Procedural Fairness in Lotteries Assigning Initial Roles in a Dynamic Setting

Gianluca Grimalda[*], Anirban Kar[†] and Eugenio Proto[‡]

September 29, 2015

Abstract: We extend the study of procedural fairness in three new directions. Firstly, we focus on lotteries determining the initial roles in a two-person game. One of the roles carries a potential advantage over the other. All the experimental literature has thus far focused on lotteries determining the final payoffs of a game. Secondly, we modify procedural fairness in a dynamic – i.e. over several repetitions of a game – as well as in a static – i.e. within a single game - sense. Thirdly, we analyse whether assigning individuals a minimal chance of achieving an advantaged position is enough to make them willing to accept substantially more inequality. We find that procedural fairness matters under all of these accounts. Individuals clearly respond to the degree of fairness in assigning initial roles, appraise contexts that are dynamically fair more positively than contexts that are not, and are generally more willing to accept unequal outcomes when they are granted a minimal opportunity

[*]Universitat Jaume I of Castelló, Spain. Institute for the World Economy, Kiel, Germany. Centre for the Study of Global Cooperation Research, University of Duisburg-Essen, Germany; Corresponding author: grimalda@uji.es

[†]Delhi School of Economics, University of Delhi

[‡]University of Warwick, IZA and CESifo

to acquire the advantaged position. Unexpectedly, granting full equality of opportunity does not lead to the highest efficiency.
JEL Classification: C92, C78, D63

Keywords: Procedural fairness; Equality of Opportunity; Experiments.

1 Introduction

The idea that individuals' sense of justice encompasses not just inequality in final outcomes, but also the fairness of the process leading to such outcomes is now widespread in the social sciences and political philosophy. Procedural fairness is made possible by *"impartial rules ensuring that each of the agents involved in an interaction enjoys an equal opportunity to obtain a satisfactory outcome"* (Krawczyk, 2011). Experimental evidence generally shows strong individual preferences for fair procedures and that individuals are willing to accept more unequal final allocations, the fairer the procedures determining such allocations.

In this paper, we extend the study of fairness in three directions that have not been explored so far. Firstly, we focus on lotteries determining the initial roles in a two-person game. To the best of our knowledge, all the experimental literature has thus far focused on lotteries determining the final payoffs of a game. Secondly, we study the effect of modifying procedural fairness in a *dynamic* as well as in a *static* sense. Finally, we analyse whether assigning individuals a minimal chance of achieving an advantaged role is enough to make them willing to accept substantially more inequality.

We take the Ultimatum Game (UG henceforth) as our basic interaction. This game is suitable for our experiment because one role has a clearly identifiable advantage over the other. The proposer has greater bargaining power and hence can expect a greater share of the surplus than the receiver (Oosterbeek et al., 2004). When asked to bid on the two roles of a UG before playing the game, players offer twice as much to occupy the proposer's role as they do for the receiver's role (Guth and Tiez, 1986). A lottery giving one player higher chances of being assigned the proposer's role than another player can conceivably be seen as *unfair*.

The main novelty of our experimental design is to make the access to the two UG roles subject to the outcome of various lotteries, and to manipulate their degree of fairness. The baseline case is that both players have *equal opportunities*, as the lottery assigns both individuals a 50% chance of acquiring the proposer role. The initial lottery becomes increasingly biased in favour of one of the two players in three other treatments. The *favoured* player has, respectively, 80%, 99%, and 100% probability of becoming the proposer, while the *unfavoured* player has the residual probability. Receivers who are only concerned with the outcomes of the game should behave in the same way in these treatments. Consequently, we interpret differences in rejection rates across treatments as caused by receivers' preferences over procedural fairness, as in Bolton et al. (2005). We run 20 interactions of the stage game with random rematching of subjects before each interaction. We define the probability that the unfavoured player has of becoming the proposer within each round as p, where $p \in \{0\%; 1\%; 20\%;\ 50\%\}$. We define the bias in this lottery as measuring *static unfairness*, because it refers to the unfairness of the lottery within each round.

In addition, we also study procedural fairness in a *dynamic* perspective. In a subset of our treatments, which we call Variable Position Conditions ($VPCs$), we introduce another lottery preceding the lottery assigning the game roles. This first lottery gives each player equal chances to acquire either the probability p or the probability $(1-p)$ of becoming proposer in the second lottery. In $VPCs$ this unbiased first lottery is run at the beginning of each round. In the alternative subset of treatments, which we call Fixed Position Conditions ($FPCs$), this first lottery is *only* run in the very first round. Therefore, in $FPCs$ an unfavoured player keeps the same probability p of being assigned the proposer's role in each of the 20 rounds. We argue that $VPCs$ guarantee *dynamic* fairness, because the expected probability of being assigned the advantaged role in the game in *future* rounds is always the same - namely, one half - for every player, regardless of the round of the experiment. This cannot be said for $FPCs$ (except for the very first round). Again, individuals who are only concerned with final outcomes should be indifferent to this manipulation, while receivers who are concerned with procedural fairness will respond to it. Overall, we have six treatments (in addition to the baseline 50% condition): one treatment for each of the three $p \in \Pi \equiv \{0\%, 1\%, 20\%\}$, run under either the $FPCs$ or $VPCs$.

We find that procedural fairness matters under all of the accounts outlined above. A general trend exists such that receivers reject less, *ceteris paribus*, when p is higher (static fairness), and when the meta-lottery reassigning p at every round is run (dynamic fairness). They also reject less when they have been granted a minimal opportunity to acquire the advantaged position in comparison to being given no chance. However, this result in $VPCs$ is not robust to the introduction of sample demographic controls. We also find unexpected results, in that granting full equality of opportunity does not lead to the lowest rejection rates. Rather, these are obtained in two of the $VPCs$. Our research enables us to speculate on the shape of individuals' demand for opportunities' and suggests that individuals' subjective perceptions of when the 'playing field' is 'level' may radically differ from the objective distribution of chances. This calls for more in-depth research on the interaction between the various dimensions of fairness that people perceive. It also suggests that policies aiming at maximizing the opportunities for disadvantaged groups in society should take into account individuals' preferences over procedural as well as outcome fairness.

The paper is organised as follows. We review the existing literature in the next section. We enumerate the main hypotheses and describe the experimental protocol in section **??**. Section **??** reports the results. Section **??** discusses the results and concludes the paper.

