 between high and low help levels. Figure 2 shows the relation between the number of changes in help
30 and both the minimum effort levels of low ability types and the payoffs of high ability types. Firms from
31 the Endogenous Help treatment are broken into categories by how many times the level of help was
32 changed in Rounds 11 – 30. There is an obvious negative relationship between instability in help levels
33 and either minimum effort or profit levels.
34
35

36
37
38
39
40
41
42 [Figure 2 around here]
43

44 Frequent changes to the level of help imply both frequent increases and frequent decreases.
45 Figure 3 illustrates why the decreases are particular harmful. It shows the effect of the high ability type
46 changing the level of help on the average minimum effort level of the three low ability types in a firm.¹²
47 Data is taken from Rounds 12 – 30. Round 11 is excluded since the effect of help is confounded with the
48 bonus rate increase. Changes in the amount of help unambiguously affect the minimum effort of the three
49
50
51

52
53
54 ¹¹ Regression analysis finds that the differences in low ability type payoffs between the Symmetric costs treatment
55 and the Endogenous Help and No Help treatments are statistically significant in Rounds 11 – 20 at the 5% level and
56 at the 1% level in Rounds 21 – 30.

57 ¹² An observation here is the minimum effort by the three low ability types in a firm for a single round.
58
59
60

1
2
3 low ability types. When the lagged minimum effort of the low ability types is 0, an increase is the only
4 possible change. Such improvement mainly occurs when help is increased. If the lagged minimum effort
5 for the three low ability types is 40, only decreases are possible. This is far more likely after a decrease in
6 the level of help. The most interesting case is when the lagged minimum effort level for low ability types
7 is 10, 20, or 30. Given that most groups which successfully coordinate do so gradually, these cases are
8 critical ones in determining a firm's success. Increasing help or holding it constant only has a moderate
9 positive effect on minimum effort levels for low ability types, but decreasing help has a strong negative
10 effect on their minimum effort levels.¹³
11
12

13
14
15
16
17 [Figure 3 around here]
18

19 Ordered probit regressions (see Table B.2 in the appendix) provide formal evidence that the
20 response of low ability types' minimum effort to changes in help is asymmetric. These regressions
21 include controls for the lagged minimum effort of low ability types, time effects, and individual effects.
22 Summarizing the results, either increasing or decreasing help leads to statistically significant changes in
23 the minimum effort of low ability types. By itself, the current level of help does not have a significant
24 relationship with changes in effort. Changing help, not the level of help, causes changes in effort. The
25 negative effect of decreasing help is estimated to be almost three times as large as the positive effect of
26 increasing help.
27
28

29 Given the large negative consequences of decreasing help, it is puzzling that high ability types
30 frequently change the level of help. Two features of the data suggest an explanation for this behavior.
31 First, decreases in help primarily occur when lagged minimum effort is high.¹⁴ Just as things start going
32 well, high ability types tend to throw a wrench in the works by decreasing help. This is consistent with
33 premature attempts at profit taking by high ability types. We conjecture high ability types are too
34 optimistic about having gotten the firm out of its performance trap and reduce the help they provide,
35 failing to anticipate a drop in the effort of others. A second feature of our data supports this
36 interpretation. Consider observations in Rounds 12 – 30 where the minimum effort of others in the
37 previous round was strictly positive (10, 20, 30, or 40). When help is cut, 93% of high ability types either
38 choose the same effort level as the previous period or *increase* it. Increased effort is more than four times
39 as likely as decreased effort.¹⁵ This tendency by high ability types to couple decreases in help with
40
41
42
43
44
45
46
47
48
49

50
51 ¹³ The response of individual effort levels to changes in help are also asymmetric, but less so than minimum effort of
52 the low ability types. The minimum function accentuates the asymmetry.

53 ¹⁴ A decrease of help is 70% more likely than an increase when the lagged minimum effort of low ability types is
54 strictly positive.

55 ¹⁵ There are 23 increases and 5 decreases out of 69 total observations. If we limit the sample to cases where the
56 lagged effort was neither 0 nor 40, so changes in both directions are possible, these figures become 14 increases and
57 3 decreases out of 27 total observations.
58
59
60

1
2
3 increases in their own effort only makes sense if the low ability workers are not expected to react to the
4 decrease in help by cutting their effort.
5
6

7 *Conclusion 2: Our analysis of the dynamics in the Endogenous Help treatment suggests that a lack of*
8 *commitment by high ability types, especially a tendency to slash help when things are going well, leads to*
9 *poor performance in the treatment.*
10

11 Conclusion 2 is largely based on the negative relationship between frequent changes in help and
12 the minimum effort achieved by a firm, as illustrated by Figure 2, but the evidence provided by the
13 Endogenous Help treatment is not sufficient to establish a causal relation between these two variables.
14 Frequent switches might be caused by poor performance rather than the other way around. The Forced
15 Commitment treatment, the final treatment in Experiment 1, is designed to establish a causal relationship
16 between commitment and increased minimum effort.¹⁶ In this treatment, the high ability (low cost) type
17 chooses how much help to provide only in Rounds 11 and 21. In these two rounds, she sets the level of
18 help in exactly the same fashion as in the Endogenous Help treatment, choosing from $H \in \{1, 2, 3, 4\}$.
19 The effort cost of all three low ability (high cost) types is lowered by H and the high ability type's effort
20 cost increases by $3H$. As in the Endogenous Help treatment, the same amount of help must be provided
21 to all three low ability types. Unlike the Endogenous Help treatment, where help is chosen for each
22 round, in the Forced Commitment treatment the high ability type must stick to his decision for the entire
23 ten-round block. This prevents the frequent changes to the level of help observed in the Endogenous Help
24 treatment. Effort levels are still chosen round-by-round as in the other treatments.
25
26
27
28
29
30
31
32
33

34 Hypothesis 2 follows from our interpretation of behavior in the Endogenous Help treatment.
35 Without forced commitment, overly optimistic high ability types tend to decrease help too soon. In the
36 presence of forced commitment this premature reduction of help is not possible, facilitating successful
37 equilibration at a higher effort level.¹⁷
38
39
40

41 *Hypothesis 2: Minimum effort will be higher for the Forced Commitment treatment than for either the No*
42 *Help or Endogenous Help treatments.*
43
44
45

46 ¹⁶ The Symmetric Costs treatment also introduces an element of commitment, since workers are told in Rounds 11
47 and 21 that effort costs will be set equal across all four workers for the next ten rounds. This is mathematically
48 equivalent to imposing a commitment to provide two units of help. However, the comparison with the Endogenous
49 Help treatment is not clean due to multiple confounds: help is forced to be at a high level, help cannot be adjusted in
50 reaction to the initial response of the low ability types, and help is imposed exogenously rather than being set by one
51 of the four workers.

52 ¹⁷ When looking at the data from the first three treatments of Experiment 1, we noticed that groups which performed
53 poorly were often held back by low effort choices from a single worker. This suggested that allowing for
54 asymmetric provision of help, with different levels of help provided to different workers, might be more effective
55 than the symmetric help allowed in Endogenous Help. We tested this hypothesis with a treatment that allowed
56 asymmetric provision of help. This improved minimum effort levels, but the effect was small and not statistically
57 significant. For details, see Brandts, Cooper, and Fatas (2011).
58
59
60

Returning to Figure 1, we now focus on the Forced Commitment treatment. This treatment yields obviously higher minimum efforts than either the No Help or Endogenous Help treatments, a difference that grows between the first and second block. Minimum efforts in the Forced Commitment treatment are also moderately higher than in the Symmetric costs treatment for both blocks. Looking at effort rather than minimum effort leads to similar conclusions – the Forced Commitment treatment leads to a substantial and persistent increase in effort levels over the No Help and Endogenous Help treatments.

Payoffs improve for both types in the Forced Commitment treatment relative to the Endogenous Help treatment, with average payoffs over Rounds 11 – 30 of 413 vs. 350 ECUs for high ability types and 348 vs. 278 ECUs for low ability types. Over the final ten rounds, average payoffs for both types are (roughly) as high in the Forced Commitment treatment as in the best of the other three treatments.¹⁸

The regression analysis described in Appendix B (Table B.1) supports our conclusions about Forced Commitment. Compared with either the No Help or Endogenous Help treatments, the Forced Commitment treatment leads to significantly higher minimum effort and effort levels in Rounds 21 – 30 than in either the No Help or Endogenous Help treatments. Payoffs for both types in the final block are also significantly higher in the Forced Commitment treatment than in the Endogenous Help treatment. For low ability types, payoffs in both blocks are also significantly higher than in the No Help treatment. As should be expected since help is costly, the difference in high ability payoffs between the Endogenous Help with Commitment and No Help treatments is never statistically significant.

Conclusion 3: The Forced Commitment treatment has a large persistent effect on effort. Compared with the No Help treatment, Forced Commitment makes low ability types better off without harming the high ability types. The data supports Hypothesis 2.

The positive effect of Forced Commitment is probably not due to an increase in the level of help offered, as the average level of help over rounds 11 - 30 is higher in the Forced Commitment treatment than in the Endogenous Help treatment (1.40 vs. 1.21) but not dramatically so.¹⁹ As in the Endogenous Help treatment, high ability types in the Forced Commitment treatment often cut help for successfully coordinated firms – of twelve firms coordinated at 40 in Round 20, four cut help for the second block. The difference between the Endogenous Help and Forced Commitment treatments is that forced commitment gives time for groups to strongly converge to a new equilibrium in the first block with help (Rounds 11 – 20). In the groups where play was coordinated at 40 in Round 20 and help decreased for the final block, all four employees had been choosing 40 since at least Round 14. These groups had a

¹⁸ For high ability types, average payoffs over Rounds 21 – 30 are 476 ECUs, 381 ECUs, 377 ECUs, and 467 ECUs in the No Help, Symmetric Costs, Endogenous Help, and Forced Commitment treatments. Analogous figures for low ability types are 288 ECUs, 383 ECUs, 302 ECUs, and 377 ECUs.

¹⁹ Running a t-test on the average help provided low ability types, the difference is not statistically significant at even the 10% level (59 obs; $t = 0.98$; $p = .33$).

1
2
3 firmly established norm of coordinating on the efficient outcome. Two of the groups saw a brief decrease
4 in Round 21, in both cases due to a change by a single employee, but every employee in these four groups
5 chose effort level 40 in Rounds 22 – 30. Increasing help for the final block leads to increased minimum
6 effort, but, unlike the Endogenous Help treatment, *a decrease in help in the Forced Commitment*
7 *treatment does not cause a decrease in effort by low ability types*. For the seven groups where help
8 decreased, the average minimum effort *increased* slightly between blocks (23.7 for Rounds 11 - 20 vs. 24
9 for Rounds 21 - 30). In the five groups where help increased for Rounds 21 – 30, the average minimum
10 effort increased strongly between the two blocks (from 5.2 to 25.6). With forced commitment, changes in
11 help push effort up rather than down. We conjecture that forcing commitment gives the time for strong
12 convergence to the efficient equilibrium to occur before the high ability types can disrupt things by
13 cutting help, a conjecture formally examined by the model contained in Section 6.
14
15
16
17
18
19
20
21
22

23 **6. An Structural Model of Behavior in Experiment 1**

24 We now introduce a structural model designed to explore the causes of the two key observations
25 from Experiment 1: (1) Forced Commitment leads to an increase in minimum effort relative to No Help
26 and Endogenous Help and (2) high ability types in the Endogenous Help treatment cut help prematurely,
27 undermining the positive effects of having provided help. We develop a model that reproduces these
28 features. In Section 5 we conjectured that over-optimism by high ability types plays an important role in
29 generating these regularities by encouraging high ability types to abandon help prematurely. Going
30 beyond intuition, the fitting exercise directly tests whether allowing for over-optimism significant
31 improves the model's fit to the data and improves its ability to track differences between the treatments.
32 Reciprocity could also play a role in driving the dynamics, a proposition tested through the fitting
33 exercise.
34
35
36
37
38
39

40 *Model Specification:* Our basic model is a simplified version of the sophisticated EWA model of
41 Camerer, Ho, and Chong (2002). The technical details are given in Appendix C. The model includes two
42 types of individuals, unsophisticated and sophisticated learners. Individuals are randomly assigned types
43 with the probability of being a sophisticated type given by the parameter α . Note that both high and low
44 ability types are equally likely to be sophisticated learners, as the ability type (high or low) is assigned
45 exogenously by the experimenter while the individuals learning type (sophisticated or unsophisticated) is
46 an exogenous characteristic of the individual subject.
47
48
49
50
51

52 Unsophisticated players follow a simple adaptive learning rule closely related to Cournot
53 learning. The choice of Cournot learning rather than a more flexible learning rule like noisy fictitious
54 play, reinforcement learning, or EWA is driven by the structure of information in the experiment.
55 Sophisticated learners have to be able to model the learning of unsophisticated types. In our experiment,
56
57
58
59
60

1
2
3 subjects can observe the choices of the other three individuals in their group, but not in a way that is
4 identifiable. This means that it is impossible to track an individual's choices across periods.²⁰ Modeling
5 Cournot learning by others does not require this information, but models like noisy fictitious play,
6 reinforcement learning, and EWA require the ability to track individuals across periods. We therefore use
7 a model that does *not* require sophisticated types to use information they don't possess.

