UNIVERSITY OF AMSTERDAM

UvA-DARE (Digital Academic Repository)

Laboratory tests of theories of strategic interaction

Yang, Y.

Publication date
2014

Document Version
Final published version

Link to publication

Citation for published version (APA):
Yang, Y. (2014). *Laboratory tests of theories of strategic interaction*. [Thesis, fully internal, Universiteit van Amsterdam]. Tinbergen Institute.

General rights
It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations
If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: https://uba.uva.nl/en/contact, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

UvA-DARE is a service provided by the library of the University of Amsterdam (https://dare.uva.nl)

Download date:05 Oct 2025

2

Inequity Aversion Revisited

2.1 Introduction

Why are you reading yet another discussion on inequity aversion (IA)? One would think that the plethora of papers that have studied this phenomenon since the seminal work by Fehr and Schmidt (1999) and Bolton and Ockenfels (2000) would have said it all.[1] In our view, this is not the case. We believe that this extensive literature has at least two important gaps. First, the literature has failed to adequately distinguish between inequity aversion *per se* and (preferences for) reciprocity. This distinction is potentially important, for instance, because the literature on ultimatum games has revealed that the responder's choice may depend heavily on the actions available to the proposer, a finding that is not captured by Fehr and Schmidt's (1999) IA model ("the IA model" henceforth).[2] Second, a systematic test of the model's robustness to systematic variations in its premises has not been undertaken for settings where players cannot reciprocate others choices.[3] This chapter aims to fill these gaps. This is important, because both aspects pertain directly to the empirical applicability of the model. Given that the IA model has been used to explain people's behavior in many experimental and real-life environments and is even seen as a preferred first approach for studying behavior in many areas (Camerer, 2003, :472), a systematic approach to its strengths and weaknesses seems of utmost importance.

Such a 'scientific' approach to critically assessing theories was recently advocated by Binmore and Shaked (2010) for economics in general and for the IA model in particular. Their main focus is on the methodology used by Fehr and Schmidt, however.[4] Though we agree that a critical assessment of methods is important, it is also necessary to stress test a model's robustness, especially when it is so widely used: Binmore and Shaked report that as of 2010 Google Scholar listed 2390 works citing Fehr and Schmidt; by June 2014 this had increased to over 6500. Only a few of these (to be reviewed below) investigate the model's predictive power, none systematically studies its robustness in reciprocity-free settings.

[1]See Cooper and Kagel (2009) for a review of the literature on inequity aversion.

[2]See, e.g., Brandts and Solà (2001) and Falk, Fehr, and Fischbacher (2003).

[3]In such settings, subjects could still reciprocate based on beliefs about an opponent's action in the sense of Rabin (1993). Slightly abusing terminology, we will label a game reciprocity-free if subjects cannot condition their choices on the actual actions by others throughout the chapter.

[4]Importantly, we do not intend to 'take sides' in the specific criticisms that Binmore and Shaked put forward with respect to the Fehr-Schmidt model. See Fehr and Schmidt (2010) for a concise reply to these critiques.

This chapter addresses both issues in a laboratory experiment. We will test the robustness of the IA model's estimated parameters to the inclusion of efficiency concerns and to variations in stakes. Moreover, we will introduce a new game, which we call the 'production game'. This has various desirable properties (to be discussed below) that will allow us to straightforwardly assess the IA model's predictive power at the individual level. We will do so in two environments, one in which subjects can reciprocate others choices and one in which they cannot.

More specifically, we use a two-step experiment to test the IA model. The first step is to estimate each subject's IA preferences by using a set of choice menus. We estimate such preferences using two different models, one with and one without efficiency concerns, and check the robustness to the IA model by comparing the two estimates. Moreover, we check their robustness to scaling (i.e., we test the linearity assumption underlying the model) by varying the stakes in the menus. In the second step, we let subjects play the production game. By comparing their decisions in this part to the theoretical prediction derived from the estimates of their individual IA levels as obtained in the first step, we test the predictive power of the IA model.[5]

Of course, there have been other attempts to test the IA model. In fact, economists have conducted various laboratory experiments to test the IA model and, typically, to compare it to other social preferences models (for example, see Charness and Rabin 2002, Engelmann and Strobel 2004, Brandts, Fatás, Haruvy, and Lagos 2014).[6] Typically, such studies develop a series of simple games for which various models (including IA) offer distinct predictions. Subjects' choices in these games subsequently provide evidence in favor or against specific models. In this way, it has been argued that efficiency concerns (Charness and Rabin 2002) and maximin preferences (Engelmann and Strobel 2004) are better predictors of individual choice than IA (though the latter result has been disputed by Fehr, Naef, and Schmidt 2006). Notwithstanding the elegance and usefulness of such horse races between models, they do have drawbacks. For one thing, they are ultimately only informative about the models concerned and for the games chosen. Moreover, not the

[5]Formally, our test of the role of reciprocity may also be interpreted as a robustness check of the IA parameters. In this chapter, we prefer to view IA with and without reciprocity as separate models, however, because reciprocity directly affects the interpretation of the IA parameters. This is discussed in our concluding section.

[6]Stakes effects in where participants can reciprocate others actions have been studied by Andersen, Ertaç, Gneezy, Hoffman, and List (2011) and Fehr, Tougareva, and Fischbacher (2013) among others.

models themselves are tested, but their comparative predictions. This is different in our approach. We will directly test the model's premises with respect to preferences and its predictions, independently of other models. We allow heterogeneity of preference in IA in our test.

Originally, many of the experimental tests (including the supporting evidence in Fehr and Schmidt 1999) were implemented at an aggregate-level, by checking the consistency of the distribution of subjects' behavior across different games with the distribution predicted by the model. This method has two drawbacks. First, because the optimal strategy for the players in several of the games used depends on the extent to which the opponents care about inequality, the validity of the test is based on an assumption that subjects hold correct beliefs on the distribution of IA preferences in the population. This seems quite demanding, especially in one-shot games. Second, even if the model passes the test at the aggregate level, this is not informative about its predictive power for each individual. More recent studies, including Engelmann and Strobel (2004) and Blanco, Engelmann, and Normann (2011), test predictions at the individual level. Such an individual-level test can directly check the within-subject consistency for each individual. In this way, our method is also an individual-level test of the IA model.

Aside from directly measuring IA in various circumstances (as opposed to comparing its predictions in a horse race with other models), there are other advantages to our approach, compared to previous individual-level tests of the IA model. First, we consider efficiency concerns as an addition to the IA model, so that the original model is nested in the extended version with such concerns. This allows us to directly test the robustness of measured IA parameters to allowing for such efficiency concerns.[7] Second, we introduce the production game, in which a player's behavior predicted by the IA model is a continuous function of her IA levels. In the games traditionally used to test the model (e.g., the prisoners' dilemma game or the ultimatum game), the models prediction is a binary function, e.g. cooperate/defect or accept/reject, of the player's IA levels. The production game we introduce facilitates a sharper test of the model because it avoids such bang-bang predictions. Third, subjects' risk attitudes are irrelevant in our experimental environment. This is in contrast to most of the games traditionally used for the analysis of social preferences

[7]We are aware of Engelmann's (2012) critical discussion of introducing efficiency concerns in this way. We will discuss the relevance of this criticism below, when introducing the theoretical model.

(e.g., the public good game, prisoners dilemma, or ultimatum game; see Fehr and Schmidt 1999, Blanco, Engelmann, and Normann 2011). These traditional games are characterized by strategic uncertainty; hence, a subjects decisions may be affected by her risk attitudes, which could yield biased estimates of her IA level.[8] In contrast, none of the decisions used in our experiments involve any strategic uncertainty. Therefore, risk attitudes play no role, allowing for a more accurate test of the IA model. Finally, we will present two versions of the production game that only differ in the simultaneity of decisions. A comparison will allow us to isolate the role of 'explicit reciprocity' (Charness and Rabin 2002) from IA-preferences in a way not previously done. The approach that is probably closest to ours is the one applied by Blanco, Engelmann, and Normann (2011).[9] As our chapter, their work is an individual-level study of the IA model that tests the internal consistency of the IA model across different games. Also, their method to measure the individual guilt level is based on a modified dictator game, which has a structure similar to the test menu we use to measure advantageous inequity aversion (the disutility derived from feeling guilty about having a higher payoff than others). There are three main differences that distinguish our work from Blanco, Engelmann, and Normann (2011), however. The first is that we test the IA model's robustness to allowing for efficiency and variations in stakes, which we believe are two important issues in the specification of the model. Second, the methods we use to measure individual IA levels and to test the internal consistency of the IA model are different from Blanco, Engelmann, and Normann's (2011). Instead of deriving the disadvantageous IA parameters indirectly from the ultimatum game, we use a set of simple menus to directly measure IA preferences (not only the guilt parameter). Third, instead of testing the IA model using classic games such as public good games and two-stage prisoner's dilemmas we introduce a novel game, the production game, for this test. As we will see, a major advantage of our methods is that several important factors (e.g., reciprocity concerns and risk attitudes) that may affect behavior and may play important roles in the games used in Blanco, Engelmann, and Normann (2011) are

[8]For instance, a risk averse proposer in an ultimatum game may choose a high offer because she fears rejection by a receiver with a high disutility of disadvantageous IA. If this is not taken into account when analyzing her choices (e.g., Fehr and Schmidt 1999; Blanco, Engelmann, and Normann 2011, her high offer may yield an erroneously high estimate of disutility of advantageous IA.

[9]Dannenberg, Riechmann, Sturm, and Vogt (2007) present an application of the Blanco et al. method.

well-controlled here. As will be shown in Section 4, our results differ from those reported in Blanco, Engelmann, and Normann (2011), indicating the importance of controlling for such factors. Our findings show, first, that the estimates of the inequity aversion model are robust to the inclusion of efficiency in the model and -to a large extent- also robust to variations in the stakes. Second, the IA model predicts individual subjects' behavior in the production game quite well. Adding the possibility of reciprocity to the game strongly reduces the model's predictive power, however. Finally, our estimate of the distribution of the guilt parameter is roughly similar to the one reported in previous studies, including Fehr and Schmidt (1999), but we find much lower levels of the envy parameter (measuring disutility of disadvantageous inequality). However, when deriving estimates of envy from individuals' choices in the version of the production game that allows for reciprocity, our estimates of subjects' envy levels are higher and correspond better to the previous estimates. All in all, we conclude that distributional preferences measured by the IA model are remarkably robust but that previously measured disadvantageous inequity aversion is strongly influenced by negative reciprocity. Note that Fehr and Schmidt (1999) already point out that their parameters may be interpreted as a combination of concerns for both inequity and intentions; e.g., a rejection of a positive but low ultimatum offer may be interpreted as aversion to the resulting inequity or a reciprocal response to the perceived bad intention of the proposer. We believe that our methods are the first to facilitate a clear separation of the intention-based concern from the preference for equality.

The remainder of the chapter is organized as follows. In Section 2.2 we will introduce the models, including the details and predictions for the production game. We will explain the experimental design in Section 2.3 and the results in Section 2.4. Finally, in Section 2.5, we offer a discussion of the results and a conclusion of the chapter.

2.2 Theory

We will derive hypotheses using the classic Fehr and Schmidt (1999) model of inequity aversion and an extension of this model that includes efficiency concerns. In the current section, we present these models. We also introduce here the production game that we will use below and analyze it in the light of the two models.

Consider a population of $n+1$ individuals. In the Fehr and Schmidt (1999) model, an individual derives positive utility from own earnings and (dis)utility from inequality. More specifically individual is utility is given by:

Inequity aversion model:

$$U_i(x, y_1, y_2, ..., y_n) = x - \alpha_i \frac{\sum_j \max\{y_j - x, 0\}}{n} - \beta_i \frac{\sum_j \max\{x - y_j, 0\}}{n} \tag{2.1}$$

where $(x, y_1, ...y_n)$ denotes a payoff bundle, x is a player's own payoff, and $y_j, j = 1, ...n$ denote the other n players' payoffs. $\alpha_i \geq 0$ is an envy parameter (Engelmann and Strobel 2004) measuring the marginal disutility of disadvantageous inequality. β_i measures the marginal disutility (guilt) related to advantageous inequality, with $\beta_i \in [0, \alpha_i] \bigcap [0, 1)$.[10]

Efficiency model:

$$U_i(x, y_1, y_2, ..., y_n) = x + \gamma_i (x + \sum_j y_j), \tag{2.2}$$

where $\gamma_i \geq 0$ measures the marginal utility of aggregate earnings (efficiency).

For the robustness test, we will use a general model that incorporates both inequity aversion and efficiency concerns.

General model:

$$U_i(x, y_1, y_2, ..., y_n) = x - \alpha_i \frac{\sum_j \max\{y_j - x, 0\}}{n} - \beta_i \frac{\sum_j \max\{x - y_j, 0\}}{n} + \gamma_i (x + \sum_j y_j), \tag{2.3}$$

where $\alpha_i, \beta_i, \gamma_i$ still satisfy the above mentioned assumptions in the IA and efficiency models.[11]

[10] It turns out that the restriction of β_i to lie between 0 and α_i is not needed for our purposes. We will return to this point in Section 2.4.

[11] Note that our formulation of efficiency concerns in (2.3) implies that player i's utility increases with n. This means, for example, that i is better off if more people join her reference group (with the same payoffs as hers). This reflects our interpretation of efficiency concerns. Alternatively, one could consider a formulation such as $\gamma_i \frac{(x + \sum_j y_j)}{n+1}$, where only the average earnings per reference group member affect i's utility. In our view, (2.3) provides a more intuitive description of efficiency concerns. Consider the following two situations, for example. In one, an individual may give up

The incorporation of efficiency concerns in a general model like (2.3) is elegantly analyzed in Engelmann (2012). The author illustrates that such a general model with distinct parameters for envy, guilt and efficiency, can be fully captured by a two-parameter inequity aversion model, unless one simultaneously considers games with different numbers of players.[12] It will become clear below that our experiment does vary the number of players. An implication of disregarding efficiency concerns is then, that when directly measuring the two IA parameters in the IA model - as is common in the existing literature- the measured parameters may represent a reduced-form model simultaneously reflecting subjects' IA and efficiency concerns. In contrast, our method facilitates measuring the IA parameters in both the IA model and the general model. A lack of robustness in the measurement of the IA parameters across the two models would then indicate that efficiency concerns are present.

Next, we introduce the production game that we will use to test the predictive power of the two models. This game is played by two players, Worker A and Worker B. At the start of the game, each receives a basic salary (s_i, $i = A, B$). Each worker is in charge of a department's production (departments are also denoted by A and B).

The production of each department will be equally distributed (as a 'bonus') between the two workers. Worker i chooses effort $e_i \in [0, e_{max}], i = A, B$. Department i's production p_i depends on the effort exerted by worker i in the following way:

$$p_i(e_i) = 4e_i - \frac{e_i^2}{100}, i = A, B \tag{2.4}$$

1 dollar to increase average income by 0.5 in a group of four; in the other, she faces the option to pay 1 dollar to increase the average income by 0.5 for a group of 1000 people. Though this is, in the end, an empirical question, we believe that one would likely make different choices in these two situations. The alternative formulation of efficiency concerns would predict the same choice in both cases because it fails to capture the fact that people care not only about social average income, but also about the number of people in their reference group. This consideration underlies our choice of the formulation of efficiency concerns in (2.3).

[12] Blanco, Engelmann, and Normann (2011) also note that the parameter describing a concern for efficiency is not identifiable if the environments used to measure are restricted to two-player games. Engelmann (2012) goes on to argue that a comparison of games with distinct numbers of players yields rather implausible predictions in the sense that a player will typically be inequality averse for some group sizes and inequality-loving for others. It goes too far to address this point in detail, here, but it suffices to say that we disagree that such an outcome is a priori implausible. For our purpose, this discussion is irrelevant, however, since we are solely interested in the robustness of the IA parameters.

Effort is exerted at constant marginal costs $c_i \geq 0$ for Worker i. Worker is payoff, π_i, is then given by:

$$\pi_i(e_A, e_B) = s_i + 1/2 \sum_{j=A,B} p_j(e_j) - e_i c_i, \quad i = A, B \tag{2.5}$$

From here onward, we consider the parameters $e_{max} = 100$, $s_A = 200$, $s_B = 0$, $c_A = 2$, $c_B = 1$ that we used in our experiment. Hence, A starts with a higher basic salary but faces higher marginal costs than B. These parameters ensure that, for any possible pair of effort levels up to the maximum of 100, $\pi_A \in [150, 350]$ and $\pi_B \in [0, 150]$, implying that Worker A always earns more than B (see Appendix 2.6.1 for details). Therefore, when applying model (2.1) to the production game, only the guilt parameter β for A and the envy parameter β for B are relevant. When applying (2.3), the efficiency parameter is also relevant.