2 Literature Review

Economic theory has traditionally held that individual preferences are consequentialist (Hammond, 1988; Machina, 1989). This means that individuals attach value only to the final outcomes of an interaction, disregarding the process leading to such outcomes. Consequentialist models allow for preferences to be either purely self-interested, or other-regarding (Fehr and Schmidt, 1999; Bolton and Ockenfels, 2000; Charness and Rabin, 2002). *Outcome fairness* looks either at the level of equality in final allocations, or to the degree to which final allocations reward individual contributions (Leventhal, 1980; Konow, 2003).

The idea that individuals are concerned not just with the *outcomes* of a certain social interaction, but also with the *process* leading to that outcome, has gained increased consensus in fields as disparate as legal studies (Thibaut and Walker, 1975; Tyler, 2006), political philosophy (Rawls, 1999), social choice (Elster, 1989), and more recently, economics. Comprehensive survey evidence demonstrates that individuals who believe that fair opportunities to make progress in their lives are available, also demand less redistribution from their governments (Fong, 2001; Corneo and Gruner, 2002; Alesina and La Ferrara, 2005; Benabou and Tirole, 2006). Procedural fairness is also vital in many other aspects of economic decisions, such as company wage structure, worker productivity (Bewley, 1999; Erkal *et al.*, 2010; Gill *et al.*, 2012) and institutional mechanisms to allocate scarce resources (Anand, 2001; Keren and Teigen, 2010).

Within experimental economics, Bolton *et al.* (2005) showed in a pioneering study that procedural fairness - modelled as equal chances of achieving unequal outcomes - is a substitute for outcome equality. Other studies replicated this result, but nevertheless showed that equality of opportunity is not a *full* substitute for equality of outcomes (Becker and Miller, 2009; Krawczyk and Le Lec, 2010). It was also shown that many people are willing to sacrifice money to reject allocations that are brought about by procedures that are extremely biased (Bolton et al., 2005; Karni *et al.,* 2008). In subsequent studies, however, changing the fairness of procedures only generated limited effects on individual behaviour (Krawczyk, 2010; Cappelen *et al.*, 2013). This suggests that procedural fairness may lose part of its prominence when other factors, such as individual merit in the determination of final outcomes (Krawczyk, 2010) or individual responsibility in risk-taking (Cappelen *et al.*, 2013), affect individuals' decisions.

Other experiments have contrasted individual merit with luck as a determinant of the initial positions in UGs, or Dictator Games (see e.g. Hoffman and Spitzer, 1985; Burrows and Loomes, 1994; Hoffman *et al.,* 1994; Cappelen *et al.,* 2007 or Schurter and Wilson, 2009). The unequivocal conclusion of this literature is that outcome inequality is more accepted when first-movers 'earn' their position by performing better than their counterparts. A second class of experiments has manipulated the relative advantage of positions in UGs, making players' initial endowments unequal (Guth and Tietz, 1986; Armantier, 2006), or modifying the UGs final outside options (Binmore *et al.*, 1991; Suleiman, 1996; Schmitt, 2004; Buchan *et al.,* 2004; Handgraaf *et al.,* 2008). Generally, players take advantage of their increased power in the game. None of these studies, however, examine random assignments of initial positions. We are the first, to the best of our knowledge, to do so.

Mathematical models of procedural fairness have recently been developed. Individuals' preferences are assumed to depend on the impartiality of the procedure determining the final outcomes, as well as on the outcomes proper. The greater the fairness of the process, the greater individuals' utility. Karni and Safra (2002) offer an axiomatic account based on Diamond's (1967) idea that individuals prefer fair procedures to biased ones, even when the final outcomes are unequal. Andreozzi et al. (2013) provide an axiomatic approach based on the notion of separability between self-interested on one hand and other-regarding or procedural preferences on the other. Bolton et al. (2005) extend Bolton and Ockenfels's (2000) consequentialist model by defining the "fairest" available allocation in a game as the closest possible - in expected value - to an equal divide. Individuals then condition their social preferences on the distance between the actual allocation and the fairest allocation. Trautmann (2009) uses the ex-ante *expected* payoff difference between individuals as a proxy for the unfairness of the procedure, with a completely fair processes being characterised by expected payoff differences equal to zero. He then applies Fehr and Schmidt's (1999) model of inequality aversion to the *expected* payoffs rather than the *final* payoffs. Krawczyk (2011) posits an assumption of negative interpendence between procedural unfairness, also proxied by the expected payoff differences, and *outcome* inequality aversion *a là* Bolton and Ockenfels (2000). Accordingly, the greater the unfairness of the procedure, the greater an individual's desire for low inequality in ex-post earnings.

3 Experimental design

3.1 Round 1 of the experiment

This section describes the procedures followed in the first round of the experiment. The tree of the game is displayed in Figure 1. £10 is at stake. First, two randomly matched players are assigned the position of either Player 1 or Player 2 by means of an unbiased random draw. We call this initial lottery $_1$. In the second phase, players are informed of the result of $_1$ and make an offer to their counterpart. An offer is a proposal of how to divide the £10 between the pair. We define x_i as the amount that Player i requires for herself, while $10 - x_i$ is the remainder being offered to the counterpart, for $i \in \{1, 2\}$. In this phase players do not know their counterpart's offer.

In the third phase, one of the two offers is randomly selected through a lottery that we call $_2$. The key aspect of the design is that the treatments differ in the probability with which each player's offer is randomly selected. This is denoted by probability p for Player 2 and $(1 - p)$ for Player 1. This probability has a maximum at $p = 0.5$ for Player 2 in the 50% treatment. In the three remaining treatments, this probability equals $p = 0.2$ in the 20% treatment, $p = 0.01$ in the 1% treatment, and $p = 0$ in the 0% treatment.

Finally, in the fourth phase, the player whose proposal has *not* been selected has to decide whether she accepts or rejects the other player's offer. Suppose that Player i is drawn. x_i is then communicated to Player j, who can either accept or reject that offer. If Player j accepts, the payoffs are x_i

for Player i and $10 - x_i$ for Player j. If Player j rejects, both players' payoff is 0. Player i is informed that her offer has been selected, but Player j's offer is not communicated to Player i. After each round, each pair is informed of the outcome of their own interaction and of their respective payoffs.

