FinanzArchiv

Public Finance Analysis

1
Volume 66
March
2010

Mohr Siebeck

Editorial Note

How does behavioral and experimental economics transform the way we think about public economics? How do notions of social preferences and bounded rationality shape analysis in public economics? These were the guiding questions in our call for papers for a workshop at the Venice Summer Institute 2009, which was jointly organized by CESifo Institute and FinanzArchiv.

Among the 48 submissions we got, 17 were selected for presentation at our two-day workshop. Among the topics covered, to name just a few, were the individual perception of tax burdens and tax incentives, the behavioral nonequivalence of theoretically equivalent taxes, the evaluation of public goods and the design of voting systems for public-goods provision, new views on the concepts of merit wants and social preferences, the consumption behavior of policymakers, and the voting behavior of ordinary citizens. Two keynote lectures, by Arno Riedl and Joel Slemrod, complemented the presentations. Together, the presentations reflect the broad scope of fruitful applications of behavioral and experimental economics to important issues in public economics.

All presentations are available as CESifo working papers or will be available shortly. Three of the manuscripts, which were submitted to FinanzArchiv, have been accepted for publication and are published in this issue. We hope that the papers will stimulate the discussion on experimental and behavioral public finance, and we welcome future submissions of excellent research papers in this field to FinanzArchiv.

Bernd Genser
(Editor)

Jean-Robert Tyran
(Guest Editor)

FinanzArchiv 66 (2010).
ISSN 0015-2218

doi: 10.1628/001522110X508131

Tax Evasion, Investment, and Firm Activity

Florian Baumann and Tim Friehe*

Received 18 June 2009; in revised form 24 February 2010; accepted 4 March 2010

This paper establishes within an intertemporal framework that firm activity is influenced by tax evasion if firms can invest in their long-term competitiveness. Higher investment raises the firm's survival probability, which in turn reduces tax-evasion payoffs, since evaded taxes may also be detected and penalized later on. At the same time, tax evasion in future periods increases future expected profits, making higher investment desirable. Consequently, tax evasion distorts investment, while the latter determines firm activity. Investment in the event of tax evasion may be higher or lower, depending on parameters of tax enforcement.

Keywords: tax evasion, activity level, investment

JEL classification: H 26

1. Introduction

In many circumstances, the activity level of a firm is separable from the firm's decision on whether to evade taxes. This finding is of interest because it implies, for instance, that the neutrality of profit taxes is upheld, even if taxes can be evaded.

In this paper it will be shown that a firm's decisions on activity and tax evasion can no longer be viewed separately if the firm can determine long-term investment, and thereby firm size, in order to enhance competitiveness. We use an intertemporal framework in which investment affects the probability that the firm can continue operating in the next period and increases the legitimate net profit per period. The assumption that investment increases the survival probability of firms is in line with the empirical fact that exiting firms are on average smaller than firms that remain in a given industry (Plehn-Dujowich 2009). The activity level in a given period depends on investment. For example, if the investment reduces marginal costs, this will tend to increase the level of firm activity. Furthermore, our framework is characterized by the fact that unpaid tax liabilities accumulate over time if

* We are indebted to Laszlo Goerke, Alfons Weichenrieder, and two anonymous referees for very helpful suggestions. Florian Baumann gratefully acknowledges financial support by the German Research Foundation (DFG).

ISSN 0015-2218
doi: 10.1628/001522110X503352

they remain undetected. The fact that such a link between reporting in the past and detection exists is asserted, for instance, by Engel and Hines (1999). The analysis of this dynamic setting yields our central result, namely that a firm's activity levels are affected by the firm's decision on tax evasion. The intuition for our result runs as follows. The activity level in a given period is dependent on the level of investment. The level of investment in turn is dependent on tax evasion by the firm. The latter relationship is due to, first, the fact that tax evasion increases future payoffs, arguing for higher investment to increase the survival probability, and second, the fact that a statement deviation in a prior period may be detected later on, arguing for a decrease in investment.

In addition to establishing our central result, we also detail how the possibility of evading taxes changes investment incentives. We find that, depending on the circumstances, the optimal level of investment, given the possibility to evade taxes, may be higher or lower than the optimal level of investment without that possibility. Given the possibility of tax evasion and the profitability thereof, we find that comparative statics for the policy variables tax rate, detection probability, and penalty rate are ambiguous at the general level. To illustrate our findings, we supply a numerical illustration.

There is a significant body of literature to which the present contribution can be related. In an early contribution, Kreutzer and Lee (1986) dismissed the neutrality of a profit tax on output if monopolists can reduce their tax liability by overstating production costs, as this would induce them to increase output. Wang and Conant (1988), building on Kreutzer and Lee (1986), introduced a probability of detection and a penalty proportional to the taxes evaded. Allowing firms to choose the level of cost overstatement, they reestablished that firms' optimal output levels will be independent of the cost-overstatement decision.[1] Wang (1990) generalized the model suggested by Wang and Conant (1988), allowing for probability of detection and penalty rate as functions of cost overstatement.[2] He argued that the neutrality of the profit tax hinges on whether the probability of detection and penalty vary with cost overstatement. However, Yaniv (1996), using the same framework, established that this argumentation is unfounded, i.e., profit taxes remain neutral irrespective of the monopolist's output decision. Lee (1998) revisited the aspect of endogenous detection probability and penalty already touched upon by Wang (1990) and Yaniv (1996). In contrast to the latter contributions, which assumed that the cost overstatement determines the probability and penalty, Lee considered various determinants for the audit probability

1 Kreutzer and Lee (1988) later commented on Wang and Conant's (1988) variation of their hypothesis.

2 See also Marrelli (1984) and Marrelli and Martina (1988) for the implications of assumptions with respect to the detection probability in the context of indirect taxes.

and the penalty rate. Allowing for tax evasion, profit taxes are neutral if, for example, the probability of detection and the penalty are functions of reported profit. Yaniv (1995) established a generic model in which firm activity is always separable from the decision to evade taxes.[3]

Two more recent contributions, however, deviate somewhat from the papers discussed so far. The approach taken by Panteghini (2000) bears some similarity to that of our paper. He too analyzes tax evasion and investment, but he develops a different argument, using real option theory to establish that the decision regarding whether to invest today or in the future is affected by the possibility to evade taxes. From this observation follows the result that tax evasion influences firm activity. In contrast, we do not focus on the point in time at which investment occurs, but rather on the way in which the possibility of evading taxes changes the optimal extent of investment. Another branch of the literature on tax avoidance and evasion discusses the possibility of affecting the tax burden of multinational companies by transfer pricing, and the potential repercussions thereof for investment. Weichenrieder (2009) presents evidence for the existence of profit shifting via transfer pricing as a function of tax differentials. Such opportunities to shift profits affect the profitability of foreign direct investment and, as a result, the scale of foreign direct investment (see, e.g., Amerighi 2008, Azémar and Corcos 2009, Weichenrieder 1996, and Gresik 2001 for a survey).

Recently, Goerke and Runkel (2006) have shown in a Cournot oligopoly setting that output and evasion decisions will not be independent from one another if the number of firms is endogenous. The channel through which tax evasion influences output is an indirect one, as a firm's optimal output does not depend on evasion. Yet, since evasion increases expected profits, more firms enter the market, with consequences for individual and aggregate output. In our paper, it is similarly assumed that tax evasion influences activity only indirectly, but via its effect on investment. To the best of our knowledge, this central aspect has not yet been discussed in the literature.[4]

The structure of this paper is as follows. Section 2 describes the model, section 3 presents the analysis, and section 4 concludes.

3 This will always hold true if the firm decides on the amount evaded; if the firm decides on the fraction, however, the separability result requires an interior optimum for the evasion decision.

4 The model is designed to deliberately exclude other aspects that might make the activity and the evasion decision nonseparable.

2. The Model

The model follows that established by Yaniv (1995), albeit extended to a two-period framework.[5] This model is specific neither to a certain tax base nor to the structure of the market in which the firm is active. The firm faces a proportional tax rate θ on its tax base, and decides on both the activity level A_t and the statement deviation from the true value S_t in period t, where $t = 1, 2$. The statement deviation causes concealment costs $g(S_t)$, $g' > 0$, $g'' > 0$, and is restricted by the true value of the tax base, $0 \leq S_t \leq b_t(A_t)$.[6] The firm is audited with exogenous probability p in each period. In the event of an audit, the firm will have to pay evaded taxes and a penalty which, following Yitzakhi (1974), is a multiple of evaded taxes according to the penalty rate $\tau - 1 > 0$. In our dynamic setting, total evaded taxes in period 2 comprise taxes withheld in that period and undetected evaded taxes from period 1, where the latter are inflated using the discount factor δ, $0 < \delta < 1$, and discounted by the rate of oblivion γ, $0 \leq \gamma \leq 1$. Whilst the discount rate takes account of accrued interest, the rate of oblivion allows for the possibility that law enforcement may have greater difficulty in proving past offenses completely. The firm decides on investment I at the beginning of period 1. Investment increases legitimate net profits in period t, $\pi_t(A_t, \theta, I)$, at a diminishing rate, $\frac{\partial \pi_t}{\partial I} > 0 > \frac{\partial^2 \pi_t}{\partial I^2}$, given that the kind of investment considered improves the firm's competitiveness. Furthermore, investment also improves the long-term competitiveness of the firm, represented by the probability $q(I)$, $q' > 0 > q''$, that the firm will still be active in period 2. One might think of investment that improves production processes, and thereby decreases marginal costs. This not only increases profits in a given period, but also makes it more likely that the firm can stay active in the market should demand drop. Schmidt (1997), for instance, assumes that the investment of managerial effort increases the probability that a cost reduction is successful, which increases expected profit per period but also decreases the probability that the firm will be liquidated. With regard to net profits, we further assume that (i) an increase in investment makes an increase in the activity level more profitable, $\frac{\partial^2 \pi_t}{\partial A_t \partial I} > 0$, (ii) an increase in the tax rate does not increase net profits, $\frac{\partial \pi_t}{\partial \theta} \leq 0$, and (iii) net profits are concave in the level of activity, $\frac{\partial^2 \pi_t}{\partial A_t^2} < 0$. The first assumption will be

5 The focus on two periods is for expositional simplicity only. The results obtained can be generalized to a multiperiod setting.

6 The introduction of concealment costs is analogous to the treatment in, e.g., Virmani (1989) or Cremer and Gahvari (1992). We make these assumptions with respect to concealment costs because it is our objective to highlight one argument against the neutrality of tax evasion for activity choice. This neutrality would already not hold if concealment costs were a function of the firm's activity level. Note that the specification we use has the implication that concealment costs do not depend on firm size.

satisfied for investment that decreases marginal costs, irrespective of whether firms set prices or quantities. The second assumption is natural; the third is applied in order to ensure second-order conditions for the firm's choice of activity level. An example satisfying these properties would be the net profits of a monopolist who faces a demand function with a constant price elasticity equal to $1/\epsilon$, $0 > \epsilon > -1$, and whose marginal costs of production, $c(I) > 0$, decrease at a diminishing rate with increasing initial investment level: $c'(I) < 0 < c''(I)$. With θ being the profit tax rate, net profits in period t for given investment I amount to $\pi_t = (1-\theta)\left(A_t^{1+\epsilon} - c(I)A_t\right)$. Firms maximize the present value of net profits less after-tax investment costs.

3. The Analysis

In the following, we solve the firm's decision problem with regard to the activity levels, statement deviations, and initial investment. The time structure of the analysis is as follows. At the beginning of the first period, the firm decides on the investment level. Next, it chooses the period 1 activity level and statement deviation. Tax authorities may audit the firm at the end of period 1. Firms again determine the activity level and statement deviation in period 2 and run the risk of being audited by tax authorities. The investigation starts with the latest decision made by the firm, i.e., the decisions on the activity level and statement deviation in period 2.

Period 2 expected profits can be stated as

$$\Pi_2(A_2, S_2) = \pi_2(A_2, \theta, I) + \theta S_2 - g(S_2) - p\theta\tau(S_2 + B_1)\,. \tag{1}$$

The term B_1 allows for the fact that the statement deviation of period 1 may still be relevant in period 2. B_1 may take one of two values:

$$B_1 = \begin{cases} 0 & \text{if the firm has been audited in period 1}\,, \\ \delta^{-1}(1-\gamma)S_1 & \text{otherwise}\,. \end{cases} \tag{2}$$

This expresses the idea that the total amount of statement deviation taken into consideration by tax authorities in period 2 differs from the actual statement deviation in period 2 only if the firm was not audited in period 1. The difference is made up of the statement deviation in period 1 adjusted according to the rate of oblivion γ and the discount rate δ, where the latter adjustment corrects for interest accrued by the firm.

Solving the maximization problem posed by (1) and assuming interior solutions, we obtain

$$\frac{\partial \pi_2}{\partial A_2} = 0 \tag{3}$$

and

$$\theta(1 - p\tau) = g'(S_2^*) \tag{4}$$

as conditions that characterize optimal activity and statement deviation in period 2. Consequently, we obtain the result that the activity choice is independent of the firm's attempt to evade taxes, i.e., activity is chosen so that legitimate net profits in period 2 are maximized. The optimal statement deviation in period 2 is given by the equalization of marginal net profits from evading and marginal concealment costs, and is independent of the activity level.

The novel feature of this analysis lies in taking an intertemporal setting into account, in which periods are linked by an initial investment. We will now consider how this investment affects choices in period 2. We find that activity in period 2 increases with investment, while tax evasion is independent of the investment level: $\frac{dA_2^*}{dI} = -\frac{\partial^2 \pi_2/(\partial A_2 \partial I)}{\partial^2 \pi_2/\partial A_2^2} > 0 = \frac{dS_2^*}{dI}$.

Lemma 1 The activity choice in period 2 is not directly dependent on the extent of tax evasion, but is a function of the investment level.

Proof. This statement follows from (3). ■

Next, we move back in time and analyze activity and evasion decisions in period 1. After investment has been undertaken, the present value of expected profits can be stated as

$$\begin{aligned} \Pi_1(A_1, S_1) = {} & \pi_1(A_1, \theta, I) + \theta S_1 - g(S_1) - p\tau\theta S_1 \\ & + q(I)\delta\big[\pi_2(A_2^*, \theta, I) + \theta S_2^* - g(S_2^*) \\ & - p\theta\tau(S_2^* + (1-p)\delta^{-1}(1-\gamma)S_1)\big], \end{aligned} \tag{5}$$

which reflects the fact that tax evasion in period 1 may also be detected in period 2 if the firm is not audited in period 1. Equation (5) also implies that if the firm is no longer active in period 2, no audit will take place.[7]

The activity and the evasion levels that maximize expected profits are characterized by

$$\frac{\partial \pi_1}{\partial A_1} = 0 \tag{6}$$

and

$$\theta(1 - p\tau(1 + q(I)(1-p)(1-\gamma)) = g'(S_1^*)\,. \tag{7}$$

The condition (7) brings to the fore that the multiperiod context implies lower net benefits of tax evasion, due to the fact that tax authorities may call the firm to account even after the passing of the current period.