8
9
10
11 There are two types of unsophisticated learners which vary only in their initial beliefs about
12 others' behavior. Optimists initially believe that all other players will choose 40, the highest possible
13 effort level, while pessimists initially believe that all other players will choose 0.²¹ The initial beliefs of an
14 unsophisticated type are randomly drawn, with parameter θ giving the probability of a pessimist. After
15 each decision round, unsophisticated players update their beliefs, with the new beliefs being a weighted
16 average of the other group members' effort levels in the previous round and initial beliefs given by (1).
17
18
19

$$20 \quad q_{it}(m) = \omega(t)\eta_{it}(m) + (1 - \omega(t))q_{ui}(m) \text{ for } u \in \{\text{optimist, pessemist}\} \text{ and } \forall m \quad \mathbf{(1)}$$

21
22 Specifically, an unsophisticated player's beliefs $q_{it}(m)$ give weights for each possible minimum
23 effort $m \in \{0, 10, 20, 30, 40\}$ for the other three players. The weight on the previous round's outcome
24 is given by $\omega(t)$. This weight changes over time according to (2), where γ is a parameter fit from the data,
25 $\omega(1) = 0$, and $\omega(2) = \bar{\omega}$ ($\bar{\omega}$ is also a parameter fit from the data). The function $\eta_{it}(m)$ gives the weight on
26 minimum effort m if the probability of each effort level is given by the observed frequency in the
27 preceding period²² and $q_{ui}(m)$ is the initial weight on minimum effort m subject to the individual's type
28 (optimist or pessimist). The model allows for a reset of beliefs in Round 11 to account for the underlying
29 game changing when the bonus rate is increased and, possibly, help is allowed. This is captured by
30 subtracting the parameter ρ from $\omega(10)$, which puts less weight on experience from Round 10 and more
31 weight on initial beliefs.
32
33
34
35
36
37
38
39

$$40 \quad \omega(t) = \omega(t - 1) + \gamma \text{ if } t \geq 3 \quad \mathbf{(2)}$$

41
42 Based on these beliefs, expected payoffs can be generated for each of the five available effort
43 levels. These are transformed into "attractions" by adding extra weight to the effort level that was chosen
44
45

46
47 ²⁰ For example, suppose in two consecutive periods, the effort choices by the other three players in my group are 40,
48 20, 30 followed by 20, 40, 30. My feedback displays choices ordered from highest to lowest without IDs. So in
49 both periods I receive feedback showing 40 30 20. I have no way of knowing whether individuals have changed
50 their effort levels between periods or not.

51
52 ²¹ Adding these types helps the model capture the bimodal distribution of behaviour in the data. Adding more types
53 (or more sophisticated versions of heterogeneous initial beliefs) complicates the model, making it more difficult to
54 fit, without improving its ability to track the experimental data.

55
56 ²² For example, suppose the other three group members chose 0, 20, and 40 in the preceding period.
57 Unsophisticated types then assign a probability of 1/3 to each of these three effort levels. A probability distribution
58 over the others' minimum effort is generated by assuming that each of the other three individual group member's
59 effort levels is independently drawn with equal probability over 0, 20, and 40.
60

1
2
3 in the previous period, where the extra weight is given by a parameter δ multiplied by the number of
4 times the same choice has been repeated. Putting extra weight on repeated choices allows the model to
5 capture the fairly obvious hysteresis in the data. Earlier papers have found that incorporating hysteresis
6 into a learning model significantly improved the fit (i.e. Cooper and Stockman, 2002; Camerer, Ho, and
7 Chong, 2002; Wilcox, 2006).

8
9
10
11 The choice of an unsophisticated type is generated from beliefs using a logit decision rule, with
12 the parameter λ determining the level of noise. In the Endogenous Help and Forced Commitment
13 treatments, we also have to model help decisions by unsophisticated high ability types. This is done in a
14 manner that parallels the choice of effort levels described above, although a separate noise parameter, λ^H ,
15 is estimated for the choice of help.

16
17
18
19
20 Sophisticated types anticipate the learning of unsophisticated types and are forward looking,
21 engaging in strategic teaching. They understand that their current actions affect the beliefs of
22 unsophisticated players in the subsequent round. This makes help attractive as a tool for pushing
23 unsophisticated types toward higher effort levels.

24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65
66
67
68
69
70
71
72
73
74
75
76
77
78
79
80
81
82
83
84
85
86
87
88
89
90
91
92
93
94
95
96
97
98
99
100
101
102
103
104
105
106
107
108
109
110
111
112
113
114
115
116
117
118
119
120
121
122
123
124
125
126
127
128
129
130
131
132
133
134
135
136
137
138
139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158
159
160
161
162
163
164
165
166
167
168
169
170
171
172
173
174
175
176
177
178
179
180
181
182
183
184
185
186
187
188
189
190
191
192
193
194
195
196
197
198
199
200
201
202
203
204
205
206
207
208
209
210
211
212
213
214
215
216
217
218
219
220
221
222
223
224
225
226
227
228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
258
259
260
261
262
263
264
265
266
267
268
269
270
271
272
273
274
275
276
277
278
279
280
281
282
283
284
285
286
287
288
289
290
291
292
293
294
295
296
297
298
299
300
301
302
303
304
305
306
307
308
309
310
311
312
313
314
315
316
317
318
319
320
321
322
323
324
325
326
327
328
329
330
331
332
333
334
335
336
337
338
339
340
341
342
343
344
345
346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377
378
379
380
381
382
383
384
385
386
387
388
389
390
391
392
393
394
395
396
397
398
399
400
401
402
403
404
405
406
407
408
409
410
411
412
413
414
415
416
417
418
419
420
421
422
423
424
425
426
427
428
429
430
431
432
433
434
435
436
437
438
439
440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485
486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539
540
541
542
543
544
545
546
547
548
549
550
551
552
553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593
594
595
596
597
598
599
600
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647
648
649
650
651
652
653
654
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701
702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809
810
811
812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863
864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917
918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971
972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000

Going into details, sophisticated types assume that all of the other group members are unsophisticated types. This departs from the sophisticated EWA model in which sophisticated types place positive weight on other players also being sophisticated types. Our model therefore does not nest QRE and is more in the spirit of a level-k model. If sophisticated types believe that others may also be sophisticated, it becomes necessary to calculate a fixed point as part of fitting the model, greatly increasing the computational complexity of the exercise. Given that our model does a good job of tracking the data and gives useful insights into the underlying processes, we opt for the simpler model.

Sophisticated types have all the information necessary to calculate beliefs and attractions for the other three members of their group (assuming the others are unsophisticated). Applying the relevant logit decision rule, they can generate a distribution over choices and minimum efforts for the other three group members. These serve as a sophisticated type's beliefs about current actions. Sophisticated types also have all the information needed generate a probability distribution over the beliefs of unsophisticated types in the upcoming period, and by extension to generate a distribution over the future choices of others as a function of their own current choice.

Based on their beliefs about current and future actions by the other three players, sophisticated types calculate expected current and future payoffs as a function of their current actions (effort and, if a high ability type, help). Choices are reached via a logit decision rule where attractions are given by the sum of expected payoffs for the current and upcoming periods.²³

²³ To simplify the model, equal weight is put on current and future expected payoffs and the noise parameter is assumed to be the same for sophisticated and unsophisticated types.

1
2
3 Based on specific features of our data, described above, we added two features to the basic model.

4
5 1) Over-optimism: The basic model assumes sophisticated types have a correctly specified model of
6 how unsophisticated types form beliefs and then make choices. We depart from this assumption by
7 allowing for “over-optimism” by sophisticated types. Sophisticated types may believe that
8 unsophisticated types learn faster from experience (are more responsive to new information) than is the
9 case in reality, and therefore may be overly optimistic about their ability to affect the beliefs of
10 unsophisticated types. This is incorporated into the model by allowing the possibility that sophisticated
11 types overestimate the weight that unsophisticated players give to observed past behavior vs. their initial
12 beliefs. Specifically, sophisticated types calculating the beliefs of unsophisticated types replace $\omega(t)$ with
13 $\omega(t) + \mu$, where μ is an over-optimism parameter fit from the data.

14
15 2) Reciprocity: Reciprocity is an important feature of many environments and it seems plausible
16 that it could play a role here as well. When a high ability type increases the level of help, she (weakly)
17 helps the low ability types and harms herself. Likewise, decreasing help (weakly) harms the low ability
18 types and helps the high ability type. The preceding implies that increasing help can be seen as a kind
19 action by the high ability type while decreasing help can be seen as unkind. If the preferences of low
20 ability types incorporate reciprocity, they should want to reward the high ability type following an
21 increase in help and punish them after a decrease. This can be done within a weak link game via changes
22 in their effort levels since, in a weak-link game with stochastic choice, the expected payoffs of the other
23 players are a strictly increasing function of my own effort. From a dynamic point of view, reciprocity
24 can affect the model’s predictions because it accentuates responses to changing levels of help. We
25 therefore wanted to test for the presence of reciprocity and determine if the addition of reciprocity to the
26 model would improve its ability to capture the main features of Experiment 1.

27
28 To incorporate reciprocity into our model, we allow the attractions for a low ability type, as a
29 function of the low ability type's effort, to include weights on the expected payoff of the high ability type.
30 Positive weight is put on the high ability type’s expected payoff if he has been kind by increasing the
31 amount of help provided. The amount of weight put on the high ability type’s expected payoff in these
32 circumstances is given by the parameter κ which is fit from the data.²⁴ This is a simple and intuitive way
33 of adding reciprocity to the model, similar to the way reciprocity is added to the Charness and Rabin
34 (2002) model, and has a natural effect on the choices of low ability types. Low ability types try to help
35 (harm) high ability types who have helped (harmed) them by choosing a higher (lower) effort than they
36
37
38

39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

²⁴ For simplicity we assume that this weight is equal for positive and negative reciprocity. We doubt these weights are truly equal given experimental results to the contrary (i.e. Offerman, 2002) but allowing for greater weight on negative reciprocity would not affect the intuition underlying our results.

1
2
3 would otherwise select. Reciprocity is applied to attractions for both unsophisticated and sophisticated
4 low ability types.