More specifically (cf. Appendix 2.6.1), the IA model (2.1) yields the following optimal strategy for Worker A in the production game:

$$e_A = \begin{cases} 0 & \text{if } \beta_A \leq 0 \\ 200\beta_A & \text{if } 0 < \beta_A < 1/2 \\ 100 & \text{if } \beta_A \geq 1/2. \end{cases} \tag{2.6}$$

Similarly, B has an optimal strategy in the production game:

$$e_B = \begin{cases} 0 & \text{if } \alpha_B \geq 1 \\ 100(1-\alpha_B) & \text{if } 0 < \alpha_B < 1 \\ 100 & \text{if } \alpha_B \leq 0. \end{cases} \tag{2.7}$$

Therefore, if the IA model describes workers' preferences, A's effort choice in the production game is a dominant strategy completely determined by the guilt parameter β, while B has a dominant effort level solely dependent on her envy parameter α.

The production game is particularly suited to test the inequity aversion model (2.1). There are at least four reasons why this is the case. First, as mentioned, both players have dominant strategies. Because their optimal choice requires no beliefs about the other worker's actions, attitudes towards risk and uncertainty are irrelevant for the models predictions. This allows for a direct test of these predictions. Second, for any combination of effort choices, Worker A earns more than

Worker B. As a consequence, each worker's predicted effort depends only on one kind of inequity aversion, again allowing for a direct test of the prediction. Third, the prediction is a continuous function of the relevant inequity aversion parameter (as opposed to the bang-bang corner predictions obtained for most games). This property facilitates a sharper test. Finally, one can distinguish between a version of the production game where both players simultaneously decide on their effort choice and a sequential version where worker B sees A's choice before deciding herself. By distinguishing between these two, one can isolate the effects of reciprocity. In the simultaneous game, the players cannot observe each other's action, and their decision is only dependent on their own preference levels, so reciprocity of the other's effort choice cannot play a role. In contrast, when Worker A chooses first, Worker B may condition her choice on A's chosen effort. The sequential game allows Worker B to respond to her perception of A's kindness, allowing for reciprocity to play a role in her decision.

2.3 Experimental Design

In the experiment, we use three choice menus to directly measure each subject's social preferences (i.e., to obtain estimates of α, β, and γ). Subsequently, we test the predictive power of the models for behavior in the production game. Transcripts of instructions are presented in the Appendix 2.6.3. We start by describing the menus used.

Choice Menus

We first introduce the three choice menus used to measure social preferences.

Menu 1

Menu 1 is used to measure parameter α of the inequity aversion model (2.1) and (together with Menus 2 and 3) the parameters of the general model (2.3). This menu consists of 10 decisions (cf. Table 2.1). In each, the decision maker (denoted by 'proposer') is asked to choose between two options (A and B). Each option allocates money to the proposer and to an anonymous other participant (denoted by 'receiver'). For each of the ten decisions, the proposer is linked to the same receiver (though at most one decision will be selected for payment, as will be explained below). Each participant decides as if she is a proposer, because roles are not (randomly) determined until the end of the experiment.

Table 2.1: Menu 1

Nr.	Option A	Option B	Choose B iff:
1	Yours: 125 ; Other's: 150	Yours: 100 ; Other's: 260	$\alpha \leq -0.19$
2	Yours: 115 ; Other's: 150	Yours: 100 ; Other's: 260	$\alpha \leq -0.12$
3	Yours: 105 ; Other's: 150	Yours: 100 ; Other's: 260	$\alpha \leq -0.04$
4	Yours: 95 ; Other's: 150	Yours: 100 ; Other's: 260	$\alpha \leq 0.05$
5	Yours: 85 ; Other's: 150	Yours: 100 ; Other's: 260	$\alpha \leq 0.16$
6	Yours: 75 ; Other's: 150	Yours: 100 ; Other's: 260	$\alpha \leq 0.29$
7	Yours: 65 ; Other's: 150	Yours: 100 ; Other's: 260	$\alpha \leq 0.47$
8	Yours: 55 ; Other's: 150	Yours: 100 ; Other's: 260	$\alpha \leq 0.69$
9	Yours: 45 ; Other's: 150	Yours: 100 ; Other's: 260	$\alpha \leq 1.00$
10	Yours: 35 ; Other's: 150	Yours: 100 ; Other's: 260	$\alpha \leq 1.44$

Notes: The final column includes the values of α for which the IA model rationalizes a choice of Option B. Of course, this column was not shown to subjects.

In each of the payoff pairs in this menu, the proposer's payoff is lower than the receiver's, i.e., the proposer is always at the disadvantageous position (which is why it is informative about the envy parameter). As a consequence, the third term on the r.h.s. of (2.1) is equal to zero for all options considered in the menu. In each of the ten decisions, Option B gives 100 points to the proposer and 260 points to the receiver; and Option A is characterized by a lower (disadvantageous) inequality. Moving down from decision 1 to decision 10, the proposer's earnings decrease and inequality increases (see Table 2.1). From decision 4 onward, the proposer's own earnings are also lower in Option A than in B.

For decision 1, compared to Option B, disadvantageous inequality is 135 lower in A, while the proposer earns 25 more. Note that for non-negative α, both remaining terms on the r.h.s. of (2.1) then imply higher utility for Option A than for B. In fact, any proposer with $\alpha \geq -0.19$ will choose A.[13] At the other extreme, consider decision 10. Here, choosing A means giving up 65 in own earnings (100-35) to decrease disadvantageous inequality from 160 (260-100) to 115 (150-35). Only individuals with strong envy ($\alpha \geq 1.44$) prefer option A.

[13] Recall the assumption in the Fehr-Schmidt model that $\alpha \geq 0$. Choosing option B in decisions 1, 2, or 3 cannot be rationalized under this assumption.

In this way model (2.1) determines for each decision question, a threshold for the envy level, above which Option A should be chosen and below which Option B should be chosen. This threshold is given in the last column of Table 1. If preferences are described by (2.1), a subject will switch when moving down from decision 1 to 10 at most once from choosing Option A to choosing Option B. It is easy to see that this switching point then identifies an interval for a proposer's envy level. More details on how subjects' envy levels are estimated can be found in Appendix 2.6.4.[14]

It is relatively easy to see from menu 1 that a downward bias would occur in the estimate of the envy parameter if an individual also cares about efficiency. Such a subject would choose the B option more often than someone without efficiency concerns because the sum of the payoffs are always greater under option B than under option A. As a consequence, decisions from menu 1 alone would underestimate her parameter. As discussed below, this bias is corrected in our joint estimate of envy and efficiency concerns using the choices in menu 3.

Menu 2

Menu 2 also consists of 10 decision questions (cf. Table 2.2), each containing an Option A and an Option B, with distinct payoff pairs for a proposer and receiver. In contrast to Menu 1, for all options the payoff of the proposer is higher than for the receiver. This means that all cases yield advantageous inequality for the proposer, which sets the second term on the r.h.s. of (2.1) equal to zero and allows us to use this menu to measure her guilt parameter β (cf. Appendix 2.6.4). Once again, each participant makes a decision as if she is a proposer because random role assignment is postponed until the end of the experiment.
Again, the payoffs for Option B remain constant across all 10 decisions, with the proposer earning 170, which is 120 more than the receiver (50). For the first decision, Option A gives the proposer more (185) and yields lower inequality (95) than B. Any non-negative β then implies higher utility for A than for B. Moving down along the table, the own earnings in Option A decrease, as does the inequality. This increases

[14] As mentioned above, previous studies have used a different approach to measure envy. They did so by eliciting responder rejection thresholds in ultimatum games. One difference with our approach is that reciprocity motives may play a role in ultimatum rejections. An anonymous referee pointed out a second difference. This is that our subjects do not have the extremely egalitarian outcome at their disposal where both players earn nothing. Under the linearity assumption in the IA model, this should not matter, but one cannot exclude that it will have a behavioral effect for boundedly rational decision makers.

Table 2.2: Menu 2

Nr.	Option A	Option B	Choose B iff:
1	Yours: 185 ; Other's: 90	Yours: 170 ; Other's: 50	$\beta \leq -0.60$
2	Yours: 175 ; Other's: 90	Yours: 170 ; Other's: 50	$\beta \leq -0.14$
3	Yours: 165 ; Other's: 90	Yours: 170 ; Other's: 50	$\beta \leq 0.11$
4	Yours: 155 ; Other's: 90	Yours: 170 ; Other's: 50	$\beta \leq 0.27$
5	Yours: 145 ; Other's: 90	Yours: 170 ; Other's: 50	$\beta \leq 0.38$
6	Yours: 135 ; Other's: 90	Yours: 170 ; Other's: 50	$\beta \leq 0.47$
7	Yours: 125 ; Other's: 90	Yours: 170 ; Other's: 50	$\beta \leq 0.53$
8	Yours: 115 ; Other's: 90	Yours: 170 ; Other's: 50	$\beta \leq 0.58$
9	Yours: 105 ; Other's: 90	Yours: 170 ; Other's: 50	$\beta \leq 0.62$
10	Yours: 95 ; Other's: 90	Yours: 170 ; Other's: 50	$\beta \leq 0.65$

Notes: The final column includes the values of β for which the inequity aversion model (2.1) rationalizes choice of option B. Of course, this column was not shown to subjects.

the level of guilt needed to prefer Option A to B. The last column in Table 2.2 gives these threshold values for β.

As was the case for the estimate of the envy parameter, the estimate of the guilt parameter is biased if we use menu 2 in isolation for an individual who also has efficiency concerns. Here the bias is non-monotonic in the actual value of β. Up to decision 6, the sum of the payoffs is higher in option A. Therefore, subjects who care about efficiency and have a guilt parameter below 0.47 may choose A more often than those with the same parameter who do not care about efficiency. This will cause an upward bias in the estimated β. The reverse is true for guilt parameters greater than 0.53. Again, such bias will not occur when we correct for menu 3 decisions in the analysis.

Menu 3

Together with Menus 1 and 2, Menu 3 (cf. Table 2.3) is used to measure the parameters of the general model (2.3). This menu, again, presents 10 decision questions to a proposer, distinguishing between Options A and B. In contrast to Menus 1 and 2, the proposer is now (anonymously) grouped with five receivers. Each option specifies one payoff amount for the proposer and another amount for each of the receivers. This allows us to separate efficiency concerns from inequity aversion. Once again, each subject makes a decision as a proposer. At the end of the experiment, one

Table 2.3: Menu 3

Nr.	Option A	Option B	Choose B iff:
1	Yours: 50 ; Five Others': 50	Yours: 25 ; Five Others': 75	$\gamma \geq 0.250$
2	Yours: 50 ; Five Others': 50	Yours: 25 ; Five Others': 85	$\gamma \geq 0.167$
3	Yours: 50 ; Five Others': 50	Yours: 25 ; Five Others': 95	$\gamma \geq 0.125$
4	Yours: 50 ; Five Others': 50	Yours: 25 ; Five Others': 105	$\gamma \geq 0.100$
5	Yours: 50 ; Five Others': 50	Yours: 25 ; Five Others': 115	$\gamma \geq 0.083$
6	Yours: 50 ; Five Others': 50	Yours: 25 ; Five Others': 125	$\gamma \geq 0.071$
7	Yours: 50 ; Five Others': 50	Yours: 25 ; Five Others': 135	$\gamma \geq 0.062$
8	Yours: 50 ; Five Others': 50	Yours: 25 ; Five Others': 145	$\gamma \geq 0.056$
9	Yours: 50 ; Five Others': 50	Yours: 25 ; Five Others': 155	$\gamma \geq 0.050$
10	Yours: 50 ; Five Others': 50	Yours: 25 ; Five Others': 165	$\gamma \geq 0.045$

Notes: The final column includes the values of γ for which the efficiency model (2.2) rationalizes a choice of option B. Of course, this column was not shown to subjects.

group of six will be randomly selected and one of these six will be randomly selected as the proposer whose decision is implemented.

In each of the ten decisions, option A is the same, giving 50 to the proposer and to each of the five receivers (for aggregate earnings equal to 300). Compared to A, Option B gives less (25 in each decision) to the proposer. Moving down the menu, it gives increasingly more to each of the receivers, thereby increasing aggregate earnings. In decision 1, aggregate earnings (400) are higher in B than in A. Hence, a strong enough preference for efficiency may give the higher utility to B. Note that from model (2.3) it follows that the threshold value of γ to switch from A to B depends on the inequity aversion measured by α and β. Restricting social preferences to only efficiency concerns (model 2.2) allows one to directly determine a lower bound for γ for each decision, however. These bounds are shown in the last column in Table 2.3. Nevertheless, our main interest lies in testing the robustness of estimates of inequity aversion parameters to allowing for efficiency concerns (i.e., model 2.3), and not in the efficiency model *per se*.

Here, we can consider the bias in the estimate of γ that results from not taking inequity aversion into account if Menu 2.3 is considered in isolation. In option B, the decision maker always earns less than the others, so only envy plays a role. Both inequity and efficiency of the B option increase as one moves down the table. Envy will postpone the switch to B and γ will be underestimated.

Combining Menus 1-3

The $\alpha-$ and $\beta-$ [$\gamma-$] cutoff points in Tables 1-3 are based on the partial model (2.1) [(2.2)]. If we consider the general model (2.3), efficiency concerns may play a role when choosing in Menus 1 and 2 and inequity aversion may influence the choices in Menu 3. By combining an individual's choices in all three menus, we can avoid the biases discussed above and jointly estimate α, β, and γ. Moreover, comparing subjects' choices in the three menus will allow us to evaluate the robustness of the Fehr-Schmidt model to the inclusion of efficiency concerns. More in particular, we will estimate the extent to which estimates obtained with the Fehr-Schmidt model alone (model 2.1) are consistent with measures obtained in the general model. Details on this procedure are given in Appendix 2.6.5.

Production Game

When introducing the production game to subjects, equations (2.4) and (2.5) were (of course) not used. Instead, participants were given a calculator to see the consequences of various effort levels by Workers A and B. Diagram 1 shows the computer screen used for this purpose. The screen is split in two halves, one for A's decision and one for B's decision. We use the strategy method (with respect to a move by nature): when deciding, the subjects do not know which role they will have and are asked to make decisions as both Worker A and Worker B. For both roles' decisions, a subject can use the calculator to try out as many decisions as they like.

For each role, the participant can try out any effort levels by moving a scroll bar. The table directly below the scroll bar shows the consequences of a decision for each worker. It shows the effort chosen, the 'bonus' (share of that department's production), the effort costs for the worker concerned, and the aggregate earnings for each worker. Note that the latter does not include the bonus to be earned from production in the other department. This is because that bonus depends solely on the other worker's efforts. The consequences of this other effort can also be tried out on the other half of the screen. After having practiced, the participant can choose a decision for both roles (A and B) and finalize by clicking an 'OK' button, after which she is asked to confirm her decisions.

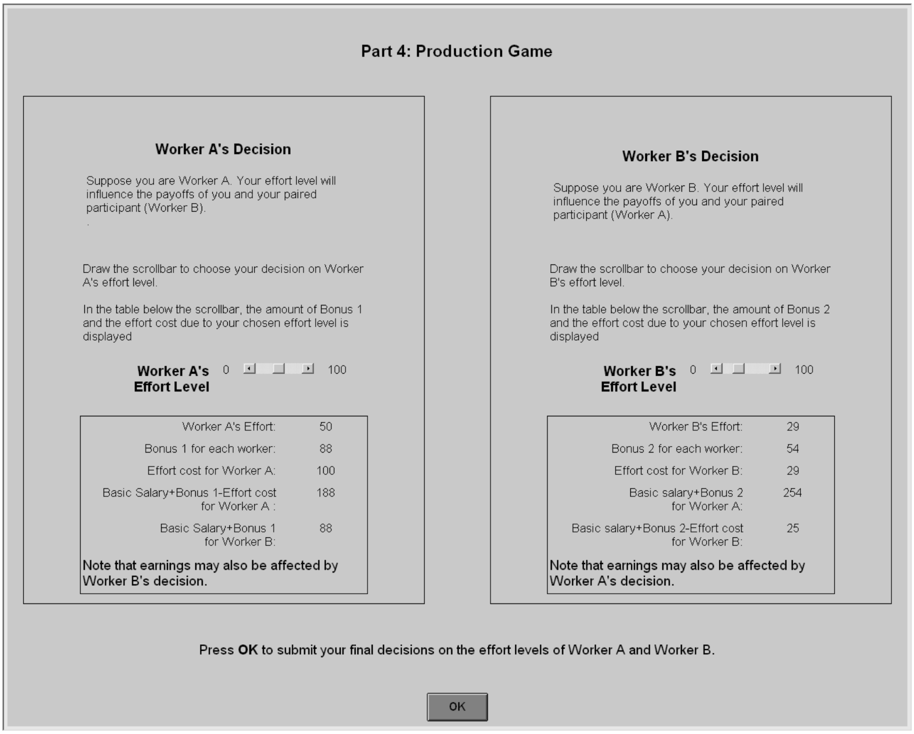

Diagram 1. Screen for SimProd

We used two versions of the production game, which we varied across subjects. In the first, subjects were not informed about their roles, and made their decisions simultaneously for both Worker A and B. We denote this simultaneous production game by SimProd. In contrast, subjects make decisions in the production game sequentially in our treatment SeqProd. For this, we again use a strategy method. Now, Worker B can condition her effort level on effort levels chosen by Worker A.