INSERT FIGURE 1 ABOUT HERE

The key difference between our extended UG and a standard UG is the introduction of lottery $_2$, which randomly selects the offer that becomes relevant for the final allocation. Players are always informed of the lottery outcomes. In particular, at the top node of the decision tree, individuals are aware of whether they are Player 1 or Player 2 at the moment they submit their proposal. As in Suleiman (1996) and Handgraaf *et al.* (1998), we do not ask Player 2s to submit an offer in 0% treatments, as this would have no possibility of being selected. We discuss the implications of this design choice in section **??**. After $_2$ has been run, the interaction proceeds exactly as in a UG. The player whose proposal has (not) been selected becomes the proposer (receiver) and the payoffs are determined as in standard UGs. All the random draws in $_1$ and $_2$ are made by a computer.

3.2 Rounds 2-20 of the experiment

After the first round is played, the ensuing 19 rounds are played, with one crucial difference between $FPCs$ and $VPCs$. In $VPCs$, Round 1 is replicated exactly as described above in each of the following 19 rounds. Both $_1$ and $_2$ are run in each round. In $FPCs$, $_1$ is not run any more from Round 2 on. Only the rest of the game - namely, Phase 2-4, or the stage game (see Figure 1) - is replicated as described above in every round. In other words, in $FPCs$ $_1$ is only run once at the beginning of the experiment, while lotteries $_2$ are run in each round. Consequently, in $FPCs$ a player remains unfavoured

(favoured) throughout the 20 rounds, while in *VPCs* each player has an even chance of being assigned the favoured or the unfavoured position in each round[1].

The different dynamic of the experiment is represented in Figure 2.

INSERT FIGURE 2 ABOUT HERE

3.3 Other aspects of the design

As already illustrated in the introduction, the experiment comprises a baseline condition where $p = 50\%$, and six treatments, one per each of three possible values of $p \in \Pi$, each played under both a VPC and an FPC treatment. Each subject participates in only one treatment, i.e. the design is between-subject. The final payoffs are given by the outcomes of two randomly-selected rounds out of the 20 to limit income effects. We prefer to pay subjects for the outcomes of two rounds instead of just one because we fear that a payment based on only one round, coupled with the relatively low show-up fee (£5), may discourage receivers from rejecting unfair offers, thus limiting the variability of our dependent variable.In fact, we show in section **??** that receivers' choices were independent from previous earnings, thus suggesting that income effects were negligible.

Experimental sessions were run at Warwick University between April and June 2007. On average, 60 students per treatment took part in our experiments, with a total of 426 students. Their demographic characteristics are reported in Table S1 in the Electronic Supplementary Material (ESM). Only subjects who had not attended courses in Game Theory were allowed to

[1]Note that we use the term 'role' to indicate whether a participant is a proposer or a receiver in the UG played in the last phase of the Stage Game (see Figure 1). We use the term 'position' to refer to whether a player is Player 1 (favoured) or Player 2 (unfavoured) in the lottery assigning UG roles - that is, $_2$ in Figure 1.

participate. We ran three sessions per treatment. Due to varying show-up rates, the number of subjects per session is not constant across sessions but varies from a minimum of 16 to a maximum of 24 subjects, with an average of around 20 subjects per session. We took care to balance the composition of the sessions in terms of gender and academic specialisation. The game was conducted using the z-tree software (Fischbacher, 2007). Supplementary details on the experimental procedures and the instructions are reported in sections 6 and 7 of the ESM.

3.4 Our Hypotheses

Firstly, we predict that a given offer is more likely to be accepted when it has been generated in a game where players had fairer initial chances. We assume that players' preferences are influenced by the procedures that determine the initial positions in the single stage-game. Consistent with procedural fairness models, it is natural to assume that agents will prefer, *ceteris paribus*, procedures providing players with a less biased distribution of opportunities. Here, we make the key assumption that players take the probability p as an index of the fairness of the procedure. This is a natural assumption because p determines, in both $FPCs$ and $VPCs$, the probability of attaining the advantaged position in the stage game. We then draw on the interaction effect proposed by Krawczyk (2011: 116). This assumes that the lower the procedural fairness, the greater the aversion to outcome inequality. This entails that a receiver who is faced with a *less* fair initial procedure is *more* inclined to reject a given allocation. We call this the 'Monotonic Fairness Hypothesis':

H_1: The higher p, the higher receivers' acceptance rates for a given split. This holds in both FPCs jointly with the 50% treatment and VPCs jointly with the 50% treatment.

As for the *dynamic* aspect of procedural fairness, in the introduction we claimed that *VPCs* will be *ceteris paribus* deemed as more fair than *FPCs* because they grant *dynamic* proceudural fairness. That is, every player always has a probability equal to 1/2 of acquiring the favoured role in the subsequent round, while in *FPCs* this probability is - with the exception of the first round - equal to $p < 1/2$. In the ESM (section 1) we also demonstrate three propositions that qualify this statement more precisely. We therefore hypothesise that $VPCs$ will be considered more procedurally fair than $FPCs$ between what we call '*corresponding treatments*'. We define corresponding treatments as pairs of treatments belonging to $FPCs$ and $VPCs$ whose $_2$ is characterised by the same p. There are three pairs of corresponding treatments, which we denote by p_VPC and p_FPC, $p \in \Pi$. For instance, $0\%\_VPC$ and $0\%\_FPC$ are corresponding treatments for $p = 0\%$. We thus posit a 'Dynamic Opportunities Hypothesis':

H_2: For any corresponding treatment, receivers' acceptance rates decrease significantly in p_FPC as compared with p_VPC, $p \in \Pi$.

Finally, we hypothesise that individuals respond to being assigned a minimal, rather than a zero, chance of success in $_2$. An extensive body of empirical and survey evidence stresses the importance for people of having a 'voice'. Frey and Stutzer (2005) show that the mere right to participate in the political process - rather than actual participation - increases individual satisfaction. Anand (2001) and Tyler (2006) report survey and experimental evidence for the importance that people place on having the right to have their opinion heard - or appropriately represented - in collective decision-making processes. The relevance of this right to a voice may be caused by the desire to express one's position, or to obtain respect for one's worth. Nozick (1994: 34) offers a rationalization of this evidence, arguing that the individual's value metric over the probability space may not be linear, and may suffer "discontinuities" in the origin of the space, i.e. when we move from full certainty to even limited uncertainty. An alternative explanation is that players may magnify the assignment of a small probability (Kahneman and Tversky, 1979). Therefore, we posit a "Discontinuity Hypothesis":

H_3: Receivers' acceptance rates decrease significantly in the 0% treatments in comparison with the 1% treatments.