5
6 *Model Results:* The full model's parameters were fit to data from the three treatments of interest, No
7 Help, Endogenous Help, and Forced Commitment, using the simulated method of moments. See
8 Appendix C for details on how the model was fit to the data, but fitting the model roughly involves the
9 following. Starting with a given set of parameters, we simulate the model for 1000 groups in all three
10 treatments. For each period in each treatment, we calculate the difference between the observed average
11 minimum effort and the simulated average minimum effort. Squaring these differences and then taking
12 the sum across periods and treatments gives us our measure of fit. The parameters are adjusted to
13 minimize the sum of squared differences. In other words, we are picking the parameters that most closely
14 track the observed minimum effort averaging across periods and treatments. The estimated parameters
15 and standard errors (in parenthesis) are presented in Table 5 below. All of the parameters are significantly
16 different from zero, except for γ , the time trend for the weight unsophisticated types' beliefs put on their
17 most recent experience, and, surprisingly, the reciprocity parameter κ . We estimate that a high proportion
18 (65.7%) of subjects are sophisticated types. The estimated fraction of pessimistic unsophisticated types is
19 quite low, 0.089, so most unsophisticated types are optimists. Finally, the over-optimism parameter is
20 significant and positive. Sophisticated types believe that unsophisticated types are putting significantly
21 more weight (about 23% more) on their most recent experience than is actually the case.
22
23
24
25
26
27
28
29
30
31
32

33 [Table 5 around here]

34
35 Using the parameters that best fit the data, we look at simulated data for each of the three
36 treatments. Figure 4 compares the average minimum efforts in the experimental (top left panel) and
37 simulated (top right panel) data. The simulated data reproduces the two main regularities of the
38 experimental data: Forced Commitment leads to an increase in minimum effort relative to No Help, and
39 but Endogenous Help does not.
40
41
42

43 [Figure 4 around here]

44
45 It is not true that *any* learning model fit to the data would reproduce the main features of the data.
46 A model's goodness of fit is being measured by the squared sum of differences between simulated and
47 observed minimum efforts *averaging across periods and treatments*. A model that fits well on average
48 can still fail to track the data from any one treatment or differences across treatments. This point is
49 illustrated by three alternative models to the data. The first only includes unsophisticated learners.
50 Mathematically, we impose a constraint that $\alpha = 0$, where α is the fraction of sophisticated types. The
51 second is the basic model with unsophisticated and sophisticated learners but no over-optimism or
52 reciprocity. Here the constraints on the full model are $\mu = 0$ and $\kappa = 0$, where μ and κ are the parameters
53
54
55
56
57
58
59
60

1
2
3 governing the strength of over-optimism and reciprocity respectively. The third adds reciprocity, but not
4 over-optimism, to the main model. The sole constraint imposed in this case is $\mu = 0$. The middle left
5 panel of Figure 4 shows simulation results for the model with only unsophisticated learners and the
6 middle right panel shows simulations for the basic model. In both cases, even though the models'
7 parameters are chosen to best fit the data, there is no discernable difference between the three
8 treatments.²⁵ The model with reciprocity but no over-optimism (lower left panel) does only slightly
9 better, with only tiny differences between the treatments.
10

11
12 Not only does reciprocity receive little weight in the fitted version of the full model, but
13 artificially adding reciprocity to the model harms its ability to track differences between treatments. To
14 show this, we fit a model constrained to have a higher level of reciprocity ($\kappa \geq 1$) than the low fitted value
15 ($\kappa = .012$). Simulated data for this model is shown in the lower right panel of Figure 4. It correctly has
16 average minimum effort higher in Forced Commitment than No Help, but has equally high average
17 minimum effort for Endogenous Commitment. In the experimental data, the average minimum effort for
18 Endogenous Commitment is significantly lower than in Forced Commitment.
19
20

21
22 To understand why sophistication and over-optimism receive significant weight in the fitted
23 version of the full model but reciprocity is not, consider how changing the learning model affects use of
24 help. Unsophisticated learners treat their group members as a fixed statistical distribution, failing to
25 recognize that changes in incentives will change their behavior. Since high ability types only benefit
26 from help if it affects the behavior of low ability types, unsophisticated learners do not anticipate any
27 benefits from help. Virtually no help is provided in the fitted model with only unsophisticated types --
28 only 0.2% of the high ability types ever use help in simulations of the Endogenous Help treatment. Given
29 that help doesn't get used, it follows that changing how help can be used won't matter. This leads to the
30 lack of differences between treatments shown in the middle-left panel of Figure 4.
31
32

33
34 In the basic model (i.e. no reciprocity or over-optimism) help still gets almost no use. In the
35 simulations, only 0.5% of the high ability types ever use help in the Endogenous Help treatment. Once
36 again, it follows that changing how help can be used will be irrelevant. Adding over-optimism gets two
37 important features. High ability types think low ability types will respond more to the provision of help
38 than they actually do, making them more willing to use help. The percentage of high ability types who
39 use help at least once in simulations of the Endogenous Help treatment rises to 65% in the full model.
40 But, because they over-estimate how much they have affected low ability types' beliefs after providing
41 help, over-optimistic high ability types are too willing to abandon help. Switches in help are almost
42 evenly split between increases and decreases (high ability types average 0.81 increases and 0.75
43
44
45
46
47
48
49
50
51
52
53
54
55

56
57 ²⁵ It may not appear that there are three lines in these two panels, but this is because they are on top of each other.
58
59
60

1
2
3 decreases) as high ability types who adopt help tend to rapidly abandon it, undercutting its positive effect.
4
5 With Forced Commitment, the high ability types can no longer abandon help, allowing minimum effort to
6
7 converge to higher levels in simulation of the full model. The mechanism underlying the differences
8
9 between treatments in simulations of the full model is similar to the mechanism in the experimental data.

10 Use of help also explains reciprocity plays little role in the ability of the model to track the
11
12 experimental data. Like over-optimism, reciprocity makes help more attractive since the positive
13
14 response to increasing help is larger. In the model with artificially high reciprocity ($\kappa \geq 1$) the percentage
15
16 of high ability types who use help at least once in simulations of the Endogenous Help treatment is still
17
18 relatively high at 37.5%. However, reciprocity also accentuates the negative response of low ability types
19
20 to a decrease in help. Since sophisticated learners model the decision process of other players and
21
22 therefore anticipate this effect, reciprocity makes decreasing help less attractive. In the model with
23
24 artificially high reciprocity there are far fewer decreases than increases (an average of 1.34 increases
25
26 versus 0.52 decreases). Reciprocity gets little weight in the full model because it undermines a critical
27
28 feature, premature decreases of help. It is worth noting that other modifications to the model that
29
30 accentuate the negative response to decreased help, such as adding inequality aversion to the utility
31
32 function, would likewise fail to improve the model's performance.

33 A pair of observations, one about the data and one about the simulations, makes it less surprising
34
35 that the full model puts little weight on reciprocity. The strong negative reaction to reductions in help in
36
37 Endogenous Help was an important piece of evidence for reciprocity, but this reaction is not present in the
38
39 Forced Commitment treatment. In groups where help is reduced prior to Round 21, the mean (and
40
41 median) change in effort by low ability types between Rounds 20 and 21 is 0. Given that reducing help
42
43 for the final ten periods is an irreversible action that harms the low ability types, reciprocity should be
44
45 highly relevant in this particular case. The data from Forced Commitment is more consistent with groups
46
47 having converged to equilibrium and hence not being destabilized by reduced help than reciprocity.

48 In the experimental data from the Endogenous Help treatment we documented an asymmetric
49
50 response of minimum effort by the low ability types to changes in the level of help (see Figure 3 and
51
52 accompanying text). Generating asymmetric responses to changes of help does not require reciprocity.
53
54 In simulations of the full model, where reciprocity is largely absent in the fitted parameters, the minimum
55
56 effort of the three low ability types increases by an average of 4.54 in response to an increase in the help
57
58 level and falls by an average of 10.00 in response to a decrease in the help level.

59 We do not claim that the structural model presented above perfectly tracks the experimental data
60
or provides the best possible fit. We have simplified the model in many ways to make it easier to

1
2
3 implement, and a more complex model could no doubt improve the fit.²⁶ Indeed, in future work we plan
4 to explore ways the model can be improved as well as experimental tests of the learning model.
5

6 To summarize, we learn three important things from fitting the learning model. (1) A relatively
7 simple structural model is able to track the most important features of the data from Experiment 1: adding
8 help does not increase minimum effort when high ability types are free to adjust help in Endogenous
9 Help, but does yield an increase when they are forced to commit in Forced Commitment. (2) It is not
10 true that any learning model fit to the data would track these features. Models without sophisticated
11 learners or with sophisticated learners but lacking over-optimism cannot reproduce the relative effects of
12 Endogenous Help and Forced Commitment.²⁷ (3) Adding reciprocity to the model, either alone or
13 coupled with over-optimism, does not help the model's ability to explain the main treatment effects. This
14 stems from the interaction between reciprocity and sophistication; reciprocity makes high ability types
15 who are sophisticated learners less willing to help since they correctly anticipate a strong negative
16 response.
17

18 *Conclusion 4: A model containing unsophisticated learners and sophisticated learners with over-*
19 *optimism captures key regularities of the experimental data. Over-optimism is necessary to capture the*
20 *main features of the data, but adding reciprocity to the model is not necessary*
21
22
23

24 **7. Experiment 2: Direct Help**

25 In this section we model help as being direct rather than indirect: high ability types provide help
26 by doing some of the work of low ability types. The new treatments parallel the Endogenous Help and
27 Endogenous Help with Commitment treatments. A low bonus rate ($B = 8$) is used in Rounds 1 – 10 to
28 trap groups at low effort levels, and then raised to $B = 14$ for Rounds 11 – 30. High ability types can help
29 low ability types in Rounds 11 – 30 by choosing a level of help before the four workers make effort
30 choices. The difference between Experiments 1 and 2 is how help is provided. In addition to working on
31 his own task (up to a maximum of 40 hours), the high ability type can also allocate time in 10 hour units
32 to the low ability types' tasks. To keep help efficiency neutral, the effort cost of a high ability type doing
33 the work of a low ability type is the same (9 ECU/hr) as it would be for a low ability type. Intuitively, we
34 model costs as reflecting the difficulty of the tasks rather than the ability of the workers, but high ability
35 types can work more hours than low ability types.²⁸ The same amount of help must be given to each low
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50

51 ²⁶ Alternatives we are exploring include making sophisticated types better at predicting the future, using more data
52 from past periods, and changing the specification to capture hysteresis.

53 ²⁷ The importance of sophisticated learners adds to existing evidence (notably Camerer, Ho, and Chong, 2002) of the
54 importance of sophisticated learners in tracking experimental data.

55 ²⁸ To keep the basic structure of the game fixed, we let the high ability type work more than 40 hours but restricted
56 the amount he could work on his own task.
57
58
59
60

ability type. When help is provided, the “effective” number of hours provided by a low ability worker is the sum of the hours received as help plus the hours worked directly by the low ability worker.

Table 6 shows the payoff tables for high and low ability types if one unit (10 hours) of help is provided.²⁹ In the payoff table for high ability types (left panel), the minimum effort level by *other* employees goes from 10 to 50 (instead of 0 to 40), since the high ability type has provided 10 hours of work to each of the low ability types. Similarly, in the payoff table for low ability types (right panel) the effective hours provided by a worker go from 10 to 50 reflecting the 10 hours of help. The minimum effective effort by other employees can still be zero since the high ability worker has the (perverse) option to not work at his own job.

[Table 6 around here]

Coordinating at effective effort levels 10, 20, 30, and 40 are Nash equilibria of the subgame shown in Table 6. There are four equilibria rather than five as previously, but the equilibria remain Pareto ranked, the efficient equilibrium remains 40, and the basic properties of a weak link game are preserved.

The two treatments with direct help, Direct Help and Direct Help with Forced Commitment, parallel the Endogenous Help and Forced Commitment treatments in Experiment 1. The mechanism through which forced commitment is expected to have a positive effect on average minimum effort is the same as for indirect help. Forced commitment prevents over-optimistic high ability (low effort cost) types from decreasing help before the equilibration process has finished. Note that the No Help treatment from Experiment 1 continues to serve as the baseline.

H4: In the long run, average minimum effort will be higher in Direct Help than in the No Help treatment and in Direct Help with Forced Commitment than in Direct Help.

We ran 6 additional sessions with 136 subjects participating in the additional two treatments (56 subjects in Direct Help and 80 in Direct Help with Forced Commitment). General procedures were identical to the ones described for Experiment 1 and subjects made an average of €19.04.

