Social Preferences and Competition*

Klaus M. Schmidt[a)]

University of Munich, CESifo and CEPR

October 2009

Preliminary Version

ABSTRACT: The folk wisdom of behavioral economics says that social preferences do not matter in competitive markets. The market experiments of Smith (1962, 1964) and others show that the standard neoclassical model predicts the actual market outcomes quite well, even if these outcomes are very unfair and give all the surplus to one side of the market. This has been confirmed theoretically by Fehr and Schmidt (1999) and Dufwenberg et al. (2008) who show that competition forces all market pariticipants to behave as if they were purely self-interested. The purpose of the current paper is to qualify this view. I will discuss the experimental and theoretical literature and argue that the folk wisdom is indeed correct if two conditions are met: preferences satisfy a "separability condition" (first identified by Dufwenberg et al. 2009) and complete contingent contracts are traded. These conditions are plausible in many markets for well defined physical goods, but they fail to hold when uncertainty is an issue (e.g. on financial markets) and when incomplete contracts are traded (e.g. on labor markets). In fact, social preferences can explain many of the anomalies frequently observed on these markets.

KEYWORDS: Social preferences, competition, asset markets, labor markets.

JEL CLASSIFICATION NUMBERS: C9, D5, J0

---

*) I would like to thank Ernst Fehr, Georg Gebhardt and Monika Schnitzer for helpful comments and suggestions. Financial support by the German Excellence Initiative and the Deutsche Forschungsgemeinschaft through SFB-TR 15 is gratefully acknowledged.

a) Klaus M. Schmidt, Department of Economics, University of Munich, Ludwigstr. 28, D-80539 München, Germany, email: klaus.schmidt@LMU.de .

1 Introduction

There is a folk wisdom in behavioral economics saying that social preferences do not matter in competitive markets. Hundreds of market experiments, starting with Smith (1962, 1964), have shown that the standard neoclassical model predicts the actual market outcomes quite well, even if these outcomes are very unfair and give all the surplus to one side of the market. This has been confirmed theoretically by Fehr and Schmidt (1999) and Dufwenberg et al. (2008) who show that competition forces all market pariticipants to behave as if they were purely self-interested. The purpose of the current paper is to qualify this view. I will discuss the experimental and theoretical literature and argue that the folk wisdom is indeed correct if two conditions are met: preferences satisfy a "separability condition" (first identified by Dufwenberg et al. 2008) and complete contingent contracts are traded. These conditions are plausible in many markets for well defined physical goods. However, they fail to hold when uncertainty is an issue (e.g. on financial markets) and when incomplete contracts are traded (e.g. on labor markets). In fact, social preferences can explain many of the anomalies frequently observed on these markets. Furthermore, even if the two conditions are met and market outcomes correspond to the predictions of the standard neoclassical model, the two fundamental theorems of welfare need not hold.

The standard neoclassical model is built on the assumption that all economic agents are only interested in their own material well-being. However, there is a large body of experimental and field evidence showing that many people are not purely self-interested. Many people care about the welfare of other people. They are willing to sacrifice own resources to promote fairness, to help those who have been kind to them and to punish those who have been unkind. This has been called "social preferences" or "other-regarding preferences" in the behavioral literature.[1] In Section 2 I will offer a brief survey of this literature. I will discuss the experimental evidence for social preferences and show that this evidence is not restricted to lab experiments for small stakes with undergraduate students, but that it extends to large stakes, representative samples of the population, and field evidence. I will also briefly discuss the most prominent theoretical models that have been suggested to explain this evidence.

Section 3 considers competitive markets. I discuss the market experiments by Smith and others showing that the neoclassical model does a surprisingly good job in predicting actual market outcomes even if many of the assumptions of this model are not satsified. The

[1] The two terms are used synonymously in the literature. For brevity I will mostly use the expression "social" preferences.

competitive equilibrium outcome prevails even if there is only a limited number of buyers and sellers, if market transperency is imperfect, if there is no Walrasian auctioneer, and if some market participants have strong social preferences. The fact that social preferences do not seem to matter in these experiments has been discussed theoretically by Fehr and Schmidt (1999) who show that in competitive ultimatum games players are forced to behave as if they were purely self-interested. Dufwenberg et al. (2008) consider a general equilibrium model and allow for a large class of social preferences. They identify a necessary and sufficient condition on preferences, called "separability", under which all agents with social preferences behave as if they were "classical" (i.e. purely self-interested). However, even if the separability condition holds, the two fundamental theorems of welfare need not hold. In particular, if some agents have social preferences competitive equilibria need not be efficient.

Dufwenberg et al. (2008) point out that the separability condition is unlikely to be satisfied when there is uncertainty, but they do not explore this issue. In Section 4 I use the example of an asset market to point out why the separability condition fails to hold with uncertainty and to show that large deviations from the neoclassical prediction are possible even if all players form rational expectations. Social preferences may give rise to mulitple equilibria, to herding, and to booms and busts on asset markets.

An implicitly assumption in Dufwenberg et al. and in the market experiments by Smith and others is that complete contracts are traded. The quality of the good can be perfectly specified for all possible states of the world in a complete contingent contract that is enforced by the courts. The market price determines how the gains from trade are split, but it does not affect the efficiency of the transaction. In the real world, however, many contracts are highly incomplete.

A prime example for an incomplete contract is a labor contract. In Section 5 I briefly discuss why labor contracts are incomplete and I argue that social preferences can be used as a substitute for performance based incentive schemes. In this case the price (wage) does not only determine the distribution but also the size of the surplus. There are several recent experimental studies showing that the role of social preferences is magnified when parties interact repeatedly and form relational contracts (Brown et al. 2004, 2008) or if parties interact one-shot but can acquire a reputation for fair or trustworthy behavior (Bartling et al. 2009). These papers also show that competition may foster the role of social preferences as an enforcement device and induce agents to spend more effort. In the experiments we observe that agents earn rents that are not competed away, that prices fail to clear the market, that unvoluntary unemployment is a stable phenomenon, that changes of market conditions affect

prices much less on markets for incomplete contracts than on market for complete contracts, and that minimum wages tend to increase even those wages that are not directly affected by the minimum wage. All of these phenomena cannot be explained by the standard neoclassical model but they are consistent with models of social preferences.

We conclude with a brief summary of the main insights of this paper and their implications for macroeconomics. Most modern macroeconomic models are explicitly or implicitly based on a general equilibrium model of perfectly competitive markets assuming that all economic agents are perfectly rational and self-interested. This is justified if markets for complete contracts with little uncertainty are considered. However, if markets are incomplete or if uncertainty is an important issue, the neoclassical model of competitive markets may be misleading.

2. Experimental Evidence and Theoretical Models of Social Preferences

After more than two decades of research in experimental economics there is a large body of experimental evidence showing that

(1) Many people do not only care about their own material well-being, but are also concerned about the payoffs of other people they interact with.
(2) People are heterogeneous. Some people care a lot about other people's payoffs, while others care very little.
(3) Social preferences have a strong and systematic impact on behavior. It is impossible to explain observed economic behavior in many situations without taking (1) and (2) into account.

In the first part of this section I give a short overview on the experimental evidence of social preferences. I will show that this evidence is quite robust and not restricted to lab experiments with undergraduate students and small stakes. In the second part I briefly discuss some theoretical models of social preferences.

2.1. Evidence on Social Preferences

The first and probably most famous experiment on social preferences is the *ultimatum game* introduced by Güth et al. (1982). In this experiment player 1, the proposer, can make a proposal on how to devide, say, $10 between himself and an anonymous player 2. Player 2,

the responder, can either accept or reject the proposal. In the latter case both players get a payoff of 0. Clearly, a rational and purely self-interested player 2 should accept any positive offer. However, a robust observation is that responders frequently reject low offers (Camerer 2003, Fehr and Schmidt 2003, Roth 1995). Because this decision of the responders is so simple and transparent (accept the offer or get nothing) it is difficult to argue that it is taken by mistake. Responders seem to prefer the payoff allocation (0,0) to the allocation (10-x,x) when they reject low offers. In fact, when asked why they chose to destroy their own payoff, many subjects in the role of player 2 say that the proposer's offer was "unfair". They want the unfair behavior of player 1 punished, even if this comes at a cost to themselves. More recently, neuroscientific studies provide evidence that is consistent with the view that fairness concerns and the punishment of unfair behavior are an expression of preferences. Quervain et al. (2004) show that reward-related neural circuits are activated when subjects decide to punish unfair behavior – even if subjects have to pay to punish. These are the same neural circuits that are activitated when the subjects decide to buy a good that they value.