In this paper, we focus on receivers' behaviour. The analysis of proposers' behaviour is reported in the ESM, section 4. There we show that proposers' behaviour adapts to receivers' behaviour. That is, in treatments where acceptance rates are higher the offers are lower and *vice versa*. This is not surprising given the repeated feedback in our experiment. Nonetheless, we also show in the ESM (section 5) that proposers were on average able to anticipate the differences in the receivers' behaviour across treatments in the first round of the game. Therefore, not only was the proposers' behaviour adaptive but it was also successfully predictive.

4 Results

4.1 Results for Fixed Position Conditions

4.1.1 Descriptive Analysis

Table 1 reports descriptive statistics for proposers, and receivers' in each treatment. The overall acceptance rate in the treatment that is closest to standard UGs - namely, *0%_FPC* - is 77.58%, while the mean proposers' demand is equal to 62.8%. This is largely in line with standard results from UGs[2]. Comparing the 50% treatment and *FPCs* shows the existence of a monotonic pattern consistent with H_1. As the bias in the initial lottery increases, both the mean and the median values of rejected demands decrease (see Table 1, column 1). This means that as the initial lottery becomes more biased, receivers require larger shares of the pie to accept an offer.

INSERT TABLE 1 ABOUT HERE

Second, the acceptance rates of low offers, i.e. those lower than 20% of the pie, decreases as the bias in the initial lottery increases (see Table 1, column 3). The drop in the acceptance rate for low offers is particularly pronounced between *1%_FPC* and *0%_FPC,* consistently with H_3. This is further corroborated by Figure S1 in the ESM, which plots acceptance rates corresponding to various offer classes.

4.1.2 Econometric Analysis for FPCs

Our econometric analysis pools all the observations coming from the different treatments together. We model the repeated nature of the data through an individual-level random effects model. In this section we only report the results relative to FPCs. All the ensuing analyses also use a random effects model. We fit a probit model where the dependent variable is the dichotomic variable $ACCEPT$. This indicates with the value of 1 (0) a receiver's acceptance (rejection) of an offer. The complete regression can be

[2]In their meta-analysis, Oosterbeek *et al.* (2004) report that the weighted average acceptance rate from 66 UG studies is 84.25%, whereas average demands equal 59.5% of the pie in 75 UG experiments.

found in ESM: Table S2. Here we report the main results. In the models reported in Table 2, the key variable is *CHANCE*. This takes the value of p (see section **??**), thus providing a general measure of the bias in the initial lottery. In the first model, *CHANCE* has a strong and positive effect (P-*value* [P henceforth]$= 0.004$) (see Table 2, column 1). In accordance with hypothesis H_1, receivers were more likely to accept offers when these came after a less unbiased initial lottery. We also control for whether a subject had been assigned the favoured role in the initial lottery. Favoured and unfavoured subjects may have formed different earnings expectations and this may have affected their behaviour when drawn as receivers. However, this variable is insignificant in all the models considered. Round dummies are also included in the model to control for time trend effects or effects associated with specific rounds. The offer size is also included in the model. Not surprisingly, higher offers had a significantly higher probability of being accepted ($P < 0.001$). All regressors are exogenous. Hence, the individual-specific effects must be uncorrelated with the other regressors, thus ensuring that a between-estimator is consistent.

The subsequent models in Table 2 interacts *CHANCE* with either the $FPCs$ or $VPCs$ in order to investigate whether *CHANCE* had different effects in the two sets of treatments. We define $\overline{FPC} \equiv \{0\%\_FPC \cup 1\%\_FPC \cup 20\%\_FPC\}$ as the set including $FPCs$, while $\overline{FPC \cup 50\%} \equiv \{0\%\_FPC \cup 1\%\_FPC \cup 20\%\_FPC \cup 50\%\}$ is the set including $FPCs$ *plus* the 50% treatment. Similarly, $\overline{VPC} \equiv \{0\%\_VPC \cup 1\%\_VPC \cup 20\%\_VPC\}$ and $\overline{VPC \cup 50\%} \equiv \{\overline{VPC} \cup 50\%\}$ are the corresponding sets for $VPCs$. In the regression in Table 2, column 2, '*CHANCE X FPC* $\cup$ *50'* is the interaction term between *CHANCE* and $\overline{FPC \cup 50\%}$. '$CHANCE\ X\ VPC'$ is instead the interaction term between $CHANCE$ and $\overline{VPC}$. In this model, *CHANCE*

X FPC $\cup$ *50* and $CHANCE\ X\ VPC$ capture separately the influence of $CHANCE$ on $\overline{FPC \cup 50\%}$ and $\overline{VPC}$, respectively. This model also adds a number of additional controls as a robustness check. We introduce a control for a subject's history in the game. It is obvious that subjects' expectations of what is fair may be influenced by past interactions and the outcomes that have been observed. We include '*PREVIOUS ROUND OFFER',* which is the offer that a receiver obtained the last time she acted as a receiver prior to the current round. Some demographic controls - namely, a subject's gender, age, a dummy identifying UK citizenship, and a dummy identifying attendance of Economics degrees - are also added.

CHANCE X FPC $\cup$ *50* exerts a positive and significant effect ($P = 0.010$). This supports the monotonicity hypothesis H_1 over $\overline{FPC \cup 50\%}$. Interestingly, *PREVIOUS ROUND OFFER* is negative and strongly significantly different from zero ($P < 0.001$). This means that subjects observing a higher offer in the past period were significantly *less* likely to accept an offer in the current period, *ceteris paribus*. This suggests that subjects may have used some kind of Bayesian updating rule in their estimation of the distribution of offers, and that they used this information in assessing the fairness of an offer. Some demographic variables are also significant[3]. The next model in Table 2 substitutes '*PREVIOUS ROUND EARNINGS'* for '*PREVIOUS ROUND OFFER'* as a control for a subject's history in the game. '*PREVIOUS ROUND EARNINGS'* denotes the latest earnings obtained as a receiver prior to the current round. This variable is not significantly different from zero ($P = 0.318$). The same result would hold using accumulated past earnings (not reported). Overall, this result indicates that income effects over

[3]The probability of acceptance was significantly higher for students attending Economics degrees ($P = 0.009$), women ($P = 0.029$), and students with UK citizenship ($P = 0.033$). Note that including these variables comes at the cost of a considerable loss of observations due to missing variables.

the course of the game were negligible for receivers.