Figure 5 shows the average minimum effort for the two direct help treatments as well as the No Help treatment from Experiment 1. Note that this is *effective* minimum effort which includes any help provided by the high ability type as part of the effort for each low ability type. Compared to the No Help treatment, the effect of the Direct Help treatment is dramatic and *negative*. There is a slight increase between Rounds 10 and 11, but average minimum effort stabilizes below 5. The addition of commitment once improves performance as average minimum effort in the Direct Help with Forced Commitment

²⁹ Subjects were allowed to provide more units of help, but this is sufficiently expensive that we did not anticipate higher levels of help being used. On rare occasions we observed provision of two units of help.

1
2
3 treatment rapidly stabilizes in the neighborhood of 20. Adding forced commitment to direct help
4 overcomes the miserable performance without commitment but does not lead to an overall improvement.
5
6

7 [Figure 5 around here]
8

9 Regression analysis described in Appendix B confirms the obvious. Compared with the No Help
10 treatment, Direct Help leads to a statistically significant decrease in effective minimum effort. Direct
11 Help with Forced Commitment has no significant effect relative to No Help but significantly improves
12 effective minimum effort relative to the Direct Help treatment.
13
14

15 The extraordinarily low minimum effort levels in Direct Help are explained in part by factors
16 similar to those responsible for weak performance in the Endogenous Help treatment from Experiment 1.
17 A majority of the high ability types (8 of 14) try offering help at least once, and offering help leads to
18 improved performance. Following an increase in help, the average effective minimum help increases by
19 an average of 14.0. The problem is that the lack of commitment by high ability types is extreme. In 18 of
20 the 23 cases where help is increased, the increase is reversed in the following round. There is no case of
21 help being offered for more than three consecutive rounds. For every single case where help is decreased,
22 the minimum effort immediately returns to 0. The inability of high ability types in the Direct Help
23 treatment to stick with help undoes any positive effects of the initial increases. Average effective
24 minimum effort is actually slightly higher in the groups where help is never offered (4.2) than in groups
25 where help is offered at least once (3.0).
26
27

28 Cutting help does *not* cause effective minimum effort to collapse in Direct Help with
29 Commitment. In this treatment, 12 of the 20 high ability types offer positive levels of help in Round 11.
30 For Rounds 16 – 20, average minimum effective help for these twelve groups is 23.8 as compared to 7.5
31 when no help is provided. Help is eliminated in ten of the twelve groups for Rounds 21 – 30. The
32 response is negative but small. Minimum effort falls for six of the ten groups in Round 21, but only two
33 collapse to a minimum effort of zero. Comparing the average effective minimum efforts for Rounds 16 –
34 20 and Rounds 26 – 30, there is only a slight decrease from 22.6 to 18.8 for these ten groups. As with
35 indirect help, forced commitment with direct help allows groups time to converge to equilibrium so that a
36 later reduction in help is less disruptive.
37
38

39 *Conclusion 4: The effect of Direct Help is strongly negative, reducing minimum effort to almost the*
40 *lowest possible level. Forced commitment reverses this effect, but does not lead to minimum effort levels*
41 *above the No Help treatment.*
42
43

44 **8. Final remarks**

45 The purpose of this paper is to study whether high ability employees can use help as an effective
46 tool for leading groups out of performance traps. Help turns out to be a double-edged sword that can
47
48
49
50
51
52
53
54

1
2
3 harm as much as it helps. Increasing help leads to improved performance, but prematurely decreasing
4 help in an effort to take profits undoes any positive effects from the initial provision of help. Forcing
5 high ability types to commit to a stable level of help over time consistently improves the effect of help on
6 groups' performance in the turnaround game. A relatively simple structural model shows that the
7 importance of commitment can be attributed to over-optimism, the belief by sophisticated types that
8 learning by unsophisticated types is faster than is actually the case.
9

10
11 Our work suggests that it is not sufficient for managers to encourage help among workers (for
12 example by forming work teams as in Hamilton *et al.* (2003)). Management must also encourage stability
13 in how help is provided. This need not be overly complicated. Simply holding a meeting on a monthly
14 basis where employees discuss what they will do for the upcoming month, including what help they will
15 provide others, may serve as a useful device for fostering more stable commitments to help. Keeping
16 work groups together for an extended period of time rather than re-matching workers may also have the
17 effect of creating a more stable environment where levels of help are more likely to be steady.
18

19
20 Our work does not address endogenous commitment to help. It would be useful to know if high
21 ability types would take advantage of a commitment device if one was offered but not required. We
22 believe that they would not – if subjects understood the value of commitment, they probably wouldn't be
23 undercutting successful coordination in the first place – but the question is ultimately an empirical one.
24

25
26 Finally, there is a broader point to be taken away from our work. This paper focuses on the
27 effectiveness of help in escaping performance traps, but the importance of commitment is likely to be a
28 more general phenomenon. Escaping a performance trap is just one of many cases where leaders might
29 seek to overturn an existing norm in favor of a more socially desirable one, and help is just one of a
30 multitude of tools available to potential leaders. Establishing a new more desirable norm is a gradual
31 process, and one that can be disrupted by negative changes along many dimensions. As described in the
32 introduction, the same insights that this paper generates about excessive switching of help by high ability
33 workers also seem to apply to excessive switching of bonus rates by managers acting as leaders. We
34 imagine there are many examples where a self-defeating tendency to declare victory and cease costly
35 attempts at leadership undercuts the benefits of leadership. If our basic insights generalize, a good
36 strategy for individuals to follow in attempting to establish a new socially desirable norm is to pick an
37 approach and stick with it. If effective leadership shares a common element of strategic teaching, even
38 the best of approaches will fail if used inconsistently.
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

References

- 1
2
3
4
5 Antonetti, Marilyne and Alexandra Rufini (2008), "Social Norms, Coordination and Collaboration in
6 Heterogeneous Teams," Managerial and Decision Economics, **29**, 547-554.
7
8
9 Blume, Andreas and Andreas Ortmann (2007), "The Effects of Costless Pre-play Communication:
10 Experimental Evidence from Games with Pareto-ranked Equilibria," Journal of Economic
11 Theory, **132**, 274- 290.
12
13 Brandts, Jordi and Gary Charness (2003), "Truth or Consequences: An Experiment," Management
14 Science, **49**, 116-130.
15
16 Brandts, Jordi and David J. Cooper (2006), "A Change Would Do You Good: An Experimental Study of
17 How to Overcome Coordination Failure in Organizations," American Economic Review, **96**, 669-
18 693.
19
20 Brandts, Jordi and David J. Cooper (2007), "It's What You Say Not What You Pay: An Experimental
21 Study of Manager-Employee Relationships in Overcoming Coordination Failure," Journal of the
22 European Economic Association, **5**, 1223-1268.
23
24
25 Brown, Michelle and John S. Heywood (2009), "Helpless in Finance. The Cost of Helping Effort Among
26 Bank Employees," Journal of Labor Research, **30**, 176-195.
27
28 Cachon, Gerard and Colin Camerer (1996), "Loss-avoidance and Forward Induction in Experimental
29 Coordination Games," Quarterly Journal of Economics, **111**, 165-194.
30
31 Camerer, Colin (2003), Behavioral Game Theory: Experiments in Strategic Interaction, Princeton
32 University Press, Princeton, NJ.
33
34 Camerer, Camerer, T-H Ho, and J-K Chong (2002), "Sophisticated Learning and Strategic Teaching,"
35 Journal of Economic Theory, 104, 137-188.
36
37 Camerer, Colin F., and Teck Ho (1999), "Experience-Weighted Attraction in Games," Econometrica, **67**,
38 827-874.
39
40 Capra, Monica, Jacob Goeree, Rosario Gomez and Charles Holt (1999), "Anomalous Behavior in a
41 Traveler's Dilemma?," American Economic Review, **89**, 3, 678-690.
42
43 Charness, Gary and Matthew Rabin, "Understanding Social Preferences with Simple Tests," Quarterly
44 Journal of Economics, (2002), 817-869.
45
46 Chaudhuri, Ananish, Andrew Schotter and Barry Sopher (2006), "Talking Ourselves to Efficiency:
47 Coordination Conventions in an Inter-Generational Minimum Effort Game with Private, Almost
48 Common and Common Knowledge of Advice," Economic Journal, **119**, 91-122.
49
50 Cooper, David J. and John Kagel (forthcoming), "Other Regarding Preferences: A Selective Survey of
51 Experimental Results," in The Handbook of Experimental Economics, vol. 2, (J. Kagel and A.
52 Roth, eds.).
53
54
55
56
57
58
59
60

- 1
2
3 Crawford, Vincent and Bruno Broseta (1998), "What Price Coordination? The Efficiency-Enhancing
4 Effect of Auctioning the Right to Play," American Economic Review, **88**, 198-225.
5
6 Drago, Robert and Gerald T. Garvey (1998), "Incentives for Helping on the Job: Theory and Evidence,"
7 Journal of Labor Economics, 16, 1-25.
8
9 Fehr, Dietmar, Dorothea Kübler and David Danz (2012), "Information and Beliefs in a Repeated Normal-
10 form Game," Experimental Economics, **15**(4), 622-640.
11
12 Hamilton, Barton H., Jack A. Nickerson, and Hideo Owan (2003), "Team Incentives and Worker
13 Heterogeneity: An Empirical Analysis of the Impact of Teams on Productivity and Participation,"
14 Journal of Political Economy, **111**, 465-497.
15
16 Hamilton, Barton H., Jack A. Nickerson, and Hideo Owan (2004), "Diversity and Productivity in
17 Production Teams," working paper.
18
19 Hamman, John, Scott Rick and Roberto A. Weber (2007), "Solving Coordination Failure with "all-or-
20 none" Group-level Incentives," Experimental Economics, **10**, 285-303.
21
22 Harbring, Christine and Bernd Irlenbusch (2011), "Sabotage in Tournaments: Evidence from a Laboratory
23 Experiment," Management Science, **57**, 611-627.
24
25 Hyndman, Kyle, Antoine Terracol and Jonathan Vaksman (2009), "Learning and Sophistication in
26 Coordination Games," Experimental Economics, **12**, 450-472.
27
28 Hyndman, Kyle, Erkut Y. Ozbay, Andrew Schotter and Wolf Ze'ev Ehrblatt (2012), "Convergence: An
29 Experimental Study of Teaching and Learning in Repeated Games," Journal of the European
30 Economic Association, **10**, 573-604.
31
32 Kogan, Shimon, Anthony M. Kwasnica and Roberto A. Weber (2009), "Coordination in the Presence of
33 Asset Markets," American Economic Review, **101**(2), 927-947.
34
35 Liang, Kung-Yee, and Scott L. Zeger (1986), "Longitudinal Data Analysis Using Generalized Linear
36 Models," Biometrika, **73**, 13-22.
37
38 McKelvey, Richard M. and Thomas R. Palfrey (1995), "Quantal Response Equilibria in Normal Form
39 Games," Games and Economic Behavior, 7, 6-38.
40
41 Moulton, Brent (1986), "Random Group Effects and the Prevision of Regression Estimates," Journal of
42 Econometrics, **32**, 385-397.
43
44 Offerman, Theo (2002), "Hurting Hurts More Than Helping Helps," European Economic
45 Review, 46, 1423-37.
46
47 Terracol, Antoine and Jonathan Vaksman (2009) "Dumbing down rational players : Learning and
48 teaching in an experimental game," Journal of Economic Behavior and Organization, **70**(1-2),
49 54-71.
50
51
52
53
54
55
56
57
58
59
60

- 1
2
3 Van Huyck, John B., Raymond Battalio and Richard Beil (1990), "Tacit Coordination Games, Strategic
4
5 Uncertainty, and Coordination Failure," American Economic Review, **80**, 234-248.
6
7 Weber, Roberto A. and Colin Camerer (2003), "Cultural Conflict and Merger Failure: An Experimental
8
9 Approach," Management Science, 49, 400-415.
10
11 Weber, Roberto A. (2006), "Managing Growth to Achieve Efficient Coordination in Large Groups,"
12 American Economic Review, **96**, 114-126.
13
14 Zwick, Rami and Xiao-Ping (1999), "What Price Fairness? A Bargaining Study," Management Science,
15 45, 804-823.
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

APPENDIX A
INSTRUCTIONS FOR THE ENDOGENOUS HELP TREATMENT

I N S T R U C T I O N S

The purpose of this experiment is to study how individuals make decisions in certain contexts. The instructions are simple and if you follow them carefully you will at the end of the sessions confidentially obtain a sum of money in cash, given that nobody will know the payments received by any of the other participants. You can ask us at any time any questions you may have by raising your hand. Apart from these questions, any kind of communication between you is prohibited and may lead to immediate exclusion from the experiment.