Another robust regularity of ultimatum game experiments is that most subjects in the role of player 1 offer between 40 and 50 percent of the surplus to player 2. There are two possible explanations for this behavior. Player 1 may offer a fair share to player 2 because he prefers a fair allocation with a lower payoff to himself to an unfair allocation where he gets more. Or he may offer a fair allocation because he is afraid that an unfair offer will be rejected. The *dictator game* experiment discriminates between these hypotheses. Like in the ultimatum game, player 1 can make a proposal on how to divide $10 between himself and player 2. However, in the dictator game player 2 cannot reject the proposal. The proposal is dictated by the proposer. Forsythe et al. (1994) were the first to compare offers in the ultimatum game to offers in the dictator game. They find that offers in the dictator game are considerably less generous. This shows that many proposers are generous only for strategic reasons in the ultimatum game. A significant fraction of dictators give nothing. Nevertheless, most subjects still give a positive amount (up to 50 percent of the pie) to player 2. Again, the dictator game is so straightforward that proposers must know what they are doing. Thus, the experimental evidence suggests that many subjects are willing to give up own resources to help others.

There are many other experiments in which observed behaviors are inconsistent with the self-interest assumption.[2] For example, in *public good games* many people deviate from the dominant strategy of free-riding and voluntary contribute to the public good (Ledyard

[2] See Fehr and Schmidt (2006) for a detailed survey on this literature.

1995). Furthermore, if given the opportunity, they are willing to punish noncontributors even if this is costly to themselves (Fehr and Gächter, 2000). In *trust games* player 1 can make an investment that increases the payoff of player 2 (the "trustee"). The trustee can return the favor and voluntarily pay something back. A selfish trustee will not pay back and should not be trusted. However, many subjects in the role of player 1 invest in the (often correct) expectation that the trustee will return the favor and pay back (Berg et al. 1995). In *gift exchange games* (further discussed in Section 5) subjects in the role of workers provide higher effort than contractually enforceable if their employer offers a generous wage rent (Fehr et al. 1993).

If everybody cared strongly about fairness and reciprocity it would be easy to solve public good and incentive problems. However, this is definitely not the case. While many subjects are strongly concerned about other people and are willing to spend ressources in order to achieve a fair allocation or to reciprocate kind or unkind behavior, there are also many subjects who behave very selfishly. They don't give anything in a dictator game, they free ride in public good games, and they do not reciprocate in trust and gift exchange games. To understand the outcomes of these experiments it is necessary to study the interaction of the fair-minded and the self-interested subjects (Fehr and Schmidt 1999, 2006).

Most of these experiments have been conducted with modest amounts of money at stake and with subject pools of undergraduate students at Western universities. Thus, a natural question is how robust these results are. Several papers examine high stakes experiments. Perhaps surprisingly, even large increases in the monetary stakes did little or nothing to change behavior. Cameron (1999) conducted ultimatum games in Indonesia. In one treatment subjects could earn the equivalent of three months' income. Nevertheless she finds no effect of the stake level on proposers' behavior and only a small reduction of the rejection probability of the responder when stakes are high. Fehr, Fischbacher and Tougareva (2002) conduted gift exchange games in Russia. In one treatment their subjects earned, on average, the income of one week, in another treatment the income of ten weeks. Despite this large increase in monetary payoffs there are no significant differences in behavior across conditions.

The experimental evidence is not confined to student populations. Several studies conducted experiments with subject pools that are representative of whole countries such as Germany (Fehr et al. 2002, Dohmen et al. 2009) or the Netherlands (Bellemare and Kröger 2007, Bellemare et al. 2008). Differences in behavior to student subject pools are small. If

anything, students behave slightly more selfishly than a representative sample of the population.

There are also a few studies that did experiments across different cultures. Roth et al. (1991) conducted ultimatum games in Israel, Japan, Slovenia and the United States. They find somewhat lower rejection rates and lower offers in Japan and Israel compared to the US and Slovenia, but the differences are small. Large differences in behavior across cultures are observed by Henrich et al. (2001) and Henrich et al. (2004). They report the results of ultimatum game experiments conducted in 15 small-scale societies with little exposure to Western societies. For example, the average offer made by the Machiguenga (who live in the Amazon jungle of Peru) is only 26 percent, while the Lamalera (whale hunters on a remote Pacific island in East Indonesia) offer 56 percent on average. This evidence suggests that fairness norms are at least partially determined by culture, but also that the cultural differences concerning fairness between most Western countries are small.

Laboratory experiments allow the researcher to tightly control the decisions of the subjects. However, behavior observed in the artificial environment of a lab may differ substantially from behavior in natural environments. Some recent studies have addressed this issue and implemented gift exchange situations in natural environments. In these studies experimenters have exogenously manipulated the wage paid to real workers in situations with a one-shot character. For example, workers had to perform tasks such as data entry (Gneezy and List, 2006, Kube et al., 2006, Englmaier and Leider, 2008), stuffing envelopes (Al-Ubaydli et al., 2008), planting trees (Bellemare and Shearer, 2007) and newspaper promotion (Cohn et al. 2007). The general message of these studies is that significant reciprocal responses exist in these field environments. For example, a wage cut relative to the promised or expected payment reduces workers' output significantly (Kube et al., 2006). The impact of a wage increase is less pronounced. Cohn et al. (2007) find that only those workers who considered the previous wage as unfairly low respond to an increase in wages with a significantly positive effort increase, while those workers who perceive the previous wage as fair do not work harder.[3]

[3] Kube et al. (2008) show that workers who receive a non-monetary gift in gift wrap paper exhibit a large increase in effort while workers who received the monetary value of the gift increase their effort significantly less. However, most workers preferred the monetary amount to the gift. Thus, the gift wrapping has a strong effect, perhaps because it makes the gift more salient and more personal.

2.2. Theoretical Models of Social Preferences

Various theories have been developed to explain the experimental evidence. All of these theories employ "social preferences" in the sense that the utility of a subject not only depends on his own consumption level but also on the consumption of other players. Three types of models can be distinguished: (1) *outcome-based social preferences*, (2) *intention-based social preferences*, and (3) *type dependent social preferences*.

In models of *outcome-based social preferences* the utility of each player depends directly on the material payoff of other players in his reference group. For example, theories of unconditional altruism (Andreoni and Miller 2002, Charness and Rabin, 2002) assume that the utility of a player not only depends on his own material payoff but is also monotonically increasing in the monetary payoff of other players. Altruism can explain giving in dictator games and reciprocal behavior in trust and gift exchange games, but it cannot explain why subjects reject unfair offers in ultimatum games or punish free-riders in public good games. The opposite assumption is made in models of unconditional envy or spitefulness (Bolton 1991, Kirchsteiger 1994). Envy is consistent with spiteful behavior in ultimatum games and public good games with punishment, but it cannot explain positive reciprocity and generosity. A conditional form of altruism and/or envy is inequity aversion (Fehr and Schmidt 1999, Bolton and Ockenfels 2000). An inequity averse person is willing to spend own ressources in order to achieve a more equitable allocation of payoffs in his reference group. In most experimental games "equity" is defined as equality of monetary payoffs and the reference group is assumed to be the set of players a subject interacts with.[4] Thus, if an inequity averse person is worse off than the other people in his reference group his utility function decreases with the monetary payoff of his opponents, while if this person is better off, he becomes altruistic. These models are consistent with generous behavior in dictator, trust and gift exchange games and with spiteful behavior in ultimatum games and public good games with punishment.