This is implemented as follows. Worker B chooses a set of 'responding rules' that consist of lower and upper bounds for effort chosen by A and a corresponding effort that B chooses for those efforts by A. B can formulate as many such rules as she wishes before finalizing her decision. An example is given in Diagram 2. In this example, Worker B chooses effort level 1 if A chooses 8 or less, 13 if A chooses between 9 and 50 and 73 if A chooses 51 or more. The instructions in Appendix 2.6.3 show how participants were informed about using this 'Decision Box'.

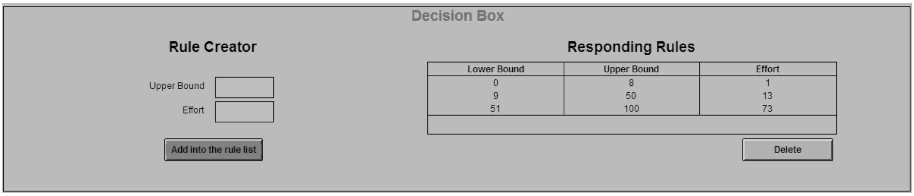

Diagram 2. Decision Box for Worker B in SeqProd

For SeqProd, we also implemented two subtreatments, SingleRole and DoubleRole. Subjects know their roles (either A or B) in the SingleRole treatment, and only make decision for their own roles. In the DoubleRole treatment, subjects are not informed about their roles and need to specify an effort level as Worker A, and give a set of responding rules as Worker B as well. Only the decisions made for their true roles, which will be only revealed at the end of the experiment, will be implemented.[15]

In summary, there were six different versions of the production game, which were varied across subjects (see Table 2.4). The main distinction is between on the one hand the simultaneous version (SimProd), and on the other hand the two sequential versions (SeqProd/SingleRole and SeqProd/DoubleRole). The other versions varied the payoff stakes in SimProd. Payoffs varied across four levels. Payoffs in the sequential version were scaled in the category "low". More details are provided below.

Table 2.4: Production Game Variations

	Simultaneous or Sequential		
Payoff level	SimProd	SeqProd/SingleRole	SeqProd/DoubleRole
lowest	88	—	—
low	34	30	30
high	30	—	—
highest	22	—	—

Note: Numbers indicate the number of subjects. All treatments presented in the table were varied across subjects.

[15] As will become clear, we used these variations of the production game (simultaneous versus sequential and single versus double roles) to isolate the effect of reciprocal concerns on inequity aversion.

Experimental Procedures

The experiments were conducted at the CREED laboratory of the University of Amsterdam (UvA). Subjects were recruited from the CREED subject pool, which consists of approximately 2000 students, mainly UvA undergraduates from various disciplines. 284 students participated and earned on average 38.60 Euro, including a 7-Euro show-up fee. All sessions lasted less than 60 minutes. At the start, participants are told that the experiment consists of several parts, and that the instructions to each part will be distributed before that part starts. Control questions are used to test understanding of these instructions. Parts 1, 2, and 3 measure social preferences using Menus 1, 2, and 3. In Part 4, one of the production games of Table 2.4 is played. Subjects do not learn about their roles and their payoffs in any part until the end of the experiment.

In Part 1, every subject is asked to make her choices for Menu 3. Payoffs are at the level "lowest". Groups of six are randomly formed. Subjects know that one of the ten decision questions will be randomly selected and that one of the six subjects in each group will be randomly assigned to be the proposer. This proposer's decision for the selected question is implemented for her group. Subjects are also told that if the number of subjects is not a multiple of six, some will not be allocated to a group, and hence are not paid for this part. Of course, when making their decision, they do not know whether or not they have been allocated to a group.

In Part 2, subjects are randomly assigned into pairs. Each subject is asked to make choices for Menus 1 and 2. Again, payoffs are at the "lowest" level. At the end of the experiment one of the in total 20 decisions is randomly selected to be paid. In each pair one of the participants is appointed proposer and the choice of this proposer for the selected question is implemented. In Part 3, Part 2 is repeated with higher payoffs. In three different treatments (varied across subjects), different payoffs are implemented. Specifically, the payoffs of Menus 1 and 2 are multiplied by 10 (denoted by "low" in Table 4), 30 ("high") and 60 ("highest"), compared to part 2. To control for income effects due to the possible earnings from the first two parts, each subject must choose whether or not to enter Part 3 (for a similar procedure, see Holt and Laury 2002). If a subject chooses to enter, she forfeits all earnings from the first two parts. If she chooses not to enter this part, all her previous earnings are kept, and she waits until this part finishes. In the experiment 234 subjects had to make this decision. Only 8 (3.4%) chose not to enter part 3.

Our focus in Parts 1-3 of the experiments is on the social preference models (2.1)-(2.3). Whereas choices in Parts 1 and 2 allow us to obtain parameter estimates for these models in the way described above, the scaling of payoffs in Part 3 allow us to test the underlying linearity assumptions. For example, in the Fehr and Schmidt (1999) inequity aversion model, the marginal disutility of inequality is independent of the payoff level. This means that measures of α and β obtained from Parts 2 and 3 should be equal.

Finally, in Part 4, we randomly paired subjects and let them play the production game described above, as SimProd, SeqProd/SingleRole, or SeqProd/DoubleRole. As described above, this will be used to test the predictive power of the estimates. For each of the payoff levels of Part 3, we ran a "lowest" payoff simultaneous production game with payoffs in the order of magnitude of Parts 1 and 2 and a second payoff version with payoffs in the magnitude of Part 3 in the session concerned (i.e., either "low", "high" or "highest"). As mentioned above, these were varied across subjects.

Table 2.5: Treatments

Treatment	Part 1 (menu 3)	Part 2 (menus 1&2)	Part 3 (menus 1&2)	Part 4 (production game)
I (n=28)	lowest	lowest	low	SimProd - lowest
II (n=34)	lowest	lowest	low	SimProd - low
III (n=30)	lowest	lowest	high	SimProd - lowest
IV (n=30)	lowest	lowest	high	SimProd - high
V (n=30)	lowest	lowest	highest	SimProd - lowest
VI (n=22)	lowest	lowest	highest	SimProd - highest
VII (n=30)	lowest	lowest	low	SeqProd/SingleRole - low
VIII (n=30)	lowest	lowest	low	SeqProd/DoubleRole - low
IX (n=24)	—	—	low	—
X (n=26)	—	—	—	SimProd - high

Notes: Numbers indicate the number of subjects participating. "lowest" indicates the benchmark payoff scale used in parts 1 and 2 (cf. Menus 1-3). These numbers were multiplied by 10 (30/60) in "low" ("high"/"highest"). All treatments presented in the table were varied across subjects.

Table 2.5 summarizes our treatment combinations. Because no information about others' choices is given until the end of the experiment, we consider each individual as an independent observation for our statistical analyses.

2.4 Results

This section starts with an overview of our estimates of the envy and guilt parameters for the IA model (2.1) as derived from the Menus 1 and 2. This includes a comparison to values estimated in previous studies. Then, we check the robustness of these estimates to allowance of efficiency considerations (model 2.3) by including Menu 3 and to the scaling of payoffs.[16] Finally, we test the ability of the Fehr-Schmidt model to predict behavior in the production game. This enables us to test the robustness of the estimates to reciprocity concerns.

Estimates of Envy and Guilt

Our first estimates of the envy parameter (α) are derived from Menu 1. Specifically, the threshold values shown in the last column of Table 1 provide intervals for the estimated value. For example, a subject who chooses option A for decisions 1-4 and B for decisions 5-10 is estimated to have $\alpha \in [0.05, 0.16)$. Note that this procedure requires a maximum of one switch when moving down from decision 1 to decision 10. In fact, 216 of the 234 subjects (over 92%) who participated in a standard treatment session (i.e., all subjects except those in the MenuOnly and ProdOnly control sessions) are "(IA-)consistent" in this way, with the remaining 18 subjects labeled as IA-inconsistent.[17] Amongst the group of 216 IA-consistent subjects, 103 are estimated to have a negative value of α either in the lowest-stakes or the high-stakes environment, which would indicate a preference for increased (disadvantageous) inequity aversion. Note that this possibility is excluded by assumption in the Fehr-Schmidt model. However, as we will show later, 66 of these 103 subjects can be rationalized to have non-negative envy level by using a model that also allows for efficiency concerns.[18] The remaining 37 subjects are denoted as "non-IA-rationalizable" since

[16]Because preferences for efficiency per se are not the main interest of this chapter, estimates of γ can be found in Appendix 2.6.7, which derives such estimates using either the efficiency model (2.2) or the general model (2.3). We basically find that little importance is attributed to efficiency. Alternatively, one could make assumptions on preferences that aid in interpreting multiple changes. Given the ad hoc nature of such assumptions and the low number of subjects involved, we have decided not to do so.

[17]Alternatively, one could make assumptions on preferences that aid in interpreting multiple changes. Given the ad hoc nature of such assumptions and the low number of subjects involved, we have decided not to do so.

[18]Charness and Rabin (2002) and Engelmann and Strobel (2004) discuss how a preference for efficiency may give a subject reason to sacrifice own payoff for an increase in the payoff of another, even if this other already has a higher payoff. In a partial inequity aversion model like (2.1),

their choices are inconsistent with the inequity aversion framework. We will exclude the 18 IA-inconsistent and the 37 non-IA-rationalizable observations form the further data analysis, leaving 179 IA-consistent-and-rationalizable (IA-C&R henceforth) observations.[19] For completeness' sake an analysis applied to the set of 216 IA consistent subjects (and thus including the non-IA-rationalizable observations) is provided in Appendix 2.6.8. Our main conclusions turn out to be unaffected by this choice.

On the other hand, we have decided not to exclude subjects with β estimates that contradict the restriction $\beta \in [0, \alpha] \bigcap [0, 1)$ proposed by Fehr and Schmidt (1999). In fact, many subjects violate this condition. 65 participants (36.3% of the IA-C&R subjects) have an estimated β level greater than their α. We see no theoretical reason for this assumed restriction, however, and the IA model can straightforwardly be applied without it. We therefore include these subjects in our analyses. Finally, one subject has a negative estimated level, which is estimated to lie in the range (-0.60, -0.14]. A reason to include this observation is that we think it in fact reasonable to exhibit (slightly) negative guilt levels, a preference indicating that being slightly better off increases one's utility. In contrast, we see no such 'reasonable interpretation' for negative envy levels.

Our estimates of the IA-C&R subjects' envy are summarized in Figure 2.1(a). For comparison, we include the distributions reported by Fehr & Schmidt (1999) (FS henceforth) and Blanco, Engelmann & Normann (2011) (BEN henceforth). Our results differ substantially from those previously found. χ^2-tests reject the null-hypotheses that the distribution of our estimates for α equals the distribution reported by FS or BEN, both at the 0.01 level. In fact, our estimates of subjects' envy parameters are almost completely (97.8%) clustered at the lowest interval ($\alpha < 0.25$). This stark difference between our estimates of envy and those in FS and BEN will be extensively discussed below.

such a subject may be perceived as having a negative envy level. For this reason, we relax the original restriction on envy ($\alpha \geq 0$) and only enforce it for the general model (2.3). In other words, non-IA-rationalizability refers to negative envy, even after correcting for efficiency concerns.

[19]Recall that each of these subjects chose to enter part 3 of the experiment.

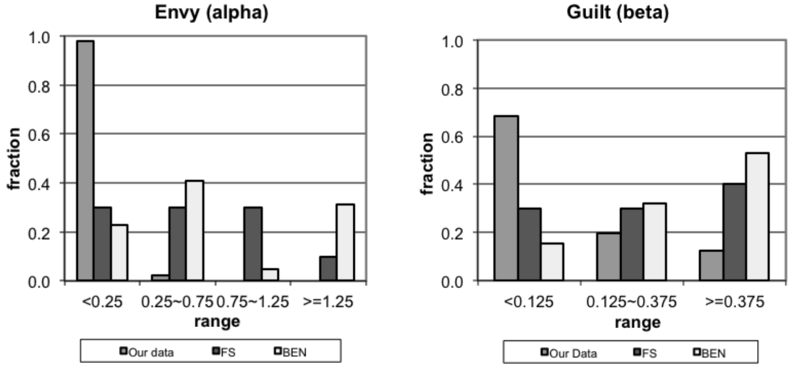

(a) (b)

Figure 2.1: Distributions of envy (a) and guilt (b) estimates

In a similar way, we use Menu 2 to provide a first estimate of the guilt parameter, β. Our estimates are presented and compared to the FS and BEN estimates in Figure 2.1(b).[20] The distribution of our guilt parameter is more comparable to those reported by FS and BEN than envy. Again, however, our estimated distribution is more skewed towards the left. The null-hypotheses of the distribution of our β estimates being the same as the distributions estimated by FS and BEN, are both rejected at the 0.01 level.

Before testing the predictive power of our estimates, we first test their robustness to stakes and efficiency concerns.

Robustness to Stakes and Efficiency concerns

Because all of the IA-C&R subjects chose to enter Part 3, we have 179 observations for which we can compare α and β estimated under different payoff scales. To start, Figure 2.2 shows the estimated distributions for the various payoff scales used. Recall from Table 2.5 that this is a within-subject comparison: each subject participated in the benchmark scale ("lowest") and in one of the higher scale treatments.

[20] Again, we center our intervals around the FS estimates, in this case β=0, 0.25 and 0.6.

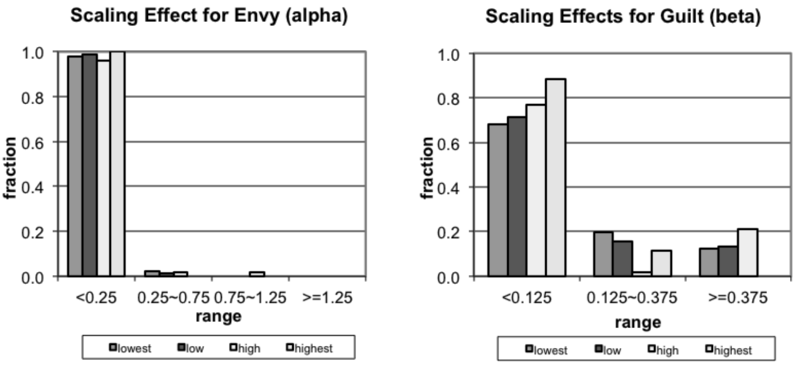

(a) (b)

Figure 2.2: Distributions of envy (a) and guilt (b) estimates for distinct payoff scale

A first impression from the figures is that estimates of the envy parameter are insensitive to changes in the stakes, though this may be due to the almost complete lack of envy in the first place. For guilt, there appears to be a shift towards lower β-values with increasing scale. To test whether this effect is statistically relevant, we use Wilcoxon sign-rank tests comparing at the individual level the IA parameters obtained from the two scale menus an individual participated in. The results are summarized in Table 2.6.

Table 2.6: Stakes Effects

	low (84 obs.)		high (52 obs.)		highest (43 obs.)	
	z	p-Value	z	p-Value	z	p-Value
α	0.527	0.598	0.495	0.620	1.613	0.107
β	0.543	0.587	-1.463	0.144	-3.155	0.002***

Notes. Results are presented from Wilcoxon Sign-rank tests of H0: $\alpha_{x,i} = \alpha_{lowest,i}$, vs. H1: $\alpha_{x,i} \neq \alpha_{lowest,i}$ and H0: $\beta_{x,i} = \beta_{lowest,i}$, vs. H1: $\beta_{x,i} \neq \beta_{lowest,i}$, $x \in \{low, high, highest\}$. *** indicates statistical significance at the 0.01 level.

For the estimates of the envy parameter, the test results do not reject the null hypothesis that they are invariant to changes in stakes. Therefore, we do not reject the hypothesis that the disutility from the payoff difference is linear in disadvantageous inequality. For the estimates of the guilt parameter, the test also does not reject the robustness of the estimates when payoffs are scaled up by a factor 10 or 30. We do, however reject the null of no stakes effect when payoffs are 60 times higher than in the benchmark case (i.e., in "highest"). This result implies that subjects feel less guilt about receiving more money than others when the amount earned is (much) more. In other words, the marginal disutility of advantageous inequity is decreasing in income. The linearity assumption for advantageous inequity is warranted for moderate increases in payoff levels, but not for major 60-fold increases.

Next, we consider the robustness of our results to allowing for efficiency concerns. We do so by deriving estimates of α and β from the general model (2.3) (as explained in Appendix 2.6.5) and comparing these to the estimates from the standard model (2.1). To provide a first measure of robustness, Figure 2.3 compares the derived distributions of α and β across the categories derived from FS estimates, which is dominant in the literature. At first sight, the introduction of efficiency concerns does not affect the distribution of the envy parameters. For the guilt parameter, the General model estimates put slightly more weight in the lowest range than the partial model estimates. The difference is not significant, however. According to the given categories as shown in the figures, χ^2 tests give a p-value of 0.268 for α and 0.684 for β, so we cannot reject the null hypothesis that the distributions of the IA estimates are the same as the General model estimates.