1. For participating in this experiment you obtain an initial payment of 2500 ECUs (a virtual money unit). This experiment consists of various *blocks* of 10 *rounds* each. The conditions of the experiment are identical within each block and although some of the conditions can vary between on block and the next, you will at each point informed about whatever changes are introduced. In each round you are part of the same group of 4 participants. The composition of these groups which are called *firms* will not vary during the experiment. Given that nobody will know the identity of the members of each group, all the actions you take during the experiment will be absolutely anonymous.
2. As as worker of your firm you have to decide how to split your working week of 40 working hours between two activities: Activity A and Activity B. Given that the hours that you don't assign to one of the activities is automatically assigned to the other we will just ask you that you decide how many weekly hours you spend on Activity A. Given that the available options are to spend 0, 10, 20, 30 or 40 hours on Activity A, this automatically implies spending 40, 30, 20, 10 or 0 on Activity B (respectively).
3. In each round you will receive a base salary of 200 ECUs. Your payoff will depend on this base salary, on your decisions and on a *bonus* that depends on the *minimum* number of hours spend on Activity A by any of the members of your group. This *bonus* is chosen by the central server of the lab, which simulates the decisions of the firm's management. Before making any decision, you will receive information about the value of this *bonus*.
4. Your earnings in each round depend on the hours you spend on Activity A, on the number of hours by the other workers in your firm on Activity A, on the *bonus* B selected by management and on C_e a variable that represents your cost of effort at work. Your earnings are given by the following formula:

$$\text{Earnings of worker} = \begin{array}{rcccl} \text{Base salary} & - & \text{Effort} & + & \text{Bonus} \\ 200 & - & (C_e \times H_i) & + & (B \times \min(HA)) \end{array}$$

where H_i is the number of hours that you spend on Activity A and $\min(HA)$ is the minimum number of hours that a worker of your firm spends on Activity A. You don't need to memorize this formula: the computer will show you an earnings table whenever you have to make a decision.

5. In each firm there are two types of workers: 1 worker of type 1 and 3 workers of type 2. At the beginning of the session you will be informed of what type of worker during the whole experiment, since types will not vary. The two types of workers will differ exclusively in their *initial* effort cost (C_e), which is higher in on the types relative to the other type. During the first 10 rounds:
 - a. Workers of type 1 have an effort cost of $C_e = 1$ and their earnings are calculated according to the following formula:

$$\text{Earnings of type 1} = 200 - 1 \times H_i + B \times \min(HA)$$
 - b. Workers of type 2 have an effort cost of $C_e = 9$ and their earnings are calculated according to the following formula:

$$\text{Earnings of type 2} = 200 - 9 \times H_i + B \times \min(HA)$$
6. Although none of the participants will receive earnings as manager of the firm, you will be shown information about its earnings, which depend on the hours spent on Activity A by the workers of the firms and on the *bonus* B , as shown in the following formula (which again you don't need to memorize):

$$\text{Firm profit} = 100 + (52 - 4 \times B) \times \min(HA)$$
7. In each round, the computer will show you a table similar to those that are shown below. The payoffs that are shown adjust to the (variable) values of B . In the examples that follow we use $B = 8$. Note that this

information is shown in the upper part of the payoff table. For players of **type 1** with an effort cost equal to 1 the payoff table is the following:

$$B = 8$$

$$\text{Firm profit} = 100 + 20 \times \min(HA)$$

		Minimum of hours spent on Activity A by the other workers				
		0	10	20	30	40
My hours spent on Activity A	0	200	200	200	200	200
	10	190	270	270	270	270
	20	180	260	340	340	340
	30	170	250	330	410	410
	40	160	240	320	400	480

For players of **type 2**, with an effort cost equal to 9, the payoff table is:

$$B = 8$$

$$\text{Firm profit} = 100 + 20 \times \min(HA)$$

		Minimum of hours spent on Activity A by the other workers				
		0	10	20	30	40
My hours spent on Activity A	0	200	200	200	200	200
	10	110	190	190	190	190
	20	20	100	180	180	180
	30	-70	10	90	170	170
	40	-160	-80	0	80	160

8. Each worker chooses the number of hours spent on Activity A using the buttons located on the right side of the screen. You can change your choice as often as you want, but once you click "OK" this decision is irreversible.
9. When you make your choice you will not know what the other workers of your firm have done, but, after each round you will be informed of the number of hours that the workers of your firm have spent on Activity, the profit of your firm, your earnings in that round and your accumulated earnings. You will also be able to see a summary of results from previous rounds.
10. At the end of the experiment, your accumulate earnings will be converted from ECUs to euros at an exchange rate of 500 ECUs = €1.

Example

Before starting the experiment, please go over the following questions. For each of the questions assume that $B = 8$ and that you are a worker of type 1 with an effort cost of 1.

Using your earnings table shown on the previous page suppose that you choose spending 10 hours on Activity A and that the other workers of your firm decide to spend 30, 20 and 40 hours on Activity A.

In this case:

The minimum number of hours that a worker of your firm spends on Activity A is 10.

Your earnings are 270 ECUs.

Now consider that you are a worker of type 2 with an effort cost equal to 9. The relevant earnings table would be the second on the previous page. Suppose that you choose to spend 20 hours on Activity A and that the other workers of your firm spend 30, 20 and 40 hours on Activity A.

In this case:

The minimum number of hours that a worker of your firm spends on Activity A is 20.

Your earnings are 180 ECUs.

A little test

Imagine that you are a worker of type 1 (effort cost = 1). Using your earnings table suppose that you choose to spend 20 hours on Activity A. The other workers of your firm choose to spend 30, 0 and 20 hours on Activity A.

The minimum number of hours that a worker of my firm spends on A is _____.

My earnings in ECUs are _____.

Suppose you decide to spend 0 hours on Activity A. The other workers of your firm choose to spend 20, 30 and 10 hours on Activity A.

The minimum number of hours that a worker of my firm spends on A is _____.

My earnings in ECUs are _____.

I will be in the same firm during the whole experiment: True. False.

My actions and earnings will be confidential: True. False.

Imagine that you are a worker of type 2 (effort cost = 9). Using your earnings table suppose that you choose to spend 20 hours on Activity A. The other workers of your firm choose to spend 30, 0 and 20 hours on Activity A.

The minimum number of hours that a worker of my firm spends on A is _____.

My earnings in ECUs are _____.

Suppose you decide to spend 0 hours on Activity A. The other workers of your firm choose to spend 20, 30 and 10 hours on Activity A.

The minimum number of hours that a worker of my firm spends on A is _____.

My earnings in ECUs are _____.

I will be in the same firm during the whole experiment: True. False.

My actions and earnings will be confidential: True. False.

From this round on there are two changes in the experiment: (i) the value of the *bonus* becomes 14 and (ii) the worker of type 1 in each group has the possibility of reducing the effort costs of the workers of type 2 of his group by increasing his own effort cost. The type 1 worker will make this decision before the effort decisions for the round are made. For each unit by which he reduces the effort costs of the three workers of type 2, the effort cost of the type 1 worker will increase by 3 units. The reduction will have to be the same for all workers of type 2 in the group, so that the concrete options available to the type 1 worker will be the following:

	Effort cost for type 1	Effort cost for type 2
Help level 0	1	9
Help level 1	4	8
Help level 2	7	7
Help level 3	10	6
Help level 4	13	5

The possibilities of modifying the effort costs lead to different payoff tables. Given that the possible effort costs are 1, 4, 5, 6, 7, 8, 9, 10 and 13 below we show the payoff tables for each of the possible effort costs:

$$C_e = 1$$

$$B = 14$$

$$\text{Firm profit} = 100 + -4 \times \min(HA)$$

Minimum of hours spent on Activity A by the other workers

		0	10	20	30	40
My hours spent on Activity A	0	200	200	200	200	200
	10	190	330	330	330	330
	20	180	320	460	460	460
	30	170	310	450	590	590
	40	160	300	440	580	720

$$C_e = 4$$

$$B = 14$$

$$\text{Firm profit} = 100 + -4 \times \min(HA)$$

Minimum of hours spent on Activity A by the other workers

		0	10	20	30	40
My hours spent on Activity A	0	200	200	200	200	200
	10	160	300	300	300	300
	20	120	260	400	400	400
	30	80	220	360	500	500
	40	40	180	320	460	600

$C_e = 5$

$B = 14$

Firm profit = $100 + -4 \times \min(HA)$

Minimum of hours spent on Activity A by the other workers

		0	10	20	30	40
My hours spent on Activity A	0	200	200	200	200	200
	10	150	290	290	290	290
	20	100	240	380	380	380
	30	50	190	330	470	470
	40	0	140	280	420	560

$C_e = 6$

$B = 14$

Firm profit = $100 + -4 \times \min(HA)$

Minimum of hours spent on Activity A by the other workers

		0	10	20	30	40
My hours spent on Activity A	0	200	200	200	200	200
	10	140	280	280	280	280
	20	80	220	360	360	360
	30	20	160	300	440	440
	40	-40	100	240	380	520

$C_e = 7$

$B = 14$

Firm profit = $100 + -4 \times \min(HA)$

Minimum of hours spent on Activity A by the other workers

		0	10	20	30	40
My hours spent on Activity A	0	200	200	200	200	200
	10	130	270	270	270	270
	20	60	200	340	340	340
	30	-10	130	270	410	410
	40	-80	60	200	340	480

$C_e = 8$

$B = 14$

Firm profit = $100 + -4 \times \min(HA)$

Minimum of hours spent on Activity A by the other workers

		0	10	20	30	40
My hours spent on Activity A	0	200	200	200	200	200
	10	120	260	260	260	260
	20	40	180	320	320	320
	30	-40	100	240	380	380
	40	-120	20	160	300	440

$C_e = 9$

$B = 14$

Firm profit = $100 + -4 \times \min(HA)$

Minimum of hours spent on Activity A by the other workers

		0	10	20	30	40
My hours spent on Activity A	0	200	200	200	200	200
	10	110	250	250	250	250
	20	20	160	300	300	300
	30	-70	70	210	350	350
	40	-160	-20	120	260	400

$C_e = 10$

$B = 14$

Firm profit = $100 + -4 \times \min(HA)$

Minimum of hours spent on Activity A by the other workers

		0	10	20	30	40
My hours spent on Activity A	0	200	200	200	200	200
	10	100	240	240	240	240
	20	0	140	280	280	280
	30	-100	40	180	320	320
	40	-200	-60	80	220	360

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

$C_e = 13$

$B = 14$

Firm profit = $100 + -4 \times \min(HA)$

Minimum of hours spent on Activity A by the other workers

		0	10	20	30	40
My hours spent on Activity A	0	200	200	200	200	200
	10	70	210	210	210	210
	20	-60	80	220	220	220
	30	-190	-50	90	230	230
	40	-320	-180	-40	100	240

APPENDIX B

Treatment Effects: The regressions shown in Table B.1 provide a formal statistical backing for our conclusions about the effects of the various treatments on subject choices. Models 1 and 1a look at firm-level data with the minimum effort as the dependent variable while Models 2 and 2a look at employee-level data with effort as the dependent variable. Given that these dependent variables are ordered categories, use of an ordered probit specification is natural.³⁰ Both data sets include all observations from all seven treatments for Rounds 11 – 30. Since there are four employees per firm, Model 2 has four times as many observations as Model 1. The independent variables include a dummy for Rounds 21 – 30 and interaction terms between treatment dummies and dummies for Rounds 11 – 20 and Rounds 21 – 30. These interaction terms capture differences between the treatment and the base for the time period in question. Models 1 and 1a differ only in what treatment is used as the base – the No Help treatment in Model 1 and the Endogenous Help treatment in Model 1a. Models 2 and 2a differ in the same manner.