A particularly simple formalization of inequity aversion is the Fehr-Schmidt (1999) model that assumes that the utility function of player *i* is given by:

$$U_i(x_i, x_{-i}) = x_i - \alpha_i \frac{1}{N-1}\sum_{j=1}^{N} \max\left\{x_i - x_j, 0\right\} - \beta_i \frac{1}{N-1}\sum_{j=1}^{N} \max\left\{x_j - x_i, 0\right\}$$

[4] In the real world the definition of equity and of the reference group is often less obvious and depends on the specific application.

with $0 \le \beta_i \le \alpha_i$ and $0 \le \beta_i < 1$. Thus, inequity aversion is assumed to be linear in payoff differences, and people are assumed to care more strongly about inequity that is to their disadvantage than about inequity to their advantage ($\beta_i \le \alpha_i$). Because of its simplicity this model can allow for heterogeneous preferences and still be easily applied to any experimental game. If all agents were inequity-averse to the same degree it would be impossible to explain the wide spectrum of fair-minded and unfair behavior in many experiments. Fehr and Schmidt show that these outcomes are often driven by the interaction between self-interested and inequity-averse players. For example, a few strongly inequity averse players are sufficient to induce full cooperation by all subjects in a public good game with punishment while a few selfish players ($\alpha_i = \beta_i = 0$) are sufficient to induce a very unfair outcome in an ultimatum game with responder competition.

The Fehr-Schmidt model is consistent with observed behavior in many experimental games but it assumes that people care only about outcomes and not about intentions. In some situations this is problematic. For example, Falk et al. (2003) consider a mini-ultimatum game in which the proposer is restricted to choose between two different allocations. In one treatment he can choose between (8,2) and (5,5), in another treatment he can choose between (8,2) and (10,0). In the first treatment the choice of (8,2) may be considered greedy, because he could have offered the equal split (5,5). In the second treatment the choice of (8,2) may be considered generous, because the only alternative would have been (10,0). Models of outcome-based social preferences predict that if the second mover is offered (8,2) he should behave the same way no matter whether the alternative had been (5,5) or (10,0). In the experiment about 45 percent of the responders reject (8,2) if the alternative is (5,5), while only 20 percent reject this proposal if the alternative is (10,0). Thus, in addition to outcome based inequity aversion intentions clearly matter.

Models of *intention-based social preferences* try to capture this effect (Rabin 1993, Dufwenberg and Kirchsteiger 2004, Falk and Fischbacher 2006). In these models preferences are defined not only over outcomes but also over beliefs about why an agent has chosen certain actions. This requires the use of psychological game theory (Genakoplos, Pearce, Stacchetti, 1989). Unfortunately, these models often give rise to multiple equilibria with self-fulfilling beliefs: If everybody believes that everybody else acts kindly it is optimal to be kind, but it is also an equilibrium that everybody is hostile because everybody believes everybody else to be hostile as well. Thus, these models are less suitable to predict behavior.

Finally, there are models of *type-based social preferences* (Levine, 1998). In these models a player behaves kindly to a “good” person and hostile towards a “bad” person. Note

that it is the type of a person and not the intention behind an action that affects preferences. Thus, type-based reciprocity can be modeled by using conventional game theory. However, if the type of player is private information, the game turns into a signalling game which is again plagued by multiple equilibria.

3. Are Social Preferences Irrelevant in Competitive Markets?

3.1. Market experiments

The first experiments on competitive markets were conducted by Edward H. Chamberlin in Harvard in the 1940s. On the first day of his beginning graduate course on monopolistic competition he used to divide his students into buyers and sellers of a fictitious good. Each buyer was assigned a different reservation price and each seller a different cost to deliver the good. Reservation prices and costs were private information. Then he gave his students a few minutes to find a trading partner and to haggle about the price. When he compared the actual trades with the prediction of neoclassical price theory, the typical result was that prices fluctuated widely and that the traded quantity was often larger than the competitive equilibrium quantity (Chamberlin 1948).

In the late 1950s Vernon Smith, a student of Chamberlin at Harvard, conjectured that the problem with Chamberlin's market experiment was the lack of public information about available bids and offers. Smith (1962, 1964) conducted a series of market experiments that differed from Chamberlin's experiments in two dimensions: First, all bids and offers were publicly recorded in order to improve market transparency. Because buyers and sellers can make bids and offers simultaneously, this market design is called a "double auction". It resembles the trading rules and procedures of the traditional trading floor of most financial markets before the introduction of computerized trading. Second, each experimental session had several rounds, so that his subjects could learn by experience. The experimental results are striking: Prices quickly converge to equilibrium prices and the traded quantity is very close to the efficient quantity predicted by the competitive equilibrium.

The experimental results match the predictions of neoclassical price theory even under extreme conditions. For example, Holt, Langan and Villamil (1986) conducted a double auction in which all buyers have the same reservation price and all sellers have the same cost. If there is excess supply, the theory predicts that all surplus goes to the buyers; if there is

excess demand, all surplus goes to the sellers. Many market participants consider these outcomes as very unfair. Nevertheless, after a few trading periods these are exactly the market outcomes observed in the experiments.

Roth et al (1991) introduced competition into the ultimatum game. In their experiment there are n proposers who simultaneously propose a share $s_i \in [0,1]$, $i \in \{1,...,n\}$, to one responder. Then the responder can either accept or reject the *highest* offer $\bar{s} = \max\{s_i\}$. They conducted experimental sessions in four different countries. In all experiments the maximum offer converged quickly to one and was accepted by the responder. Thus, the responder received all the gains from trade while the proposers got nothing. Similarly, Güth, Marchand and Rulliere (1997) conducted an ultimatum game with responder competition. One proposer offers a share $s \in [0,1]$ to n responders. At the same time each responder decides on his acceptance threshold. If several responders have an acceptance threshold that is smaller than the proposed share, one of them is selected at random. After five periods the average acceptance threshold is below five percent, with 71 percent of the responders stipulating a threshold of exactly zero. Thus, even though most subjects reject unfair offers in the standard ultimatum game, they are willing to accept them if there is competition.

3.2 Irrelevance Theorems

Why is it that many people seem to have strong social preferences when they interact in small groups, while they seem to be purely self-interested when there is competition? Fehr and Schmidt (1999) addressed this puzzle. Using their model of inequity aversion they analyse the ultimatum games with responder and proposer competition. They show that even if all people are highly inequity averse, competition forces them to behave as if they were purely self-interested (Fehr and Schmidt 1999, Propositions 2 and 3). The intuition for this result is easy to understand in the ultimatum game with proposer competition. Note first that if $\beta < 1$ the responder will accept any offer. Consider now an equilibrium candidate where $\bar{s} < 1$. It cannot be the case that one proposer offers less than $\bar{s}$ in equlibrium, because then he would get a monetary payoff of 0 for sure and suffer from the inequality to his disadvantage, while by offering $\bar{s}$ he would get $1-\bar{s}$ and less inequality to his disadvantage with positive probability. But if he offers $\bar{s}+\varepsilon$, he gets $1-\bar{s}-\varepsilon$ with probability one, which is even better. Thus, as in a game with Bertrand competition, each proposer has an incentive to overbid the other proposers, and the only equilibrium outcome has at least two proposers offering $s = 1$. With

competition it is impossible for any one party to enforce a fair or equal outcome. If a player insists on his fair share, he only hurts himself, but he cannot prevent the other market participants from trading. [5]

More recently, Dufwenberg et al. (2008) have looked at this question from the perspective of general equilibrium theory. They consider a standard general equilibrium model and allow for a very general class of social preferences. The preference relation of each market participant can be represented by a utility function that may depend not only on his own consumption vector, x_i, but also on the consumption vectors x_{-i} of all other market participants. Furthermore, it may depend on the vector of budget sets, *B*, of all market participants. Thus, they allow for the possibility that a consumer does not care about the consumption bundle chosen by another consumer, but he does care about the consumption possibilities (the budget set) of this consumer.