When we turn to the effect that investment has on activity and evasion, we find that $\frac{dA_1^*}{dI} > 0 > \frac{dS_1^*}{dI}$. With regard to the latter, it should be noted that higher

7 Alternatively, without altering our main result, it could be assumed that the audit probability is lower for firms that have left the market or that the rate of oblivion is higher for these firms, as it becomes more difficult for the tax authority to verify past tax evasion.

investment increases the survival probability. This implies that the compound probability that tax evasion undertaken in period 1 will be detected is higher, which reduces incentives for evasion.

Lemma 2 The activity choice in period 1 is not directly dependent on the extent of tax evasion, but is a function of the investment level.

Proof. This statement follows from (6). ■

The level of investment is endogenously determined at the beginning of the first period. Taking into account the expected payoffs in periods 1 and 2, the firm considers $\Pi(I) = \Pi_1(I) - I(1-\tau)$, which can be stated more extensively as

$$\begin{aligned}\Pi(I) = {} & \pi_1(A_1^*, \theta, I) + \theta S_1^* - g(S_1^*) - p\tau\theta S_1^* - I(1-\theta) \\ & + q(I)\delta\big[\pi_2(A_2^*, \theta, I) + \theta S_2^* - g(S_2^*) \\ & - p\tau\theta(S_2^* + (1-p)\delta^{-1}(1-\gamma)S_1^*\big]. \end{aligned} \tag{8}$$

Here we make use of the functional dependences established above and assume that investment expenses are immediately tax-deductible.

The optimal investment I^* fulfils the following condition, which was derived by simplifying the first-order condition using the envelope theorem:

$$\begin{aligned}\frac{d\Pi}{dI} = {} & \underbrace{\frac{\partial \pi_1}{\partial I} + q(I^*)\delta\frac{\partial \pi_2}{\partial I} + q'(I^*)\delta\pi_2(A_2^*, \theta, I^*)}_{\text{A}} - \underbrace{(1-\theta)}_{\text{B}} \\ & \underbrace{- q'(I^*)(1-p)p\tau\theta(1-\gamma)S_1^*}_{\text{C}} \\ & \underbrace{+ q'(I^*)\delta(\theta S_2^* - g(S_2^*) - p\tau\theta S_2^*)}_{\text{D}} = 0\,. \end{aligned} \tag{9}$$

This condition may be contrasted with that obtained in the absence of tax evasion, given by

$$\frac{d\widetilde{\Pi}}{dI} = \frac{\partial \pi_1}{\partial I} + q(I)\delta\frac{\partial \pi_2}{\partial I} + q'(I)\delta\pi_2(A_2^*, \theta, I) - (1-\theta) = 0\,. \tag{10}$$

It can easily be observed that the expression labeled A, which describes the effects of an increase in investment on the legitimate net profits in respective periods and the survival probability, is common to both conditions. Similarly, the second term, B, which gives marginal net investment costs, is present both in the condition with and that without tax evasion.

However, the condition given in (9) deviates from the one given in (10) in that it contains terms C and D. These terms thus capture the effects tax evasion bears on the investment decision. Two effects can be distinguished, both arising from the fact that more investment raises the probability that the

firm will still be active in period 2: (i) A higher survival probability implies that taxes evaded in period 1 will be detected more often. This decreases the expected benefits of tax evasion in period 1 according to term C and therefore argues against higher investment. (ii) A higher survival probability also implies that the gain due to tax evasion in period 2, represented by term D, becomes more likely. The expected profit in period 2 is higher due to tax evasion, which makes reaching period 2 more attractive for the firm, i.e., this term argues for more investment.

In summary, tax evasion distorts the investment choice by introducing additional marginal benefits and costs.

The above allows the following result to be deduced:

Lemma 3 The optimal investment of the firm depends on statement deviations in period 1 and period 2.

Proof. This statement follows from (9). ■

Given the fact that the investment choice is a function of the firm's decision on tax evasion, we can establish that the activity choice is likewise nonseparable from the choice regarding tax evasion.

Proposition 1 The firm's activity choice is a function of the extent of tax evasion.

Proof. This statement is a direct consequence of Lemmas 1–3. ■

The above establishes our central result: that the possibility of evading taxes influences a firm's activity choice. We have shown that tax evasion influences investment, while investment influences activity choice. In the following, we elaborate on the way in which tax evasion influences investment and therefore turn to comparative statics of the optimal level of investment I^*. We are concerned with the policy variables tax rate θ, detection probability p, and penalty rate $\tau - 1$. We first turn to the tax rate and obtain

$$\frac{dI^*}{d\theta} = -\frac{\frac{\partial^2 \pi_1}{\partial I \partial \theta} + q(I^*)\delta \frac{\partial^2 \pi_2}{\partial I \partial \theta} + q'(I^*)\delta \frac{\partial \pi_2}{\partial \theta} + 1}{\partial^2 \Pi / \partial I^2} + \frac{q'(I^*)\left[(1-p)p\tau(1-\gamma)\left\{S_1^* + \theta \frac{\partial S_1^*}{\partial \theta}\right\} - \delta S_2^*(1-p\tau)\right]}{\partial^2 \Pi / \partial I^2}, \tag{11}$$

where $\frac{\partial S_1^*}{\partial \theta} > 0$ follows from (7). The change in net profits from an increase in investment is usually smaller, the higher the tax rate, i.e., $\frac{\partial^2 \pi_t}{\partial I \partial \theta} < 0$. The first line gives the terms that also occur in the absence of the possibility of evading taxes, whereas the terms in the second line are additional effects due to tax evasion. In general, we find that the sign of the $\frac{dI^*}{d\theta}$ cannot be determined unambiguously. Referring to (9), a higher tax rate increases term C, which represents the effect of the expected penalty for tax evasion in period 1

detected in period 2. Likewise, a higher tax rate increases term D, due to higher profits resulting from tax evasion in period 2.

Next, we turn to the detection probability and find

$$\frac{dI^*}{dp} = -\frac{q'(I^*)\tau\theta}{\partial^2\Pi/\partial I^2}\left[-(1-p)(1-\gamma)\left\{S_1^* + p\frac{\partial S_1^*}{\partial p}\right\} + p(1-\gamma)S_1^* - \delta S_2^*\right], \tag{12}$$

where $\frac{\partial S_1^*}{\partial p} < 0$ follows from (7). Referring to (9), the first two terms in the brackets in (12) indicate the change in term C, i.e., the variation in the expected fine for tax evasion in period 1 detected in period 2. *Inter alia*, the effect depends on the sign of $S_1^* + \theta\frac{\partial S_1^*}{\partial p}$, which itself cannot be stated. The second term stems from the fact that the probability of tax evasion being detected in period 1 increases, thereby reducing the possibility of first-period tax evasion being detected belatedly. Finally, the third term indicates lower profits out of tax evasion in period 2. All in all, the sign of $\frac{dI^*}{dp}$ cannot be determined unambiguously.

As a next step, we inquire into the change of optimal investment when the penalty rate is varied. We obtain

$$\frac{dI^*}{d\tau} = -\frac{q'(I^*)p\theta}{\partial^2\Pi/\partial I^2}\left[-(1-p)(1-\gamma)\left\{S_1^* + \tau\frac{\partial S_1^*}{\partial \tau}\right\} - \delta S_2^*\right], \tag{13}$$

where $\frac{\partial S_1^*}{\partial \tau} < 0$ follows from (7). The reasoning with regard to the respective terms and the sign of $\frac{dI^*}{d\tau}$ is similar to that for $\frac{dI^*}{dp}$, as both τ and p raise the expected fine. However, with respect to the detection probability, there is also the effect of a lower probability of belated detection (see the second term in the brackets in (12)).

In summary, we find that comparative statics of the optimal level of investment are ambiguous. In order to illustrate the possible effects more concretely, we will analyze a numerical example. We specify (i) $q(I) = \frac{I}{1+I}$ with $q' > 0 > q''$, (ii) $g(S_t) = S_t^2$ with $g', g'' > 0$, and (iii) $\pi_t(A_t, \theta, I) = (1-\theta)(A_t^{1+\epsilon} - \frac{c}{V^{1/2}}A_t)$ with $\frac{\partial \pi_t}{\partial I} > 0$, $\frac{\partial^2 \pi_t}{\partial A_t \partial I} > 0$, $\frac{\partial^2 \pi_t}{\partial A_t^2} < 0$, $\frac{\partial \pi_t}{\partial \theta} < 0$, $\frac{\partial^2 \pi_t}{\partial I^2} < 0$. Given these functional relationships, we arrive at

$$A_t^* = \left(\frac{\sqrt{I}}{c}(1+\epsilon)\right)^{-1/\epsilon}, \tag{14}$$

$$S_1^* = \frac{\theta\left(1 - p\tau(1 + \frac{I}{1+I}(1-p)(1-\gamma))\right)}{2}, \tag{15}$$

$$S_2^* = \frac{\theta\left(1-p\tau\right)}{2}. \tag{16}$$

First, we calculate the level of investment that results without the possibility of evading taxes, by fixing $c = 0.3$, $\epsilon = -0.5$, and $\delta = 0.9$. Note that the optimal investment in the evasion-free benchmark is independent of the

Figure 1
Investment as a Function of the Tax Rate

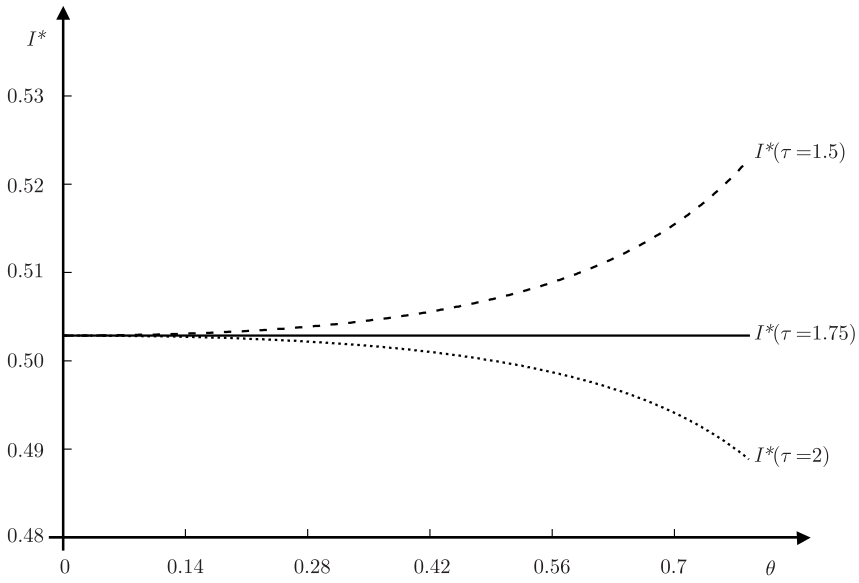

tax rate, given our specifications. This optimal level of investment is $I_B^* = 0.503$. Next, we create a configuration of parameters describing the tax system including enforcement that yield this same level of investment. This configuration is $\theta = 0.4$, $\tau = 1.75$, $p = 0.252$, and $\gamma = 0.3$. The tax rate θ roughly corresponds to the combined corporate income tax rate in the United States in 2009 according to the OECD Tax Database.[8] The level of τ is in line with figures presented by Andreoni et al. (1998). Hindriks and Myles (2006, p. 526) agree that a figure for τ on the order of 1.5 to 2 is realistic. With respect to the detection probability p, note that firms are normally audited more often than individuals. For large firms, the audit probability in the United States was almost 30% in 2003 (Slemrod 2004). At these parameter levels, we obtain tax evasion in both periods. To pursue our interest in the effects of the policy variables θ, p, and τ, we subsequently vary their levels, keeping all else equal.

Figure 1 illustrates the variation of the optimal level of investment for different tax rates and three values for the penalty rate $\tau - 1$. Remember that I_B^* is independent of the tax rate. As tax evasion becomes irrelevant for $\theta = 0$, all three curves in figure 1 originate from the same locus on the ordi-

8 See the section "Taxation of Corporate and Capital Income," available at www.oecd.org/ctp/taxdatabase.

Figure 2
Investment as a Function of the Penalty Rate

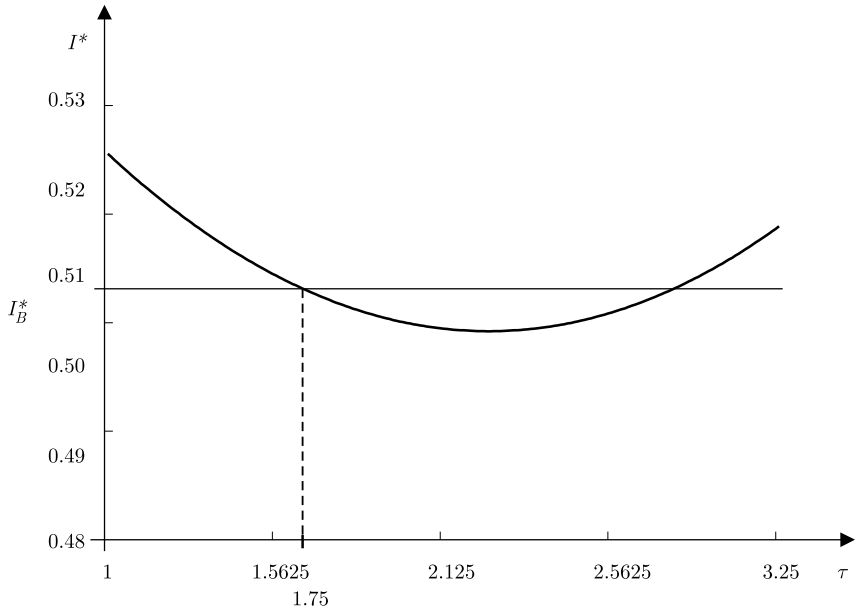

nate, viz., I_B^*. For the benchmark configuration of parameters with $\tau = 1.75$, the investment level is independent of the tax rate even though tax evasion occurs. However, as the curve for $\tau = 1.5$ ($\tau = 2$) indicates, investment can also increase (decrease) with increasing tax rate. A higher tax rate raises profits from tax evasion in period 2, which argues for higher investment. However, the higher tax rate also increases the expected fine resulting from undetected tax evasion committed in period 1, arguing for lower investment. In our specification, optimal investment is a monotonous function of the tax rate, where the slope depends on parameter values.