[Table B.1 about here]

Since there are multiple observations from each firm, correcting for firm effects is a central issue in the regressions. We do this in a couple of ways. First, standard errors in both regressions are corrected for clustering (Moulton, 1986; Liang and Zeger, 1986) at the firm level.³¹ Second, even though effort level zero is the unique equilibrium choice in Rounds 1 – 10, the majority of employees (60%) choose positive effort at least once in this first phase. We therefore controls for behavior in Rounds 1 – 10 with the idea that this will pick up an individual and/or group tendency to choose higher effort levels. Specifically we include the firm's average minimum effort in Rounds 1 – 10 and, in the regressions on employee-level data, the employee's average effort over Rounds 1 – 10.

We start by looking at the statistical support for Conclusion 1. Looking at Models 1 and 1a, few significant differences are observed at the firm-level between the three initial treatments (No Help, Symmetric Costs, and Endogenous Help). The only significant difference is between the Endogenous Help and Symmetric Costs treatments in Rounds 21 – 30, and this is only significant at the 10% level. Looking at employee-level data, in Models 2 and 2a, even this difference fails to achieve statistical significance. The regression results therefore support Conclusion 1 – Endogenous Help does little to improve (minimum) effort levels and the positive effect of Symmetric costs is modest.

³⁰ The cutoffs are not reported on Table B.1 since they are of little economic interest. Copies of all regression output are available from the authors upon request.

³¹ In Models 2 and 2a, observations from individuals in the same firm are obviously not independent. We therefore cluster at the firm level rather than the employee level.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50

Turning to Conclusion 3, the regression analysis reported in Table B.1 also supports this conclusion. Compared with either the No Help or Endogenous Help treatments, the Forced Commitment treatment leads to significantly higher minimum effort and effort levels in Rounds 21 – 30. This effect is significant at the 5% level for all four models.³² If the regressions are altered so the base is the Symmetric Costs treatments, no significant differences are found between this treatment and Forced Commitment.

The regressions provide strong support for Conclusion 4. Compared with the No Help treatment, Direct Help without Commitment leads to significantly lower minimum effort and effort throughout Rounds 11 – 30. This effect is always significant at the 1% level. In contrast the effect of Direct Help with Commitment (as compared with the No Help treatment) is always small and never approaches statistical significance at the 10% level. The difference between the two direct help treatments is always statistically significant at the 1% level.

Effects of Changing Help: Table B.2 shows the results of ordered probit regressions capturing the relation between changes in help and changes in the effort level. These provide backing for Conclusion 2. The dataset is all observations where help was available (Rounds 11 – 30) for the Endogenous Help treatment. Model 1 uses firm-level data with the dependent variable being the minimum effort chosen by the firm's three low ability types in the current round. Model 2 uses employee-level data from low ability types only. The dependent variable is the chosen effort level for the current round. As should be clear from Figures 2 and 3, the current round and the lagged (minimum) effort both play an important role in determining how the (minimum) effort changes beyond any change in the level of help provided. To control for these effects, both regressions include dummies for the current round as well as dummies for each possible value of the lagged dependent variable.³³ Parameter estimates for these dummies and the cutoffs are not reported in Table B.2 since they are not of direct interest. Copies of the full regression output are available from the authors upon request. With the inclusion of lagged dependent variables, the estimated parameters capture how the independent variables affect *changes* in the (minimum) effort. The primary independent variables are the current level of help, the *increase* in help from the previous round interacted with a dummy for a positive change in help, and the *decrease* in help from the previous

51
52
53
54
55
56
57
58
59
60

³² As a simple alternative to running ordered probits, we can calculate non-parametric tests where the unit of observation is average minimum effort for a group in Rounds 26 – 30 (e.g. the late rounds where convergence has taken place). Using a Mann-Whitney test, a relatively low powered test, the difference between No Help and Forced Commitment is significant at the 10% level and the difference between Endogenous Help and Forced Commitment is significant at the 5% level.

³³ Inclusion of round dummies also controls for the effect of the bonus rate increase prior to Round 11. To avoid colinearity, the dummies for Round 11 and lagged (minimum) effort = 0 are dropped.

1
2
3 interacted with a dummy for a negative change in help.³⁴ Standard errors are corrected for clustering at
4 the firm level.
5

6
7 **[Table B.2 about here]**

8 Looking at Model 1, the two coefficients for changes in the level of help are both strongly
9 significant, while the coefficient for the current level of help is not. The negative coefficient for negative
10 help changes is almost three times as large – in absolute terms - as the positive one for positive help
11 changes. The difference between the absolute values of these two parameters is statistically significant at
12 the 1% level. The main results for Model 2 are much the same: the level of help does not have a
13 significant effect and the negative effect of decreasing help is roughly three times as large as the positive
14 effect of increasing help with the difference being statistically significant at the 1% level. One possible
15 explanation for the asymmetric response of minimum effort to changes in help is strictly mechanical –
16 three individuals need to respond positively to increased help to raise the minimum effort for low ability
17 types, but only one needs to to respond negatively to decreased help to lower the minimum effort. The
18 results of Model 2 indicate that the disproportionate response to decreases is not caused by the minimum
19 function emphasizing decreases.
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55

56 ³⁴ To be clear, decreases are being measured as positive numbers. For example, if a high ability type decreases help
57 from 2 to 0, the variable “Decrease in Help * (Change in Help < 0)” equal 2, not -2.
58
59
60

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

Table 1: Employee i's Payoff Tables for B=8

$C_i = 1$

$C_i = 9$

		Minimum Effort by Other Employees				
		0	10	20	30	40
Effort by Employee i	0	200	200	200	200	200
	10	190	270	270	270	270
	20	180	260	340	340	340
	30	170	250	330	410	410
	40	160	240	320	400	480

		Minimum Effort by Other Employees				
		0	10	20	30	40
Effort by Employee i	0	200	200	200	200	200
	10	110	190	190	190	190
	20	20	100	180	180	180
	30	-70	10	90	170	170
	40	-160	-80	0	80	160

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

Table 2: Employee i's Payoff Tables for B=14

$C_i = 1$

$C_i = 9$

		Minimum Effort by Other Employees				
		0	10	20	30	40
Effort by Employee i	0	200	200	200	200	200
	10	190	330	330	330	330
	20	180	320	460	460	460
	30	170	310	450	590	590
	40	160	300	440	580	720

		Minimum Effort by Other Employees				
		0	10	20	30	40
Effort by Employee i	0	200	200	200	200	200
	10	110	250	250	250	250
	20	20	160	300	300	300
	30	-70	70	210	350	350
	40	-160	-20	120	260	400

Table 3: List of Treatments of Experiment 1

TREATMENT NAME	NO HELP	ENDOGENOUS HELP	SYMMETRIC COSTS	FORCED COMMITMENT
Player 1 Cost Block 1	1	1	1	1
Player 1 Cost Blocks 2 and 3	1	$1 + 3H$ $H \in \{1, 2, 3, 4\}$	7	$1 + 3H$ (H fixed for block) $H \in \{1, 2, 3, 4\}$
Players 2, 3 & 4 Cost Block 1	9	9	9	9
Players 2, 3 & 4 Cost Blocks 2 and 3	9	$9 - H$ $H \in \{1, 2, 3, 4\}$	7	$9 - H$ (H fixed for block)
Number of Firms (Number of Sessions)	20 (2 Sessions)	39 (5 Sessions)	20 (2 Sessions)	20 (2 sessions)

Table 4: Employee Payoff Tables with $B = 14$ and $C_i = 7$ for all Employees

		Minimum Effort by Other Employees				
		0	10	20	30	40
Effort by Employee i	0	200	200	200	200	200
	10	130	270	270	270	270
	20	60	200	340	340	340
	30	-10	130	270	410	410
	40	-80	60	200	360	480

Table 5: Parameter Estimates of the Structural Model

PARAMETER	EXPLANATION	ESTIMATE
μ	Over-optimism	0.046 (0.004)
κ	Reciprocity	0.012 (0.045)
α	Fraction of sophisticated players	0.657 (0.057)
θ	Fraction of type 1 players	0.089 (0.005)
$\bar{\omega}$	Initial belief weight	0.202 (0.008)
δ	Hysteresis	0.303 (0.023)
γ	Belief weight slope term	0.178×10^{-4} (1.175×10^{-4})
ρ	Reset of belief weight	0.086 (0.006)
λ	Logit function parameter	0.072 (0.002)
λ^H	Logit parameter for help decisions	0.410 (0.137)
-	Sum of squared errors	523.861

Table 6: Payoff Tables, High Ability Type Provides One Unit (10 Hours) of Help

		$C_i = 1$					$C_i = 9$					
		Minimum Effective Effort by Other Employees					Minimum Effective Effort by Other Employees on Own Job					
		10	20	30	40	50	0	10	20	30	40	
Effort by Employee 1 at Own Job	0	-70	-70	-70	-70	-70	10	200	340	340	340	340
	10	60	60	60	60	60	20	110	250	390	390	390
	20	50	190	190	190	190	30	20	160	300	440	440
	30	40	180	320	320	320	40	-70	70	210	350	490
	40	30	170	310	450	450	50	-160	-20	120	260	400

Table B.1: Ordered Probit Regressions on Treatment Effects

	Model 1	Model 1a	Model 2	Model 2a
Dependent Variable	Minimum Effort	Minimum Effort	Effort	Effort
Unit of Observation	Firm	Firm	Employee	Employee
Base Treatment	No Help	Endogenous Help	No Help	Endogenous Help
Rounds 21 – 30	.197 (.129)	.265** (.106)	-.030 (.138)	.053 (.105)
No Help * Rounds 11 – 20		.103 (.262)		.035 (.236)
No Help * Rounds 21 – 30		.035 (.315)		-.048 (.292)
Endogenous Help * Rounds 11 – 20	-.103 (.262)		-.035 (.236)	
Endogenous Help * Rounds 21 – 30	-.035 (.315)		.048 (.292)	
Symmetric Costs * Rounds 11 – 20	.185 (.288)	.288 (.258)	.102 (.287)	.137 (.255)
Symmetric Costs * Rounds 21 – 30	.496 (.365)	.532* (.315)	.486 (.351)	.438 (.304)
Forced Commitment * Rounds 11 – 20	.371 (.329)	.474 (.297)	.339 (.313)	.375 (.278)
Forced Commitment * Rounds 21 – 30	.781** (.396)	.816** (.343)	.799** (.371)	.751** (.322)
Direct Help * Rounds 11 – 20	-1.032*** (.261)	-.929*** (.218)	-.818*** (.246)	-.783*** (.197)
Direct Help * Rounds 21 – 30	-1.556*** (.377)	-1.521*** (.319)	-1.380** (.310)	-1.428*** (.244)
Direct Help with Commitment * Rounds 11 – 20	-.003 (.253)	.100 (.220)	.010 (.232)	.045 (.192)
Direct Help with Commitment * Rounds 21 – 30	-.070 (.328)	-.034 (.273)	.094 (.298)	.046 (.244)
Ave. Min. Effort Rounds 1 – 10	.170 (.109)	.170 (.109)	.176* (.105)	.176* (.105)
Average Effort, Rounds 1 – 10			.008 (.006)	.008 (.006)
Log-Likelihood	-3686.89	-3686.89	-14584.60	-14584.60
Number of Observations	2660	2660	10640	10640

Notes: Standard errors are corrected for clustering at the firm level. Three (***) , two (**), and one (*) stars indicate statistical significance at the 1%, 5%, and 10% respectively.

Table B.2: The Effect of Changes in Help on Effort by Low Ability Types

	Model 1	Model 2
Dependent Variable	Firm-Level Data Min. Effort, Low Ability Types	Employee-Level Data Effort, Low Ability Types
Current Level of Help	.044 (.062)	.045 (.066)
Increase in Help * (Change in Help > 0)	.218*** (.082)	.124** (.060)
Decrease in Help * (Change in Help < 0)	-.604*** (.137)	-.377*** (.059)
Log Likelihood	-674.65	-2416.30
Number of Observations	780	2340

Notes: Standard errors are corrected for clustering at the firm level. Three (***), two (**), and one (*) stars indicate statistical significance at the 1%, 5%, and 10% respectively.