In a classical general equilibrium model the utility of a consumer depends on his own consumption vector only. Thus, a consumer behaves "as if classical", if his demand function is independent of the consumption bundles and budget sets of all other consumers. Dufwenberg et al. (2008) offer a necessary and sufficient condition for a consumer to behave "as if classical", i.e. for his demand function to be independent of the consumption and income of all other consumers. A consumer behaves "as if classical" if and only if his preferences are separable, i.e. if and only if his preference relation can be represented by a utility function of the form

$$V_i(m_i(x_i), x_{-i}, B)$$

and if $V_i(\cdot)$ is strictly increasing in its first argument. The function $m_i(x_i)$ is called the consumer's *internal utility function*. The idea of the separability condition is simple. Because V_i is strictly increasing in the internal utility m_i, the consumer wants to maximize m_i. However, m_i depends only on x_i and is independent of (x_{-i}, B). Thus, for any (x_{-i}, B) the consumer will choose the same consumption bundle x_i^*, i.e., he will behave as if classical. If all consumers have preferences that are separable, then all consumers will behave as if classical, and social preferences do not affect market behavior. Thus, an economy in which agents have social preferences is observationally equivalent to an economy in which each agent only cares about his own consumption.

[5] Bolton and Ockenfels (2000) consider Bertrand and Cournot Games and show that the standard Nash equilibria of these games are also equilibria if some of the players suffer if they have ERC preferences, i.e., if their payoff differs from the average payoff in the group.

In an economy with separable social preferences Walrasian equilibria must be internally efficient, i.e., it is impossible to increase the internal utility m_i of some consumers without reducing the internal utility of some other consumers. However, an internally efficient allocation need not be Pareto efficient. The reason is that social preferences give rise to externalities that will not be internatized if each agent chooses a consumption bundle that maximizes his internal utility. Thus, the two fundamental welfare theorems need not hold in this economy. A Walrasian equilibrium may be Pareto-inefficient, and a Pareto-efficient allocation need not be a Walrasian equilibrium.

Dufwenberg et al. (2008) show that if a social monotonicity condition holds, i.e. if it is possible to distribute any amount of money z in a such way among consumers that all consumers are strictly better off, then the set of Pareto-efficient allocations is a subset of the set of internally-efficient allocations. In this case the second welfare theorem trivially holds, because any internally efficient allocation can be implemented as a Walrasian equilibrium by using suitable lump-sum transfers. However, the first welfare theorem need not hold. There may be Walrasian equilibria that are internally efficient but not Pareto-efficient. Thus, if people have social preferences, redistribution may be necessary to achieve a Pareto-efficient allocation. For example, it there are two groups of consumers, rich and poor, and if the rich have sufficiently strong altruistic preferences, then a Pareto improvement can be achieved by transferring wealth from the rich to the poor. Note that this cannot be achieved with bilateral transfers if each of the rich is small. If one rich person gives some of his wealth to the poor, he suffers from the loss of his own consumption while the gains of the poor are negligible. However, if all of the rich give up some of their wealth there is a strong impact on the welfare of the poor and everybody is better off. A solution to this collective action problem is redistributive taxation. Note that this argument for redistribution is based only on individual preferences.

The separability condition offers an explanation for why we observe the competitive equilibrium outcomes predicted by neoclassical price theory in the market experiments by Smith and others and in the ultimatum games with competition. This condition is satisfied by the most prominent models of outcome based social preferences (such as Fehr-Schmidt 1999, Bolton-Ockenfels 2000 and Charness-Rabin 2002).[6] Nevertheless, it is important to keep in mind that the separability condition is restrictive. It requires that the marginal rate of substitution between two goods is independent of the consumption vectors and budget sets of all other agents in the economy. In a world without uncertainty this assumption is plausible as

[6] However, it is non-generic in the class of all possible outcome-based social preferences.

a first approximation.[7] With uncertainty, however, separability is very unlikely to hold. In the next section I will provide a simple example of an asset market illustrating this point.

4. Competition under Uncertainty

Consider an economy with aggregate uncertainty in which each consumer has to decide how to allocate his wealth across different states of the world. If the consumer has social preferences, his utility in one state of the world depends not only on his own consumption level in this state but also on the consumption of other agents in his reference group. For example, a consumer may experience a larger disutility from a sudden drop in consumption if he is the only one who is suffering as compared to a situation in which everybody in his neighborhood is affected by the same shock. Thus, the marginal rate of substitution between consumption in state 1 and in state 2 depends on how much other people consume in these two states. This violates separability.

If there is aggregate uncertainty and people have social preferences there is a natural tendency for herding and multiple equilibria on asset markets even if these markets are competitive and everybody is perfectly rational. This has been pointed out by Gebhardt (2002, 2004). Let me illustrate this point with a simple example.

Suppose that there is a continuum of identical consumers indexed by i, $i \in [0,1]$, each of whom has to decide on how much of his wealth to invest in a safe and in several risky assets. Consumers are inequity averse in the sense of Fehr and Schmidt (1999), i.e., their utility function is given by

$$V(x) = m(x_i) - \alpha \int_0^1 \max\{x_j - x_i\}\, dj - \beta \int_0^1 \max\{x_i - x_j\}\, dj$$

where x_i is the level of wealth enjoyed by the consumer after the state of the world has materialized and $m(x_i)$ is an internal (indirect) utility function with $m'(x_i) > 0$ and $m''(x_i) < 0$. Consumers want to maximize the expected value of $V(x)$.

If $\alpha = \beta = 0$ this is just a von Neumann-Morgenstern utility function. In this case, under standard assumptions, there exists a unique optimal portfolio choice for any vector of asset prices and a unique asset market equilibrium. If α and/or β are strictly positive,

[7] To be sure, even in a world without uncertainty the separability condition need not always hold. For example, my marginal utility from conspicious goods (expensive cars, jewelry, fashion goods, etc.) may depend on how much of these goods my neighbors consume and/or on how wealthy they are.

however, the consumer is not only concerned about his consumption risk, but also about the "social" risk that he may fall behind (or move ahead) of everybody else. Therefore, agents dislike taking risks that are not taken by their reference group. For example, suppose that there are two assets that have the same return and risk profiles but are not perfectly correlated. A self-interest agent wants to buy a convex combination of the two assets in order to reduce his total risk exposure. In contrast, an agent who is sufficiently inequity averse wants to hold the same asset allocation that is held by his reference group. If everybody holds only the first asset, he will do so as well. The reason is that if he buys some of the other asset he may fall behind his reference group in some states of the world and suffer a utility loss from inequity aversion in addtion to the monetary loss. By buying the same portfolio as his reference group he insures against this risk.

Thus, an agent may be willing to hold a portfolio with a risk-return profile that is strictly dominated by some other portfolio provided that enough other people in his reference group do so as well. Herding is an optimal strategy if investors have social preferences. If everybody buys internet stocks or large new houses an investor with social preferences may do so as well even if he believes that the expected return is negative. If prices continue to go up, he does not want to be left behind. If prices fall, he is not the only one whose wealth is reduced ("Two in distress make sorrow less"). This gives rise to multiple asset market equilibria. Gebhardt (2002, 2004) shows that these effects can be used to explain time varying risk premia, stock market bubbles and crashes even if all market participants have rational expectations and behave optimally.

5. Competition for Incomplete Contracts

An implicit assumption in the general equilibrium model of Dufwenberg et al. (2008) as well as in the market experiments by Smith and others is that complete contracts are traded. The quality of the good is observable by the trading parties and verifiable by the courts. The parties can write a complete contingent contract that specifies all aspects of their relationship in every possible state of the world. Any deviations from the terms of the contract are deterred by the threat of high damage payments that are enforced by the courts. The market price determines how the gains from trade are split, but it does not affect the efficiency of the transaction.