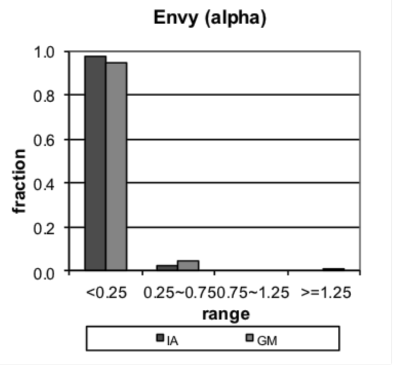

(a)

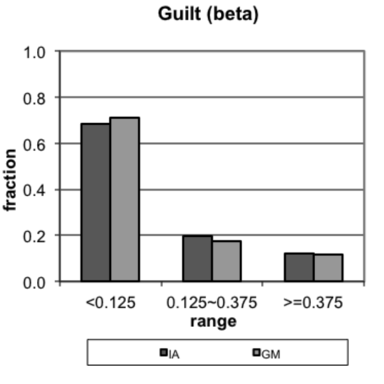

(b)

Figure 2.3: Distributions of envy (a) and guilt (b) estimates after correcting for efficiency concerns

The FS categorization is very crude, however. Our data allow for a more fine-tuned comparison. As explained in Appendix 2.6.6, a direct comparison of point estimates for α and β derived using the two models (2.1) and (2.3) has drawbacks. For example, a point estimate for the IA model is derived from the intervals for α and β implied by the last columns in Menus 1 and 2. These intervals change if we allow for efficiency concerns in the general model. A particular switching point in Menu 1 then yields a (slightly) different interval for α. Any method used to derive a point estimate from an interval thus yields distinct parameter estimates for the two models, which strongly influences the test results. An alternative way to check the robustness of the IA estimates to the allowance of efficiency concerns is to use the Overlap Ratio (see Appendix 2.6.6). This ratio, denoted by Ω, measures the average overlap of the intervals measured using models (2.1) and (2.3) as a percentage of the interval measured by the IA model (2.1). $\Omega = x\%$ then indicates that on average $x\%$ of the estimated α (or β) values that are consistent with the IA model based on Menu 1 (or 2) are also consistent with the general model jointly based on Menus 1 and 3 (or 2 and 3). Weighted by the number of observations[21], we find $\Omega = 82.1\%$ for α and $\Omega = 88.5\%$ for β.

An alternative measure for the robustness of IA parameter estimates to allowing for efficiency concerns simply determines the fraction of cases for which the interval derived from the partial IA model overlaps with the interval derived from the general model (as described in Appendix 2.6.5). An overlap means that there are values of α (β) that allow a subject's choices to be rationalized by both models. 88.8% of the IA-C&R subjects fulfill this criterion (cf. Appendix 2.6.6 of this chapter) for both parameters.

All in all, our estimates of IA parameters are quite robust to allowing for efficiency concerns. This can partly be attributed to the finding that our subjects show less concern for efficiency (see Appendix 2.6.7) than observed in some previous studies (e.g., Charness and Rabin 2002; Engelmann and Strobel 2004). A possible explanation for this difference with previous studies is that we are the first to directly measure preferences for efficiency using simple individual choice menus.[22] We

[21]We weigh in this way, because multiple observations exist for various combinations of switching points in Menus 1 (or 2) and 3.

[22]Charness and Rabin (2002) use a series of dictator choices to test subjects' preferences, but without trying to measure the degree of preference for efficiency (for which they use the term "social welfare").

will not further dwell upon possible reasons for these differences, because our main interest is in the robustness of the IA model and not in efficiency concerns *per se*.

Behavior in the Production Game

We will use the production game to test the predictive power of the IA model. Before doing so, we note that this is a novel game that has not previously been studied in the laboratory. We therefore start with a brief overview of observed behavior in this game. Table 2.7 gives an overview of average effort choices in the various treatments for this game.

Table 2.7: Effort Choices in the Production Game

Payoff level	SimProd Worker A Worker B	SeqProd/SingleRole Worker A Worker B	SeqProd/DoubleRole Worker A Worker B
lowest	39.9 96.3	—	—
low	45.3 95.7	70.8 66.2	53.5 47.8
high	42.1 98.0	—	—
highest	23.2 89.6	—	—

Notes. The first (second) number in each cell gives average effort by worker A (B) for all IA-C&R In the SeqProd treatments, worker B's effort was determined by applying her 'responding rule' to the effort actually chosen by worker A.

The results for the simultaneous version of the game (SimProd) show that A-workers exhibit less effort than B-workers. Across all payoff levels, the average (weighted by the number of observations) efforts are 38.8 and 95.5, respectively. It appears that in the simultaneous game, subjects take into account the distinct marginal costs of effort (2 for A and 1 for B). Because the benefits from production are shared equally and A has lower aggregate costs, the average production choices increase the inequality between A and B (recall that A always earns more than B) compared to the case where they only have their basic salaries. The distinct payoff scales in the game yield only small differences, with the exception of the highest payoffs (60 times the lowest). We use two-sided Kolmogorov-Smirnov (KS) tests to investigate whether the distribution of efforts (separately for A and B) in "lowest" differs significantly from that in "low", "high" or "highest". This is not the case for "low" or "high" or for B's effort in "highest" (all $p \geq 0.138$). Worker A produces significantly less effort in "highest" than in "low", however (p=0.001). Hence, only an increase in payoff scale with a factor of 60 has an effect on (A's) effort choices.

The differences between the two workers do not appear in the two sequential games. When B can condition her effort on that of A, A produces much more effort and B much less than in the simultaneous case. Comparing to the simultaneous case with the same "low" payoff scale, the increase for worker A is significant at the 10%-level for the SingleRole case (KS, p=0.077) and insignificant for DoubleRole (KS, p=0.338). The decrease compared to SimProd in worker B's effort is strongly significant for SeqProd/DoubleRole treatments (KS, p=0.013 but insignificant for SeqProd/SingleRole (KS, p=0.208). All in all, the possibility of reciprocation introduced by the sequential nature of the game thus destroys the inequity increasing effects observed in the simultaneous game by inducing B-workers to provide less effort and A-workers to provide more.

Next, we test whether using the strategy method has an effect on the effort levels chosen. We do so by comparing the SingleRole and DoubleRole treatments. It turns out that the distributions of effort levels are not significantly different (KS; p=0.800 for worker A and p=0.905 for worker B).

Finally, we consider whether the 'responding functions' that B-workers submitted in the SeqProd treatments exhibit reciprocity. This is the case if their effort choices are increasing in A's effort. We tested this in the following way. First, for each IA-C&R worker B in the SeqProd treatments (9 SingleRole and 20 in DoubleRole), we create 101 fictitious A decisions (choosing effort levels 0,1,,100). Then we used B's responding rules to determine for each B her effort in response to each of these 101 effort levels. Finally, we regressed B's effort on A's (fictitious) effort. This shows that Worker B on average increases her effort by 0.30 for each unitary increase in A's effort (with an associated p-value <0.001). We conclude that B-workers on average do reciprocate A's effort.

Predictive Power of the IA model in the Production Game

We can use the estimates of an individual's α and β to predict her behavior in the production game. We do so in two ways. First, we use the fact that the models predict specific relationships between individual envy and guilt parameters on the one hand and their effort choices on the other. We use regression equations to investigate whether we can reproduce these relationships. Second, we will use individual subjects' estimated parameters to derive for each subject an interval in which her effort level is predicted to lie and investigate whether her observed effort level lies in this interval.

We start by presenting the results from regressions of effort choices on individual envy and guilt parameters. Recall that we have two versions of the production game, SimProd and SeqProd (cf. Table 2.4). To test whether the predictions for the production game are supported by the observations from the lab, we ran the following regression:

$$e^T_{R,i} = \delta^T_{R,0} + \delta^T_{R,1}\alpha^{IA}_i + \delta^T_{R,2}\beta^{IA}_i + \varepsilon^T_{R,i} \tag{2.8}$$

where T =SimProd, SeqProd, $R = A, B$ represents the subject's role in the production game[23], and i is the index of a subject. α^{IA}_i and β^{IA}_i are the IA estimates of i's envy and guilt parameter (as derived from the lowest-stakes Menus 1 and 2 in Part 2), respectively. Table 2.8 gives the estimated coefficients of the δ's in (2.8). We present these results for the IA model using eqs. (2.6) and (2.7), because this is where our primary research question lies. For the parallel results for the general model (2.3) see Appendix 2.6.5.

Table 2.8: Coefficients of IA parameters in the production game

	IA Prediction		SimProd		SeqProd	
	e_A	e_B	e_A (obs. 144)	e_B (obs. 144)	e_A (obs.26)	e_B (obs. 21)
const.	0	100	29.70***	95.84***	58.10***	76.80***
			(3.24)	(1.36)	(7.69)	(8.05)
α	0	-100	-4.58	-61.28***	107.39	-245.91
			(31.85)	(13.35)	(119.07)	(157.10)
β	200	0	100.52***	-9.91	18.46	-140.98***
			(15.44)	(6.47)	(34.65)	(62.11)

The predicted coefficients are taken from eqs. (2.6) and (2.7). For SeqProd (DoubleRole and SingleRole) workers B do not enter a single effort level but effort as a function of worker A's effort. To obtain a single choice, as a consequence, only observations from subjects who are assigned to be worker B in SeqProd are used. Numbers in parentheses give standard errors. *** indicates statistical significance at the 1%-level.

Comparing the results in Table 2.8 to the theoretical predictions of the IA model

[23] Recall that in most treatments, subjects played both roles.

shows that in the simultaneous production game (SimProd) the relationships between envy/guilt and effort go in the direction of the predicted coefficients. Worker A's effort level only relies significantly (and positively) on her guilt parameter. The estimated coefficient is below the predicted level of 200, however. Worker B's effort is only affected by the envy parameter, which has the predicted negative impact on her effort level, though again the marginal effect falls short of what is predicted. For the two treatments of the sequential production game, the regression outcome differs substantially from the theoretical prediction, however. We find either that none of the predicted effects exists (for Worker A's effort), or that the significance is exactly opposite to the theoretical prediction (only the guilt levels have significant impact, for Worker B's effort).[24]

Next, we directly compare for each subject the predicted and observed effort levels. Using the estimated interval of the envy and guilt parameters derived from each subject's choices in Menu 1 and 2, eqs. (2.6) and (2.7) yield for each subject a prediction for the interval in which her effort as Worker A or B will lie. We derive these intervals for each subject and simply check whether her observed efforts in each worker role lie within the predicted intervals. This shows that in SimProd, 58.3% (84 out of 144) of the observed efforts as Worker A and 86.8% (125 out of 144) as Worker B lie in the predicted interval. The success rate is much lower in SeqProd, where only 19.23% (5 out of 26) of Worker A's efforts and 47.6% (10 out of 21) of Worker B's efforts lie in the predicted intervals.[25] We will return to this difference between the two games when discussing our results below.

The weak predictive power of the IA model in SeqProd strongly suggests that the possibility to reciprocate has an impact on decision makers' preferences for envy and guilt. Therefore, we conclude this section by presenting direct estimates of individuals' envy and guilt parameters for both SimProd and SeqProd. We use the inverse relationships of (2.6) and (2.7) to do so. The resulting distributions are displayed in Figure 2.4. For comparison, we include our estimates from Menus 1 and 2 as well as the FS and BEN distributions. For SeqProd (both the DoubleRole and

[24]For the regression in SeqProd, for Worker B's effort we use the realized levels, as obtained by applying their responding rules to their paired Worker A's effort levels. We use all the data on Worker A efforts from all the subjects who have made that decision (i.e. all subjects in DoubleRole and the Worker A's in SingleRole).

[25]For SeqProd, the data on Worker A's effort are from all IA-C&R subjects in SeqProd-DoubleRole and those who have been assigned the role of Worker A in SeqProd-SingleRole. The data used for Worker B are selected in an analogous way.

SingleRole treatments) worker B gives responding rules instead of a specific effort level. To obtain a single choice, we again use the effort level that is realized after the chosen responding rule has been applied to A's chosen effort level. As a consequence, only observations from subjects who are assigned to be worker B in SeqProd (21 observations in total) are used to estimate α; in contrast, all effort levels by workers A (also those that ex post were allocated to be a worker B) are used to estimate β.

Figure 2.4(a) shows that the simultaneous production game yields a distribution of envy that is very much like the distribution estimated with the IA model from the choices for Menu 1 (cf. Figure 2.1(a)). Hence, in SimProd, where workers B cannot reciprocate worker A's choice, the envy parameters are very similar to those estimated from Menu 1 (where reciprocating the other's choice is not feasible either). This explains why the regression results reported for SimProd in Table 2.8 are quite consistent with the theoretical prediction.

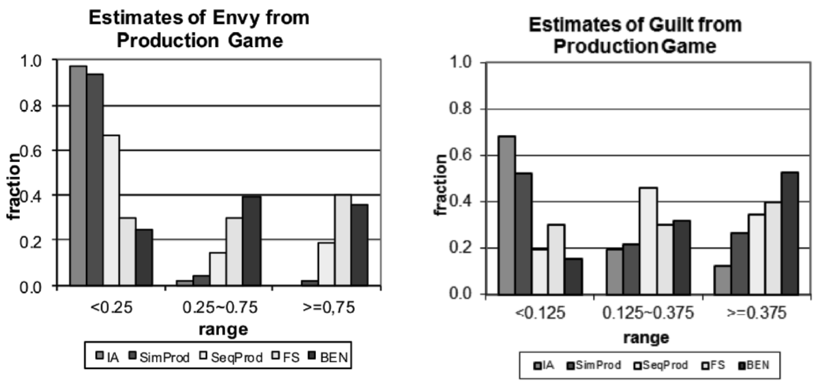

(a) (b)

Figure 2.4: Distributions of envy (a) and guilt (b) estimates derived from the production game.

The results are different for the sequential production game, where reciprocity is possible. Here, far fewer subjects (66.7% as opposed to 93.8% in SimProd) are estimated to have an envy parameter in the lowest category. We observe a substantial number of subjects showing envy levels at higher values. This seems to imply that high envy levels are triggered by (negative) reciprocity. As a result, the distribution we derive from SeqProd is closer to those reported in FS and BEN. Recall that their envy parameters were also derived form an environment where subjects may be exposed to negative reciprocity.

In Figure 2.4(b), the distribution of guilt levels derived from Worker A's effort, in SimProd and SeqProd, are compared to those in FS and BEN. Similar to envy, we observe higher guilt levels in SeqProd than in SimProd. This may be attributed to a strategic choice by A, anticipating B's negative reciprocal response if her effort level is too low.

2.5 Concluding Discussion

In this chapter, we have examined the robustness of parameter estimates and the predictive power of the IA model. Our main findings are the following. First, our estimates of disadvantageous inequity aversion (envy) and advantageous inequity aversion (guilt) are robust to allowing for efficiency concerns and also reasonably robust to increases in the stakes (though guilt seems to become less important when payoffs are scaled up very strongly). Second, while our results indicate that guilt is important for some subjects we find far less evidence of envy than has been observed in previous studies. Only when we introduce the possibility to reciprocate others choices, we observe more subjects for whom envy plays a role. Third, we observe that in settings where explicit reciprocity is ruled out, the model's predictive power is higher than what results from previous studies suggest.

The effect of reciprocation in our production game shows up in changes in both workers' behavior in the sequential production game compared to the simultaneous one. In the sequential game, the second-mover may condition her choices on her perception of the opponent's kindness or meanness, which may trigger her motivation to reciprocate or retaliate, something that is impossible in the simultaneous game. In other words, the second-mover may have reciprocal preferences that are distinct from her feelings of envy or guilt (see also Falk and Fischbacher 2006 and Charness and Rabin 2002, pp. 824-825). If this is the case, these will surface in the sequential game and choices made there may be mistakenly interpreted as evidence of inequity aversion. Similarly, the first-mover may behave strategically differently in the sequential game than in the simultaneous game, anticipating changes in the second-mover's decisions in the former case.

Fehr and Schmidt (1999) acknowledge the possibility that the parameters of their model can be interpreted in two ways, when they argue that "positive α_i's and β_i's can be interpreted as a direct concern for equality as well as a reduced-form concern for intentions. [...] As a consequence, our preference parameters are compatible with the interpretation of intentions-driven reciprocity." In our view, strategic and reciprocal tendencies should be distinguished from inequity aversion per se, however. In other words, we favor the view that reciprocal preferences should be distinguished from preferences with respect to equality. The alternative (that preferences about inequality vary with the environment, i.e., with the possibility to reciprocate) requires allowing for endogenous preferences, which would substantially reduce the predictive power of the model. In our preferred interpretation, a comparison between the estimates of envy and guilt estimates from the simultaneous and the sequential production games reveals that in an environment with explicit reciprocity, subjects' behavior will yield higher estimates of inequity aversion. Hence, the current literature (where estimates of the envy parameter are traditionally derived from responder behavior in the ultimatum game) may provide biased estimates of the envy parameter.