Figure 2 depicts optimal investment as a function of the penalty rate $\tau - 1$. Our simulation shows that the relation between investment and the penalty rate is nonmonotonic, with investment being higher than the benchmark level I_B^* for low and high levels of the penalty rate. Note that for the range of penalty rates used in figure 2, statement deviations are positive in both periods.[9] A higher penalty rate discourages investment, as the profits from tax evasion in the second period are reduced. However, this could be com-

9 For very high levels of the penalty rate, given a specific level of the detection probability, tax evasion is no longer profitable. Therefore, optimal investment will approach the benchmark level I_B^* for high penalty rates.

pensated if a higher penalty rate causes lower tax evasion in period 1, thereby lessening the (potential) threat of the fine imposed in case of detection in period 2.

Finally, in figure 3 we depict how optimal investment is affected by the detection probability p. As outlined above, the effect of the detection probability bears some similarity to that of the penalty rate, in that both increase the expected fine. However, varying the detection probability also entails the additional effect due to a higher detection probability in period 1 making the event of being not detected in period 1 but detected in period 2 less likely. This latter fact encourages investment. We find that investment first decreases with increasing detection probability but increases for larger values of this probability, as was the case for the penalty rate. For $p = 0$, only the effect of higher profits in period 2 exists, so that investment surpasses the benchmark value without tax evasion. In our simulation, the investment level then falls below this benchmark value but eventually exceeds it again.

All in all, our numerical results show that (i) tax evasion can provoke either a higher or a lower level of investment than results without the possibility of evading taxes, and (ii) the deviation in the optimal investment level responds to changes in policy parameters. For all the policy parameters considered, the optimal investment can be higher or lower than the benchmark value

Figure 3
Investment as a Function of the Detection Probability

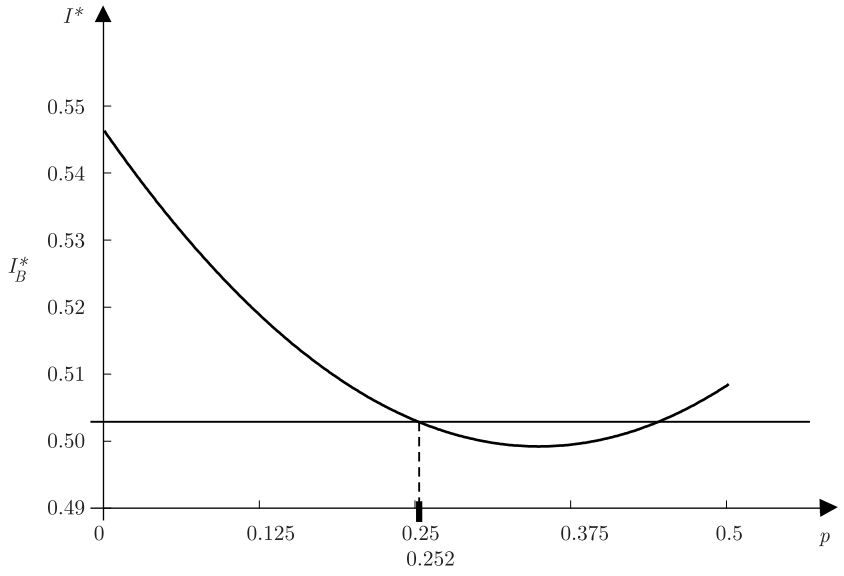

without tax evasion. The numerical illustration often suggests nonmonotonic relations between policy parameters and optimal investment.

4. Conclusion

This paper, using an intertemporal framework, proves that a firm's activity choice is no longer separable from tax evasion if investment by the firm is taken into account. The investment considered improves the firm's competitiveness and thereby links different periods. The firm's activity levels are determined by investment, which in turn is affected by tax evasion considerations. Consequently, we establish an indirect link between tax evasion and the firm's activity levels and thereby contribute to the literature that casts doubt on the separability of activity levels and tax evasion. Moreover, we establish that the possibility of evading taxes implies no clear effect on the optimal level of investment.

Our analysis bears the following policy implications. Most importantly, we establish that the policy on tax evasion has decisive repercussions on a firm's activity and investment decisions. Consequently, deterrence of tax evasion has consequences that are beyond its direct effect on tax revenues. We highlight that harsher tax enforcement might have the unintended consequence of reducing firm investment. From this follows, for example, a link between tax policy and employment at the firm level. Furthermore, we find that privately optimal tax evasion is dependent on firm size. This dependence may allow enforcement resources to be employed more effectively.

References

Amerighi, O. (2008), Transfer Pricing and Enforcement Policy in Oligopolistic Markets, in: Brakman, S. (ed.) Foreign Direct Investment and the Multinational Enterprise, MIT Press, Mass, Cambridge, 117–154.

Andreoni, J., Erard, B., and J. Feinstein (1998), Tax Compliance, Journal of Economic Literature 36, 818–860.

Azémar, C., and G. Corcos (2009), Multinational Firms' Heterogeneity in Tax Responsiveness: The Role of Transfer Pricing, World Economy 32, 1291–1318.

Cremer, H., and F. Gahvari (1992), Tax Evasion and the Structure of Indirect Taxes and Audit Probabilities, Public Finance/Finance Publiques 47, 351–365.

Engel, E.M.R., and J.R. Hines (1999), Understanding Tax Evasion Dynamics, NBER Working Paper No. 6903.

Gresik, T.A. (2001), The Taxing Task of Taxing Transnationals, Journal of Economic Literature 39, 800–838.

Goerke, L., and M. Runkel (2006), Profit Tax Evasion under Oligopoly with Endogenous Market Structure, National Tax Journal 59, 851–857.

Hindriks, J., and G.D. Myles (2006), Intermediate Public Economics, MIT Press, Mass, Cambridge.

Kreutzer, D., and D.R. Lee (1986), On Taxation and Understated Monopoly Profits, National Tax Journal 39, 241–243.

Kreutzer, D., and D.R. Lee (1988), Tax Evasion and Monopoly Output Decisions: A Reply, National Tax Journal 41, 583–584.

Lee, K. (1998), Tax Evasion, Monopoly, and Nonneutral Profit Taxes, National Tax Journal 51, 333–338.

Marrelli, M. (1984), On Indirect Tax Evasion, Journal of Public Economics 25, 181–196.

Marrelli, M., and R. Martina (1988), Tax Evasion and Strategic Behaviour of the Firms, Journal of Public Economics 37, 55–69.

Panteghini, P.M. (2000), Tax Evasion and Entrepreneurial Flexibility, Public Finance Review 28, 199–209.

Plehn-Dujowich, J.M. (2009), Entry and Exit by New versus Existing Firms, International Journal of Industrial Organization 27, 214–222.

Schmidt, K.M. (1997), Managerial Incentives and Product Market Competition, Review of Economic Studies 64, 191–213.

Slemrod, J. (2004), The Economics of Corporate Tax Selfishness, National Tax Journal 57, 877–899.

Virmani, A. (1989), Indirect Tax Evasion and Production Efficiency, Journal of Public Economics 39, 223–237.

Wang, L.F.S. (1990), Tax Evasion and Monopoly Output Decisions with Endogenous Probability of Detection, Public Finance Quarterly 18, 480–487.

Wang, L.F.S., and J.L. Conant (1988), Corporate Tax Evasion and Output Decisions of the Uncertain Monopolist, National Tax Journal 41, 579–581.

Weichenrieder, A.J. (1996), Transfer Pricing, Double Taxation, and the Cost of Capital, Scandinavian Journal of Economics 98, 445–452.

Weichenrieder, A.J. (2009), Profit Shifting in the EU: Evidence from Germany, International Tax and Public Finance 16, 281–297.

Yaniv, G. (1995), A Note on the Tax Evading Firm, National Tax Journal 48, 113–120.

Yaniv, G. (1996), Tax Evasion and Monopoly Output Decisions: Note, Public Finance Quarterly 24, 501–505.

Yitzhaki, S. (1974), A Note on Income Tax Evasion: A Theoretical Analysis, Journal of Public Economics 3, 201–202.

Tim Friehe
University of Konstanz
Department of Economics
Box 136
78457 Konstanz
Germany
tim.friehe@uni-konstanz.de

Old George Orwell Got It Backward: Some Thoughts on Behavioral Tax Economics

Joel Slemrod*

Received 24 August 2009; in revised form 6 January 2010; accepted 10 January 2010

It is entirely appropriate that the study of public finance take seriously ''behavioral'' inconsistencies with traditional models of individual and collective decision-making. This raises the question of whether the state should play a role in protecting individuals from themselves, and whether individuals are susceptible to manipulation, or even exploitation, by the people who comprise the state. In this essay I two aspects of this issue – tax complexity and tax compliance. In addressing these issues I ask, and offer some tentative answers to, what is distinctive about behavioral tax economics as a sub-field of behavioral economics and as a sub-field of tax economics.

Keywords: complexity, compliance

JEL classification: H 20, H 27

"Old George Orwell got it backward. Big Brother isn't watching. He's singing and dancing. He's pulling rabbits out of a hat. Big Brother's busy holding your attention every moment you're awake. He's making sure you're always distracted. He's making sure you're fully absorbed."
Chuck Palahniuk (2002), in Lullaby

1. Introduction

It is entirely appropriate that the study of public finance take seriously "behavioral" inconsistencies with traditional models of individual and collective decision-making. A central tension in political economy – the extent to which people need to be protected from the state versus whether the state is needed to protect people from each other and the vicissitudes of life – takes on new dimensions once one recognizes that people often act irrationally and in ways that are contrary to their own long-term interests, and are cognitively bounded. This raises the question of whether the state should play a role in

* I am grateful to Dan Silverman, Jean-Robert Tyran, Neslihan Uler, conference attendees and an anonymous referee for helpful comments on an earlier draft.

FinanzArchiv 66 (2010), 15–33 doi: 10.1628/001522110X503361
ISSN 0015-2218

protecting individuals from themselves, and whether individuals are susceptible to manipulation, or even exploitation, by the people who comprise the state.

In this essay I address one aspect of this issue – how it affects an economic analysis of tax systems. In addressing this task I ask, and offer some tentative answers to, what is distinctive about behavioral tax economics as a sub-field of behavioral economics and as a sub-field of tax economics. In his review of behavioral economics, Camerer (2006, p. 196) refers to the "franchising" of behavioral economics into sub-fields, listing finance, labor economics, and public finance as notable examples, but he does not elaborate on the distinctive problems that the public finance franchise must confront. The papers collected in McCaffery and Slemrod (2006) address some of these questions, and Kirchler (2007) provides a nice overview of some as well. In this essay I offer some thoughts on two of these distinctive problems. The first is tax complexity and how it relates to the framing of the fiscal environment and the bounded rationality of taxpayers and citizen-voters. The second is tax compliance and the roles played by altruism and reciprocity.

2. Tax Complexity

Observers of tax systems have noted that many tax systems are extraordinarily complicated, although the standard for judging excess is generally not made explicit. Consider the U.S. income tax. The resource cost of collecting income taxes – including both the administrative cost and compliance cost – has been estimated to be about 10 percent of revenue collected.[1] In 2005 the number of words in the federal income tax code was 1,286,000, and 5,778,000 words in federal tax regulations.[2] In tax year 2006, 62 percent of individual income taxpayers in the United States paid someone to help them file; 72 percent did so in Australia.

Tax complexity can affect the private decisions made by taxpayers, as well as the voting and other social choice behavior of citizens. Each of these issues has been recognized, although the connection between the two has not been given much attention.

2.1. Taxpayers as Decision Makers

Accurately calculating one's tax liability in self-assessment income tax systems is difficult for many taxpayers. To be sure, this is heterogeneous. People differ in their cognitive ability, and the cognitive (and energy) process re-

1 See Slemrod (2004) for the source of this figure.
2 http://www.taxfoundation.org/research/show/1961.html, accessed on April 7, 2009.

quirements of tax compliance vary greatly across taxpayers. For those with complicated financial affairs, especially with respect to capital income, it can be very difficult. For those whose income is mostly wages and salary, the process can be very straightforward.

Calculating one's tax liability is tantamount to calculating the distance between a pre-tax budget line and a post-tax budget line at a given set of choices about labor supply and consumption basket. Of more interest to economists, though, is how complexity affects individuals' perceptions of the whole budget set, and in particular the (relative) prices as reflected by its gradients.

There is a long strand of research about taxpayer perceptions of their average and marginal tax rate, a component of some central relative prices such as between labor and leisure. Sheffrin (1994) reviews studies of American, British, and Canadian taxpayers that find that taxpayers generally underestimate both their total tax liability and their marginal tax rates. What is not clear is the connection between knowledge of one's marginal tax rate and the complexity of the tax system. Bartolome (1995) showed in an experiment that at least as many individuals used the average tax rate as if it was the marginal tax rate as used the true marginal tax rate in making marginal decisions. Moreover, how the tax table was represented mattered a lot, as almost all subjects used the true marginal tax rate when the tax table was redesigned to stress it. Liebman and Zeckhauser (2004) argue that it is because of cognitive limitations that taxpayers presume that their marginal tax rate is the easier-to-calculate average tax rate; they call this rule-of-thumb behavior "ironing," one of two examples of what they dub "schmeduling," defined as an inaccurately perceived price schedule. They (as did Bartolome, 1995) show that ironing behavior eliminates some of the deadweight loss from high marginal taxes, so that when the optimal tax schedule with non-ironing taxpayers would be convex, superior outcomes are available. An empirical analysis of the introduction in 1998 of the child care credit uncovers evidence that is consistent with "schmeduling," but is not conclusive.

That taxpayers have cognitive limitations has many implications for tax analysis. One is that the distribution of tax burden may depend on cognitive ability in addition to the intended characteristics of taxpayers. Another is that taxpayers who are not so good at addressing tax matters may avoid certain types of employment status, such as self-employment, that require or reward this kind of savvy. Finally, some equivalences taken for granted by tax theory – for example, that between a labor income tax and a consumption tax – may not obtain because they are perceived differently.[3] Blumkin, Ruffle,

3 The textbook equivalences may also fail because of differences in the administrative and compliance costs. See Slemrod (2008a).

and Ganun (2008) find, in an experimental setting, that subjects reduce their labor supply significantly more in response to an income tax than they do in response to an equivalent consumption tax, and speculate that this occurs because subjects underestimate the present value of the taxes levied on future consumption. Sausgruber and Tyran (2005) show, also in an experimental setting, that buyers systematically underestimate the tax burden of an indirect tax levied on sellers but do not do so with an equivalent direct tax. Moreover, Sausgruber and Tyran (2008) demonstrate that this perception bias can distort voting on taxes in a way that is reduced by experience but not by deliberation.