In the real world, however, many contracts are highly incomplete. A prime example for an incomplete contract is the employment relationship. A labor contract cannot specify in detail what the worker is supposed to do. Most occupations involve multiple and complex tasks, the optimal action to be taken depends on many contingencies that are unknown in advance, and it is impossible to constantly monitor the employee. This gives rise to severe incentive problems. In this section I will argue that in markets for incomplete contracts concerns for fairness and reciprocity may mitigate or exacerbate incentive problems. In these markets the market price has two functions. It not only determines how the gains from trade are split, it also affects the incentives of the involved parties and thereby the size of the surplus that is generated. Social preferences can explain why markets sometimes do not clear and why prices are more rigid than predicted by the neoclassical general equilibrium model. They can also explain why the contracting parties choose not to contract on some variables that could easily be contracted upon and to leave their contract more incomplete than necessary.

5.1 Incomplete Contracts and Incentives

Incentive problems arise when there are informational assymmetries that make it impossible to enforce the efficient allocation with a complete contingent contract. The question how to provide optimal incentives in an environment with asymmetric information has been addressed by the theory of "mechanism design" or "contract theory" which is based on the standard model of rational and purely self-interested agents. A general insight of contract theory is that the optimal contract is "comprehensive" (Hart 1995, p. 22), i.e. it has to be contingent on all verifiable information that contains statistical information about the agent's actions (Holmström 1982). A comprehensive contract offers performance based incentives in order to better align the interests of the agent and the principal. The optimal contract cannot implement the first best efficient allocation that would be implemented if there were no problems of asymmetric information, but it will be optimal subject to the informational constraints. Once the second best optimal contract has been found, contract theory predicts that this contract will be traded on competitive markets like any other contract.

However, this is not what we observe in real labor markets. Performance based compensation contracts are prevelant for chief executives and other employees at the high end of the hierarchy, but even these contracts are not comprehensive because they leave out many contingencies that could be contracted upon. For most employees, however, incentive pay is

the exception. The large majority of workers receive fixed hourly wages or monthly salaries without any explicit monetary incentives (MacLeod and Parent 1999). It may be impossible to contract on the employee's total performance, but there are always some performance measures that could be used to give at least some incentives to the employee. However, firms deliberately choose not to contract on these performance measures but to leave the contract highly incomplete.

An explanation for this puzzle has been offered by Holmström and Milgrom (1991). They argue that most employees work on multiple tasks. Typically, performance on some tasks is easy to measure while for other tasks there are no accurate performance measures. Rewarding performance on those tasks that are easy to measure may induce the agent to distort his effort allocation and to concentrate his efforts on the tasks that are rewarded and away from those that are not. To avoid a distorted effort allocation it may be necessary not to give any explicit incentives at all.

The problem with this argument is that a rational and purely self-interested agent will not spend any costly effort if he receives a fixed wage. This does not hold for an agent with social preferences. A fair-minded worker will reciprocate to a generous wage by spending more effort. Furthermore, if he is offered a generous fixed wage, he will allocate this effort efficiently across tasks in order to benefit the principal as much as possible at minimal cost to himself. Thus, preferences for fairness and reciprocity may act as an enforcement device. They induce the agent to work even if there are no explicit incentives. It may be optimal for a firm not to rely on performance based incentives schemes but rather to offer a generous fixed wage, i.e. a highly incomplete contract appealing to the reciprocity of the employee.

This has been confirmed in a laboratory experiment by Fehr and Schmidt (2004). In their experiment an agent has to engage in two complementary tasks. One task can be measured perfectly, the other not at all. When principals offer a piece-rate contract that makes pay a linear function of the performance in the measurable task almost all agents put all their efforts into this task and ignore the non-measureable task - which is highly inefficient. However, when principals offer a generous fixed wage (complemented by a discretionary bonus payment after the agent has taken his actions) most agents allocate their efforts efficiently and both parties, principals and agents, are much better off.

Employers and managers seem to be well aware of the multi-task problem and often stress that "workers have many opportunities to take advantage of employers so that it is not wise to depend on coercion and financial incentives alone as motivators. … other motivators are necessary, which are best thought of as having to do with generosity" (Bewley 1995).

5.2 Social Preferences as a Contract Enforcement Device

How effective are social preferences in motivating costly effort provision? It turns out that the answer to this question depends crucially on the strategic environment in which people interact and the reference point that they use to evaluate the fairness of a wage offer.

There is substantial experimental evidence from the lab and the field showing that many people withdraw effort when they feel treated unfairly and that some people are willing to work harder when they are treated generously. The work horse of the experimental literature is the gift exchange game introduced by Fehr, Kirchsteiger and Riedl (1993). This two stage game captures the basic strategic features of an incomplete employment contract. At the first stage, the subject in the role of a "firm" can offer a fixed wage to the subject in the role of a "worker". At the second stage the worker chooses an effort level. Effort increases the payoff of the firm and the gains from trade, but it is personally costly to the worker. Standard theory predicts that a self-interested worker who is offered a fixed wage will always choose the minimum effort level. In contrast, models of social preferences predict that workers who are concerned about fairness will reciprocate to higher wages by spending more effort.

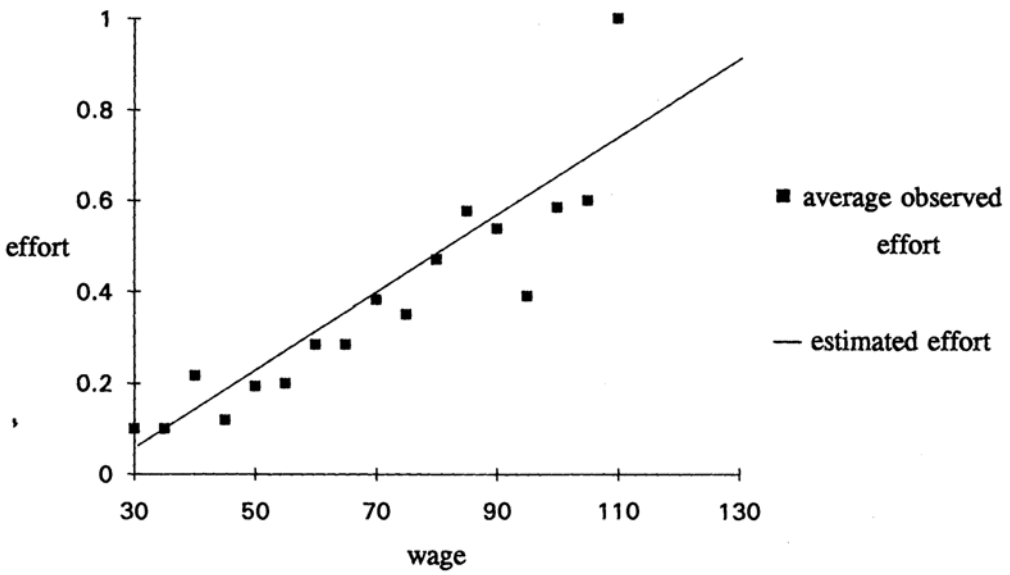

Figure 1: The Effort-Wage Relation in a Gift Exchange Game
Source: Fehr, Kirchsteiger and Riedl (1993, Fig. I)

Figure 1 shows what happened in the experiment: For low wages most workers choose the minimum effort level. However, if higher wages are offered, many workers choose higher effort levels. There is a lot of heterogeneity among workers. Many workers choose the

minimum effort level no matter what wage they are offered, but other workers are willing to reciprocate to high wages by working harder. The average observed effort level is increasing in wage.

This effect has been corroborated by many other experimental studies including Charness (2004), Charness et al. (2004), Gächter and Falk (2002), Fehr, Klein and Schmidt (2007), Hannan et al (2002). More recently, there have also been several field experiments confirming this result, including Gneezy and List (2006), Al-Ubaydli et al. (2008), Kube et al. (2006) and Kube et al. (2008). [8] However, in all of these experiments the average effort level is far below the efficient effort level for two reasons. First, only some of the workers reciprocate to high wages by choosing high effort levels. Second, because of this, it is often not profitable for employers to offer high wages.