In short, we believe that the control for explicit reciprocity offered by our design (both in the choice menus and in the production game) is important for isolating preferences for equality and therefore for accurately measuring pure inequity aversion levels. We believe that our design provides stronger evidence for the robustness and predictive power of the IA model than was found in any previous research. Our understanding of this finding is that the IA model does a good job in explaining and predicting subjects' behavior in environments where explicit reciprocity is ruled out. However, when reciprocity is involved, the model should be augmented with a reciprocity term (e.g., as in Charness and Rabin 2002). Much of the previous literature seems to have taken an alternative 'as if' approach, in the sense that any choice that simultaneously yields lower inequity and lower own payoff is reflected in the measured inequity aversion parameters, irrespective of whether other motivations could be involved. Whether or not this is a problem in practice depends on one's goals. If one is interested in applying the model to an environment where reciprocity is deemed to be important, one can measure inequity aversion in a situation that also allows for reciprocity and interpret the resulting parameters 'as if' they measure envy and guilt *per se*. Though one may question the interpretation of the parameters of the model, the model's predictive power need not be affected. Nevertheless, from a scientific point of view the distinction between various kinds of preferences seems important.

2.6 Appendix

2.6.1 IA Predictions for the Production Game

In this section, we derive the predictions for the production game for the case where preferences are given by the IA model. Recall that the preferences (given in (2.1)) are described by:

$$U_i(x, y_1, y_2, ..., y_n) = x - \alpha_i \frac{\sum_j \max\{y_j - x, 0\}}{n} - \beta_i \frac{\sum_j \max\{x - y_j, 0\}}{n} \tag{2.9}$$

where x denotes i's own monetary payoffs and y_j denotes the monetary payoffs of others with whom i interacts.

Furthermore, recall that payoffs in the production game, are given by (see (2.5)):

$$\pi_i(e_A, e_B) = s_i + 1/2 \sum_{j=A,B} p_j(e_j) - e_i c_i, \quad i = A, B \tag{2.10}$$

where e_A and e_B denote the two workers' (A and B) effort levels and (cf. (2.4)):

$$p_i(e_i) = 4e_i - \frac{e_i^2}{100}, \quad i = A, B. \tag{2.11}$$

Then, the payoff difference between Worker A and B, denoted by Δ, is given by:

$$\Delta(e_A, e_B) \equiv \pi_A(e_A, e_B) - \pi_B(e_A, e_B) = s_A - s_B + e_B c_B - e_A c_A. \tag{2.12}$$

Given our chosen parameters ($s_A = 200$, $s_B = 0$, $c_A = 2$, $c_B = 1$), this gives

$$\Delta(e_A, e_B) = 200 + e_B - 2e_A. \tag{2.13}$$

Two conclusions follow straightforwardly from (2.13):

1. $\Delta(e_A, e_B) \geq 0$, because A and B's effort levels are limited between 0 and 100. This means that A always earns at least as much as B. Therefore, in the production game, A is always at the advantageous position compared to B, in terms of monetary payoffs.

2. The amount that A's payoff exceeds B's, $\Delta(e_A, e_B)$, satisfies, $\frac{\partial^2 \Delta(e_A, e_B)}{\partial e_A \partial e_B} = 0$, $\frac{\partial \Delta(e_A, e_B)}{\partial e_A} \leq 0$, and $\frac{\partial \Delta(e_A, e_B)}{\partial e_B} \geq 0$.

Because of property 1, we know that in A' s utility function (2.9), the envy term (i.e., the term in α) plays no role in the production game. Then, A's utility function in the production game can be simplified to:

$$U_A(\pi_A(e_A, e_B), \pi_B(e_A, e_B)) = \pi_A(e_A, e_B) - \beta_A \Delta(e_A, e_B) \tag{2.14}$$

Similarly, it follows from property 1 that the guilt term in B's utility function is irrelevant for the parameter settings in our production game, reducing B's utility function to:

$$U_B(\pi_A(e_A, e_B), \pi_B(e_A, e_B)) = \pi_B(e_A, e_B) - \alpha_B \Delta(e_A, e_B) \tag{2.15}$$

Taking the derivative of $U_A(\pi_A(e_A, e_B), \pi_B(e_A, e_B))$ w.r.t. e_A and setting this equal to zero gives:

$$\frac{\partial \pi(e_A, e_B)}{\partial e_A} = \beta_A \frac{\partial \Delta(e_A, e_B)}{\partial e_A} \tag{2.16}$$

The solution to (2.16) gives A's best response to B's effort choice eB. Substituting (2.10), (2.11), and (2.13), the LHS and RHS of (2.16) reduce to $-\frac{e_A}{100}$ and $-2\beta_A$, respectively. Therefore, (2.16) gives $e_A = 200\beta_A$. We conclude that A's best response function e_A is independent of B's choice e_B, i.e., A has a dominant strategy that only depends on her preference parameter β_A. Taking into account the restriction $e_A \in [0, 100]$, this dominant strategy is given by (cf. eq. (2.6)):

$$e_A = \begin{cases} 0 & \text{if } \beta_A \leq 0 \\ 200\beta_A & \text{if } 0 < \beta_A < 1/2 \\ 100 & \text{if } \beta_A \geq 1/2. \end{cases} \tag{2.17}$$

A similar exercise leads to B's best response. Given A's effort level , B maximizes utility specified by (2.15), yielding the FOC:

$$\frac{\partial \pi(e_A, e_B)}{\partial e_B} = \alpha_B \frac{\partial \Delta(e_A, e_B)}{\partial e_B} \tag{2.18}$$

Substituting (2.10), (2.11) and (2.13) gives $1 - \frac{e_B}{100} = \alpha_B$, which, taking into account

$e_B \in [0, 100]$, gives (cf. eq. (2.7))

$$e_B = \begin{cases} 0 & \text{if } \alpha_B \geq 1 \\ 100(1 - \alpha_B) & \text{if } 0 < \alpha_B < 1 \\ 100 & \text{if } \alpha_B \leq 0. \end{cases} \tag{2.19}$$

Because B's best response is independent of A's effort choice, we once again have a dominant strategy (for B) that only depends on her own preference parameter (α). Since each worker's dominant strategy only depends on the own preference level, beliefs about the other player's preference levels are irrelevant. In fact, if players' utility is given by the IA model (2.9), any choice of parameters in the production game that satisfies properties 1 and 2, ensures that A has a dominant strategy only dependent on her guilt level and B has a dominant strategy only dependent on her envy level.

2.6.2 General Model Predictions for the Production Game

In this section, we derive the predictions for the production game for the case where preferences are given by the general model. Recall that these preferences (given in eq. (2.3)) are described by:

$$U_i(x, y_1, y_2, ..., y_n) = x - \alpha_i \frac{\sum_j \max\{y_j - x, 0\}}{n} - \beta_i \frac{\sum_j \max\{x - y_j, 0\}}{n} + \gamma_i (x + \sum_j y_j), \tag{2.20}$$

In this mode, three factors affect a worker's utility: the own payoff, the difference between A and B's payoffs ($\Delta(e_A, e_B)$), and the aggregate payoff of the two workers, which is given by:

$$T(e_A, e_B) \equiv \pi_A(e_A, e_B) + \pi_B(e_A, e_B). \tag{2.21}$$

As explained in Appendix 2.6.1, Worker A always earns more than B in the production game and, thus, the utility functions of the two workers can be written as:

$$U_A(\pi_A(e_A, e_B), \pi_B(e_A, e_B)) = \pi_A(e_A, e_B) - \beta_A \Delta(e_A, e_B) + \gamma_A T(e_A, e_B) \tag{2.22}$$

$$U_B(\pi_A(e_A, e_B), \pi_B(e_A, e_B)) = \pi_B(e_A, e_B) - \alpha_B \Delta(e_A, e_B) + \gamma_B T(e_A, e_B) \tag{2.23}$$

Then, A's best response to B's effort level e_B is the solution to the FOC obtained from (2.22):

$$\frac{\partial \pi_A(e_A, e_B)}{\partial e_A} - \beta_A \frac{\partial \Delta(e_A, e_B)}{\partial e_A} + \gamma_A \frac{\partial T(e_A, e_B)}{\partial e_A} = 0 \tag{2.24}$$

Analogously, B's best response to A's effort e_A, is determined by

$$\frac{\partial \pi_B(e_A, e_B)}{\partial e_B} - \alpha_B \frac{\partial \Delta(e_A, e_B)}{\partial e_B} + \gamma_B \frac{\partial T(e_A, e_B)}{\partial e_B} = 0 \tag{2.25}$$

Using the chosen production function (2.11), profits (2.10) and the marginal effort costs ($c_A = 2$, $c_B = 1$), (2.24) gives after some straightforward algebra: $e_A =$

$100\left(1+\frac{-1+2\beta_A}{1+2\gamma_A}\right)$ which, given that $e_A \in [0,100]$, yields:

$$e_A = \begin{cases} 0 & \text{if } \beta_A \leq -\gamma_A \\ 100\left(1+\frac{(-1+2\beta_A)}{1+2\gamma_A}\right) & \text{if } -\gamma_A < \beta_A < 1/2 \\ 100 & \text{if } \beta_A \geq 1/2. \end{cases} \tag{2.26}$$

Considering next worker B, (2.25) yields in a similar manner: $e_B = 100\left(1.5 - \frac{0.5+\alpha_B}{1+2\gamma_B}\right)$. Using $e_B \in [0,100]$, this gives:

$$e_B = \begin{cases} 0 & \text{if } \alpha_B \geq 1+3\gamma_B \\ 100\left(1.5-\frac{0.5+\alpha_B}{1+2\gamma_B}\right) & \text{if } \gamma_B < \alpha_B < 1+3\gamma_B \\ 100 & \text{if } \alpha_B \leq \gamma_B. \end{cases} \tag{2.27}$$

As was the case for the IA model (Appendix 2.6.1), neither workers' best response depends on the other's effort level, so (2.26) and (2.27) again define dominant strategies for A and B in the production gamewhich only depend on the own preference levels, now assuming the workers' utility functions follow the general model. Moreover, envy does not enter A's optimal strategy and guilt does not affect B's. Note also that the effort of both workers is increasing in γ: the more workers care about efficiency, the more effort they will exert. This is intuitive because the greater γ the more a worker will take into account that her effort increases both players' payoffs.

In a way analogous to the analysis presented in Table 2.8 in Section 2.4, we can test the predictive power of the general model. We start with a regression, which in this case is nonlinear in the model's parameters (cf. (2.26) and (2.27)). Specifically, for Worker A/B's effort, we run the nonlinear regression as shown in (2.28).

$$e_{R,i} = 100\left(c_1 + \frac{c_2 + a\alpha^{Gen}_{lowest,i} + b\beta^{Gen}_{lowest,i}}{1+g\gamma^{Gen}_{lowest,i}}\right) + \varepsilon_i \tag{2.28}$$

where i denotes the index of each subjects and $R = A, B$ represents the role of the worker. Notice that in the regressions, the constant terms in the denominators are fixed at the theoretically predicted value, 1, for the sake of identifiability. We run the regressions separately for SimProd and SeqProd. The results are as displayed in Table 2.9. The theoretical prediction of the coefficients c_1, c_2, a, b and g listed in the table are derived by a comparison of (2.28) with (2.26) or (2.27). The estimation using the SimProd observations show that the (in)significant influence of envy and guilt is found to be consistent with the model's predictions. The results on the influence of the efficiency concerns is in contrast to what the general model suggests, however, as the estimated coefficient of γ has exactly the opposite sign than predicted. In the regression using SeqProd data, none of the preference parameters show significant influence on Worker A/B's chosen effort levels. This result is similar to what we found in the linear regression of IA model's prediction (c.f. Table 2.8 in Section 2.4).

Table 2.9: Coefficients of general model's parameters in the production game

	GEN Prediction		SimProd		SeqProd	
	e_A	e_B	e_A (144 obs)	e_B (144 obs)	e_A (26 obs)	e_B (21 obs)
c_1	1	1.5	$1.28***$	$1.46***$	$1.50***$	$1.31***$
			(0.03)	(0.01)	(0.08)	(0.10)
c_2	-1	-0.5	$-0.96***$	$-0.50***$	$-1.00***$	$-0.49***$
			(0.02)	(0.00)	(0.02)	(0.02)
$a \quad (\alpha)$	0	-1	0.07	$-0.44***$	2.32	-2.44
			(0.21)	(0.15)	$(2.80E+6)$	(2.30)
$b \quad (\beta)$	2	0	$0.62***$	-0.06	0.13	-0.97
			(0.15)	(0.06)	(0.61)	(0.66)
$g \quad (\gamma)$	2	2	$-0.92**$	$-0.33***$	0.77	-0.72
			(0.39)	(0.10)	$(3.40E+6)$	(2.74)

Note: The value displayed in the parentheses below each estimate is the standard error calculated by bootstrap.

There is no added value to apply to the general model the second method we used to test the predictive power of the IA model in the production game. This directly compares for each subject the predicted and the observed effort levels. Doing so for the general model would yield exactly the same outcome as reported in Section 2.4. This is because the production game involves a two-player environment, as do the tasks related to Menus 1 and 2. For a detailed theoretical explanation, we refer to Engelmann (2012), Section 2.

2.6.3 Experimental Instructions

Instructions (Standard Treatments)

Welcome to this experiment on decision making. You have already earned 7 Euro for showing up on time. You may make more money in today's experiment. How much more you make depends on your decisions and the decisions of other participants. You will be paid privately at the end of the session. This is an anonymous experiment: your identity will not be revealed to any other participant. This experiment has several parts. You will receive a specific instruction before each part starts.

Please read the instructions carefully. During the experiment, you are not allowed to talk to other participants or to communicate with them in any other way. If you want to ask any questions, please raise your hand.

All payments in the experiment are denoted in points. At the end of the experiment, points will be exchanged to Euros at a rate of 200 points=1 Euro.

Instructions for Part 1

In this part, participants are randomly matched into groups of 6. Within each group, one group member will be randomly selected as the proposer, while the other five will be the receivers. Because group and role assignment will not be made until the very end of the experiment, every participant in the group is asked to make a series of proposer decisions. In this way we will have your decisions in case you are appointed proposer at the end.

For each decision question, you will face two options, each involving two amounts: one is the amount of payoff for yourself, the other is the amount of the payoff for each of the 5 receivers in your group. For example, if you are proposed and choose Option B in the example below, this means that you get 10 points while all 5 receivers in you group earn 0 points.

Example:

Option A		**Option B**	
Your Payoff:	20	Your Payoff:	10
Five Others' Payoff:	20	Five Others' Payoff:	0

At the end of the experiment, we will first randomly form groups of 6. If there are not enough participants to form a new group, some of you will not be put in a group and earn nothing from this part of the experiment. For example, if there are 26 participants, we will form 4 groups of 6 (24 participants) and two randomly chosen participants will not make money from this part. After groups have been formed, we will select only one of the six members as the proposer whose decision will be implemented. Other participants' choices in the group will therefore not affect payoffs in any way. We then select one of the decisions made by the selected proposer to be paid.

Before Part 1 starts, a few questions will be asked to make sure that you have fully understood these instructions.

Please press **Continue** on the screen to proceed to these questions after you have finished reading these instructions.

Instructions for Part 2

In this part of the experiment, you will be asked to make 20 decisions (in two sets of 10). Each decision involves a choice between an Option A and an Option B, taking the form of:

Option A		**Option B**	
Your Payoff:	...	Your Payoff:	...
Others' Payoff:	...	Others' Payoff:	...

The options refer to payments to you and one of the other participants in this experiment. For each option, two amounts will be displayed: one amount that you will receive yourself, and one amount that the "Other" will receive. This other participant is one that we will randomly choose. Of course, this other participant will remain anonymous to you.

You will be given a table with 10 choices like this. After you have made those choices, you will be given a table with 10 new choices to make.

Your decision in this part will determine your payoff and the payoff of the other participant in the following way.

At the end of the entire experiment, all participants will be randomly matched into pairs. In each pair, one participant will be randomly chosen, with a probability of 50%, to be the **Proposer**, and the other will be the **Receiver**.

If you are chosen to be the Proposer, one out of your 20 decisions made in Part 2 will be randomly selected with equal probability. You will receive the amount you decided to receive in that chosen decision and your paired Passive Receiver will receive the amount you decide to give to "Other". Your paired Passive Receiver's decisions in Part 2 will have no influence at all.

If you are chosen to be the Receiver, one out of the 20 decisions made by your paired Proposer in Part 2 will be randomly selected with equal probability. Your paired Proposer will receive the amount he/she decided to receive in that chosen decision and you will receive the amount he/she decided to give to "Other". Your decisions in Part 2 will have no impact at all.