I believe that the public finance community needs to give more thought to the appropriate econometric methodology for measuring taxpayer responses to fiscal instruments in the presence of cognitive biases and rules of thumb that economize on cognitive resources. A first but necessary step is to understand how taxpayers map these instruments into relative prices. This step includes, but is not limited to, understanding which changes are ignored as not salient, possibly because rules of thumb conserve cognitive resources and focus only on changes above a certain level of materiality. The burgeoning literature on salience and taxation has provided many insights, but has not yet been placed in an adequate dynamic context; after all, a series of small tax changes, each ignored, can add up to a major displacement from an optimal choice.[4]

2.2. Taxpayers as Voters

Adult workers and consumers are also potential voters, and as such must decide which candidates' tax positions to favor, including how complex tax systems should be. In their role as voters they are confronted by the difficulty of figuring out what the consequences of alternative tax policies are, both for themselves and for aggregate economic outcomes.

Some context is appropriate. Political science research is fairly persuasive that voters know very little about the details of government generally. Delli Carpini and Keeter (1996), in a comprehensive survey of the political knowledge of voters covering several decades and hundreds of surveys, show that majorities of voters are ignorant of many key aspects of the U.S. political system, such as who has the power to declare war, the respective functions of the three branches of government, and who controls monetary policy. In contrast, the policy implications of this lack of political knowledge are highly controversial among political scientists. For example, Lupia (2001) argues that political knowledge, as commonly measured by scales that count the

4 See Chetty, Looney, and Kroft (2009).

number of correct responses to a small number of questions about public affairs, represent neither necessary nor sufficient conditions for voter *competence* in making choices, where a choice is defined to be competent if it is the same choice that would be made given the most accurate information about its consequences. There is little systematic evidence that voters are misled to support policies that are almost certainly not in their interest, although Slemrod (2006) argues that this is the case with respect to Americans' support for a flat income tax: because many (mistakenly) believe the current U.S. income tax system is regressive, they view, and support, the move to a flat tax as a move toward a more progressive distribution of the tax burden.

Why some jurisdictions' tax systems, or to be exact some taxes levied in some jurisdictions, become complex while others do not has not been widely studied, hampered by the absence of reliable, comparable measures of complexity across countries. Slemrod (2005) uses the variation in U.S. state income tax systems, and their differential change over time, to examine what engenders tax complexity, measured simply by the number of lines in the tax forms and the number of pages in the instruction booklets, and reveals some intriguing patterns. The analysis shows that, in 2000, states with more professional legislatures, as measured by the salaries paid, tended to have more complex tax systems, as did states with a less active voting population. The former relationship suggests that complexity is one of the things that professional legislatures *do*, although it may also be that states that want more activist policy want professional legislatures and choose more complex tax systems. The latter relationship suggests that a more politically involved citizenry acts as a deterrent to tax complexity.

Professional economists cannot be too smug about voter confusion, of course, because there is much we do not understand ourselves in all areas of economics, with the economics of taxation being no exception. Central questions such as the incidence of the corporate tax and deficit financing, and the long-term growth implications of alternative tax systems, are unresolved and controversial.

Politicians have an incentive and often the ability to take advantage of taxpayers' behavioral quirks and cognitive limitations.[5] McCaffery (1993) and Krishna and Slemrod (2003) argue that the U.S. income tax has many features that take advantage of cognitive biases to reduce the perceived tax burden, and do so by applying well-known features of what in marketing

5 To be sure, there are other explanations for tax complexity. Hettich and Winer (1999) argue that complex tax structures emerge as a by-product of the struggle for political office, in the course of which political parties are forced to propose and implement policies that discriminate or distinguish as carefully as possible among heterogeneous voters. In their view it is administrative costs that limit the desire of governments to discriminate fully among taxpayers.

science is known as "price presentation," such as the use of discounts (as in deductions from a broad measure of income) and of small frequent disbursements (as in employer withholding).[6] The laboratory experiments of Baron and McCaffery (2003) provide some support that such a strategy can be successful, as they demonstrate that people tend to underestimate the total tax burden when it is spread among multiple taxes. The concern that some taxes, such as the value added tax or corporation income tax, are "hidden" from taxpayers is a major reason why some conservatives oppose these taxes: they feel that the hiddenness causes voters to underestimate the true cost of government.[7] As an example, Finkelstein (2009) argues that the switch from manual, per-trip, remittance of traffic tolls to automatic electronic charging facilitated toll increases because the act of remittance became less salient to driver/voters.

The analogy of tax design to price presentation raises the issue of what is different between the public finance setting and a market environment. Although in democracies there is some degree of political competition, it seems likely that in social choices the intermediation of the market is less relevant, whereas in many cases involving behavioral economics, markets might plausibly arbitrage away, or exploit, irrationalities. But even this is not obvious. As Mullainathan and Thaler (2001) emphasize, many decisions, such as with regard to retirement savings, are made infrequently and so learning by doing is not likely to be very important; less-than-rational people "survive" and influence market outcomes.

Scitovsky (1950) observed sixty years ago that ignorance can be a source of oligopoly power because it limits price and quality competition among established firms and protects them from potential entry, thus facilitating collusion among established firms. Garrod (2007) remarks that obfuscation is widespread in several markets including, somewhat surprisingly, Internet retailing, and retail financial products such as index funds, money market funds, credit cards, conventional fixed-rate mortgages, life annuities, and term life insurance. Indeed, recent theoretical research has shown how even competitive markets might not drive out private firms' obfuscation about prices. As a baseline, Milgrom (1981) had shown that, if consumers have the cognitive ability to infer that they should avoid firms with hidden information, then competing firms will fully inform consumers of product information if doing so is feasible and costless. But a more recent literature has established that

6 Kim and Kachersky (2006) critically review the marketing science literature on price salience.

7 Although see Galle (2009, pp. 94–98), who questions whether hidden taxes are necessarily lower, arguing that if sophisticates are aware that naives are unaware of the tax and hence do not lobby, sophisticates know that they can no longer free ride. Thus, total political opposition to hidden taxes may be higher.

obfuscation can be profitable in equilibrium with competitive constraints. For example, in the "shrouded attributes equilibrium" of Gabaix and Laibson (2006), firms can obfuscate their prices for a complementary, avoidable add-on to a good, and may optimally do so when there are a sufficient fraction of consumers that myopically do not consider the add-on. Of particular interest for the topic at hand is the result that in equilibrium the sophisticated consumers actually benefit by taking advantage of the lower price for the basic good and not purchasing the add-ons.[8]

One insight from this literature is that firms may benefit from complicating, or obfuscating, the available information and that this process may benefit sophisticated consumers at the expense of unsophisticated ones. I have argued that the people in an incumbent government may try to, and may be able to, take advantage of framing difficulties to benefit themselves by reducing the perceived burden of what they do. In the process the more sophisticated of taxpayers may actually benefit. But the people in the government are also "just" people, and so may themselves be subject to framing issues. One important and understudied issue is how to model the behavior of policy makers subject to cognitive limitations: are they subject to the same kind of heuristic biases as taxpayers/voters?

3. Tax Compliance

Once a tax system is in place, influenced by the voting decisions of citizens, then these same citizens must decide whether and to what extent to comply with the tax rules in place. Can the psychology of attitudes toward authority shed light on the question of whether taxpayers generally free ride, or under some circumstances look beyond their cost-benefit calculus of risk and reward to be influenced by, for example, the fairness of the distribution of tax burden and the process that determines the burden?[9] Camerer (2006, p. 9) places this issue squarely within the realm of behavioral economics, remark-

8 Carlin (2009) generalizes the standard assumption in this literature that the fraction of uninformed consumers is constant, and allows that firms may influence how informed the consumer population is by affecting the quality of information they are given; in this model an increase in competitive pressure generates an increase in price complexity. Garrod (2007) discusses other models that address how firms strategically set prices in equilibrium in response to the cognitive shortcomings of their consumer population.

9 Of course, behavioral game theory may also apply to the strategic interaction between a taxpayer (and perhaps a professional tax preparer) and government auditors. For example, Slemrod, Blumenthal, and Christian (2001) found in a field experiment that sophisticated, high-income taxpayers report *less* taxable income when informed their tax return would be audited; one explanation is that the taxpayers understand that an audit is a negotiation in which the auditor has imperfect information, and that in such a negotiation a low initial bid (i.e., a low reported income) may be part of an optimal strategy.

ing that "the idea that people care only about their own monetary or goods payoff is not a central tenet of rational choice theory, but it is a common simplifying assumption."

3.1. Beyond Deterrence

Although the deterrence framework introduced by Allingham and Sandmo (1972) has dominated the economics literature addressing tax evasion, some have argued that it misses important elements of the tax evasion decision because it predicts a compliance rate much lower than what we actually observe. For example, Feld and Frey (2002, p. 5) assert that it is "impossible to account for tax compliance in terms of expected punishment." The dismissive argument goes as follows: given the low average probability of audit (in the United States recently less than 1 percent for individual returns with no business income), the low penalties generally assessed for noncompliance (typically 10 percent of the amount underpaid in the U.S.), and what we know about the degree of risk aversion from other contexts, noncompliance should be much higher than it apparently is.

But this dismissive argument is not persuasive, because the low average audit coverage rate vastly understates the chances that a typical dollar of unreported net income would be detected. A wage or salary earner whose employer submits the employee's taxable income and Social Security number electronically to the Internal Revenue Service, but who does not report that income on his own personal return, will be flagged for further scrutiny with a probability much closer to 100 percent than to 1 percent. Thus, the low rates of noncompliance for labor income (about 1 percent) calculated as part of the IRS tax gap study (U.S. Department of Treasury, 2005) by no means patently contradict the deterrence theory. Whether the 57 percent noncompliance rate of non-farm sole proprietors the IRS calculates is less than the deterrence theory predicts is less clear, and Andreoni, Erard, and Feinstein (1998, pp. 821–822) argue that it is.

Nonetheless, there is considerable experimental (and anecdotal) evidence that there is more to the story of tax evasion than an amoral cost-benefit calculation. Frey (1997) argues that it is important to differentiate between the intrinsic motivation under which taxpayers comply with tax liabilities because of "civic virtue" and extrinsic motivation in which they pay because of threat of punishment. He suggests that increasing extrinsic motivation – say with more punitive enforcement policies – may "crowd out" intrinsic motivation by making people feel that they pay taxes because they have to, rather than because they want to. Gneezy and Rustichini (2000) argue that this explains why parent tardiness *increased* after an Israeli day care center instituted monetary fines for late pick-up of children. In an experimental

setting, Scholz and Lubell (2001) find that the level of cooperation in certain settings declines significantly when penalties are introduced, suggesting that the increased level of deterrence did not compensate for the change in how people frame their decision brought about by the higher penalties.

Some laboratory experiments have found that subjects respond not only to the probabilities and stakes of a tax evasion game, but also to the context provided to them, as in Spicer and Becker (1980) and Alm, Jackson, and McKee (1992).[10] Alm, Jackson, and McKee (1993) found that (1) experimental subjects are willing to pay more in taxes when they first choose the use of their taxes by voting than when the identical use is imposed upon them, (2) compliance is somewhat greater when the vote is decisive compared to when the vote is close, and (3) tax compliance is significantly lowered by the imposition of an unpopular program.

It may be that tax evasion decisions depend on perceptions of the fairness of the tax system. If, the argument goes, perceived tax equity strengthens the social norm against evasion, then evasion becomes more costly in terms of bad conscience (if not caught) or bad reputation (if caught). Note also that an individual may find unfairness in what the government uses tax revenues for - a person with some of the spirit of Henry David Thoreau[11] may avoid taxes because that person thinks government (non-tax) policy wrong (Andreoni, Erard, and Feinstein, 1998). But such individual judgments can be complex; for example, expenditures on warfare might contribute to a sense of fairness tolerated in a patriotic period, but rejected during another period characterized by anti-militarism.[12]

These patterns suggest that a form of reciprocal altruism may be at work, in which the taxpayer's behavior depends on the behavior, motivations, and intentions not of any subset of particular individuals, but of the government itself. Levi (1997, p. 91) argues that when citizens believe that the government will act in their interests, that its procedures are fair, and that their trust of the state and others is reciprocated, then people are more likely to become "contingent consenters" who cooperate in paying taxes even when their short-term material interest would make free riding the individual's best option. Some survey evidence is consistent with this hypothesis. Torgler (2003) and Slemrod (2003) show there is a positive relationship across countries between survey-based attitudes toward tax evasion on the one hand and professed trust in government, and Slemrod (2003) finds that the same

10 Alm and Jacobson (2007) critically review the use of laboratory experiments in public economics.

11 Thoreau, the author of the influential 1849 book *Civil Disobedience* advocating resistance to unjust forms of authority, in 1846 refused to pay delinquent poll taxes because of his opposition to the Mexican-American War and slavery.

12 This argument is made by Daunton (1998).

relationship holds across individuals within the United States and Germany. Of course attitudes and actions are not the same.[13] A 2002 poll in the Czech Republic indicated that a person would be more likely to evade taxes if that person believed government services were substandard (Hanousek and Palda, 2004). None of these studies, though, establishes a causal connection between the two attitudes, and some of the observed correlation might be due to an ex post rationalization of tax-noncompliant behavior.

If perceptions matter for tax compliance, a natural question is to what extent tax compliance behavior can be manipulated by the government to lower the cost of raising resources. Appeals to conscience go back at least to Hammurabi's reign in ancient Babylon, when the tax collector sent the following notice when payments were late: "Why have you not sent to Babylon the 30 lambs as your tax? Are you not ashamed of such behavior?"[14] Wartime appeals to patriotism to induce citizens to pay their taxes (and, often, buy war bonds) are common; the U.S. Secretary of Treasury during World War I, William Gibbs McAdoo, referred to these campaigns as "capitalizing patriotism." Kang and Rockoff (2006) discuss the World War I experience, while Jones (1988/1989) discusses fiscal propaganda during World War II. Feldman and Slemrod (2009), using cross-country data on interstate conflicts from 1970 to the present and on attitudes toward tax evasion from the World Values Survey, find that positive attitudes towards tax compliance increase with the number and length of conflicts that a country faces, but decrease in the number of fatalities incurred in these conflicts. Konrad and Qari (2009) find a positive cross-country and within-country correlation between professed patriotism and tax compliance attitudes, although it is difficult to establish causality with their data.

That such campaigns are successful during ordinary (non-war) times has not been compellingly demonstrated. In a randomized field experiment with Minnesota taxpayers in a peacetime setting, Blumenthal, Christian, and Slemrod (2001) find no evidence that either of two written appeals to taxpayers' consciences had a significant effect on compliance. One letter stressed the beneficial effects of tax-funded projects, while the other conveyed the message that most taxpayers were compliant. Torgler (2004), using a controlled field experiment in Switzerland, also found that moral suasion had hardly any effect on taxpayers' compliance behavior. Fellner, Sausgruber, and Traxler (2009) find that similar written appeals had no discernible impact on compliance with Austrian television registration fees.