FIGURE 1. MINIMUM EFFORT




FIGURE 2. FREQUENCY OF CHANGING HELP, MINIMUM EFFORT, AND HIGH ABILITY TYPE PROFITS

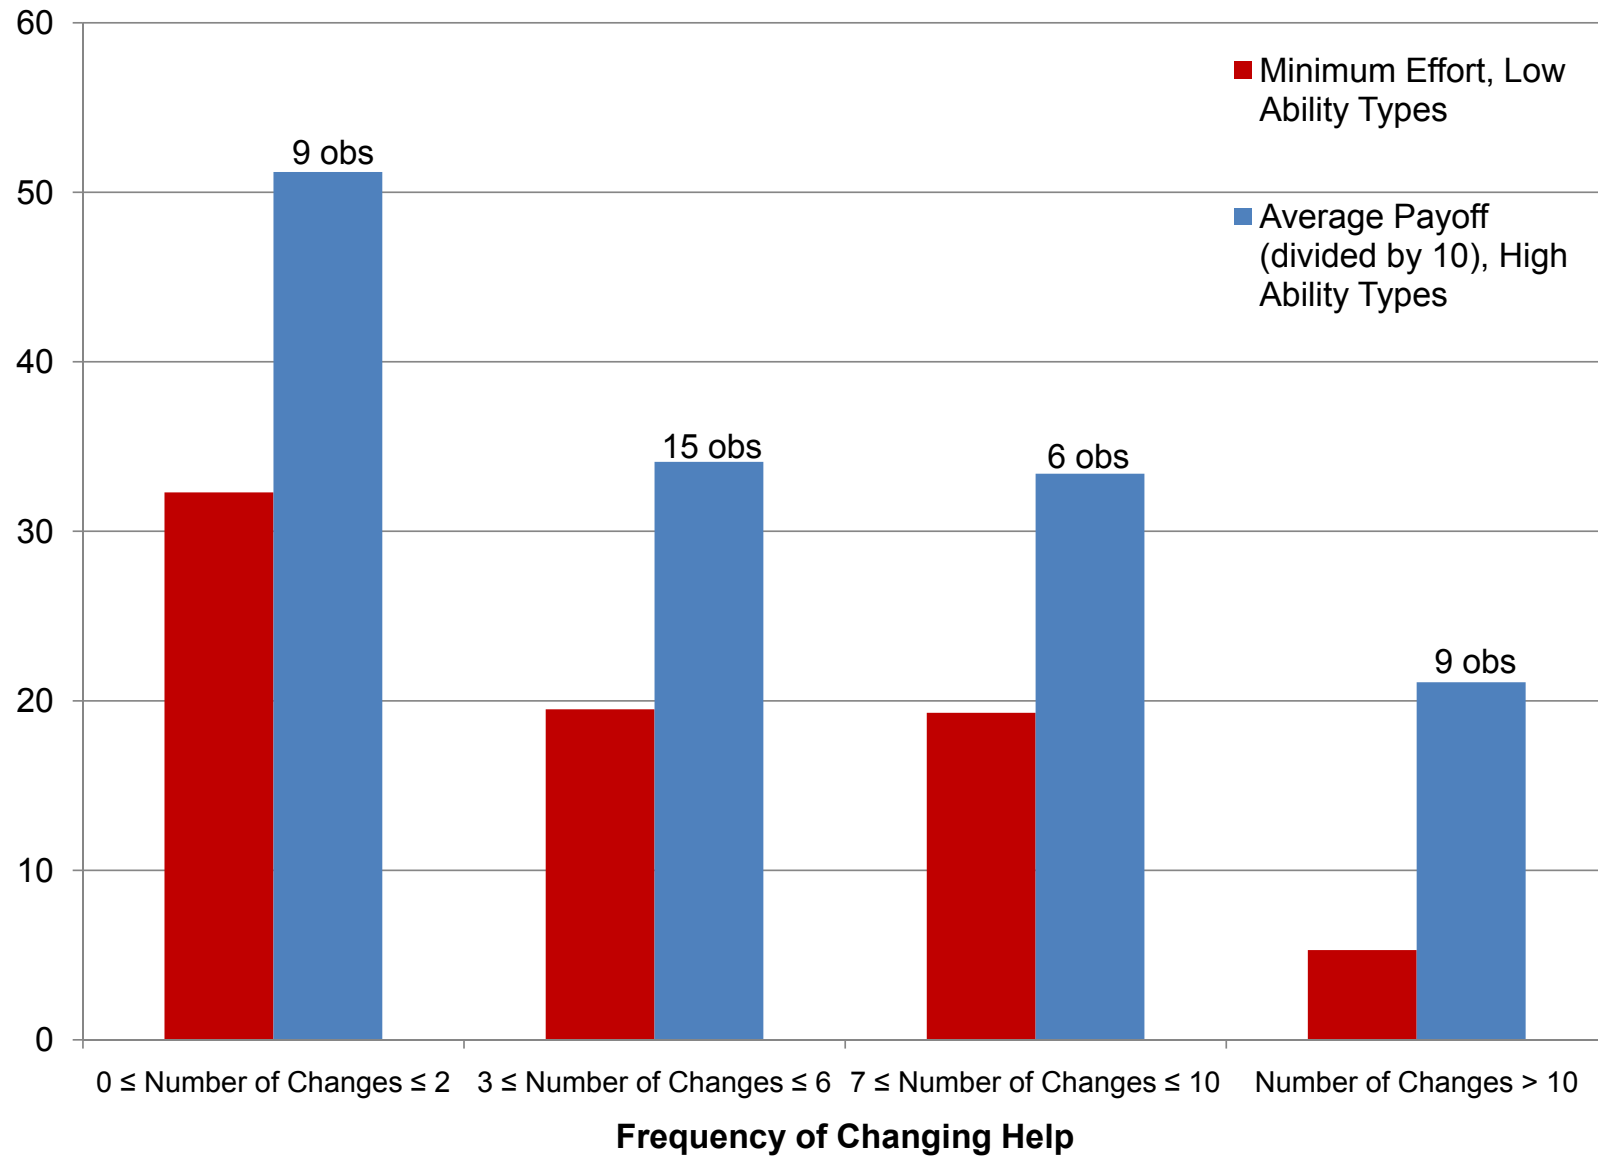


FIGURE 3. EFFECT OF CHANGING HELP ON LOW ABILITY TYPE MINIMUM EFFORT (ROUNDS 12-30)