The effectiveness of social preferences as a contract enforcement device increases substantially if parties interact repeatedly so that relational contracts are feasible (MacLeod and Malcomson 1989, Baker et al, 1994). Brown et al. (2004) conducted an experiment where employers can hire the same workers repeatedly. The employer can identify the worker he is matched with by an identification number and make a private contract offer to him in the next period if he wants to rehire him. If the offer is accepted, the employer and the worker stay together for another period, if it is rejected, the firm can still hire another worker on the market. The experiment ends after 15 periods. In the treatment with one-shot interaction, the average effort is 3.3, significantly more than the minimum effort of 1, but far less than the efficient effort of 10. With repeated interaction the average effort increases to 6.9.

The possibility of repeated interaction greatly amplifies the role of social preferences, even if all players interact for only a limited number of periods. If it is common knowledge that all players are selfish, a simple backward induction argument shows that workers will shirk in all periods and employers will offer the lowest possible wage. However, if there is asymmetric information and employers do not know whether they face a selfish or a fair-minded worker, high effort and high wages can be sustained. The intuition is familiar from the work of Kreps et al. (1982) on finitely repeated games with asymmetric information. If a worker is identified as selfish, he will be offered the lowest possible wage in all future periods. Therefore a selfish worker has an incentive to choose a high effort level in order to maintain a good reputation. In equilibrium all employers offer high wages to workers with a good reputation and low wages to workers who shirked, and all workers work hard except for the last few periods when the selfish workers start shirking with positive probability while the

[8] See Fehr, Goette and Zehnder (2009) for a survey of this literature.

fair-minded workers continue to spend high effort. Thus, a small fraction of fair-minded workers can induce purely self-interested workers to work hard and employers to pay generous wages.

Reputation does not require relational contracting. Bartling, Fehr and Schmidt (2009) conduct an experiment in which employers and workers interact one-shot, but the employer can observe the agent's effort level in the last three periods. Furthermore, the employer can choose whether to control or to trust his agent. If the employer trusts, the worker is free to choose any effort level between 1 and 10. If he controls, the worker must choose at least an effort level of 3, but he is slightly less productive. Bartling et al. (2009) find that most workers spend high effort if and only if they are offered a generous wage. If they are offered a low wage, most workers choose low effort (or reject the contract), even if this damages their reputation. The optimal strategy for employers is to trust and to offer generous wages to workers with a good reputation and to control and offer low wages to workers with a poor reputation.

In the experiment, Bartling et al. (2009) observe a clustering of two types of job offers. Employers either offer "good jobs" that involve full discretion, high wages and substantial rents that are left on the table to induce the worker to work hard. Or they offer "bad jobs" that involve control, low wages and hardly any rents for workers. Good job offers induce much higher effort and are more profitable for employers if and only if they are offered to high reputation workers. If workers have a low reputation or if reputation building is impossible because employers do not observe how workers behaved in the past then bad job offers outperform good job offers.

The clustering of job attributes that we observe in the lab is confirmed by field evidence that indicates that high wages, high rents, and high disretion are highly correlated. In the Socio-economic Panel (SOEP) a representative sample of the German population is asked a wide range of questions including questions on workplace characteristics, job satisfaction, earnings, education, work experience, etc. The 2001 wave of the survey asked 22,351 individuals the following questions: "Can you decide yourself how to complete your work tasks?" and "Is your work performance strictly monitored?" The answers to these questions (either "applies completely", "applies partly" or "does not apply at all") can be taken as a measure of job discretion. After controlling for occupation, industry, education, labor market experience, tenure, firm size, hours of work, gender and many other factors we still find that there is a positive, highly significant correlation between job discretion and earnings and between job discretion and job satisfaction. Jobs with high discretion (full autonomy and no

monitoring) are associated with 10 percent higher wages than jobs with no discretion. Furthermore, jobs with high discretion are associated with significantly higher job satisfaction, indicating that these jobs offer rents to employees.

5.3 Competitive Markets for Incomplete Contracts

Brown, Falk and Fehr (2004) implemented the gift exchange game in a competitive labor market. In each period firms make wage offers. Every employer can hire at most one worker and every worker can accept at most one job offer. There is an excess supply of workers, so not all workers can get a job. The experimental results confirm that higher wages induce workers to choose higher effort levels (on average). Furthermore, even though unemployed workers are eager to accept job offers with lower wages, employers are very reluctant to offer lower wages because they are afraid that lower wages will induce less effort. In fact, when some of them do reduce wages, workers accept these contracts but shirk. Thus, because effort is endogenous it is indeed more profitable to pay a rent rather than to hold workers down to their reservation utility. This implies that the price mechanism may fail to clear the labor market. If there is a demand shock there may be persistent unvoluntary unemployment because firms are unwilling to cut wages.

This has been confirmed by the Bewley (1995, 1999). Based on interview evidence with CEOs and personell managers of hundreds of companies he concludes that firms are very reluctant to cut wages because they expect a wage cut to hurt morale. It is more profitable to fire some workers rather than to reduce wages for all workers because the former strategy causes less resistance among the remaining workforce. This explains why wages do not fall in a recession.

When there is unvoluntary unemployment, relational contracts are very valuable to employed workers. They are willing to provide high effort in order to keep their jobs and not to become unemployed. The question arises whether relational contracts also work when the labor market is tight and workers can always find a new job. Brown, Falk and Fehr (2008) implemented an experimental labor market with an excess demand for workers. They find that effective relational contracts emerge even if there is full employment. If a worker works hard the incumbent firm offers a high wage in the next period that exceeds the going market wage that is paid to workers who change employers. Thus, workers who stay with one firm and work hard receive a rent that induces them to be loyal to their current employer. The lower

wage payed to workers who change employers is based on the expectation that these workers are more likely to shirk.

Comparing the market outcomes with an excess supply of workers in Brown et al (2004) to the market outcomes with an excess demand for workers in Brown et al (2008) the authors find that there is more turnover if there is excess demand. However, the average effort is roughly the same under both conditions, presumably because the negative effect on effort of a higher turnover is compensated by the positive effect of higher wages that are paid if there is an excess demand for workers. It is interesting to note that the wage difference between the excess supply and the excess demand condition is positive but rather small. In two control treatments the authors also implemented markets for complete contracts in which effort could be specified in the contract and enforced by the courts. They find that the wage difference between excess supply and excess demand is much larger if complete contracts are traded than if incomplete contracts are traded. This indicates that prices are more rigid when contracts are incomplete. A possible explanation is that with incomplete contracts worker are paid a rent not only when they are in excess demand but also when they are in excess supply.

Bartling et al. (2009) consider a competition treatment where employers compete for workers and workers compete for jobs. In their experiment the number of workers equals the number of jobs, so there is neither excess demand nor excess supply. They find a new effect of competition. If each firm is matched with exactly one worker (bilateral monopoly) only half of the contracts are trust contracts with high wages. In this case the average effort is 4.7, significantly more than the minimal enforceable effort of 3, but less than half of the efficient effort of 10. If firms compete for workers and workers compete for jobs the fraction of generous trust contracts increases to almost 80 percent and average effort rises to 7.3. The authors show that competition makes reputation building more attractive. Firms compete for workers with the best reputations driving up their wages. This induces workers to spend more effort in order to gain a high reputation. Those workers who fail to do so get control contracts with very low wages. Thus, competition increases the quasi-rents paid for good performance and thereby fosters trust and trustworthiness. It also increases the reliance on incomplete contracts and on reciprocity as an enforcement device. This contradicts the folk wisdom that competition drives out the role of social preferences.

Whether a certain wage is considered fair strongly depends on the reference group to which workers compare themselves. When a new worker is hired for a job the going market wage for this job is a natural benchmark. On the other hand, an incumbent worker is more likely to compare a proposed change in the employment relationship to the status quo. Thus,

theories of social preferences give rise to the following predictions: First, entry-level wages respond more strongly to changes in labor market conditions than the wages of incumbent workers. Second, there are cohort effects. A worker who entered the firm with a high entry-level wage (because the labor market was tight) considers this high wage as the reference point for the wage negotiations in the next period and will demand and get a higher wage than another worker who entered the firm with a low entry-level wage.