13 Kirchler (2007, p. 55) concludes from a review of the literature that most studies find a statistically significant, but weaker, relationship between attitudes toward taxation and self-reported compliance behavior, and goes on to suggest that this implies that the relationship between attitudes and actual behavior "is expected to be even weaker."

14 This quotation is cited in Webber and Wildavsky (1986, p. 58).

Survey evidence also suggests that attitudes about the acceptability of tax evasion vary considerably across countries. In the World Values Surveys done between 1999 and 2002, respondents were asked whether, given the chance, tax evasion is never, sometimes, or always justified, where a value of 1 corresponds to "never justifiable" and a value of 10 corresponds to "always justifiable." These attitude measures of the World Values Survey across countries are associated, holding other factors constant, with already-discussed measures of the shadow economy and widely used survey measures of actual evasion (Torgler, 2004). But, again, attitudes are not behavior.

The difficulties of separating out whether people pay their taxes because they feel they "ought to" or whether they fear the penalties attendant to not doing so is well illustrated by some evidence from a recent survey sponsored by the Internal Revenue Service Oversight Board (U.S. Department of Treasury, 2006). While 96 percent of those surveyed in 2005 mostly or completely agreed that "It is every American's civic duty to pay their fair share of taxes," 62 percent also said that "fear of an audit" had a great deal or somewhat of an influence on whether they report and pay their taxes "honestly."

Behavioral game theory and laboratory experiments may shed some light on the conditions under which taxpayers may be willing to deviate from their Allingham-Sandmo optimal level of evasion. For example, much research about the ultimatum game suggests that people are willing to take costly actions that express their concerns for fairness. Many people express "negative reciprocity," meaning that they will take actions that lower the welfare of the *person* who treated them in a way that they perceive to be unfair, and will do so at a cost to themselves. The experimental results reported in Blount (1995) suggest that beliefs about what motivated another person and judging the appropriateness of the motives, their "intentionality," is critical to explaining behavior toward that person.[15] According to Cooper and Kagel (forthcoming, p. 49), the Blount result "makes it completely obvious why outcome-based preferences are not enough."

Some observers have interpreted behavior reflecting intentionality as the vestigial expression of a behavior that had survival value in a setting where people repeatedly interacted with the same people in a small group. Also

15 The role of intentionality is nicely illustrated, although not resolved, by a famous interchange in Joseph Heller's (1961) novel, *Catch-22*, between the protagonist John Yossarian, an Army Air Force bombardier, and another member of the bomb crew, the naïve Clevinger:

"They're trying to kill me," Yossarian told him calmly.

"No one's trying to kill you," Clevinger cried.

"Then why are they shooting at me?" Yossarian asked.

"They're shooting at everyone," Clevinger answered. "They're trying to kill everyone."

"And what difference does that make?"

of interest is the laboratory result that inducing a sense of entitlement, by allowing the ultimatum game proposer to be the winner of a contest, lowers offers; the sense of entitlement leads people to give away less of what is theirs. Ultimatum games with multiple players suggest that responders care about whether proposers are unfair to *them*, but do not care much about how the proposer treats others. This is an important distinction for understanding individuals' attitudes toward government, because government policies do not generally single out particular individuals other than through enforcement actions, but may single out groups of people defined by income, geography, demographics, tastes, or choices.

Note also that the concern for fairness that is evident in two-player games tends to disappear in large markets, where even those who care about fairness behave self-interestedly either because they are not sure whether others are being fair or they cannot easily punish those that are acting unfairly. As Camerer (2006) remarks, "a competitive market is simply a place in which it is hard to express your concern for fairness." It is indeed hard, but not impossible. Consumer boycotts date back as far as the fourteenth century, and have had both sociopolitical objectives (as with the U.S. civil rights bus boycotts) and objectives of changing corporate practices (as with the Nike boycott designed to stop their use of "sweatshop" labor).

There is an active controversy about what exactly fairness means. Is it an aversion to inequality, where people dislike both getting less than a fair share and getting more than a fair share? Or is a concern for reciprocity, where how people feel about others depends on how they expect to be treated? As mentioned, the research suggests that people care about the *intentions* of other players.

3.2. The Psychology of Authority

We know little about to what extent the psychological dynamics of individuals' relations with other individuals may be different than the psychological dynamics of individuals versus an agency of the government. For example, there is evidence from laboratory experiments that many people are willing to reciprocate what they perceive to be kindness in other individuals, and to not reciprocate – or even punish – perceived meanness in others. In addition, Falk and Kosfeld (2006) show that implementing a minimum performance requirement causes most agents to reduce their overall performance in response; when asked how they perceived the minimum performance requirement, most of those who reacted negatively said that they perceived it as a signal of distrust and as a limitation on their choice autonomy. But how do individuals ascribe human qualities like kindness, meanness, or distrust to a government? For example, do such feelings change with a change of

government? Certainly we should be sensitive to an "anthropomorphic fallacy" of attributing human thoughts and emotions to inanimate objects or animals, but that does not imply that related attributions, and reactions to those attributions, do not occur in interactions with governments.

Here again we may learn by looking at similar situations. Although government is not exactly like any other organization or institution – in particular, it has a monopoly on coercive power, including the power to tax – government is not the only organization or institution that individuals interact with. For example, people interact on a regular basis as employees and customers with firms, some of which are as large as some governments. As employees they have to decide whether to give maximal effort or slack off, whether to pilfer or even embezzle. Indeed, in his survey of behavioral economics, Camerer (2006, p. 177) asks whether angry workers consider "management to be a single monolithic player and get angry the same way that they get angry at a spouse who threatens to leave them or a driver who cuts them off on the LA Freeway." As customers, people have to make decisions about shoplifting, insurance fraud, and the like. Firms invest resources in deterring employee crime (and encouraging effort), with accounting systems and hidden cameras. Many companies try to instill identification with the company, so as to achieve both goals. Konrad (2008) argues that, similarly, countries must make decisions about how many resources to invest in instilling identification with the country, what he refers to as patriotism, in order to increase tax compliance.

Because governments have much more power than any other organization, individual psychological attitudes toward them might be fundamentally different than toward other organizations. Because they purport to serve their interests, individuals might feel more dutiful, and even *obedient*, toward government. Invocation of the word obedience, though, invokes a darker side of the relationship between individuals and government as an authority figure. I am speaking of the controversial, indeed notorious, experiments conducted by the Yale University psychologist Stanley Milgram (1963), which showed that unwitting subjects were willing to deliver what they thought were substantial electric shocks when instructed to, and encouraged to, by authority figures. This research ignited several controversies, one of which centered on the nature and influence of authority figures – in Milgram's experiments the men in white coats who were urging the subjects to continue the apparent shock treatments. Commentators such as Morelli (1983) differentiated between a person who is "in authority" and "an authority," where the former refers to legitimate coercive power and the latter refers to a presumption of expert knowledge. Although reaction to the former may be characterized as "obedience," response to the latter might be better denoted as "deference."

Authority, obedience, and deference are central to many important questions in public economics. For example, the extensively demonstrated[16] effect of defaults on individual choices may be due to the decider's presumption of expertise on the part of an authority, or to the decider's cost of obtaining information. Distinguishing between "in authority" and "an authority" is a worthwhile research objective. It will not likely be easy to address in an experimental setting, though, in part because both internal validity and external validity questions arise. Internal validity issues, also known as experimenter demand effects (EDE), arise when behavior by subjects depends on cues about what constitutes appropriate behavior. As Zizzo (forthcoming) states, "it is unavoidable that the experimenter is in a position of authority relative to subjects," having both legitimacy and expertise. Indeed, Zizzo (forthcoming, p. 6) says that the Milgram experiment is "an extreme case of EDE at work in an experiment where the effect of such social EDE was itself the objective of the experiment." External validity questions arise when the authority for the real-world tax enforcer (often referred to as the "tax authority") is crucial to behavioral response, and whose attributes are not (and maybe cannot) be replicated in a lab.

Before concluding, it is interesting to note that complexity and compliance are not completely independent concerns. It may be that a tax system is so complex that it is not worth the taxpayer's time and expense to accurately calculate tax liability. In that case, there will be more capriciousness in the assignment of tax burden, but the errors should be symmetric. It may also be that complexity engenders a particular strong negative assessment of the government and the tax process, which overcomes the free-rider calculus and leads to noncompliance. On the other hand, Scotchmer and Slemrod (1989) suggest that uncertainty about true tax liability may in some settings cause risk-averse taxpayers to be less aggressive than otherwise in their tax reporting behavior.

4. Final Thoughts

In this essay I have offered some observations about what insights behavioral economics can bring to bear on the issues of tax complexity and tax compliance. I ponder why the government, in the words of the epigraph, is often "singing and dancing" – presenting taxpayers with a tax system of cacophonic complexity. In part it may be, as the epigraph's author Chuck Palahniuk suggests, to distract taxpayers from the magnitude of the tax burden imposed or to reward those who can manage and even profit from the complexity. After all, firms operating in competitive markets have discovered that this can be

16 See, for example, Choi, Laibson, Madrian, and Metrick (2004).

a profitable strategy, and the constraints of political markets are likely to be smaller. It is, though, these same taxpayers in their role as voters that help set the tax system and its level of complexity, marking a difference between government and Procter & Gamble offering coupons it knows most buyers never cash in. In an environment of complicated tax systems and differing opinions among economists about the ultimate consequences of even simple tax systems, we must ask who is fooling whom.

Governments and firms differ in another important way. Most firms are selling private goods and therefore can, instances of theft notwithstanding, generally withhold their product absent payment. Governments are in part offering non-excludable services, and so must rely on remittances that are either made dutifully or "encouraged" by an enforcement system based on withholding, information reporting, and audits backed by penalties for noncompliance.[17] Behavioral economics, largely observing the results of laboratory experiments, has shown that some people will deviate in some situations from the choices that would seem appropriate for rational, purely self-interested individuals. For example, their decisions depend on the perceived reciprocity or intentionality of those they interact with, as well as their perceived meanness or kindness. But by their nature laboratory experiments involve the decisions of individuals interacting with other individuals, or perhaps a random-number-generating computer, and not with a largely impersonal government or its tax agency, an authority figure imbued with a rich and complicated history. Field experiments offer more promise for capturing how real people react to real policy changes made by real governments, although they come with their own limitations.

References

Allingham, M., and Sandmo, A. (1972), Income Tax Evasion: A Theoretical Analysis, Journal of Public Economics 1, 323–338.

Alm, J., Jackson, B. R., and McKee, M. (1992), Estimating the Determinants of Taxpayer Compliance with Experimental Data, National Tax Journal 45, 107–114.

17 Another part of what government does is effect transfers from one group to another. This suggests that lobbyists, along with legislators, might work together to contribute to tax complexity, and that the objective is not only to reduce the perceived average burden of taxation but also to obfuscate the distribution of that burden. This raises the issue of the cognitive biases and limitations of the lobbyists and the interests they represent. Slemrod (2008b) argues that business associations' positions on tax reform often make arbitrarily simple and internally inconsistent assumptions about incidence, especially by distinguishing between "business taxes" (presumed to negatively affect the return to business operations) and non-business taxes (not likely to do so) that are subject to the framing of the taxes and the remittance system (i.e., who "writes the check" to the tax authority).

Alm, J., Jackson, B. R., and McKee, M. (1993), Fiscal Exchange, Collective Decision Institutions, and Tax Compliance, Journal of Economic Behavior & Organization 22, 285–303.

Alm, J., and Jacobson, S. (2007), Using Laboratory Experiments in Public Economics, National Tax Journal 60, 129–152.

Andreoni, J., Erard, B., and Feinstein, J. (1998), Tax Compliance, Journal of Economic Literature 36, 818–860.

Baron, J., and McCaffery, E. J. (2003), The Humpty Dumpty Blues: Disaggregation Bias in the Evaluation of Tax Systems, Organizational Behavior and Human Decision Processes 91, 230–242.

Bartolome, C. A. M. (1995), Which Tax Rate Do People Use: Average or Marginal?, Journal of Public Economics 56, 79–96.

Blount, S. (1995), When Social Outcomes Aren't Fair: The Effect of Causal Attributions on Preferences, Organizational Behavior and Human Decision Processes 63, 131–144.

Blumenthal, M., Christian, C., and Slemrod, J. (2001), Do Normative Appeals Affect Tax Compliance? Evidence from a Controlled Experiment in Minnesota, National Tax Journal 54, 125–138.

Blumkin, T., Ruffle, B. J., and Ganun, Y. (2008), Are Income and Consumption Taxes Ever Really Equivalent? Evidence from a Real-Effort Experiment with Real Goods, CESifo Working Paper No. 2194.

Camerer, C. (2006), Behavioral Economics, in: Blundell, R., Newey, W. K., and Persson, T. (Eds.), Advances in Economics and Econometrics: Theory and Applications. Econometric Society Monographs, Cambridge University Press, Cambridge, 181–214.

Carlin, B. I. (2009), Strategic Price Complexity in Retail Financial Markets, Journal of Financial Economics 91, 278–287.

Chetty, R., Looney, A., and Kroft, K. (2009), Salience and Taxation: Theory and Evidence, American Economic Review 99, 1145–1177.

Choi, J. J., Laibson, D., Madrian, B., and Metrick, A. (2004), For Better or For Worse: Default Effects and 401(k) Savings Behavior, in: Wise, D. (Ed.), Perspectives on the Economics of Aging, University of Chicago Press, Chicago.

Cooper, David J., and Kagel, J. H. (forthcoming), Other Regarding Preferences: A Selective Survey of Experimental Results, in: Kagel, J. H., and Roth, A. (Eds.), Handbook of Experimental Economics, Volume 2.

Daunton, M. (1998), Trusting Leviathan: British Fiscal Administration from the Napoleonic Wars to the Second World War, in: Levi, M., and Braithwaite, V. (Eds.), Trust and Governance, Russell Sage Foundation, New York.

Delli Carpini, M. X., and Keeter, S. (1996), What Americans Know About Politics and Why It Matters, Yale University Press, New Haven.

Falk, A., and Kosfeld, M. (2006), The Hidden Costs of Control, American Economic Review 96, 1611–1630.

Feld, L., and Frey, B. (2002), Trust Breeds Trust: How Taxpayers Are Treated, Economics of Governance 3, 87–99.

Feldman, N. and Slemrod, J. (2009), War and Taxation: When Does Patriotism Overcome the Free-Rider Impulse?, in: Martin, I. W., Mehotra, A. K., and Prasad, M. (Eds.), The New Fiscal Sociology: Taxation in Comparative and Historical Perspective, Cambridge University Press, Cambridge, 137–154.