Because all the random pairing and role assignments will be done at the very end of the experiment, no one knows what role he/she will play in this part of the experiment. Therefore, when deciding between options there is a 50% chance that your decisions will actually be used to determine your payoff and that of a randomly chosen other participant.

Instructions for Part 3[26]

Your task in Part 3 is very much like in Part 2. You will again be given two sets of 10 choices between Option A and Option B. The numbers in these sets are the numbers from Part 2 multiplied by t, so all outcomes are now t times higher.[27]

Participation in Part 3 is optional. You have to decide whether or not you want to participate.

However, you may participate in Part 3 only if you give up the earnings that you may have made in Parts 1 and 2. If you choose not to participate in Part 3, you will have to wait until all the Part 3 participants have finished, and then continue to the next part of the experiment together with them.

If the number of the Part 3 participants is odd, the system cannot combine all of them in pairs at the end of the experiment. In this case, the system will randomly select one. This person will not make money from decisions made in this part, but will instead receive 500 points as his/her payoff in this part.

Instructions for Part 4 (SimProd)

[26] For this part of the experiment, we did not provide handouts; subjects read the text from their computer monitor only.

[27] We put in the instructions t=10, 30, 60, for the "low", "high", and "highest" treatments, respectively.

This is the last part of the experiment. In this part, you will participate in a Production Game.

The production game involves two workers, Worker A and Worker B, who are in charge of Department 1 and 2, respectively. Each worker chooses an effort level (an integer between 0 and 100), which will determine the production of the department he/she is in charge of. A worker's total income from this production game consists of four parts: (1) Basic salary; (2) A bonus dependent on the production of Department 1; (3) A bonus dependent on the production of Department 2; (4) Effort cost, which is dependent on one's own effort level. We discuss each of these in turn.

1. **Basic salary**. The basic salary is x points for Worker A and 0 points for Worker B.[28]

2. **Bonus 1**. The production from Department 1 will be equally divided between Worker A and Worker B as Bonus 1. The production is totally determined by Worker A's effort. The higher the effort level Worker A chooses, the more Department 1 produces, and, hence the more Bonus 1 both Worker A and Worker B will get.

3. **Bonus 2**. The production from Department 2 will be equally divided between Worker A and Worker B as Bonus 2. The production is totally determined by Worker B's effort level. The higher the effort level Worker B chooses, the more Department 2 produces, and, hence the more Bonus 2 both workers will get.

4. **Effort cost**. A worker bears a cost for each unit of effort input into the department's production. Each unit of effort in Department 1 costs c_A points for Worker A. Each unit of effort in Department 2 costs c_B points for Worker B.

For a worker, the total payoff from the production game is

Total income = Basic salary + Bonus 1 + Bonus 2 - Effort cost.

Please note that, because Worker A's basic salary is x while Worker B's is 0, the total income for Worker A is always higher than Worker B, no matter what effort

[28]In the printouts, we replaced x, c_A, c_B with the value of $200t$, $2t$, and t, respectively, with t=1, 10, 30 and 60 for the "lowest", "low", "high", and "highest" treatments, respectively.

levels Worker A and B choose. Of course, the difference varies as Worker A and Worker B choose different effort levels.

During the experiment, after you enter an effort level, you will immediately see the corresponding potential Bonus amounts and effort cost displayed. You can try out different effort levels. When you finalize your decisions, make sure the numbers in the blanks are your decisions, and press "Submit" at the bottom of the page.

In this part, you will be randomly paired and assigned the role of Worker A or Worker B. The result of the random pairing and role assignment will not be revealed until Part 4 has finished. For this reason, every participant is asked to make a decision as a Worker A and as a Worker B. At the end of the experiment, only your decision on Worker A's effort level will apply if you are assigned the role of the Worker A, otherwise, you are assigned the role of Worker B, and only your decision on Worker B's effort level will be taken into account.

Instructions for Part 4 (SeqProd-DoubleRole)

This is the last part of the experiment. In this part, you will participate in a Production Game.

Production Game

The production game involves two workers, Worker A and Worker B, who are in charge of Department 1 and 2, respectively. Each worker chooses an effort level (an integer between 0 and 100), which will determine the production of the department he/she is in charge of. A worker's total income from this production game consists of four parts:

(1) Basic salary;

(2) A bonus dependent on the production of Department 1;

(3) A bonus dependent on the production of Department 2;

(4) Effort cost, which is dependent on one's own effort level.

We discuss each of these in turn.

1. Basic salary. The basic salary is 2000 for Worker A and 0 for Worker B.

2. Bonus 1. The production from Department 1 will be equally divided between Worker A and Worker B as Bonus 1. The production is totally determined

by Worker A's effort. The higher the effort level Worker A chooses, the more Department 1 produces, and, hence the more Bonus 1 both Worker A and Worker B will get.

3. Bonus 2. The production from Department 2 will be equally divided between Worker A and Worker B as Bonus 2. The production is totally determined by Worker B's effort level. The higher the effort level Worker B chooses, the more Department 2 produces, and, hence the more Bonus 2 both workers will get.

4. Effort cost. A worker bears a cost for each unit of effort input into the department's production. Each unit of effort in Department 1 costs 20 points for Worker A. Each unit of effort in Department 2 costs 10 points for Worker B.

For each worker, the total payoff from the production game is

Total income = Basic salary + Bonus 1 + Bonus 2 - Effort cost.

Please note that, because Worker A's basic salary is 2000 while Worker B's is 0, the total income for Worker A is always higher than Worker B, no matter what effort levels Worker A and B choose. Of course, the difference varies as Worker A and Worker B choose different effort levels.

In this game, a Worker A decides on the effort level she/he would like to put into Department 1. Meanwhile, a Worker B decides on a set of **'responding rules'**. These specify her/his effort level, in response to different possible efforts chosen by Worker A.

A set of responding rules consists of a number of rules. Each rule has three elements: a **Lower Bound (LB)**, an **Upper Bound (UB)**, and an **Effort Level (E)**, meaning that "When the paired Worker A's effort level is between **LB and UB**, my effort level is set to be **E**".

Worker B is free to choose any number of responding rules. At one extreme, if LB=UB for every rule, he/she chooses 101 different effort levels, one for each possible effort by Worker B (which could be any integer between 0 and 100). At the other extreme, if LB=0 and UB=100, Worker B can choose one single effort level, independently of Worker A's effort choice. Any set of rules between these two extremes is possible.

For example, Worker B specifies the responding rules as shown in Figure 1. According to Figure 1, Worker B will choose effort 1 if Worker A chooses an effort level between 0 and 8. Similarly, Worker B's effort will be 13 if Worker A's effort lies between 9 and 50. Finally, Worker B's effort level will be equal to 73 for all Worker A's effort levels between 51 and 100.

Meanwhile, the paired Worker A chooses an effort level of 10. Therefore, Worker A's effort level lies within the range of the second rule set by the paired Worker B, [9, 50], so the Worker B's effort in Department 2 is determined as the effort level indicated in the second rule, 13.

Responding Rules

Lower Bound	Upper Bound	Effort
0	8	1
9	50	13
51	100	73

Figure 1

In this part, you will be randomly paired and assigned the role of either Worker A or Worker B. The result of the random pairing and role assignment will not be revealed until Part 4 has finished. For this reason, every participant is asked to make a decision as a Worker A and as a Worker B. If you are assigned the role of the Worker A at the end of the experiment, only your decision on Worker A's effort level will apply. If you are assigned the role of Worker B, only your decision on Worker B's effort level will be taken into account.

How to specify your decisions

Worker A's Decision: When at the top of the page it says 'Stage of Worker A's Decision', please specify the effort level you would like to choose in Department 1, if you are assigned the role of Worker A. In this stage, a 'Decision Box' and a 'Trial Box' will be displayed on the screen. In each box, you will find a scroll bar which you can use to try different effort levels of Worker A or Worker B. You will immediately see the corresponding (potential) Bonus amounts and effort costs displayed. When you finalize your decision, make sure the number displayed as 'Worker A's effort level' in the Decision Box is your decision on Worker A's effort, and press "Submit" at the bottom of the page.

Worker B's Decision: When at the top of the page it says 'Stage of Worker B's Decision', please specify a set of responding rules, as if you are assigned the role of Worker B. You will see two 'Trial Boxes', each containing a scroll bar which you can use to try different effort levels of Worker A and Worker B. You will immediately see the corresponding (potential) Bonus amounts and effort costs displayed.

The lower part of the screen, as shown in Figure 2, is a 'Decision Box', which is composed of two parts. The left part is a Rule Creator, which you can use to create rules. You create a rule by specifying two numbers, an Upper Bound and an Effort Level. You do not need to specify the Lower Bound for a rule. Based on your other rules, the lower bound will be automatically generated. If applicable, it will also be updated as you create other rules or make further changes to existing rules.

For example, assume that the existing responding rules are as shown in Figure 2. If you create a new rule of which the Upper Bound is 50 and the Effort level is 13, the new set of responding rules will be as shown in Figure 3: the Lower Bound of the new rule is automatically set to be 9. In addition, the Lower Bound of the second rule in the old list shown in Figure 2 (the third rule in the new list, in Figure 3) is automatically updated to be 51. We suggest using pencil and paper for trying out rules before you enter your responding rules into the computer. You may, however change any previous rule until you have finalized a complete set of responding rules.

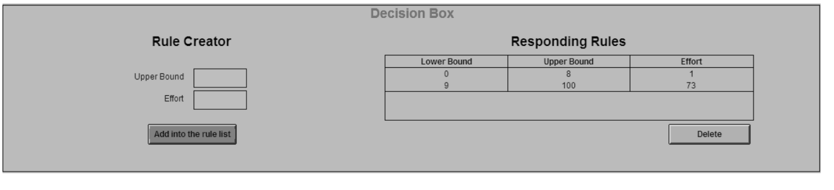

Figure 2

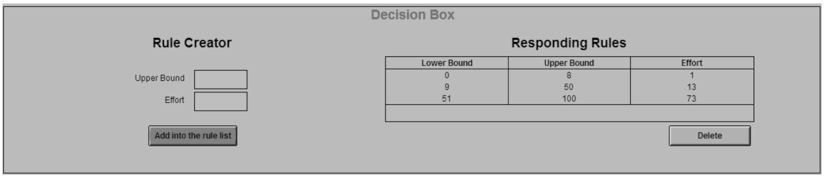

Figure 3

The right part of the Decision Box displays all existing responding rules that you have set. As long as you have not finalized a complete set, you can delete any rule from the list by first selecting it and then pressing the "Delete" button. For example, if the existing responding rules are as shown in Figure 3, and you delete the second rule, then the responding rules will be updated to be as shown in Figure 2.

When you make your final decision, make sure the responding rules displayed in the Decision Box are your decisions, and press "Submit" at the bottom of the page.

Instructions for Part 4 (SeqProd-SingleRole, Worker A)

This is the last part of the experiment. In this part, you will participate in a Production Game.

Production Game

The production game involves two workers, Worker A and Worker B, who are in charge of Department 1 and 2, respectively. Each worker chooses an effort level (an integer between 0 and 100), which will determine the production of the department he/she is in charge of. A worker's total income from this production game consists of four parts:

(1) Basic salary;

(2) A bonus dependent on the production of Department 1;

(3) A bonus dependent on the production of Department 2;

(4) Effort cost, which is dependent on one's own effort level.

We discuss each of these in turn.

1. **Basic salary**. The basic salary is 2000 for Worker A and 0 for Worker B.

2. **Bonus 1**. The production from Department 1 will be equally divided between Worker A and Worker B as Bonus 1. The production is totally determined by Worker A's effort. The higher the effort level Worker A chooses, the more Department 1 produces, and, hence the more Bonus 1 both Worker A and Worker B will get.

3. **Bonus 2**. The production from Department 2 will be equally divided between Worker A and Worker B as Bonus 2. The production is totally determined by Worker B's effort level. The higher the effort level Worker B chooses, the more Department 2 produces, and, hence the more Bonus 2 both workers will get.

4. **Effort cost**. A worker bears a cost for each unit of effort input into the department's production. Each unit of effort in Department 1 costs 20 points for Worker A. Each unit of effort in Department 2 costs 10 points for Worker B.

For each worker, the total payoff from the production game is

Total income = Basic salary + Bonus 1 + Bonus 2 - Effort cost.

Please note that, because Worker A's basic salary is 2000 while Worker B's is 0, the total income for Worker A is always higher than Worker B, no matter what effort levels Worker A and B choose. Of course, the difference varies as Worker A and Worker B choose different effort levels.

In this game, a Worker A decides on the effort level she/he would like to put into Department 1. Meanwhile, a Worker B decides on a set of 'responding rules'. These specify her/his effort level, in response to different possible efforts chosen by Worker A.

A set of responding rules consists of a number of rules. Each rule has three elements: a **Lower Bound (LB)**, an **Upper Bound (UB)**, and an **Effort Level (E)**, meaning that "When the paired Worker A's effort level is between LB and UB, my effort level is set to be E".

Worker B is free to choose any number of responding rules. At one extreme, if LB=UB for every rule, he/she chooses 101 different effort levels, one for each possible effort by Worker B (which could be any integer between 0 and 100). At the other extreme, if LB=0 and UB=100, Worker B can choose one single effort level, independently of Worker A's effort choice. Any set of rules between these two extremes is possible.

For example, Worker B specifies the responding rules as shown in Figure 1:

According to Figure 1, Worker B will choose effort 1 if Worker A chooses an effort level between 0 and 8. Similarly, Worker B's effort will be 13 if Worker A's effort lies between 9 and 50. Finally, Worker B's effort level will be equal to 73 for all Worker A's effort levels between 51 and 100.

Meanwhile, the paired Worker A chooses an effort level of 10. Therefore, Worker A's effort level lies within the range of the second rule set by the paired Worker B, [9,50], so the Worker B's effort in Department 2 is determined as the effort level indicated in the second rule, 13.

Responding Rules

Lower Bound	Upper Bound	Effort
0	8	1
9	50	13
51	100	73

Figure 1

You have been assigned the role of Worker A.

In this stage, a 'Decision Box' and a 'Trial Box' will be displayed on the screen. In each box, you will find a scroll bar, which you can use to try different effort levels of Worker A or Worker B. You will immediately see the corresponding (potential) Bonus amounts and effort costs displayed. When you make your final decision, make sure the number displayed as 'Worker A's effort level' in the **Decision Box** is your decision, and press "Submit" at the bottom of the page.

Instructions for Part 4 (SeqProd-SingleRole, Worker B)

This is the last part of the experiment. In this part, you will participate in a Production Game.

Production Game

The production game involves two workers, Worker A and Worker B, who are in charge of Department 1 and 2, respectively. Each worker chooses an effort level (an integer between 0 and 100), which will determine the production of the department he/she is in charge of. A worker's total income from this production game consists of four parts:

(1) Basic salary;

(2) A bonus dependent on the production of Department 1;

(3) A bonus dependent on the production of Department 2;

(4) Effort cost, which is dependent on one's own effort level.

We discuss each of these in turn.

1. **Basic salary**. The basic salary is 2000 for Worker A and 0 for Worker B.

2. **Bonus 1**. The production from Department 1 will be equally divided between Worker A and Worker B as Bonus 1. The production is totally determined by Worker A's effort. The higher the effort level Worker A chooses, the more Department 1 produces, and, hence the more Bonus 1 both Worker A and Worker B will get.

3. **Bonus 2**. The production from Department 2 will be equally divided between Worker A and Worker B as Bonus 2. The production is totally determined by Worker B's effort level. The higher the effort level Worker B chooses, the more Department 2 produces, and, hence the more Bonus 2 both workers will get.

4. **Effort cost**. A worker bears a cost for each unit of effort input into the department's production. Each unit of effort in Department 1 costs 20 points for Worker A. Each unit of effort in Department 2 costs 10 points for Worker B.

For each worker, the total payoff from the production game is

Total income = Basic salary + Bonus 1 + Bonus 2 - Effort cost.

Please note that, because Worker A's basic salary is 2000 while Worker B's is 0, the total income for Worker A is always higher than Worker B, no matter what effort levels Worker A and B choose. Of course, the difference varies as Worker A and Worker B choose different effort levels.

In this game, a Worker A decides on the effort level she/he would like to put into Department 1. Meanwhile, a Worker B decides on a set of 'responding rules'. These specify her/his effort level, in response to different possible efforts chosen by Worker A.

A set of responding rules consists of a number of rules. Each rule has three elements: a **Lower Bound (LB)**, an **Upper Bound (UB)**, and an **Effort Level (E)**, meaning that "When the paired Worker A's effort level is between LB and UB, my effort level is set to be E".

Worker B is free to choose any number of responding rules. At one extreme, if LB=UB for every rule, he/she chooses 101 different effort levels, one for each possible effort by Worker B (which could be any integer between 0 and 100). At the other extreme, if LB=0 and UB=100, Worker B can choose one single effort level, independently of Worker A's effort choice. Any set of rules between these two extremes is possible.