The next econometric specifications replace the variable *CHANCE* with dummy variables identifying individual treatments. This enables us to study the differential effects of pairs of treatments on the propensity to accept, thus performing both a more stringent test of H_1 and a direct test of H_3. The whole model is reported in the ESM: Table S3. Figure 3a reports the probabilities of acceptance in each treatment for various offers. For any offer, as the bias in the initial lottery decreases, the probability of acceptance increases. For intermediate offer values, sizable differences emerge across treatments. For instance, for offers equal to 20% of the pie, the predicted probability of acceptance is equal to 0.89 in the baseline case, drops to 0.86 in 20%_*FPC* and to 0.68 in 1%_*FPC*, and drops to a mere 0.25 in 0%_*FPC*.

INSERT FIGURE 3 ABOUT HERE

Results of two-tailed Wald tests of the null hypothesis that pairs of treatment coefficients are equal to each other are reported in Table 3a. The signs are always positive, consistent with H_1, except in one case. In four cases out of six the null is rejected, denoting significant treatment differences in line with H_1. The acceptance rate of 0%_*FPC* always results as being significantly smaller than other treatments' coefficients. In particular, hypothesis H_3 of a symbolic value of opportunity is supported. The difference between $\beta_{0\%\_FPC}$ and $\beta_{1\%\_FPC}$ (where β is the value of the treatment dummy indicated in the subscript) is negative and significant ($P = 0.013$). The receivers in 1%_*FPC* had, *ceteris paribus,* a significantly higher probability of accepting a given offer than the receivers in the 0%_*FPC*. The results are virtually unchanged when demographic controls and *PREVIOUS ROUND OFFER* are introduced in the regression (see ESM: Table S3, column 2, and ESM: Table S4a).

INSERT TABLE 3 ABOUT HERE

On the basis of this analysis, we conclude:

Conclusion 1 *Descriptive and econometric analysis supports H_1 and H_3 in FPCs.*

4.2 Results for Variable Position Conditions

4.2.1 Descriptive Analysis

Table 1 shows that the monotonic pattern linking bias in the initial lottery and rejection rates still holds moving from *0%_VPC* to *20%_VPC,* but is reversed between *20%_VPC* and *50%.* Looking at the mean and median values of rejected demands, we note that receivers' hostility decreases from *0%_VPC* up to *20%_VPC,* but then rises again (see Tables 1, colum 1; and ESM: Figure S1). A similar trend can be detected with respect to the acceptance rate of low offers (see Tables 1, column 3), and also for mean and median offers (see Tables 1, column 4).

4.2.2 Econometric Analysis for VPCs

We modify the models formulated above to study the impact of the variable $CHANCE$ limitedly to $\overline{VPC \cup 50\%}$ (see section **??**). *CHANCE X VPC* $\cup$ *50* is the interaction term between $CHANCE$ and $\overline{VPC \cup 50\%}$, while $CHANCE\ X\ FPC$ is $CHANCE$ interacted with $\overline{FPC}$. The descriptive statistics suggest the existence of a non-linearity so the model in Table 2, columm 4 also includes a quadratic term for the interaction between $CHANCE$ and $\overline{VPC \cup 50\%}$. Indeed, both terms are strongly significantly different from zero. The predicted probability of acceptance shows an inverted-U pattern, reaching a maximum for $CHANCE = 0.29$. This supports the monotonicity hypothesis limitedly to $\overline{VPC}$, but not $\overline{VPC \cup 50\%}$. Again, this result is robust to the inclusion of demographic controls (see ESM: Table S2, column 6), even when they are interacted with $FPCs$ (see ESM: Table S2, column 7). We find no significant effect for any of such interaction terms. We conclude that the control variables we used do not appear to have different effects in $FPCs$ compared to $VPCs$.

Figure 3b plots the predicted probability of acceptance for $VPCs$. $VPCs$ follow a monotonic trend. For instance, for offers equal to 15% of the pie, the predicted probability of acceptance is equal to 0.42 in 0%_VPC, rises to 0.70 in 1%_VPC, and to 0.88 in 20%_VPC. However, the probability of acceptance drops to 0.67 in *50%*. Pairwise comparisons of treatment coefficient differences confirm the existence of a non-linearity in how receivers responded to variations in p (see Table 3b). All the three signs of the z-statistics, limitedly to $\overline{VPC}$, are positive and statistically significant. As far as H_3 is concerned, the Wald test rejects the null hypothesis that the treatment dummies are equal in *0%_VPC* and *1%_VPC* ($P = 0.047$).

Introducing demographic controls and *PREVIOUS ROUND OFFER* into the model somewhat attenuates the individual treatments effects (see ESM: Table S4b). The difference in the coefficients for *0%_VPC* and *1%_VPC* is no longer significant ($P = 0.24$), and some other pairwise comparisons lose significance. This is partly due to the fact that the $VPCs$ are most severely affected by missing observations. Moreover, introducing *PREVIOUS ROUND OFFER* absorbs some of the treatments effects, because offers are on average significantly higher in some treatments than in others (see ESM, section 4).

We conclude:

Conclusion 2 *Descriptive and econometric analysis supports H_1 in $\overline{VPC}$. The monotonic pattern breaks between* $20\%\_VPC$ *and the* 50% *treatment however. The treatment effects are somewhat attenuated when demographic controls and PREVIOUS ROUND OFFER are included.*

Conclusion 3 *Descriptive and econometric analysis supports H_3. However, this result is not robust to the inclusion of demographic controls and PREVIOUS ROUND OFFER.*

4.3 Comparing VPCs and FPCs

Descriptive statistics from Table 1 support H_2. For each pair of corresponding treatments (see section **??**), the mean and median value of rejected offers, and the acceptance rate of low offers, are all lower in $FPCs$ than VPCs. Secondly, Table 3c reports the results of Wald tests conducted on pairs of coefficient differences. The acceptance rates are *ceteris paribus* significantly lower in $FPCs$ than in $VPCs$ in all the corresponding treatments. The difference is highly significant between *0%_FPC* and *0%_VPC* ($P = 0.002$), and significant both between *20%_FPC* and *20%_VPC* ($P = 0.014$), and between *1%_FPC* and *1%_VPC* ($P = 0.037$). Introducing demographic controls and *PREVIOUS ROUND OFFER* into the analysis leaves the results unchanged except for the comparison between *1%_FPC* and *1%_VPC,* which loses significance (see ESM: Table S4c). This may be due to the reasons outlined in section **??**.