Fellner, G., Sausgruber, R., and Traxler, C. (2009), Testing Enforcement Strategies in the Field: Legal Threat, Moral Appeal, and Social Information, CESifo Working Paper No. 2787.

Finkelstein, A. (2009), EZ Tax: Tax Salience and Tax Rates, Quarterly Journal of Economics 124, 969–1010.

Frey, B. (1997), A Constitution for Knaves Crowds Out Civic Virtues, Economic Journal 107, 1043–1053.

Gabaix, X., and Laibson, D. (2006), Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets, Quarterly Journal of Economics 121, 505–540.

Galle, B. (2009), Hidden Taxes, Washington University Law Review 89, 59–113.

Garrod, L. (2007), Price Transparency and Consumer Naivety in a Competitive Market, Centre for Competition Policy Working Paper 07-10, University of East Anglia.

Gneezy, U., and Rustichini, A. (2000), A Fine is a Price, Journal of Legal Studies 29, 1–17.

Hanousek, J., and Palda, F. (2004), Quality of Government Services and the Civic Duty to Pay Taxes in the Czech and Slovak Republics, and Other Transition Countries, Kyklos 57, 237–252.

Heller, J. (1961), Catch-22, Simon & Schuster, New York.

Hettich, W., and Winer, S. L. (1999), Democratic Choice and Taxation: A Theoretical and Empirical Analysis, Cambridge University Press, Cambridge.

Jones, C. (1988/1989), Class Tax to Mass Tax: The Role of Propaganda in the Expansion of the Income Tax during World War II, Buffalo Law Review 37, 685–737.

Kang, S. W., and Rockoff, H. (2006), Capitalizing Patriotism: The Liberty Loans of World War I, NBER Working Paper No. 1191.

Kim, H. M., and Kachersky, L. (2006), Dimensions of Price Salience: A Conceptual Framework for Perceptions of Multi-Dimensional Prices, Journal of Product & Brand Management 15, 139–147.

Kirchler, E. (2007), The Economic Psychology of Tax Behaviour, Cambridge University Press, Cambridge.

Konrad, K. (2008), Mobile Tax Base as a Global Common, International Tax and Public Finance 15, August, 395–414.

Konrad, K., and Qari, S. (2009), The Last Refuge of a Scoundrel? Patriotism and Tax Compliance, IZA Discussion Paper No. 4121, Bonn.

Krishna, A., and Slemrod, J. (2003), Behavioral Public Finance: Tax Design and Price Presentation, International Tax and Public Finance 10, 189–203.

Levi, M. (1997), Consent, Dissent, and Patriotism, Cambridge University Press, Cambridge.

Liebman, J. B., and Zeckhauser, R. J. (2004), Schmeduling, unpublished paper, Harvard University.

Lupia, A. (2001), What We Should Know: Can Ordinary Citizens Make Extraordinary Choices?, unpublished paper, University of Michigan.

McCaffery, E. J. (1993), Cognitive Theory and Tax, UCLA Law Review 41, 1861–1947.

McCaffery, E. J., and Slemrod, J. (2006) (Eds.), Behavioral Public Finance, Russell Sage Foundation, New York.

Milgram, S. (1963), Behavioral Study of Obedience, Journal of Abnormal and Social Psychology 67, 371–378.

Milgrom, P. (1981), Good News and Bad News: Representation Theorems and Applications, Bell Journal of Economics 12, 380–391.

Morelli, M. F. (1983), Milgram's Dilemma of Obedience, Metaphilosophy 14, 183–189.

Mullainathan, S., and Thaler, R. H. (2001), Behavioral Economics, in International Encyclopedia of Social and Behavioral Sciences, Pergamon Press, 1094–1100.

Palahniuk, C. (2002), Lullaby, Doubleday & Co, Inc, New York.

Sausgruber, R., and Tyran, J.-R. (2005), Testing the Mill Hypothesis of Fiscal Illusion, Public Choice 122, 39–68.

Sausgruber, R., and Tyran, J.-R. (2008), Tax Salience, Voting, and Deliberation, University of Copenhagen Department of Economics Discussion Paper 08-21.

Scholz, J. T., and Lubell, M. (2001), Cooperation, Reciprocity, and Collective-Action Rhetoric, American Journal of Political Science 45, 160–178.

Scitovsky, T. (1950), Ignorance as a Source of Oligopoly Power, American Economic Review 40, 48–53.

Scotchmer, S. and Slemrod, J. (1989), Randomness in Tax Enforcement, Journal of Public Economics 38, 17–32.

Sheffrin, S. (1994), Perceptions of Fairness in the Crucible of Tax Policy, in: Slemrod, J. (Ed.), Tax Progressivity and Income Inequality, Cambridge University Press, Cambridge, 309–334.

Slemrod, J. (2003), Trust in Public Finance, in: Cnossen, S., and Sinn, H.-W. (Eds.), Public Finance and Public Policy in the New Century, MIT Press, Cambridge, MA, 49–88.

Slemrod, J. (2004), Written Testimony Submitted to the Committee on Ways and Means, Subcommittee on Tax Simplification, Washington D.C., June 15.

Slemrod, J. (2005), The Etiology of Tax Complexity: Evidence from U.S. State Income Tax Systems, Public Finance Review 33, 279–299.

Slemrod, J. (2006), The Role of Misconceptions in Support for Regressive Tax Reform, National Tax Journal 59, 57–75.

Slemrod, J. (2008a), Does It Matter Who Writes the Check to the Government? The Economics of Tax Remittance, National Tax Journal 61, 251–275.

Slemrod, J. (2008b), Is Tax Reform Good for Business? Is a Pro-Business Tax Good for America?, in: Diamond, J. W., and Zodrow, G. R. (Eds.), Fundamental Tax Reforms: Issues, Choices and Implications, MIT Press, Cambridge, MA, 143–170.

Slemrod, J., Blumenthal, M., and Christian, C. (2001), Taxpayer Response to an Increased Probability of Audit: Evidence from a Controlled Experiment in Minnesota, Journal of Public Economics 79, 455–483.

Spicer, M. W., and Becker, L. A. (1980), Fiscal Inequity and Tax Evasion: An Experimental Approach, National Tax Journal 33, 171–175.

Torgler, B. (2003), Tax Morale, Rule-Governed Behaviour and Trust, Constitutional Political Economy 14, 119–140.

Torgler, B. (2004), Moral Suasion: An Alternative Tax Policy Strategy? Evidence from a Controlled Field Experiment, Economics of Governance 5, 235–253.

U.S. Department of the Treasury, Internal Revenue Service (2005), Preliminary Update of the Tax Year 2001 Individual Income Tax Underreporting Gap Estimates, 8-16, Washington, D.C., Office of Research, Analysis and Statistics, (Online access http://www.irs.gov/taxstats/compliancestats/article/0,,id=97177,00.html)

U.S. Department of the Treasury, Internal Revenue Service, Oversight Board (2006), Taxpayer Attitude Survey, February, Washington, D.C.

Webber, C., and Wildavsky, A. (1986), A History of Taxation and Expenditure in the Western World, Simon and Schuster, New York.

Zizzo, D. J. (forthcoming), Experimenter Demand Effects in Economic Experiments, Experimental Economics.

Joel Slemrod
University of Michigan
Stephen M. Ross School of Business
701 Tappan Street, Rm. R5396
Ann Arbor, MI 48109-1234
United States
jslemrod@umich.edu

Voting on Thresholds for Public Goods: Experimental Evidence

Julian Rauchdobler, Rupert Sausgruber, and Jean-Robert Tyran*

Received 15 December 2009; in revised form 23 February 2010; accepted 19 March 2010

Introducing a threshold in the sense of a minimal project size transforms a public-good game with an inefficient equilibrium into a coordination game with a set of Pareto-superior equilibria. Thresholds may therefore improve efficiency in the voluntary provision of public goods. In our one-shot experiment, we find that coordination often fails and exogenously imposed thresholds are ineffective at best and often counterproductive. This holds over a range of threshold levels and refund rates. We test whether thresholds perform better if they are endogenously chosen, i.e., whether a threshold is approved in a referendum, because voting may facilitate coordination due to signaling and commitment effects. We find that voting does have signaling and commitment effects, but they are not strong enough to significantly improve the efficiency of thresholds.

Keywords: Provision of Public Goods, Threshold, Voting, Experiments

JEL classification: H 41, D 72, C 92

1. Introduction

In some cases, public goods have a threshold, i.e., a minimal project size, for technological reasons (building half a bridge does not make much sense), but in other cases, the public good can be provided continuously (think of donations to start a community library). This paper investigates whether introducing a threshold to a public good that does not have a specific threshold value for technological reasons can increase efficiency.[1] More specifically, we

* Rauchdobler: University of Innsbruck, Department of Public Economics, Universitätsstr. 15, A-6020 Innsbruck. E-mail: Julian.Rauchdobler@uibk.ac.at. Sausgruber: University of Innsbruck, Department of Public Economics, Universitätsstr. 15, A-6020 Innsbruck. E-mail: Rupert.Sausgruber@uibk.ac.at. Tyran: University of Copenhagen, Department of Economics, Øster Farimagsgade 5, DK-1353 Copenhagen K. E-mail: Jean-Robert.Tyran@econ.ku.dk. We thank Charles Noussair, Rudolf Kerschbamer, and Wolfgang Höchtl for helpful comments. We are grateful for financial support by the Austrian Science Fund (FWF) under Projects No. 17029 and S10307-G14 and to the Danish Science Foundation FSE.

1 Threshold public goods are also called step-level, binary, discrete, lumpy, or provision-point goods.

ISSN 0015-2218
doi: 10.1628/001522110X503370

ask whether such thresholds are more effective if they have been approved in a referendum rather than just imposed by some authority.

The issue we investigate is of potential relevance in both the small and the large. For example, consider a fundraising drive to start the community library. Will citizens donate more if the charity commits to a minimal size of the library than if the fundraising drive is started without any lower bound for the size of the library? To provide an example from the other end of the scale, consider international agreements to abate greenhouse gases. Would nations be willing to contribute more to prevent global warming if the international agency commits to a minimum target for abatement? Are nations more likely to reach the threshold if they approved of the threshold in a referendum than if the threshold was imposed by some international organization? These examples illustrate that our study refers to situations in which potential contributors to the public good can be identified and, therefore, contributors can vote on the threshold, but the central authority is weak and does not have the power to enforce contributions to the public good. In the examples above, the charity cannot force citizens to donate, and the international organization cannot force sovereign states to reduce greenhouse gases. The mechanism we investigate is therefore a potential remedy to the free-rider problem when no strong enforcement institution exists and when the players keep the full decision sovereignty on whether or not to contribute to the public good. All that is required is that an institution can credibly commit not to provide the public good if voluntary contributions are insufficient to meet the threshold.

The theoretical rationale for expecting that introducing a threshold increases efficiency is that a threshold transforms a public-good game with a unique and inefficient equilibrium into a coordination game with a set of additional equilibria in which all players are better off (Bagnoli and Lipman 1989). The intuition for why this transformation may increase efficiency is that a player's contribution can be essential for whether the public good is provided at all.[2] If a single player makes the difference between providing the public good at the threshold level and providing none at all, he may find

2 Note that we investigate the effect of introducing a *threshold* in the sense of a minimal project size. In a situation without a threshold, the provision of the public good is continuous and proportional to the aggregate contributions. In contrast, in a situation with a threshold, the public good is provided if contributions meet or exceed the threshold, but is not provided at all if the threshold is not met. If contributions exceed the threshold, the public good is again continuous and proportional to aggregate contributions, as in a situation without a threshold. This property is also called the "extended benefits rule"; it tends to induce high contributions compared to rules in which contributions beyond the threshold are simply wasted (no rebate) or returned to the participants (full rebate). See Marks and Croson (1998) for a comparison of these rules, and Spencer et al. (2008) for a systematic discussion of alternative rebate rules.

it in his best interest to contribute. However, this situation typically occurs if others make contributions such that the threshold is *almost* reached, i.e., if it is feasible (and not too costly) for a particular player to make the difference. If others already contribute more than the threshold (and the public good is provided anyway) or if they contribute too little for a single player to make the difference between provision and no provision (i.e., the public good is not provided anyway), free riding is optimal for a self-interested player. Thus, thresholds improve incentives to cooperate only if players manage to solve a coordination problem.

The main focus of this paper is to investigate the effect of voting on thresholds, i.e., of choosing the level of the threshold in a vote. In particular, we ask whether efficiency is higher (because coordination is more successful) when a threshold has been approved in a referendum than when it is exogenously imposed. Our hypothesis is that strong electoral support for a threshold improves coordination because it serves to signal a willingness to contribute, i.e., widespread support is a reduced form of communication that shapes expectations. To illustrate, suppose that in a referendum all N voters support threshold level T. It seems plausible in a symmetric game (for reasons of focality and inequality aversion) to assume that, upon learning the result of the referendum, all players expect T to be reached and that all players contribute T/N.[3] In contrast, if only a small majority votes for T, players may hold the belief that some voters will contribute less than T/N each. This may result in a contribution of more than T/N by player i (if he thinks he is essential) or in zero contributions (if he thinks aggregate contributions are too low for his contribution to make a difference). Thus, the signaling effect is theoretically ambiguous, and the efficiency-enhancing effect of voting on thresholds is a fundamentally empirical issue, which can be systematically investigated only in the experimental laboratory (e.g., Falk and Heckman 2009).

We test the hypothesis that approving of a threshold in a referendum improves the effectiveness of thresholds in a two-stage game for different threshold levels (low, intermediate, or high) and for different refund rates (no, partial, or full refund). In stage 1, the players vote on implementing a threshold. The players know that if they accept, the public good will be provided only if their total contributions in stage 2 reach or exceed the threshold. In stage 2, upon learning the outcome of the vote, the players make contribution decisions. If a positive threshold is approved, there is a set of equilibria in pure strategies where the threshold is exactly met.

3 This conjecture is supported by a substantial body of research on social dilemmas, showing that under conditions of strategic uncertainty, group members tacitly coordinate their choice behavior by anchoring their decisions on rules of fairness (e.g., Allison et al. 1992, Suleiman et al. 2001).

Because these equilibria Pareto-dominate the equilibrium of the standard public-good game, it is a weakly dominant strategy to vote for the threshold in stage 1. Variation of the threshold level is interesting because both benefits and costs of coordination increase with the threshold. Higher thresholds are more efficient if they are reached, but players may be less confident that they can be reached, which increases the risk of wasting their contribution on a project that turns out not to be implemented.[4] Because of this trade-off, intermediate thresholds may be more efficient than high thresholds, and our experiment allows us to investigate if more ambitious thresholds are more or less efficient than low or intermediate thresholds. We also test the hypothesis at various refund rates (no, partial, and full refund) in case the project is not realized. A refund provides partial or full insurance against the risk of wasting one's contribution on a project that is not realized. Clearly, such insurance reduces the cost of miscoordination, and we expect better coordination at higher refund rates.