For example, Worker B specifies the responding rules as shown in Figure 1:

Responding Rules

Lower Bound	Upper Bound	Effort
0	8	1
9	50	13
51	100	73

Figure 1

According to Figure 1, Worker B will choose effort 1 if Worker A chooses an effort level between 0 and 8. Similarly, Worker B's effort will be 13 if Worker A's effort lies between 9 and 50. Finally, Worker B's effort level will be equal to 73 for all Worker A's effort levels between 51 and 100.

Meanwhile, the paired Worker A chooses an effort level of 10. Therefore, Worker A's effort level lies within the range of the second rule set by the paired Worker B, [9,50], so the Worker B's effort in Department 2 is determined as the effort level indicated in the second rule, 13.

You have been assigned the role of Worker B.

On the upper part of the screen, you will find two trial boxes, each containing a scroll bar which you can use to try different effort levels of Worker A or Worker B. You will immediately see the corresponding (potential) Bonus amounts and effort costs displayed.

The lower part of the screen, as shown in Figure 2, is composed of two parts. The left part is a Rule Creator, which you can use to create rules. You create a rule by specifying two numbers, an Upper Bound and an Effort Level. You do not need to specify the Lower Bound for a rule. Based on your other rules, the lower bound will be automatically generated. If applicable, it will also be updated as you create other rules or make further changes to existing rules.

For example, assume that the existing responding rules are as shown in Figure 2. If you create a new rule of which the Upper Bound is 50 and the Effort level is 13, the new set of responding rules will be as shown in Figure 3: the Lower Bound of the new rule is automatically set to be 9. In addition, the Lower Bound of the second rule in the old list shown in Figure 2 (the third rule in the new list, in Figure 3) is automatically updated to be 51.

We suggest using pencil and paper for trying out rules before you enter your responding rules into the computer. You may, however, change any previous rule until you have finalized a complete set of responding rules.

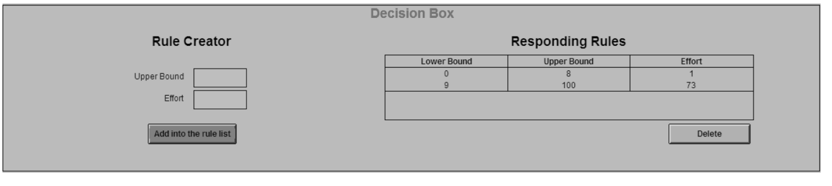

Figure 2

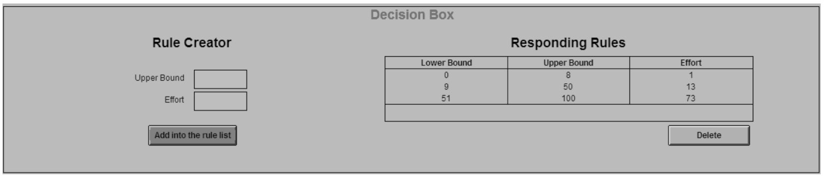

Figure 3

The right part of the figures displays all existing responding rules that you have set. As long as you have not finalized a complete set, you can delete any rule from the list by first selecting it and then pressing the “Delete” button. For example, if the existing responding rules are as shown in Figure 3, and you delete the second rule, then the responding rules will be updated to be as shown in Figure 2.

When you make your final decision, make sure the responding rules displayed in the **Decision Box** are your decisions, and press “Submit” at the bottom of the page.

2.6.4 Estimating the IA Model's Parameters from Menus 1 and 2

Using Subject j's decisions in Menus 1 and 2, we estimate her parameters α_j and β_j from the IA model (2.1) in the following three consecutive steps. Before we present these steps, it is useful to define the values listed in the last column of Menu i for decision k as v_{ik}. To do so formally, denote by $o_{ikA} = (x_{ikA}, y_{ikA})$, $o_{iB} = (x_{ikB}, y_{ikB})$ the payoff pairs to individuals x (the decision maker) and y in Menu i $(= 1, 2)$, decision $k(= 1, ..., 10)$ if option A or B, respectively, is chosen. Then, for Menu 1 we have:

$$v_{1k} \equiv \min_{\alpha : U(o_{1kA}) \geq U(o_{1kB})} \alpha \tag{2.29}$$

For Menu 2:

$$v_{2k} \equiv \min_{\beta : U(o_{1kA}) \geq U(o_{1kB})} \beta \tag{2.30}$$

Step 1: Define for each Menu $i = 1, 2$ the 'switching point', σ_{ij} as the first decision for which the Subject j chooses option B in the menu. If j always chooses A, her switching point equals 11. Hence $\sigma_{ij} \in \{1, ..., 11\}$. If j chooses A at some point in the menu after choosing B earlier in the menu, define $\sigma_{ij} = -\infty$.

Step 2: All j for whom $\exists i$, such that $\sigma_{ij} = -\infty$ are categorized as 'IA-inconsistent', implying that no α, β exist that allow her choices to be rationalized by the IA model. Denote the set of 'IA consistent' subjects by C $(= \{j | \sigma_{ij} > 0, \forall i\})$.

Step 3: Let $j \in C$, Any $\sigma_{ij} \in \{2, ..., 10\}$ can be rationalized by a finite range of parameter values in the IA model. In particular, an upper bound for the relevant parameter (α in Menu 1 and β in Menu 2) is given by v_{ik}.[29] This is denoted by $\bar{\tilde{\alpha}}_j = v_{i\sigma_{ij}}$,if $i = 1$; and $\bar{\tilde{\beta}}_j = v_{i\sigma_{ij}}$if $i = 2$. A lower bound of the IA estimate is given by $\underline{\tilde{\alpha}}_j = v_{i\sigma_{ij}-1}$, if $i = 1$ and $\underline{\tilde{\alpha}}_j = v_{i\sigma_{ij}-1}$ if $i = 2$. As our estimate of the envy and guilt parameter for Subject j we choose the average of the upper and lower bounds: $\alpha_j^{IA} = \frac{\bar{\tilde{\alpha}}_j + \underline{\tilde{\alpha}}_j}{2}$ and $\beta_j^{IA} = \frac{\bar{\tilde{\beta}}_j + \underline{\tilde{\beta}}_j}{2}$, respectively.

If $\sigma_{ij} = 1$, we define $v_{i\sigma_{ij}-1} = -1$, $i = 1, 2$. $v_{i\sigma_{ij}}$ still determines the upper bound for the IA estimates, $\bar{\tilde{\alpha}}_j$ $[\bar{\tilde{\beta}}_j]$. Because the lower bound $\underline{\tilde{\alpha}}_j$ $[\underline{\tilde{\beta}}_j]$ is $-\infty$, for $i = 1[2]$, we cannot take the mean as an estimate. Instead, we choose the upper bound as the IA estimate in this case, i.e., $\alpha_j^{IA} = \bar{\tilde{\alpha}}_j$ $[\beta_j^{IA} = \bar{\tilde{\beta}}_j]$.

[29] For example, v_{13} gives the lowest for which someone with inequity averse preferences will choose option A in decision 3 of Menu 1. It therefore gives the highest α for which she will choose B.

Finally, if σ_{ij}=11, then we still have lower bound $\underline{\tilde{\alpha}}_j$ $[\underline{\tilde{\beta}}_j]$=$v_{i\sigma_{ij}-1=v_{i10}}$, but the upper bounds $\bar{\tilde{\alpha}}_j$ $[\bar{\tilde{\beta}}_j]$=$+\infty$, if $i = 1[2]$. Hence, we choose the value of the lower bound as the IA estimate in this case, that is, $\alpha_j^{IA} = \underline{\tilde{\alpha}}_j$ $[\beta_j^{IA} = \underline{\tilde{\beta}}_j]$.

2.6.5 Estimating the General Model's Parameters from Menus 1-3

Using Subject j's decisions in Menus 1-3, we first estimate her parameters α_j and γ_j from the general model (2.3) in six consecutive steps. These steps separate the envy and efficiency parameters from the guilt parameter. We can do so because menus 1 and 3 both involve no advantageous inequity. Hence, of the three pairs of inequalities for the parameters α, β and γ that can be derived from the three switching points in menus 1-3, two pairs are only in α and γ and can be solved independently before deriving an interval for β from the other pair of inequalities.

Step 1: Define for each Menu $i = 1, 2, 3$ the 'switching point', σ_{ij}, as the first decision for which the Subject j chooses option B in the menu. If j always chooses A, her switching point equals 11. Hence the $\sigma ij \in \{1, ..., 11\}$. If j chooses A at some point in the menu after choosing B earlier in the menu, define $\sigma_{ij} = -\infty$.

Step 2: All j for whom $\exists i$, such that $\sigma_{ij} = -\infty$ are categorized as 'GM-inconsistent', implying that no α, β, γ exist that allow her choices to be rationalized by the general model. Denote the set of 'GM-consistent' subjects by $\mathcal{G}$ ($\equiv \{j | \sigma ij \geq 0, \forall i\}$).

Step 3: Let $j \in \mathcal{G}$. Find the set of (α, γ)-pairs that rationalizes $(\sigma_{1j}, \sigma 3j)$ if j's utility is given by eq. (2.3). If $(\sigma_{1j}, \sigma 3j) \in \{2, ..., 10\} \times \{2, ..., 10\}$, this yields four inequalities that $(\sigma_{1j}, \sigma 3j)$-pairs have to satisfy. Analogously, if $\sigma_{1j} \in \{1, 11\}$ or $\sigma_{3j} \in \{1, 11\}$, there are two or three inequalities that $(\sigma_{1j}, \sigma 3j)$-pairs need to satisfy. For all these cases, we denote Θ_j as the set that contains all the $(\sigma_{1j}, \sigma 3j)$-pairs that satisfy these inequalities, i.e., the pairs that rationalize j's choices using the general model. Denote by $\underline{\alpha}_j$ ($\bar{\alpha}_j$) the lowest (highest) value of α that satisfies these inequalities, i.e. $\underline{\alpha}_j = \min_{(\alpha,\gamma)\in\Theta_j} \alpha$, $\bar{\alpha}_j = \max_{(\alpha,\gamma)\in\Theta_j} \alpha$ and define $\underline{\gamma}_j$ ($\bar{\gamma}_j$) in an analogous way.

For $\sigma_{3j} < 11$, there are four possible cases:

1. If for j's choices, all bounds are finite, define $\alpha_{j0} \equiv \frac{\underline{\alpha}_j + \bar{\alpha}_j}{2}$; $\gamma_{j0} \equiv \frac{\underline{\gamma}_j + \bar{\gamma}_j}{2}$

2. If exactly one of the bounds is infinite, take the value of the related finite bound as the initial estimate, i.e. $\alpha_{j0} = \bar{\alpha}_j$, if $\underline{\alpha}_j = -\infty$ [$\gamma_{j0} = \bar{\gamma}_j$, if $\underline{\gamma}_j = -\infty$]; $\alpha_{j0} = \underline{\alpha}_j$, if $\bar{\alpha}_j = +\infty$ [$\gamma_{j0} = \underline{\gamma}_j$, if $\bar{\gamma}_j = -\infty$].

3. If both bounds are infinite, we have no information about the parameters,[30] and we set $\alpha_{j0} = \gamma_{j0} = -\infty$.[31]

4. If no (α, γ)-pair satisfies the set of inequalities, i.e. if $\Theta_j = \varnothing$, mark j as 'non-GM-rationalizable', (i.e., j's choices are behaviorally inconsistent with the general model). Denote the set of subjects with GM-rationalizable choices by $\mathcal{R}$. α_{j0} and γ_{j0} are not defined for $j \notin \mathcal{R}$. In our experiments, all subjects' choices were GM-rationalizable.

If $\sigma_{3j} = 11$ (j chooses only A in Menu 3), we set $\gamma_{j0} = 0$, which is a value that is always within the range that satisfies the set of inequalities.[32] We let $\underline{\alpha}_j$ ($\bar{\alpha}_j$) denote the lowest (highest) value of α that satisfies two inequalities rationalizing j's choice in Menu 1, given preferences (2.3) and $\gamma_j = 0$: $\underline{\alpha}_j = \min_{(\alpha,\gamma)\in\Theta_j} \alpha$, , $\bar{\alpha}_j = \max_{(\alpha,\gamma)\in\Theta_j} \alpha$. Again let $\alpha_{j0} \equiv \frac{\underline{\alpha}_j + \bar{\alpha}_j}{2}$. Note that this ensures that the α-estimates from the general model are the same as for the IA model, when $\gamma_j = 0$.

<u>Step 4</u>: $\forall j \in \mathcal{R}$, if $\alpha_{j0} < 0$, set $\alpha_{j1} = 0$. For all other j, set $\alpha j1 = \alpha_{j0}$. $\forall j \in \mathcal{R}$, set $\gamma_{j1} = \gamma_{j0}$. This step is taken to satisfy the Fehr-Schmidt (1999) condition that envy is non-negative ($\alpha \geq 0$).

<u>Step 5</u>: Next, we adjust the estimates to ensure monotonicity in the relationship between switching point σ_{ij} and parameter estimate (e.g., to ensure that a choice to give up private earnings and decrease inequity does not yield a decreased estimate of inequity aversion). Formally, if $\sigma_{1j} > \sigma_{1k}$ and $\alpha_{j1} < \alpha_{k1}$, set $\alpha_{j2} = \alpha_{k2} = \alpha_{k1}$; and if $\sigma_{3j} > \sigma_{3k}$ and $\gamma_{j1} > \gamma_{k1}$, set $\gamma_{k2} = \gamma_{j2} = \gamma_{j1}$. For all other j, set $\alpha_{j2} = \alpha_{j1}$, $\gamma_{j2} = \gamma_{j1}$.

<u>Step 6</u>: If $(\alpha_{j2}, \gamma_{j2})$, applied to general model (2.3) rationalizes σ_{1j} and σ_{3j}, set $\alpha_j^{GM} = \alpha_{j2}$ and $\gamma_j^{GM} = \gamma_{j2}$. If not, define $\Xi_j = \{(\alpha, \gamma) | \alpha \geq \alpha_{j2}; \gamma \geq \gamma_{j2}\}$, and set $\alpha_j^{GM} = \min_{(\alpha,\gamma)\in\Theta_j \bigcap \Xi_j} \alpha$ and $\gamma_j^{GM} = \min_{(\alpha,\gamma)\in\Theta_j \bigcap \Xi_j} \gamma$,[33] that is, use the lowest values of (α, γ) that are consistent with both the general model and the monotonicity of the estimates, if $\Theta_j \bigcap \Xi_j \neq \varnothing$. Otherwise, i.e., when $\Theta_j \bigcap \Xi_j = \varnothing$, we denote Subject j

[30] We do have information on the area in which the (α, γ) pairs lie in the $\alpha - \gamma$ space, but with one of them unfixed, we cannot draw any restriction on the other.

[31] At this stage, the precise values of α, γ are not important, for these values will be changed in later steps.

[32] If $\sigma_{3j} = 11$, it holds for any σ_{1j} that $\underline{\gamma}_j = 0.5$, or $\underline{\gamma}_j = -\infty$, and $\bar{\gamma}_j > 0$.

[33] Notice there is always a unique solution to these two problems, because the boundaries of $\Theta_j \bigcap \Xi_j$ are always monotonically increasing functions of γ in α.

as non-IA-rationalizable. We impose $\alpha_j^{GM} = 0$ and $\gamma_j^{GM} = 0$ for these subjects, but exclude these subjects' observations in the data analysis in Section 2.4.

Finally, we estimate j's guilt parameter β_j^{GM} from the general model (2.3) in a seventh step:

Step 7: For any $j \in \mathcal{G}$, for given γ_j, j's switching point in Menu 2, determines a lower bound $\underline{\beta}_j$ and an upper bound $\bar{\beta}_j$). For $\sigma_{2j} \in \{2, ..., 10\}$, both bounds are finite and we set $\beta_j^{GM} \equiv \frac{\underline{\beta}_j + \bar{\beta}_j}{2}$, For $\sigma_{2j} = 1$, $\underline{\beta}_j = \infty$ and we set $\beta_j^{GM} = \bar{\beta}_j$, whereas $\sigma_{2j} = 11$ gives $\bar{\beta}_j = +\infty$ and we set $\beta_j^{GM} \equiv \underline{\beta}_j$.

2.6.6 Consistency of Estimates Derived from the Two Models

To test consistency, we compare j's parameters estimated via the IA model those estimated via the general model. We believe a direct comparison of point estimates for α and β derived from the two models to be inappropriate for this purpose, however. This is because each model only identifies from the decisions for each parameter an interval of possible values. Simply comparing points derived from each interval may yield conclusions biased by the point chosen to represent the interval and not directly related to the measured inequity aversion levels. Hence, for the consistency check, we instead propose methods, which are based on the intervals that are consistent with the respective models. For the IA model, these intervals can be directly derived from the final columns in Menus 1 and 2, respectively. We define these as $[\underline{\tilde{\alpha}}, \bar{\tilde{\alpha}}]$ and $[\underline{\tilde{\beta}}, \bar{\tilde{\beta}}]$. For example, Menu 2 in Section 2.3 shows that if $\sigma_{1j} = 3$, then $[\underline{\tilde{\alpha}}, \bar{\tilde{\alpha}}] = [-0.04, 0.05]$. For the general model, these intervals are given by $[\underline{\alpha}, \bar{\alpha}]$ and $[\underline{\beta}, \bar{\beta}]$.