There is a lot of empirical evidence supporting these predictions. Several studies show that the wages of individuals who enter a firm are far more sensitive to the business cycle than wages of incumbent workers. There is also substantial evidence for cohort effects. Oreopoulos et al. (2006) show that Canadian students graduating in a boom year get wages that are about 9 percent higher than average. After five years wages are still 4 percent higher, and the effect fades only after 10 years.[9]

Finally, social preferences shed new light on some puzzling empirical findings about minimum wages. First, several empirical papers report that minimum wages have spill-over effects (Card and Krueger 1995, Teulings 2003). After the introduction of a minimum wage firms increase wages by more than necessary to comply with the new regulation. Second, firms make little use of the possibility to pay subminimum wages to some of their workers. For example, Katz and Krueger (1991) report that wages of teenage workers did not decline significantly after the introduction of an exception to the minimum wage law that allowed firms to pay subminimum wages to young workers.

Falk et al. (2006) consider the introduction of a minimum wage in an experimental labor market. They find that the minimum wage strongly affects reservation wages. This suggests that the minimum wage is perceived as a benchmark for evaluating the fairness of wage offers. Most workers are willing to accept wages significantly below the minimum wage before its introduction. After the minimum wage has been intoduced, however, many workers perceive a wage payment at the minimum wage level as unfair and reject these offers. A wage that is perceived as quite generous before the introduction of the minimum wage is perceived as greedy thereafter. This explains why firms have to pay more than the minimum wage if they want to appeal to the reciprocity of their workers. Furthermore, Falk et al (2006) observe a hysteresis effect. If the minimum wage is abolished, workers still use it as a benchmark and consider wage offers below the minimum wage as unfair. This is consistent with the observation of Katz and Krueger (1991) that firms did not use the possibility to pay

[9] See Fehr et al. (2009) for a more detailled discussion of these effects and the empirical literature.

subminimum wages to young workers after the corresponding change in the minimum wage law.

6. Conclusions

Most modern macroeconomics is explicitly or implicitly based on a general equilibrium model of perfectly competitive markets assuming that all economic agents are perfectly rational and self-interested. This is justified if markets for complete contracts with little uncertainty are considered. However, if markets are incomplete or if uncertainty is an important issue, the neoclassical model of competitive markets may be misleading. …

[To be completed]

References

Al-Ubaydli, Omar; Andersen, Steffen; Gneezy, Uri and John List. 2008. "For Love or Money? Comparing the Effects of Non-pecuniary and Pecuniary Incentive Schemes in the Workplace. George Mason University. mimeo.

Andreoni, James and John H Miller. 2002. "Giving According to GARP: An Experimental Test of the Consistency of Preferences for Altruism." *Econometrica,* Vol. 70, 737–753.

Baker, George; Gibbons, Robert and Kevin J Murphy. 1994. "Subjective Performance Measures in Optimal Incentive Contracts." *The Quarterly Journal of Economics,* Vol. 109, 1125–1156.

Bartling, Björn; Fehr, Ernst and Klaus M Schmidt. 2009. "Reputation and Job Design. University of Munich.

Bellemare, Charles and Bruce Shearer. 2007. "Gift Exchange within a Firm: Evidence from a Field Experiment." *SSRN eLibrary*.

Berg, Joyce; Dickhaut, John and Kevin McCabe. 1995. "Trust, Reciprocity, and Social History." *Games and Economic Behavior,* Vol. 10, 122–142.

Bewley, Truman F. 1999. "Why wages don't fall during a recession." Cambridge, Mass.: Harvard Univ. Press.

Bewley, Truman F. 1995. "A Depressed Labor Market as Explained by Participants." *The American Economic Review,* Vol. 85, 250–254.

Bolton, Gary E. 1991. "A Comparative Model of Bargaining: Theory and Evidence." *The American Economic Review,* Vol. 81, 1096–1136.

Bolton, Gary E and Axel Ockenfels. Mar., 2000. "ERC: A Theory of Equity, Reciprocity, and Competition." *The American Economic Review,* Vol. 90, 166–193.

Brown, Martin; Falk, Armin and Ernst Fehr "Competition and Relational Contracts: The Role of Unemployment as a Disciplinary Device." *SSRN eLibrary*.

Brown, Martin; Falk, Armin and Ernst Fehr. 2004. "Relational Contracts and the Nature of Market Interactions." *Econometrica,* Vol. 72, 747–780.

Camerer, Colin. 2003. "Behavioral Game Theory. Experiments in Strategic Interaction." Princeton, N.J.: Princeton University Press.

Cameron, Lisa A. 1999. "Raising the Stakes in the Ultimatum Game: Evidence from Indonesia." *Economic Inquiry,* Vol. 37, 47–59.

Card, David and Anne Krueger. 1995. "Math and Measurement: The New Economics of the Minimum Wage." Princeton, N.J.: Princeton Univ. Press.

Chamberlin, Edward H. 1948. "An Experimental Imperfect Market." *The Journal of Political Economy,* Vol. 56, 95–108.

Charness, Gary. 2004. "Attribution and Reciprocity in an Experimental Labor Market." *Journal of Labor Economics,* Vol. 22, 665–688.

Charness, Gary; Frechette, Guillaume R and John H Kagel. 2004. "How Robust is Laboratory Gift Exchange." *Experimental Economics,* Vol. 7, 189–205.

Charness, Gary and Matthew Rabin. 2002. "Understanding Social Preferences with Simple Tests." *The Quarterly Journal of Economics,* Vol. 117, 817–869.

Cohn, A.; Fehr, Ernst and Lorenz Goette "Gift Exchange and Effort: Evidence from a Field Experiment. University of Zürich. mimeo.

Dohmen, Thomas; Falk, Armin; Huffman, David and Uwe Sunde. 2009. "Homo Reciprocans: Survey Evidence on Behavioural Outcomes." *Economic Journal,* Vol. 119, 592–612.

Dufwenberg, Martin; Heidhues, Paul; Kirchsteiger, Georg; Riedel, Frank and Joel Sobel. 2008. "Other-regarding Preferences in General Equilibrium. CEPR Discussion Paper Series, 6815.

Dufwenberg, Martin and Georg Kirchsteiger. 2004. "A Theory of Sequential Reciprocity." *Games and Economic Behavior,* Vol. 47, 268–298.

Englmaier, Florian and Stephen Leider. 2008. "Contractual and Organizational Structure with Reciprocal Agents." *SSRN eLibrary.*

Falk, Armin; Fehr, Ernst and Urs Fischbacher. 2003. "On the Nature of Fair Behavior." *Economic Inquiry,* Vol. 41, 20–26.

Falk, Armin; Fehr, Ernst and Christian Zehnder. 2006. "Fairness Perceptions and Reservation Wages - The Behavioral Effects of Minimum Wage Laws." *Quarterly Journal of Economics,* Vol. 121, 1347–1382.

Falk, Armin and Urs Fischbacher. 2006. "A theory of reciprocity." *Games and Economic Behavior,* Vol. 54, 293–315.

Fehr, Ernst; Brown, Martin and Christian Zehnder. 2009. "On Reputation: A Microfoundation of Contract Enforcement and Price Rigidity." *The Economic Journal,* Vol. 119, 333–353.

Fehr, Ernst; Fischbacher, Urs; Rosenblatt, Bernhard von; Schupp, Jürgen and Gert G Wagner. 2002. "A Nation-Wide Laboratory - Examining Trust and Trustworthiness by Integrating Behavioral Experiments into Representative Surveys." *Schmollers Jahrbuch,* 519–542.

Fehr, Ernst; Fischbacher, Urs and Elena Tougareva. 2002. "Do High Stakes and Competition Undermine Fairness? Evidence from Russia." University of Zürich. Institute for Empirical Research in Economics, Working Paper Series, No. 120.