Our main findings are that thresholds are counterproductive if exogenously imposed and if less than full insurance is provided. When full insurance is provided, exogenous thresholds cease to be counterproductive and become merely ineffective: the efficiency is the same with no, low, or high thresholds. While higher thresholds are generally associated with higher contributions, the contributions often do not increase sufficiently to match the more ambitious thresholds. Higher thresholds induce the belief that others will contribute more generously, which increases contributions somewhat, but these beliefs are excessively optimistic. As a result, contributions often fail to meet the threshold.

When thresholds are endogenously chosen in a referendum, we find that thresholds are more popular with full insurance and that intermediate thresholds are more popular than ambitious ones. Thus, voting patterns reflect the (anticipated) risk of coordination failure when a threshold is ambitious and insurance unavailable. We do find that approval of a threshold serves as a coordination device in the sense that expected contributions for chosen thresholds are higher than for imposed ones, and this effect is more pronounced if the threshold receives stronger support among other voters. However, these effects are weak and cannot significantly reduce the massive coordination failure we observe in our design. As a result, overall efficiency is no greater when participants choose a threshold than when it has been imposed. Thus, introducing a threshold – whether by fiat or in a referendum – is not an effective cure for the inefficiency in the provision of public goods in our framework.

4 A reason might be that the set of asymmetric Pareto-efficient equilibria shrinks with rising threshold.

The paper proceeds as follows. Related literature is discussed in section 2; section 3 explains the experimental design. In section 4 we further discuss the predictions and hypotheses. Section 5 reports the results, and section 6 concludes.

2. Related Literature

The effects of exogenously imposed thresholds have been extensively studied in the experimental literature, but not, to the best of our knowledge, the effects of endogenous thresholds. The previous literature tends to find rather mixed results on imposed thresholds. For example, the literature review of Croson and Marks (2000) shows that results of previous studies have varied widely, with success rates ranging between 10% and 82%.[5] This meta-study shows that coordination tends to be more successful with a higher ratio of total benefits of the public good to its costs [6], with higher refund rates, and with better communication. Leadership contributions and other possibilities to choose sequentially also seem to increase the effectiveness of thresholds (e.g., List and Rondeau 2003, Cadsby and Maynes 1999, Coats et al. 2009). Controlled evidence from the field seems to support these findings (List and Lucking-Reiley 2002). Rondeau et al. (2005) report higher efficiency under a threshold mechanism when contributions below the target are fully refunded.

Closely related to our study are a field and lab experiment by Rondeau and List (2008) and a theoretical contribution by Gerber and Wichardt (2008). Rondeau and List (2008) investigate (among other things) the effect of introducing a threshold into a public good that does not have a threshold for technological reasons (a fundraising drive by the Sierra Club to provide environmental education) in a field experiment under conditions very similar to our lab study (contributions are fully refunded, i.e., $r = 1$ if the threshold is not met and the extended benefits rule that we use in all treatments applies). In line with our results, the authors find that a higher threshold (USD 2500 versus USD 5000) increased donations, but the increase was insignificant and donations (USD 945 versus USD 1375) were in both cases clearly insufficient to meet the threshold. The authors also implement a one-shot game in the lab with thresholds at USD 22.50 versus USD 45. Now, contributions increase significantly (USD 5.4 versus USD 7.5), but the increase is again less than proportional to the increase in the threshold (39% versus 100%). We infer

5 Our definition of a threshold as a minimum project size (allowing for project sizes exceeding the threshold) differs from that in some experimental studies surveyed in Croson and Marks (2000), which define the threshold as the only feasible project size.

6 Note that we hold this ratio, which is also called the step return, constant at 1.5 across all conditions.

from these numbers (the paper does not say) that the success rate must have fallen with a rising threshold, i.e., that higher thresholds were counterproductive in this respect. Gerber and Wichardt (2008) suggest a mechanism to provide a public good in the absence of sanctioning institutions. In their two-stage game, the players choose to pay a deposit in stage 1. The deposit is lost unless a player contributes to the public good. Hence, the deposit serves as a commitment device that renders contributing to the public good a dominant strategy.

3. Experimental Design

Our design has six treatments, which vary along two dimensions (see table 1). We vary (i) whether thresholds T are imposed (EXO) or endogenously chosen in a majority vote (END), and (ii) the refund rate r across treatments. Each subject participates only in one treatment (i.e., a subject makes choices either in EXO or in END) and experiences only one of the refund rates. The numbers in parentheses in table 1 show the number of participants in each condition. For example, we had 36 subjects participating in condition END0, which means that these subjects voted over threshold levels (END) and received no refund in case a threshold was not met ($r = 0$). The next two sections explain the parameters and procedures in EXO and END, respectively.

Table 1
Treatments of the Experiment (Number of Subjects per Cell in Parenthesis)

Decision mode	Refund rate 0%	50%	100%
END	END0 (#Subj.: 36)	END50 (#Subj.: 36)	END100 (#Subj.: 36)
EXO	EXO0 (#Subj.: 36)	EXO50 (#Subj.: 33)	EXO100 (#Subj.: 33)

3.1. Imposed Thresholds (EXO)

Treatments in EXO implement a standard version of the threshold public-good game (see Isaac et al. 1989). Subjects are randomly assigned to groups of $N = 3$, and group composition remains constant throughout the experiment. The subjects are endowed with $E = 20$ experimental points and decide how

many of these points to keep or contribute to a public good. The payoffs are determined by

$$\pi_i = \begin{cases} E - c_i + \alpha \sum_j c_j & \text{if } \sum_j c_j \geq T\,, \\ E - c_i + rc_i & \text{if } \sum_j c_j < T\,, \end{cases} \tag{1}$$

where π_i is subject i's payoff in points, c_i is i's contribution to the public good, and T is the threshold. If the sum of contributions within a group reaches or exceeds the threshold, each subject receives $\alpha = 0.5$ times this sum as payoff from the public good in addition to the amount kept, $E - c_i$. If the sum of contributions fails to meet the threshold, the public good is not provided and contributions are refunded at the rate r, with $0 \leq r \leq 1$. The parameter α is the marginal per capita return (MPCR) from the public good.

In EXO, participants make contribution choices for low ($T = 21$), intermediate ($T = 39$), and high ($T = 57$) thresholds. The case with $T = 0$ is a standard linear public-good game and serves as a control. The subjects make contribution choices for each of these thresholds in a randomized order. We provide no feedback about outcomes until the end of the experiment. We chose thresholds that are divisible by $N = 3$ to facilitate coordination, thus making equal contribution by all group members focal.

Each participant makes contribution choices given one refund rate, r. Refund rates vary the cost of contributing when the public good is not provided. For example, a value of $r = 0$ makes coordination failure costly because it implies that all contributions to the public goods are wasted if the threshold is not reached. In contrast, a value of $r = 1$ implies full insurance in the sense that contributions to the public good are fully refunded should the threshold not be met. Table 2 summarizes the parameters of the experiment.

For each threshold level, participants make only one contribution decision. The one-shot nature of the game serves to investigate if participants are able to solve the difficult coordination problem absent any opportunities for communication, learning, and experience. For each threshold level, subjects state their expectations on the contributions of others. The data on

Table 2
Parameters of the Experiment

	Variable	Value
Endowment	E	20
Group size	N	3
MPCR	α	0.5
Threshold	T	{0, 21, 39, 57}
Refund rate	r	{0, 0. 5, 1}

beliefs enables us to evaluate best-response behavior. Beliefs are elicited by rewarding a correct point prediction by an additional payment of 10 points. Incorrect beliefs were not rewarded even if they were close. Point incomes from all choices are converted into money and paid out at the end of the experiment according to the exchange rate of 10 points = 0.8 euros. Subjects were paid for each contribution choice.

3.2. Voting on Thresholds (END)

The END treatments are essentially the same as the EXO treatments except that participants vote on which threshold to implement before making contribution decisions. Voting is over pairs of thresholds T_L and T_H, with $T_L < T_H$. Participants vote on all six pairwise comparisons of thresholds, i.e., $T_L = 0$ versus $T_H = 21$, $T_L = 0$ versus $T_H = 39$, ..., $T_L = 39$ versus $T_H = 57$. To avoid sequence effects, we randomize the order of voting over subjects. Participants make conditional contribution decisions (i.e., according to the strategy method) for all possible outcomes of the vote. More specifically, subjects make contributions for the case that zero, one, or two of the others in the group vote for T_H. Obviously, the outcome of the referendum may depend on the subject's own vote. For example, if a subject has voted for T_H, then T_H is accepted for $H_{-i} \in \{1, 2\}$, while if the subject has voted for T_L, then T_H is accepted only for $H_{-i} = 2$. The decision screen in the experiment takes account of this fact (see Appendix).

Applying the strategy method has the important advantage that we observe choices for all contingencies, including the cases that are not implemented. In particular, we can analyze how the contribution behavior depends on the subject's own voting and on other group members' voting choices. This rich data allows us to investigate the effects of voting choices on contributions – the main purpose of the paper – in great detail. For example, it allows us to disentangle the signaling and commitment effects discussed below. However, the use of the strategy method has the disadvantage of making choices more complicated – participants make 18 (6 votes $\times$ 3 cases) contribution choices in END, compared to 4 contribution choices in EXO – and perhaps also more cognitively demanding.

In treatment END, participants state their beliefs about the contributions by others in the group after having made voting and contribution choices. Since the subjects make conditional contribution decisions, they also state beliefs conditional on all possible voting outcomes $H_{-i} \in \{0, 1, 2\}$. In addition, we ask subjects to state their beliefs regarding H_{-i}, i.e., the number of others' votes for the higher of the two thresholds, T_H.

In total, 210 (see table 1) undergraduate students from the University of Innsbruck participated in our computerized (z-Tree by Fischbacher 2007)

experiment. A session lasted approximately 45 minutes, and the average subject earned 8.8 euros.[7] Because subjects made four contribution choices in treatment EXO as compared to six in END, average payoffs were higher in END. The average earning in EXO was 7.2 euros (including 0.3 euro for rewarding correct point predictions). In END they earned 10.8 euros (including 0.6 euros for correct predictions).

4. Predictions and Hypotheses

The game without a threshold ($T = 0$) has a unique inefficient equilibrium in which all participants contribute zero, $\sum_j c_j = 0$. The threshold public-good game with $T > 0$ has multiple pure-strategy equilibria (see, e.g., Isaac et al. 1996). In addition to the inefficient free-riding equilibrium, there is a set of efficient equilibria that contains all feasible combinations of contributions along the mutual best response where the threshold is exactly met, $\sum_j c_j = T$. This set contains symmetric equilibria in which each participant contributes T/N and asymmetric equilibria in which participants contribute different amounts.[8] Because $\alpha N > 0$, equilibria involving positive contributions Pareto-dominate the zero-contribution equilibrium.[9]

Figure 1 illustrates individual best responses for the low ($T = 21$), intermediate ($T = 39$), and high ($T = 57$) thresholds as a function of the sum of contributions by others in the group. The figure is drawn for $\alpha = 0.5$ and $N = 3$. For example, at $T = 21$, if others' contributions are below 10 points, the individual cost to meet the threshold exceeds the individual benefit from the public good. A rational and self-interested subject therefore contributes zero to the public good. If others contribute between 11 and 20 points, the best response is to contribute just as many points as needed to reach the threshold.[10] For others' contributions above 20 points, the best response is to contribute zero, because the subject's contribution is not essential for

7 Subjects were recruited via e-mail. Those with experience in public-good experiments were excluded from the recruitment. The subjects were randomly assigned to treatments. We checked proper understanding of the instructions (see appendix for sample instructions) in a series of control questions. The sessions did not start before all participants answered these questions correctly.

8 Note that the equilibria are given by $\sum_j c_j = 0$ plus $\sum_j c_j = T$, with the additional restriction that none of the subjects contributes more than 10 points at $T = 21$ and 19 points at $T = 39$.

9 Note that for $r = 1$, there is a larger set of inefficient equilibria. These equilibria obtain when an efficient equilibrium is not feasible, i.e., $E < \sum_j c_j - c_i$, and subject i is indifferent between contributing zero and any positive amount to the public good.

10 Note that the figure only serves illustrative purposes. In the experiment subjects contribute integer numbers of points to the public good.

Figure 1
Best Response as Function of the Sum of others' Contributions for $T \in \{21, 39, 57\}$

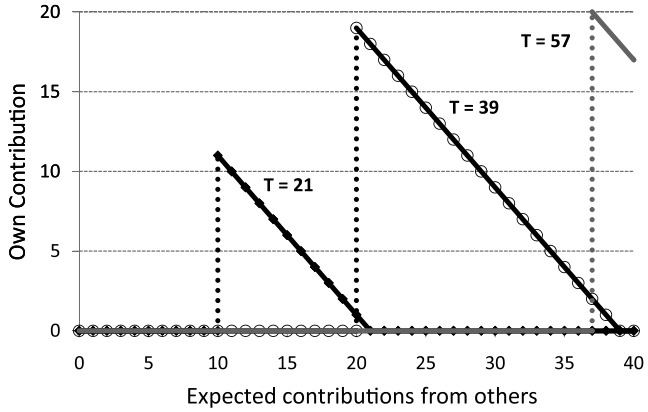

implementation of $T = 21$. The figure shows the analogous best-response functions for $T = 39$ and $T = 57$.

In treatment END, the game is solved by backward induction. If a majority in a group votes for a positive threshold rather than the zero threshold, this decision transforms a social dilemma game with a unique inefficient Nash equilibrium into a coordination game with a set of Pareto-superior equilibria. It is therefore a weakly dominant strategy to vote for a positive threshold if the alternative is a zero threshold. Matters are more complicated when the vote is between two positive thresholds $T_L > 0$ and $T_H > 0$, because (empirically debatable) assumptions about the equilibrium selection in the contribution stage of the game must be made. Unless otherwise stated, we will therefore restrict our analysis to the decisions between $T_L = 0$ and $T_H > 0$. The within-subject variation of the threshold nevertheless enables us to evaluate which threshold level is most popular and to compare this with the empirically optimal threshold.

In our one-shot design, participants cannot learn from experience, and the risk of miscoordination is therefore high. As illustrated in figure 1, higher thresholds are associated with higher critical levels triggering positive contributions by rational and self-interested players. Since $\alpha N > 1$, the set of equilibria with higher thresholds contains equilibria that Pareto-dominate all equilibria with lower thresholds. This fact may make these equilibria more focal and ease the coordination problem. At the same time, if the threshold is high, the cost of miscoordination and consequent deviation from equilibrium is also high for participants if $r < 1$. For example, a participant

who contributes 20 points when overall contributions fail to exceed $T = 57$ ends up earning zero when $r = 0$. The net effect is therefore indeterminate, and the question of which threshold level is more efficient is fundamentally empirical.