Then, j's decisions in Menus 1 [2] and 3 are marked as 'inconsistent across models' if the intersection of the intervals is empty, i.e., if $[\underline{\tilde{\alpha}}, \bar{\tilde{\alpha}}] \bigcap [\underline{\alpha}, \bar{\alpha}] = \varnothing$ or $[\underline{\tilde{\beta}}, \bar{\tilde{\beta}}] \bigcap [\underline{\beta}, \bar{\beta}] = \varnothing$. Note that this is a very 'lenient' binary measure of consistency. As long as a parameter value exists that allow both models to rationalize j's choices, the estimates are considered to be consistent.[34] It turns out that, for 92.1% of the 216 IA-consistent subjects (and 88.8% of the 179 IA-C&R subjects), the models are consistent in this way, meaning that for them, the IA model is robust to including efficiency concerns.

For a more precise measure of the consistency of the two models we define the **Overlap Ratio**, Ω_ω, for $\omega = \alpha, \beta$. This measures the average overlap of the intervals of α/β measured using models (2.1) and (2.3) as a percentage of the intervals measured by the IA model (2.1). This gives, for any Subject j:

$$\Omega_{\omega,j} \equiv \frac{|[\underline{\tilde{\omega}_j}, \bar{\tilde{\omega}}_j] \bigcap [\underline{\omega_j}, \bar{\omega}_j]|}{|[\underline{\omega_j}, \bar{\omega}_j]|} \cdot 100\%$$

where $\omega = \alpha, \beta$ indicates for which parameter the overlap ratio is calculated.

[34] On the other hand, minor errors by subjects in one of the menus will easily yield inconsistency. In this respect, the measure is less lenient.

If the interval ranges from $-\infty$ to $+\infty$, we set the interval's lower and upper bounds to be 0.833 and +106.5, respectively. These are the lowest finite lower bound and highest finite upper bound that are observed. This step is a straightforward normalization, intended to restrict the length of all intervals to finite numbers.

The overall overlap ratio from the data is then defined as

$$\Omega_\omega = \frac{\sum_{j \in \mathcal{R}} \Omega_{\omega,j}}{||\mathcal{R}||}, \quad \omega = \alpha, \beta$$

Where $||\mathcal{R}||$ is the total number of GM-rationalizable subjects.

2.6.7 Estimates for γ

Though efficiency concerns are not the main topic addressed in this paper, our methods do provide us with estimates of γ derived from the partial efficiency model (eq. (2.2) and (2.31)).

$$U(x, y_1, y_2, ..., y_n) = x + \gamma(x + \sum_j y_j) \tag{2.31}$$

To obtain estimates, we use a method analogous to the estimation of α and β described in Appendix 2.6.4, and the general model (2.20), as described in Appendix 2.6.5.

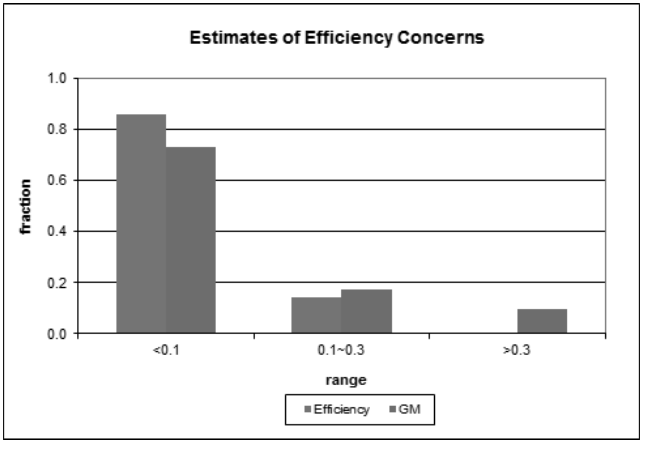

Figure 2.5: Distribution of the estimated IA-C&R subjects' efficiency concerning levels γ, from the efficiency model and the general model

Figure 2.5 compares the distributions of values derived in this way. The dominant observation from the figure seems to be that efficiency concerns play only a minor role in the choices made by our subjects, though correcting for inequity aversion does yield slightly increased estimates of γ. A χ^2-test rejects the null hypothesis that the estimates of efficiency concerns obtained from the two models have the same distribution at the 0.01 significance level. Both models yield estimates with more than 70% (85.5% by efficiency model and 73.2% by general model) of the IA-C&R subjects having efficiency concerns (measured by γ) below 0.1. Another 15% or so are within [0.1, 0.3), and 10% of the subjects, according to general model, have efficiency concerns higher than 0.3, while no one reaches this level according to the efficiency model. This is an artifact of the design, however, because Menu 3 does not facilitate the identification of higher value for this method.[35] The general model does allow us to identify γ values as high as 14. We only observe 9.5% of the subjects with efficiency estimates above 0.3, and none above 0.6, however.

2.6.8 Empirical Results Including Non-IA-Rationalizable Subjects

In this section, we represent our results obtained from the analysis of all 216 IA-consistent subjects, including the non-IA-rationalizable subjects that we excluded in the analysis reported in section 2.4. The results are qualitatively very much like those previously reported.

Estimates of Envy and Guilt

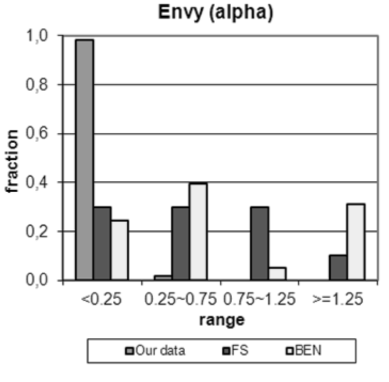

(a)

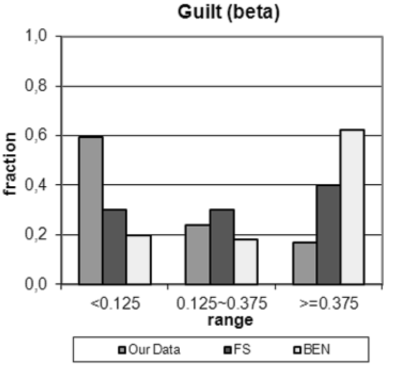

(b)

Figure 2.6: Distributions of envy (a) and guilt (b) estimates of IA-consistent subjects estimated according to IA model (2.9) by us, FS and BEN

The distributions of the IA parameters obtained from all IA-consistent subjects are very much like those obtained from only the IA-C&R subjects. For , the difference

[35]Note that this does not limit the applicability of the model. It simply means that subjects with higher values of γ would have $\sigma_{2j} = 1$ and our procedures would allocate them in the category 0.1-0.3.

in the estimated fractions is less than 0.01 for any of the categories considered. For β, the estimates from the two data sets have similar shaped distributions, but data from the expanded set show about 0.05 higher fractions of estimates lying in each of the two higher-value ranges (the fraction increases from 0.196 to 0.241 for [0.125, 0.375) and from 0.123 to 0.167 for $[0.375, +\infty)$). A χ^2-test, again, rejects the null-hypotheses that the distributions of our estimates (α and β, respectively) equals the distribution reported by FS or BEN both at the 0.01 level .

Robustness to Stakes and Efficiency concerns

Among all of the 216 IA-consistent subjects, 8 of them chose not to enter Part 3. Therefore, we have 208 observations for which we can compare α and β estimated under different payoff scales. To start, Figure 2.7 shows the estimated distributions for the various payoff scales used. Recall that each subject participated in the benchmark scale ("lowest") and in one of the higher scale treatments.

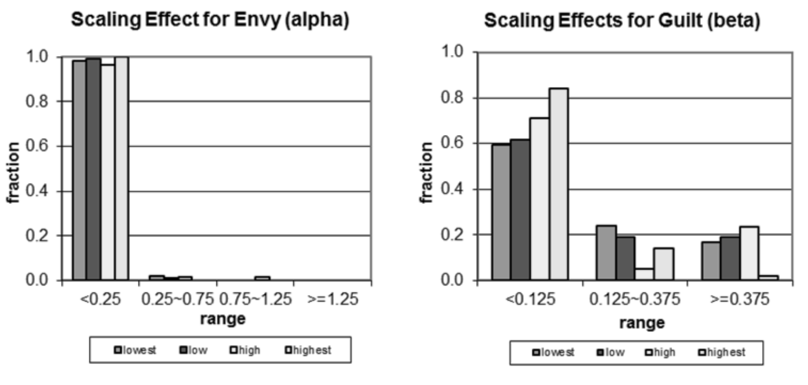

(a) (b)

Figure 2.7: Distributions of envy (a) and guilt (b) estimates for distinct payoff scales

For the test of whether the individual IA parameters estimated under the high-scaled menus remain the same as in the standard-scaled menus, Wilcoxon sign-rank tests only reject the null hypothesis when the scale is multiplied by 60 ("highest"). In contrast to the result for envy in Section 2.4, however, the envy parameters obtained from "highest" are found to be significantly higher than from "lowest".[36] These qualitative conclusions are the same as we reported in Section 2.4 for the restricted set. The statistics of the tests are summarized in Table 2.10.

Table 2.10: Stakes Effects

	low (99 obs.)		high (59 obs.)		highest (50 obs.)	
	z	p-Value	z	p-Value	z	p-Value
α	1.37	0.17	0.91	0.36	2.61	0.01***
β	1.05	0.29	-1.36	0.18	-3.98	0.00***

Notes. Results are presented from Wilcoxon Sign-rank tests of H0: $\alpha_{x,i} = \alpha_{lowest,i}$, vs. H1: $\alpha_{x,i} \neq \alpha_{lowest,i}$ and H0: $\beta_{x,i} = \beta_{lowest,i}$, vs. H1: $\beta_{x,i} \neq \beta_{lowest,i}$, $x \in \{low, high, highest\}$. *** indicates statistical significance at the 0.01 level.

Concerning the robustness of our results to allowing for efficiency concerns, the estimates of α have almost the same distribution as reported before, irrespective of whether the estimates are from the IA (2.9) or GM (2.20) model. The estimates of β from the two models both shift slightly to the higher intervals, when comparing to the smaller data set (see Figure 2.8).

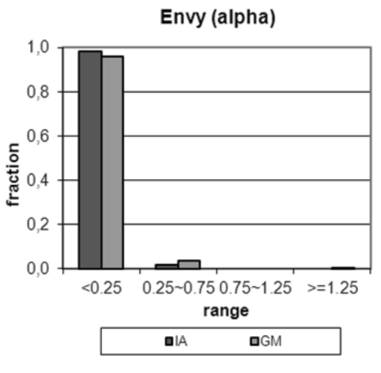

(a)

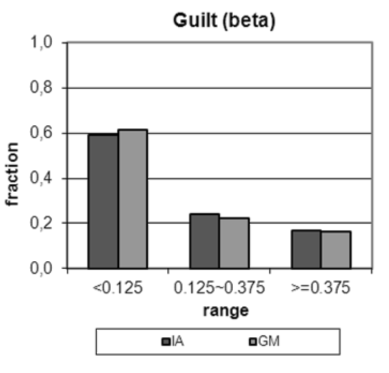

(b)

Figure 2.8: Distributions of envy (a) and guilt (b) of IA-consistent subjects after correcting for efficiency concerns

[36]This change can be attributed to some subjects with negative α estimates in the standard menu tests showing less negative or zero envy levels when the payoff scale was multiplied by 60 ("highest").

Pearson's χ^2tests do not reject the null hypothesis that the distributions of the IA model estimates are the same as the general model estimates, with a p-value of 0.373 for α and 0.762 for β. Weighted by the number of observations, we find $\Omega = 82.7\%$ for α and $\Omega = 88.5\%$ for β, virtually identical to the values obtained in Section 2.4 (82.1% and 88.5%, respectively). Hence, we conclude again that our estimates of IA parameters are quite robust to allowing for efficiency concerns.

Predictive Power of the IA Model in the Production Game

In Table 2.11, we summarize the regression results of (2.32), which is equal to 2.8 in Section 2.4, for the full IA-consistent subjects data pool.

$$e^T_{R,i} = \delta^T_{R,0} + \delta^T_{R,1}\alpha^{IA}_i + \delta^T_{R,2}\beta^{IA}_i + \varepsilon^T_{R,i} \tag{2.32}$$

where T =SimProd, SeqProd, $R = A, B$ represents the subject's role in the production game, and i is the index of a subject. α^{IA}_i and β^{IA}_i are the IA estimates of i's envy and guilt parameter, respectively.

The estimated coefficients are listed in Table 2.11 together with the theoretical prediction of the IA model. The qualitative results, including the significance and the direction of the dependence of the effort level on guilt and envy levels, are the same as the results obtained from using the IA-C&R subjects pool. More specifically, in the simultaneous production game (SimProd), A's effort only depends significantly (and positively) on her guilt level, while B's effort only significantly (and negatively) depends on her envy parameter, although the level of these dependency is not as strong as the IA model predicts. In the sequential production game (SeqProd), the predictive power of IA model strongly diminishes. Neither A nor B's effort level is significantly dependent on either of the IA levels any more.

Table 2.11: Coefficients of IA parameters in the production game

	IA Prediction		SimProd		SeqProd	
	e_A	e_B	e_A (obs. 167)	e_B (obs. 167)	e_A (obs.36)	e_B (obs. 28)
const.	0	100	30.45***	95.51***	59.62***	76.76***
			(3.22)	(1.28)	(7.74)	(8.19)
α	0	-100	-7.62	-49.56***	53.17	-139.79
			(27.93)	(11.11)	(77.84)	(112.33)
β	200	0	93.44***	-10.29*	-3.69	-64.22
			(13.83)	(5.50)	(27.79)	(44.27)

Notes: The predicted coefficients are taken from eqs (2.17) and (2.19). The number in the parentheses below each estimate is the standard error.

The distributions of the envy and guilt levels derived using the inverse of equations 2.17 and 2.19, are displayed in Figures 2.9(a) and (b). We include the FS and BEN distributions for comparison. As in Section 2.4, for SimProd, we use all the subjects' choices (for A or B, 167 observations in total) to estimate the IA parameters (either guilt or envy). For SeqProd (for both the DoubleRole and SingleRole treatments), because Worker B gives responding rules instead of a single number of the effort levels, we only use the realized effort level of B. Hence, only the observations from the subjects who are assigned to be Worker B in SeqProd (28 obs. in total) are used to estimate α.

We also applied our second individual-level test of the IA model's prediction in the production game by comparing for each subject whether her observed chosen effort levels for Worker A/B lies in the predicted intervals (derived from the α and β intervals as elicited from her choices in Menus 1 and 2. This yields the following results. In SimProd, 53.3% (89 out of 167) of the Worker A's and 87.4% (146 out of 167) of B's observed effort levels lie in the predicted intervals. In SeqProd, the rates of correct prediction drops to 22.2% (8 out of 36) for Worker A and 50% (14 out of 28) for Worker B.

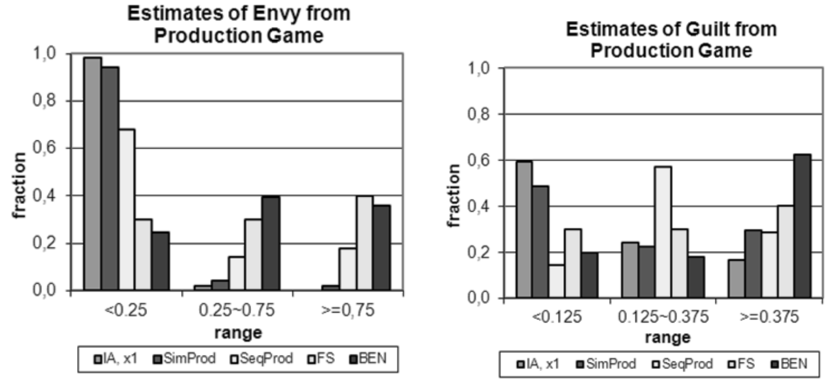

(a) (b)

Figure 2.9: Distributions of envy (a) and guilt (b) estimates derived from the production game.

Figure 2.9(a) shows that the simultaneous production game yields a distribution of envy that is very much like the distribution estimated from the IA model according to the choices in the standard-scaled Menu 1. When reciprocity is involved in SeqProd, the results are different, however. The envy levels estimated from Worker B's effort levels in the sequential production game have more fraction in the middle- and high-value ranges, with the distribution closer to that reported by FS and BEN.

Similar to envy, in Figure 2.9(b), we observe higher guilt levels in SeqProd than in SimProd. As mentioned in Section 2.4, this may be attributed to an effect of strategic choices by Worker A, fearing B's negative reciprocity.

All of these results are the same as those obtained in Section 2.4 for the restricted data set.