Fehr, Ernst and Simon Gaechter. 2000. "Cooperation and Punishment in Public Goods Experiments." *The American Economic Review,* Vol. 90, 980–994.

Fehr, Ernst; Goette, Lorenz and Christian Zehnder. 2009. "A Behavioral Account of the Labor Market: The Role of Fairness Concerns." *Annual Review of Economics,* Vol. 1, 355–384.

Fehr, Ernst; Kirchsteiger, Georg and Arno Riedl. 1993. "Does Fairness Prevent Market Clearing? An Experimental Investigation." *The Quarterly Journal of Economics,* Vol. 108, 437–459.

Fehr, Ernst; Klein, Alexander and Klaus M Schmidt. 2007. "Fairness and Contract Design." *Econometrica,* Vol. 75, 121–154.

Fehr, Ernst and Klaus M Schmidt. 1999. "A Theory of Fairness, Competition and Cooperation." *Quarterly Journal of Economics,* Vol. 114, 817–868.

Fehr, Ernst and Klaus M Schmidt. 2003. "Theories of Fairness and Reciprocity – Evidence and Economic Applications." In: Dewatripont, M. et al (Hg.): *Advances in Economics*

and Econometrics. Eighth World Congress of the Econometric Society, Vol. 1. Cambridge: Cambridge University Press, 208–257.

Fehr, Ernst and Klaus M Schmidt. 2004. "Fairness and Incentives in a Multi-Task Principal-Agent Model." *Scandinavian Journal of Economics,* Vol. 106, 453–474.

Fehr, Ernst and Klaus M Schmidt. 2006. "The Economics of Fairness, Reciprocity and Altruism - Experimental Evidence and New Theories." In: Kolm, Serge-Christophe; Ythier, Jean Mercier (Hg.): *Handbook on the Economics of Giving, Reciprocity and Altruism, Vol. 1.* Amsterdam: Elsevier, 615–691.

Forsythe, Robert; Horowitz, Joel L; Savin, N. E and Martin Sefton. 1994. "Fairness in Simple Bargaining Experiments." *Games and Economic Behavior,* Vol. 6, 347–369.

Gaechter, Simon and Armin Falk. 2002. "Reputation and Reciprocity: Consequences for the Labour Relation." *The Scandinavian Journal of Economics,* Vol. 104, 1–26.

Geanakoplos, John; Pearce, David and Ennio Stacchetti. 1989. "Psychological games and sequential rationality." *Games and Economic Behavior,* Vol. 1, 60–79.

Gebhardt, Georg. 2002. "Venture capital, Takeovers and Time-varying Returns. Three essays in financial economics. Dissertation." Berlin: dissertation.de (Dissertation.de, 632).

Gebhardt, Georg. 2004. "Inequity Aversion, Financial Markets, and Output Fluctuations." *Journal of the European Economic Association,* Vol. 2, 229–239.

Gneezy, Uri and John List. 2006. "Putting Behavioral Economics to Work: Testing for Gift Exchange in Labor Markets Using Field Experiments." *Econometrica,* Vol. 74, 1365–1384.

Güth, Werner; Marchand, Nadege and Jean-Louis Rulliere. 1998. "Equilibration et dependance du contexte. Une evaluation experimentale du jeu de negociation sous ultimatum." *Revue Economique,* Vol. 49, 785–794.

Güth, Werner; Schmittberger, Rolf and Bernd Schwarze. 1982. "An experimental analysis of ultimatum bargaining." *Journal of Economic Behavior & Organization,* Vol. 3, 367–388.

Hannan, R. L; Kagel, John H and Donald V Moser. Oct., 2002. "Partial Gift Exchange in an Experimental Labor Market: Impact of Subject Population Differences, Productivity Differences, and Effort Requests on Behavior." *Journal of Labor Economics,* Vol. 20, 923–951.

Hart, Oliver. 1995. "Firms, contracts, and financial structure." Oxford: Clarendon Press (Clarendon Lectures in Economics).

Henrich, J.; Boyd, R.; Bowles, S.; Camerer, C., et al. (eds.). 2004. "Foundations of Human Sociality." Oxford: Oxford University Press.

Henrich, Joseph and N. Smith. 2004. "Comparative experimental evidence from Machiguenga, Mapuche, Huinca, and American Populations." In: Henrich, Joseph; Boyd, Robert; Bowles, Samuel; Camerer, Colin; Fehr, Ernst; Gintis, Herbert; McElreath, Richard (Hg.): *Foundations of Human Sociality.* Oxford: Oxford University Press .

Holmstrom, Bengt. 1982. "Moral Hazard in Teams." *The Bell Journal of Economics,* Vol. 13, 324–340.

Holmstrom, Bengt and Paul Milgrom. 1991. "Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design." *Journal of Law, Economics, & Organization,* Vol. 7, 24–52.

Holt, Charles A and Loren W Langan. 1986. "Market Power in Oral Double Auctions." *Economic Inquiry,* Vol. 24, 107–123.

Kirchsteiger, Georg. 1994. "The role of envy in ultimatum games." *Journal of Economic Behavior & Organization,* Vol. 25, 373–389.

Kreps, David M; Milgrom, Paul; Roberts, John and Robert Wilson. 1982. "Rational cooperation in the finitely repeated prisoners' dilemma." *Journal of Economic Theory,* Vol. 27, 245–252.

Kube, Sebastian; Maréchal, Michel A and Clemens Puppe. 2006. "Putting Reciprocity to Work - Positive Versus Negative Responses in the Field." *SSRN eLibrary*.

Kube, Sebastian; Maréchal, Michel A and Clemens Puppe. 2008. "The Currency of Reciprocity - Gift-Exchange in the Workplace." *SSRN eLibrary*.

Ledyard, John. 1995. "Public Goods: A Survey of Experimental Research." In: Kagel, John Henry; Roth, Alvin E. (Hg.): *The handbook of experimental economics.* Princeton, NJ: Princeton Univ. Press, 111–194.

Levine, David K. 1998. "Modeling Altruism and Spitefulness in Experiments." *Review of Economic Dynamics,* Vol. 1, 593–622.

MacLeod, Bentley and Daniel Parent. 1999. "Job Characteristics and the Form of Compensation Form." *Research in Labor Economics,* Vol. 18, 177–242.

MacLeod, W. B and James M Malcomson. Mar., 1989. "Implicit Contracts, Incentive Compatibility, and Involuntary Unemployment." *Econometrica,* Vol. 57, 447–480.

Oreopoulos, Philip; Wachter, Till von and Andrew Heisz "The Short- and Long-Term Career Effects of Graduating in a Recession: Hysteresis and Heterogeneity in the Market for College Graduates." *SSRN eLibrary*.

Quervain, Dominique J F de; Fischbacher, Urs; Treyer, Valerie; Schellhammer, Melanie; Schnyder, Ulrich; Buck, Alfred and Ernst Fehr. 2004. "The Neural Basis of Altruistic Punishment." *Science,* Vol. 305, 1254–1258.

Rabin, Matthew. 1993. "Incorporating Fairness into Game Theory and Economics." *The American Economic Review,* Vol. 83, 1281–1302.

Roth, Alvin E. 1995. "Bargaining Experiments." In: Kagel, J.; Roth, Alvin E. (Hg.): *Handbook of Experimental Economics.* Princeton, N.J.: Princeton University Press .

Roth, Alvin E; Prasnikar, Vesna; Okuno-Fujiwara, Masahiro and Shmuel Zamir. 1991. "Bargaining and Market Behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An Experimental Study." *The American Economic Review,* Vol. 81, 1068–1095.

Smith, Vernon L. 1962. "An Experimental Study of Competitive Market Behavior." *The Journal of Political Economy,* Vol. 70, 111–137.

Smith, Vernon L. 1964. "Effect of Market Organization on Competitive Equilibrium." *The Quarterly Journal of Economics,* Vol. 78, 182–201.

Teulings, Coen N. 2003. "The contribution of minimum wages to increasing wage inequality." *The Economic Journal,* Vol. 113, 801–833.