We thus conclude:

Conclusion 4 *Descriptive and econometric analyses support H_2.*

Figure 4 plots the distribution of the pie between proposers and receivers in each treatment, as well as the percentage of the pie that is lost because of receivers' rejections, for the whole 20 rounds (Panel a), and for the last five rounds (Panel b). We define the *'efficiency rate'* as the proportion of the pie that does *not* go destroyed. Surprisingly, granting equality of opportunity does not lead in our experiments to the highest efficiency rates. Highest efficiency over the whole 20 rounds is in fact achieved in *20%_VPC* and *0%_VPC*, where 15% of the resources are destroyed. *50%* is only fourth in this ranking, with losses equalling 19% of the available pie. The treatments with the lowest efficiency rates are *1%_FPC* and *0%_FPC*, where 22% and 23% of the resources are lost, respectively. The same ranking holds in the last 5 rounds. In order to analyse whether these differences are statistically significant over the whole 20 rounds, we consider the average acceptance rate in each of the three sessions making up a treatment. We compute the differences in such acceptance rates in pairwise comparisons between treatments, and construct a Binomial test on the null hypothesis that session-level acceptance rates in $FPCs$ or $VPCs$ are equally likely to be higher than one another. The results of these Binomial tests are reported in Table 4. We note that the the lower acceptance rates in *1%_FPC* and *0%_FPC* with respect to other treatments result in statistically significant differences (N=9; $P = 0.039$), in eight out of ten pairwise treatment comparisons. Furthermore, a Binomial test strongly rejects the hypothesis that session-level acceptance rates in $FPCs$ or $VPCs$ are equal to each other. Sessions in $FPCs$ are significantly *more* likely to generate lower acceptance rates than sessions in $VPCs$ (N=81, P¡0.001).

INSERT TABLE 4 ABOUT HERE

We thus conclude:

Conclusion 5 *Granting equal opportunity to achieve the proposer role does not lead to the highest efficiency. This is instead attained in the 20%_VPC* *and 0%_VPC treatments. Sessions with higher efficiency rates are significantly more likely to be observed in VPCs than FPCs.*

4.4 Analysis of dynamic and session effects

Table 5 analyses the existence of dynamic effects in our experiment. We find that acceptance rates tend to grow over repetitions of the game in $\overline{VPC}$ ($P < 0.001$), in 50% ($P = 0.047$), but not in $\overline{FPC}$ ($P = 0.213$) (see Table 5, column 1). The difference in the trend effect between $\overline{VPC}$ and $\overline{FPC}$ is statistically significant ($P = 0.003$). We also analyse whether the fairness monotonicity effect tends to amplify or dampen over time. For this purpose we interact $CHANCE$ with $ROUND$. Overall, we find no effect for this interaction term (Table 5, column 2). Moreover, interacting $CHANCE$ and $ROUND$ with $\overline{VPC}$ and $\overline{FPC}$ does not highlight any significant change in the way $CHANCE$ affects probability of acceptance within each set of treatments. We obtain similar results within a model using treatment dummies instead of $CHANCE$ (see ESM: Table S3 and Table S5). We interact each treatment dummy with a dummy identifying the last 10 periods of the interaction. In no treatment do we observe a significant change in the estimated probability of acceptance between the first and the second block of 10 rounds. This implies that the monotonicity effect that we observe does not appear to either vanish or grow over time. Moreover, acceptance rates in $FPCs$ appear significantly lower than in corresponding $VPCs$ both in the first and in the second block of 10 rounds, signalling the stability of the Dynamic Opportunities Hypothesis H_2 over repetitions of the game. Finally, the Discontinity Hypothesis H_3 is strongly significantly supported in both blocks of 10 rounds in $FPCs$, while it loses significance in the second block of 10 rounds in $VPCs$.

We also perform a separate analysis for the first round (see ESM, section 5). We note that the patterns we detect over the whole experiment already occur in the first round, although in several cases they do not reach significance levels. We conclude that the repetitions strenghtened patterns of behaviour that were already present from the start of the experiment. In the ESM, section 3.2 we perform an analysis of session effects, noting that only for one treatment (*0%_VPC*) the acceptance rate in a session - the very first session we run - differ significantly from the other two sessions making up that treatment. This weakens the effects related to the *0%_VPC* treatment.

We thus conclude:

Conclusion 6 *Apart from FPCs, acceptance rates tend to increase over time. Nonetheless, the effects related to the monotonic fairness hypothesis appear to be invariant over repetitions of the game. The same is true for the differences between FPCs and VPCs in corresponding treatments. As for the Discontinuity Hypothesis, this holds in both blocks of 10 rounds in FPCs but loses significance in the second block of 10 rounds in VPCs.*

5 Discussion and conclusions

Our results confirm and extend previous results that individuals are sensitive to the procedures leading to outcomes, in addition to the outcomes themselves. We found robust support for the Monotonic Fairness Hypothesis in *Fixed Position Conditions (FPCs).* The greater the inequality in the distribution of initial opportunities, the lower the acceptance rates of a given offer. This pattern of behaviour reproduces insights from survey analyses implying that the more a society is deemed to be granting fair opportunities to its citizens, the lower the demand for redistribution, arguably because citizens are more inclined to accept the resulting income inequalities as *fair* (see section 2). Even in our experiments, inequality is more accepted when the process leading to the final bargaining is fairer. Admittedly, in *Variable Position Conditions (VPCs)* the fairness monotonicity hypothesis does not extend to the 50% treatment, but remains confined to $\overline{VPC}$ (see section **??**). This break in monotonicity is surprising. It is associated with each treatment belonging to $\overline{VPC}$ having higher efficiency rates than the baseline case of equal opportunities. A possible explanation is that the role asymmetry implicit in *Variable Position Conditions* made the possibility of achieving fairness over the whole 20 rounds of interaction salient to subjects, thus inducing them to become more lenient regarding the proposed allocations in each single round. This process is analogous to the establishment of a 'convention' in repeated coordination problems (Hargreaves-Heap and Varoufakis, 2002).

The difference that we observe between *Fixed Position Conditions* and *Variable Position Conditions* in corresponding treatments clearly points to subjects responding to the way chances are allocated dynamically, rather than just statically. However, the fact that we observe significant treatment effects across *Variable Position Conditions* - where dynamic unfairness was absent in all the treatments - clearly indicates that a sizable portion of subjects responded to static procedural fairness. Although our experiment is not designed to test for Machina's (1989) dynamic consistency hypothesis, we believe that these results point in the direction of a signficant portion of players being dynamically *inconsistent*. Roughly speaking, dynamically inconsistent individuals neglect lotteries that occurred before the current decision node and thus modify their behaviour before and after a lottery has taken place (Machina, 1989; Trautmann and Wakker, 2010; Trautmann and van de Kuilen, 2014). Since before each round in *Variable Position Condi-*

tions players always have a 50% probability of being selected as proposers, we conjecture that the change in behaviour that we observe across *Variable Position Conditions* is due to players "neglecting" $_1$, while responding to $_2$. In other words, the behaviour we observe is consistent with players redressing the unfairness associated with $_2$, while ignoring the fairness of $_1$. Future research should investigate the relevance of this conjecture, and how preferences for dynamic and static fairness interact with each other.