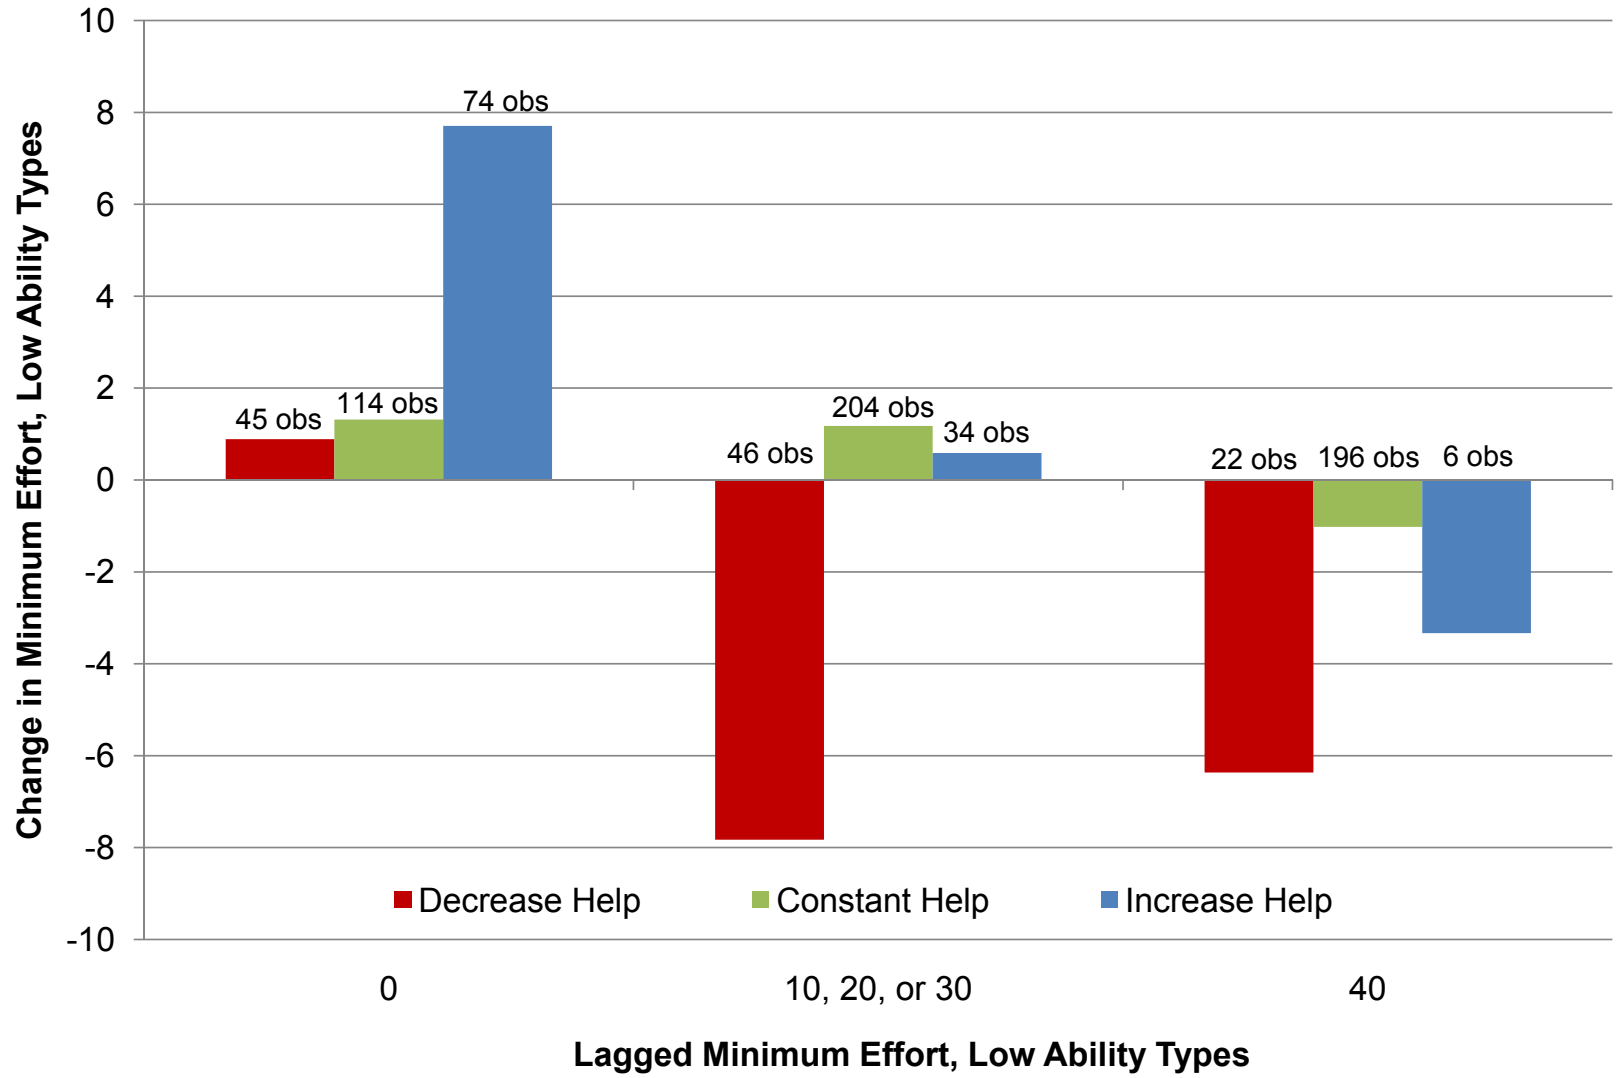


FIGURE 4. COMPARISON OF EXPERIMENTAL AND SIMULATED DATA

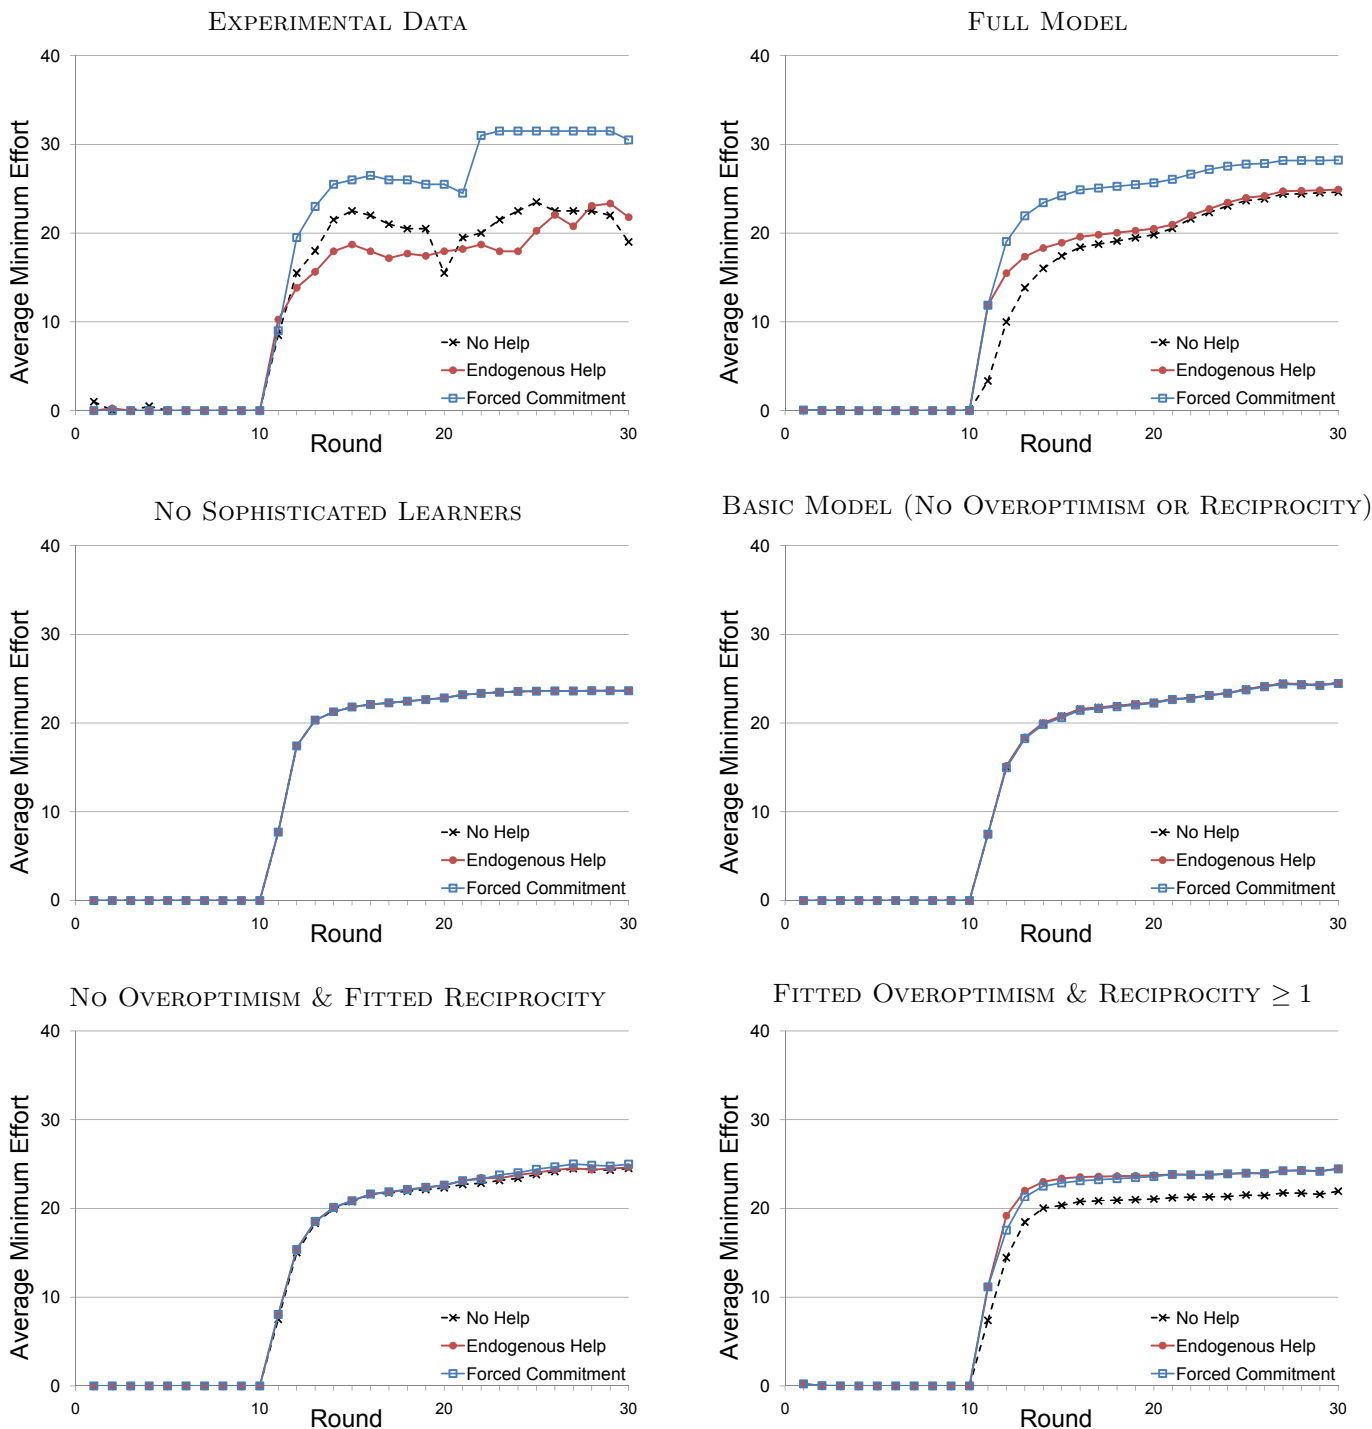



FIGURE 5. MINIMUM EFFORT IN EXPERIMENT 2



1
2
3 Responses to Associate Editor
4
5

6 Thank you for all of the effort you and the review team have put into this paper. We realize that this went
7 far beyond the ordinary and genuinely appreciate it. Please find responses to your comments below.
8

9 Note: Excerpted material from the letter is in normal print and our responses are in italics.
10

11 First, and most importantly, you should add a couple of paragraphs in Section 6 (and other appropriate
12 places in the paper) that motivates the use of your structural model. You should present *ex ante* arguments
13 why such a structural model is promising/necessary/desirable in the first place and adds value to the
14 paper. At appropriate places you should address why this is not just a fitting exercise. The fact that three
15 members of your reviewer team, who all spent considerable time on your paper, were unconvinced by the
16 value added of the structural model, should give you an incentive to sell your model better. Your letter of
17 appeal contains many good arguments that have helped me to better understand the rationale of the
18 model.
19

20
21
22
23 *We completely get your point. It doesn't matter how good we think the model if we cannot communicate*
24 *its value to readers. We have rewritten both the discussion of the model in the introduction as well as the*
25 *material in Section 6. Given your comment about the letter of appeal, much of the revision consists of*
26 *taking the arguments from the letter of appeal and incorporating them into the paper. We have tried to*
27 *improve our explanation of the model, how the fitting and simulation exercises work, and what is learned*
28 *from the model. We hope the paper now sells the model.*
29
30

31
32 Second, you should pay attention to the comments by Reviewer 2 and address them as much as possible.
33

34 *Responses to Referee 2's comments are below. The most important is that, as suggested, we have*
35 *shortened the discussion of Experiment 2. This, along with tightening the exposition elsewhere, allowed*
36 *us to explain the learning model in more detail without changing the length of the paper.*
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

Responses to Referee 2

Note: Excerpted material from the referee report is in normal print and our responses are in italics.

1) As I mentioned in my previous report the results seem not to be very pronounced. Because of the small number of observations (20 in most of the treatments) the findings should be backed by non-parametric statistics in addition to regression results. However, I could not find this backing in the results section, presumably because non-parametrically there are few differences between the treatments.

Added. Results with non-parametrics parallel the regression results and are reported with them.

2) The EWA-like model presented in the paper is not very convincing and the gained insights are limited. To me its assumptions and its modelling are so specific that it looks as if the authors selected it from a post-hoc process among other possible (maybe less supporting) models and specifications.

We have rewritten the material describing the learning model. We hope this makes it clearer what we learn from using this model, why the model is constructed as it is, and why the reciprocity term in the model is not statistically significant.

3) In my view the authors are not able to convincingly pinpoint the causal reasons why the Forced Commitment treatment is more successful in establishing higher average minimum effort levels.

We have tried to clarify why this treatment increases minimum effort level.

4) The paper could be more focused. For example, the Symmetric Cost treatment does not add much to the main story and the contributions of Experiment 2 (the two treatments with direct help) to the commitment argument could be presented in a shorter and more concise way.

We agree that the most important treatments are the No Help, Endogenous Help, and Commitment treatments and have tried to keep the focus on these three treatments. In revising, our primary goal has been to better explain the learning model, but we have also tightened the exposition (especially in the discussion of Experiment 2 as per Referee 2's comment).

In my view Conclusion 4 is far too strong: "A model of sophisticated ... best explains...". It is far from clear that the chosen model best explains the experimental data.

We have rewritten the material describing the learning model. We hope this makes it more clear what we learn from using this model, why the model is constructed as it is, and why the reciprocity term in the model is not statistically significant.

Minor comments

Footnote 1: It should be "Drago and Garvey (1998)" and not "Drago and Harvey (1998)"

Fixed.

p. 5: The sentence "Adding reciprocity to the model therefore decreases..." is very difficult to grasp.

Rewritten to, hopefully, be clearer.

1
2
3 p. 7: In my printout it reads $E_i \in \{0, 10, 20, 30, 40\}$.
4

5 *Fixed.*
6

7
8 p. 9: What does the sentence “Finally, both indirect and direct help are natural in the context of escaping
9 performance traps.” exactly mean? Is there any evidence for this?

10
11 *This has been rephrased. Hopefully it should be clearer that the types of help that can be offered in the
12 lab experiment can be easily interpreted in terms of types of help that can be offered in a field setting.*
13

14 p. 18: “are equally likely to be” instead of “are equally like to be”
15

16 *Fixed.*
17

18 p. 19: What is the meaning of “omega upper bar”
19

20 *We have clarified this point.*
21

22 p. 20: Please, explain “Occam’s Razor” or do not use the term.
23

24 *Rewritten without using this term.*
25

26 p. 25: “average minimum effort will be higher” instead of “average minimum effort will be higher”
27

28 *Fixed.*
29

30
31 Figure 5 is hard to read since the line markers are too similar
32

33 *Fixed.*
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

Management Science Data Disclosure Checklist

Please complete this form, save it as “Data Disclosure Checklist,” and upload it as a separate file after your submitted manuscript. It will be available to the entire review team, so please do not include author names. This page does not count toward the page limit on the manuscript.

Title of manuscript:

Stand By Me. Experiments on Help and Commitment in Coordination Games.

Indicate (e.g., by underlining) “Yes” or “No”:

Yes No This manuscript includes analysis of data, either field data, simulated data, or experimental data.

Yes No If our manuscript is accepted we will provide the journal with our data so that it can be posted on the journal’s website. To promote additional research and to increase the credibility of a paper’s findings, data disclosure is encouraged but not required.

Yes No A portion of our data cannot be disclosed due to a non-disclosure agreement or similar limitations on disclosure. If “Yes,” briefly explain which data cannot be disclosed and why:

For papers that report experimental data, please answer the following:

Yes No We report how we determined our sample sizes, all data exclusions (if any), all manipulations (for experimental work), and all measures collected.