It is striking that most of the observed variation in behaviour takes place as we move from 0% treatments to 1% treatments. When $_2$ is unbiased, the receivers reject offers of £2.15 on average, and when $_2$ gives people no chance of being a proposer in the *Fixed Position Conditions*, receivers reject offers of £2.96 on average. Put it differently, receivers would be willing to pay 81p on average to be in the 50% treatment rather than being in the 0%_*FPC*. By the same token, receivers would be willing to pay 43p to have a 1% chance of being proposers compared to none, and only 38p more to have equal chances compared to a 1% chance. In other words, the 'demand for opportunity' seems to be very steep near the origin of the scale, but considerably less so afterwards. In section **??** we argued that the purely symbolic opportunity of having a "voice" in a collective decision problem underpins subjects' propensity to view the situation as significantly fairer than when such an opportunity - albeit minimal - is denied. Admittedly, our design cannot rule out that what we observe is due to the tendency of individuals to overweigh small probabilities (Kahneman and Tversky, 1979). Future research may ascertain the relative importance of the purely symbolic value of 'voice' *vis-à-vis* the small probability overweighting effect. Arguably, many people who feel marginalised in societies will believe that they have neither a symbolic power of expression nor an even negligible chance of acquiring advantaged positions. We thus believe that it is important that our design has uncovered such a sizable response to procedural changes.

Our results may pave the way to refining existing theoretical models of procedural fairness. The use of expected payoffs as a proxy for procedural fairness, helpful as it may be to make models tractable, makes the predictions of both Trautmann (2009) and Kracwzyk's (2011) models unsuitable for lotteries applied to initial positions[4]. In our settings the variable p is a natural way to measure "how fair" the procedure is. In other contexts, such a clear-cut proxy for procedural fairness may not exist. Alternatives to expected payoffs, such as the *ex ante* willingness to pay to enter a game in a certain position (Stefan Trautmann, private communication) may be considered instead of expected payoffs.

Our research suggests that people are sensitive to procedures in ways and contexts that had not been explored so far. Some of our results confirm our hypotheses on procedural fairness while others are unexpected and call for more research on the topic. All of this points to the need to further investigate individuals' actual perceptions of fairness and the interaction between the various dimensions. Policies aimed at improving opportunities for disadvantaged groups should incorporate such perceptions into their design to augment their scope and maximize their efficiency.

[4] In our experiments, the average expected payoffs for receivers in the last five rounds - seemingly an appropriate measure for "equilibrium" payoffs - are the highest (£3.16) in *0%_FPC*, which is arguably the most unfair procedure in our experiments. The only unbiased procedure in our experiments, i.e. the baseline *50%* condition, only yields £2.47 to receivers and comes fifth in the ranking of expected receivers' payoffs across treatments. In our case "equilibrium" expected payoff differences are thus an inaccurate proxy for procedural fairness.

References

[1] Alesina, A., La Ferrara, E. (2005). " Preferences for redistribution in the land of opportunities", Journal of Public Economics, 89(5-6), 897-931.

[2] Anand, P. (2001). "Procedural fairness in economic and social choice: Evidence from a survey of voters", Journal of Economic Psychology, 22(2), 247-270.

[3] Andreozzi, L., Ploner, M., & Soraperra, I. (2013). Justice among strangers. On altruism, inequality aversion and fairness. Working paper No. 1304, CEEL, Trento University.

[4] Armantier, O. (2006)."Do Wealth Differences Affect Fairness Considerations?", International Economic Review, 47(2), 391-429.

[5] Becker, A., Miller, L. (2009). "Promoting justice by treating people unequally: an experimental study", Experimental Economics, 12(4), 437-449.

[6] Bewley, T. (1999). Why Wages don't Fall during a Recession. Harvard: Harvard Univ. Press.

[7] Benabou, R., Tirole, J. (2006). "Belief in the Just World and Redistributive Politics", Quarterly Journal of Economics, 699-746.

[8] Binmore, Kenneth, Peter Morgan, Avner Shaked, John Sutton (1991). "Do People Exploit Their Bargaining Power?: An Experimental Study", Games and Economic Behavior, 3, 295–322.

[9] Bolton, G. E., Brandts, J.,Ockenfels, A. (2005). "Fair Procedures: Evidence from Games Involving Lotteries", Economic Journal, 115(506), 1054-1076.

[10] Bolton, G., Ockenfels, A. (2000). "A theory of equity, reciprocity and competition", American Economic Review, 90, 166-93.

[11] Buchan, N., Croson, R., Johnson, E. (2004). "When Do Fair Beliefs Influence Bargaining Behavior? Experimental Bargaining in Japan and the United States", Journal of Consumer Research, 31, 181-190.

[12] Burrows, P., Loomes, G. (1994). "The impact of fairness on bargaining", Empirical Economics, 19(2), 201-221.

[13] Cappelen, A. W., Konow, J., Sorensen, E. O., & Tungodden, B. (2013). "Just luck: An experimental study of risk-taking and fairness", American Economic Review, 103(4), 1398-1413.

[14] Cappelen, A., Drange, A., Sørensen, E., Tungodden, B. (2007). "The Pluralism of Fairness Ideals: An Experimental Approach", American Economic Review, vol. 97(3), 818-827.

[15] Charness, G., Rabin. M. (2002). "Understanding social preferences with simple tests", Quarterly Journal of Economics, 117, 817–869.

[16] Corneo, G., Gruner, H. M. (2002). "Individual preferences for political redistribution", Journal of Public Economics, 83, 83–107.

[17] Diamond, P. (1967). "Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparison of Utility: A Comment", Journal of Political Economy, 75(5), 765-766.

[18] Elster, J. (1989). Solomonic judgements: Studies in the limitation of rationality. Cambridge University Press.

[19] Erkal, N., Gangadharan, L. Nikos Nikiforakis (2010). "Relative Earnings and Giving in a Real-Effort Experiment", American Economic Review, 101(7), 3330-48.

[20] Fehr, E., Schmidt, K. (1999). "A theory of fairness, competition and cooperation", Quarterly Journal of Economics, 114, 817-68.

[21] Fischbacher, U. (2007). "z-Tree: Zurich toolbox for ready-made economic experiments", Experimental Economics 102, 171–178.

[22] Fong, C. (2001). "Social preferences, self-interest, and the demand for redistribution", Journal of Public Economics, 82(2), 225-246.

[23] Frey, B., Stutzer, A. (2005). "Beyond outcomes, measuring procedural utility", Oxford Economic Papers 57: 90-111.

[24] Gill, D., Prowse, V., Vlassopoulos, M. (2012). "Cheating in the workplace: An experimental study of the impact of bonuses and productivity". Available at SSRN 2109698.

[25] Guth, W., Tietz, R. (1986). "Auctioning ulitmatum bargaining positions—how to act rational if decisions are unacceptable?" In R.W. Scholz (ed.), Current Issues in West German Decision Research, Frankfurt: Verlag Peter Lang, 173–185.

[26] Hammond, P. (1988). "Consequentialist foundations for expected utility", Theory and Decision, 25 (1), 25-78.

[27] Handgraaf, M., Van Dijk, E., Vermunt, R., Wilke, H., De Dreu, C. (2008). "Less power or powerless? Egocentric empathy gaps and the irony of having little versus no power in social decision making", Journal of Personality and Social Psychology 2008. 95 (5), 1136-49.

[28] Hargreaves-Heap, S., Varoufakis, Y. (2002). "Some Experimental Evidence on the Evolution of Discrimination, Co-operation and Perceptions of Fairness", Economic Journal, 112, 679-703.

[29] Hoffman, E., Spitzer, M.L. (1985). "Entitlements, rights and fairness: An experimental examination of subjects' concepts of distributive justice", Journal of Legal Studies 14, 259-297.

[30] Hoffman E., McCabe K., Shachat, K., Smith, V. (1994). "Preferences, Property Rights, and Anonymity in Bargaining Games", Games and Economic Behavior, 7 (3), 346-380.

[31] Kahneman, D., Tversky, A. (1979) "Prospect Theory: An Analysis of Decision under Risk", Econometrica, Vol. 47, No. 2. (Mar., 1979), pp. 263-292.

[32] Karni, E., Salmon, T., Sopher, B. (2008). "Individual sense of fairness: an experimental study", Experimental Economics, 11(2), 174-189.

[33] Karni, E., Safra, Z. (2002). "Individual sense of justice: a utility representation", Econometrica, 70, 263–284.

[34] Keren, G., Teigen, K. (2010). "Decisions by coin toss: Inappropriate but fair", Judgement and Decision Making, 5(2), 83-101.

[35] Konow, J. (2003). "Which is the fairest one of all? A positive analysis of justice theories", Journal of economic literature, 41(4), 1188-1239.

[36] Krawczyk, M. (2011). "A model of procedural and distributive fairness", Theory and Decision, 70(1) 111-128.

[37] Krawczyk, M. (2010). “A glimpse through the veil of ignorance: Equality of opportunity and support for redistribution”, Journal of Public Economics, 94, 131-141.

[38] Krawczyk, M., Le Lec, F. (2010). “Give me a chance! An experiment in social decision under risk”, Journal of Experimental Economics, 13, 500–511.

[39] Leventhal, G. S. (1980). What should be done with equity theory? (pp. 27-55), Springer US.

[40] Machina, M. (1989). “Dynamic Consistency and Non-Expected Utility Models of Choice Under Uncertainty”, Journal of Economic Literature, 27 (4), 1622-1668.

[41] Nozick, R. (1994). The nature of rationality, Princeton: Princeton University Press.

[42] Oosterbeek, H., Sloof, R., van de Kuilen, G. (2004). “Cultural Differences in Ultimatum Game Experiments: Evidence from a Meta-Analysis”, Experimental Economics, 7 (2), 171-188.

[43] Rawls, J. (1999). A Theory of Justice, rev. ed. Cambridge, MA: Belknap, 5.

[44] Schmitt, P. M. (2004). “On perceptions of fairness: The role of valuations, outside options, and information in ultimatum bargaining games”, Experimental Economics, 7(1), 49-73.

[45] Schurter, K., Wilson, B.J. (2009). “Justice and Fairness in the Dictator Game”, Southern Economic Journal, 76(1), 130-145.

[46] Suleiman, R. (1996). "Expectations and fairness in a modified UG", Journal of Economic Psychology, 17, 531–554.

[47] Thibaut, J. W., Walker, L. (1975). Procedural justice: A psychological analysis. Hillsdale: L. Erlbaum Associates.

[48] Trautmann, S. and van de Kuilen, G. (2014). Process fairness, outcome fairness, and dynamic consistency: Experimental evidence, mimeo.

[49] Trautmann, S., Wakker P. (2010). "Process fairness and dynamic consistency", Economics Letters, 109(3), 187-189.

[50] Trautmann, S.T. (2009). "A tractable model of process fairness under risk", Journal of Economic Psychology, 30(5), 803-813.

[51] Tyler, T. R. (2006). "Why people obey the law." Princeton University Press.

[52] Woolridge, J.M. (2002). Introductory Econometrics: A Modern Approach, Cincinnati, OH: South-Western College Publishing.

A Acknowledgements

We thank Iwan Barankay, Dirk Engelmann, Enrique Fatás, Peter Hammond, Andrew Oswald, Elke Renner, Blanca Rodriguez, Tim Salmon, Stefan Traub for useful discussion, as well as participants in the Workshop on 'Procedural fairness - theory and evidence', Max Planck Institute for Economics (Jena), the 2008 IMEBE conference, the 2008 European ESA conference, the 2010 Seminar on 'Reason and Fairness', Granada, and seminar participants at Nottingham, Royal Holloway, Trento, Warwick. We also appreciate the careful reading of the paper and the insightful suggestions by two anonymous referees and the journal editors. We especially thank Malena Digiuni for excellent research assistance, Stefan Trautmann for fruitful discussion, and Steven Bosworth for his comments on a previous version of the paper. Any errors are our sole responsibility. This project was financed by the University of Warwick RDF grant RD0616. Gianluca Grimalda acknowledges financial support from the Spanish Ministry of Science and Innovation (grant ECO 2011-23634), Bancaixa (P1·1A2010-17), Junta de Andalucía (P07-SEJ-03155), and Generalitat Valenciana (grant GV/2012/045).