We hypothesize that voting improves efficiency by reducing the risk of miscoordination. Our hypothesis is based on three arguments. First, the number of votes for a threshold may provide a signal for others' cooperativeness. Second, voting may be determined by subjects' beliefs about others' behavior as well as their personal characteristics such as social preferences or cognitive skills. If such characteristics are relevant for the behavior in the game, voting may give rise to selection effects that influence the outcome of the game. Finally, a subject who votes in favor of a threshold may feel committed to also contribute to the successful provision of the public good. It is important to note that these arguments do not univocally support a positive effect of voting on efficiency. For instance, voters may vote strategically and send misleading signals. Moreover, depending on the expected contributions with and without a threshold, a signal of contributions from others can increase, decrease, or leave unchanged the optimal contributions (see figure 1). The effect of approving a threshold in a referendum is thus theoretically indeterminate and fundamentally an empirical issue.

5. Results

Section 5.1 presents the results for exogenous thresholds. Our main findings from this analysis are that exogenous thresholds are at best ineffective (with full insurance, i.e., $r = 1$) and frequently (in all other cases) counterproductive. Thus, exogenous thresholds do not increase efficiency as measured by the sum of participants' payoffs in our experiment.[11] The counterproductive effect on efficiency is most pronounced with the most ambitious threshold ($T = 57$). While higher thresholds tend to induce higher expectations and somewhat higher contributions, the increase in contributions is usually insufficient to reach the more ambitious threshold, as shown in section 5.2. For example, on increasing the threshold by a factor of 2.7 from $T = 21$ to $T = 57$

11 Taking the sum of payoffs as a measure of efficiency assumes that (nonrefunded) contributions are simply wasted. This is a reasonable measure of efficiency in cases where contributions to the (nonprovided) public good are sunk, as when half a bridge is built (and it certainly is the appropriate measure from the perspective of the average participant in our experiment, by virtue of our design). Alternatively, we discuss the *surplus from the public good*, which measures the efficiency of public-good provision when nonrefunded contributions are not counted as waste. This measure has the advantage, as one of the referees points out, of normalizing surplus across refund rates. According to this measure, introducing intermediate thresholds does have a weakly beneficial effect.

with partial refund ($r = 0.5$), we find that expectations increase by a factor of 1.6, and contributions only increase by a factor of 1.4, falling clearly short of the required factor 2.7 in this example. The consequence is that the success rate, i.e., the percentage of cases in which the threshold is reached, falls from 73% to 27%, and efficiency measured by average payoffs falls by 23%.

Section 5.3 shows that voting on thresholds does not improve efficiency of thresholds. Again, expectations and contributions increase with chosen thresholds, but the increase falls short of what is required to meet the more ambitious threshold. To continue the example in the previous paragraph – moving from $T = 21$ to $T = 57$ with partial refund ($r = 0.5$) – the efficiency falls in END by 19%, which is in the same ballpark as the drop observed in EXO (23%). To explore the reasons for this result, we discuss voting behavior and the effects of voting on expected and actual contributions.

5.1. Results in EXO

In EXO the main results of interest are the *success rate* at a given threshold (i.e., whether subjects manage to coordinate on the Pareto-superior equilibria of the game) and how efficiency is related to various threshold levels. In addition, we are interested in how these effects interact with the refund rate.

Figure 2 shows that average contributions monotonically increase with the threshold at all refund rates. For example, in EXO0 with no refund (upper

Figure 2
Average Contributions by Threshold and Refund Rate in Treatment EXO

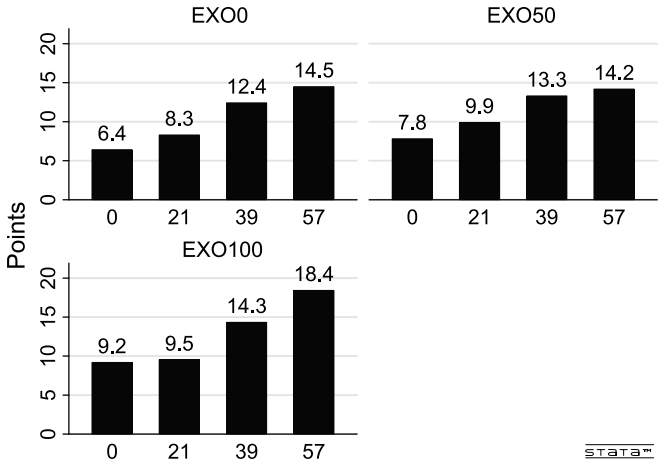

Table 3
OLS Regression Treatment EXO

Dep. Var.	(1) c_i	(2) π_i	(3) PG Surplus	(4) Belief
T21	1.060*	−0.306	0.017	0.581
	(0.582)	(0.709)	(0.025)	(1.375)
T39	1.973***	−1.517	0.102***	9.211***
	(0.716)	(0.934)	(0.034)	(1.520)
T57	2.240***	−6.984***	−0.072	14.370***
	(0.791)	(1.215)	(0.049)	(1.919)
EXO50	0.344	3.339***	0.059	1.386
	(0.787)	(1.169)	(0.055)	(1.948)
EXO100	1.640**	6.191***	0.189***	2.053
	(0.691)	(1.084)	(0.056)	(1.894)
Belief	0.396***	–	–	–
	(0.025)			
Round	−0.497***	−0.108	−0.014	−0.123
	(0.161)	(0.339)	(0.014)	(0.456)
Constant	0.824	21.073***	0.342***	18.267***
	(0.807)	(1.174)	(0.045)	(1.837)
Observations	408	408	408	408
R-squared	0.881	0.895	0.600	0.831
	$F(8, 101) = 359.60$***	$F(7, 101) = 466.06$***	$F(7, 101) = 53.53$***	$F(7, 101) = 208.81$***

Note: Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.1$

left panel of figure 2), average contributions are 6.4 without a threshold ($T = 0$), and increase to 8.3 ($T = 21$), 12.4 ($T = 39$), and 14.5 with the most ambitious threshold ($T = 57$). Figure 2 also suggests that contributions tend to be higher with a full refund (EXO100) than with partial or no refund. For example, at $T = 57$, contributions in EXO100 are 18.4 versus 14.2 in EXO50 and 14.5 in EXO0. These findings are in line with the conjecture that lower refund rates increase the cost of miscoordination, which makes the participants more reluctant to contribute.

Table 3 reports the results of an OLS regression to test the significance of these results.[12] Variables T21, T39, and T57 indicate the level of the thresh-

12 The regression uses all data from EXO, i.e., 102 subjects choosing contributions for four levels of thresholds. The regression allows for the fact that individual choices are not in-

Figure 3
Average Payoff by Threshold and Refund Rate in Treatment EXO

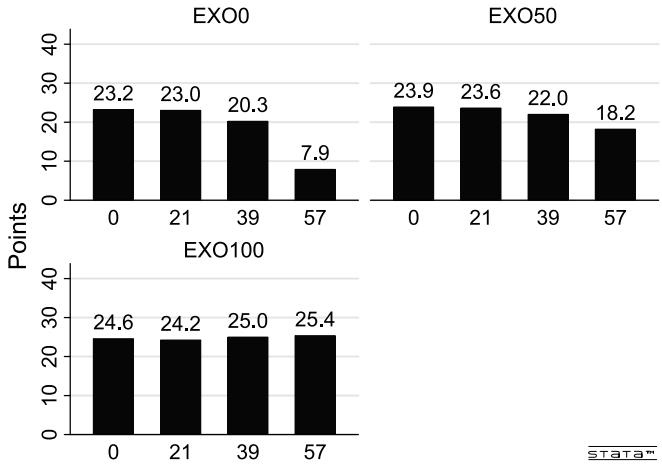

old (with $T = 0$ as the left-out category). Similarly, variables EXO50 and EXO100 capture the effect of a partial and a full refund (as compared to no refund in EXO0). The variable "Belief" stands for expected contributions by others and ranges from 0 to 40. Finally, we include a trend for the decision round to all our regressions.[13] Column (1) confirms the visual impression from figure 2. The first three coefficients show that higher thresholds significantly increase contributions, and the coefficient for EXO100 shows that a full refund increases contributions compared to no refund. In addition, we find that higher beliefs are overall strongly associated with higher contributions. For example, a player expecting others to contribute 40 points rather than 20 contributes about 8 points more on average.

Figure 3 shows the average payoff in points by threshold. The figure reveals that the payoffs tend to fall with rising thresholds when there is no (EXO0) or a partial (EXO50) refund, and this drop is particularly pronounced for the most ambitious threshold. With full refund (EXO100), average payoffs are essentially constant across thresholds, and the threshold has therefore virtually no effect. The visual impression from figure 3 is again confirmed in

dependent over multiple decision rounds by using robust standard errors adjusted for clustering of individuals. We have also estimated the effects of the thresholds by allowing for individual fixed effects. The results are robust to this modification. Note that the refund does not vary within subjects, implying that one cannot estimate the effects of different refunds with individual fixed effects.

13 Note that this variable is orthogonal to our treatment variables by design, because the order of decisions was randomized.

regression analysis. Column (2) of table 3 shows a significant negative effect of the most ambitious threshold: almost 7 points. Average payoffs are higher by about 6.2 points with full refund, which implies that overall payoffs do not drop much in EXO100 for $T = 57$, as shown in figure 3.

The average payoff is an appropriate measure of efficiency because in our design all or part of the contributions (for $r < 1$) are lost if the group fails to reach the threshold. This assumption seems reasonable for capturing real-world situations, for example, where contributions go into investments that are sunk. If we choose to ignore such sunk private cost (as has been common in some of the experimental literature), we can generate a *surplus* measure reflecting how close the provision of the public good is to the efficient level of provision. Figure 4 shows the share of the maximum surplus from the public good (30 points in our design) that was actually provided. This surplus measure clearly paints a less bleak picture. Figure 4 shows that the surplus tends to increase from $T = 0$ to $T = 39$, and to decrease again for the most ambitious threshold, independently of the refund. Column (3) of table 3 corroborates a positive effect of the threshold at level $T = 39$ for the variable "PG Surplus," but $T = 57$ is not significantly associated with negative effects. Confirming the visual impression of figure 4, regression (3) confirms that the surplus is significantly higher with a full refund (EXO100) than with no refund.

Figure 4
Average Surplus (in Percent of Maximal) of the Public Good by Threshold and Refund Rate in Treatment EXO

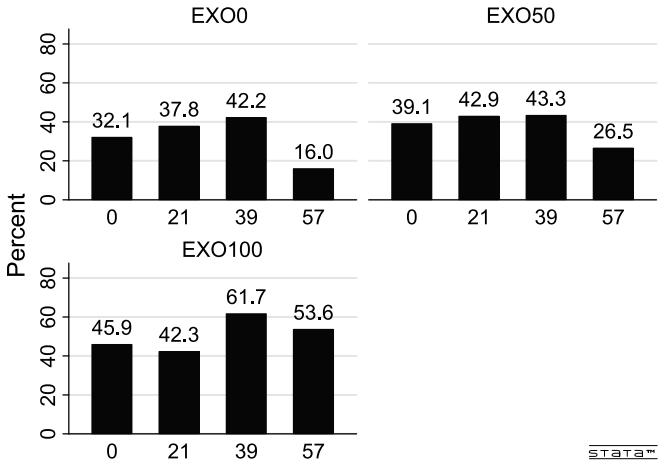

We conclude that exogenously imposing a threshold does not increase efficiency measured by overall payoffs, and ambitious thresholds are counterproductive absent full insurance in our design. The reason why thresholds fail to increase efficiency despite the fact that contributions do increase with the threshold level is that the increase is often insufficient to match the increase in the threshold. Table 4 (upper panel) shows the success rate, i.e., the share of groups who manage to reach or surpass the threshold. With zero or partial refund, low success rates are particularly costly because contributions are wasted in this case. The table shows that the average success rate decreases dramatically with the size of the threshold, whereas it clearly increases with the refund rate. Even a modest threshold under full refund does not significantly improve payoffs (or the surplus from the public good) compared to the standard linear public-good game (see EXO100 in figure 3 and figure 4) (Wilcoxon signed-rank test: $p = 0.844$ for payoffs, $p = 0.469$ for surplus).

Table 4
Success Rates by Treatment

		EXO0	EXO50	EXO100	*avg.*
Threshold	21	0.75	0.73	0.82	*0.76*
	39	0.58	0.55	0.82	*0.65*
	57	0.17	0.27	0.55	*0.32*
	avg.	*0.5*	*0.52*	*0.73*	*0.53*
		END0	END50	END100	*avg.*
Threshold	21	0.45	0.67	1.00	*0.68*
	39	0.38	0.62	0.73	*0.57*
	57	0.10	0.36	0.64	*0.46*
	avg.	*0.35*	*0.58*	*0.76*	*0.57*

5.2. Discussion of Results in EXO

Why do the subjects increase their contributions with the threshold as shown in figure 2 despite the fact that this behavior does not increase, and often reduces, their payoffs? The answer to this question comes in three parts. The first part is that the threshold level provides a signal about others' contributions. If met, a high threshold improves the efficiency of the Pareto-superior equilibria. For some of these equilibria, a subject may rationally

expect high contributions from others. At the same time, if the threshold is high, the cost of miscoordination and thus deviating from equilibrium is also high (for $r < 1$). This argument might induce lower expectations. *A priori*, the direction of the signal is therefore not clear. Column (4) of table 3 shows the results of a regression using beliefs as the dependent variable. The results show that subjects expect much higher contributions by others if the threshold is high. On average, expectations increase by about 14.4 points if the threshold is $T = 57$ rather than $T = 0$. But these expectations were excessively optimistic. Column (1) of table 3 shows that contributions only increased by about 2.2 points (see coefficient for $T = 57$).

The second part of the answer concerns the subjects' reaction given their beliefs. Table 5 shows that about a third (34% to 36%) of all subjects choose exact best responses (= BR) to their expectations. Among these subjects, a sizable fraction hold focal beliefs, i.e., beliefs about others' contribution equivalent to "cost sharing," i.e., $\frac{2}{3}T$ (see numbers in parentheses in column = BR). A majority of subjects (between 51% and 65%) contribute more than their best response (> BR). This high share of *overcontributing* subjects can be explained by three factors. First, subjects overcontribute to avert the risk of not reaching the threshold. Second, since the contributions can have positive externalities on others within the group, subjects may overcontribute due to social preferences. Finally, there may be decision errors. However, since only a few subjects contribute less than the best response (< BR), the explanatory power of unsystematic decision errors is limited.

Table 5

Number of Subjects in EXO who Contribute Exactly (= BR), less than (< BR), and more than (> BR) Best Response towards own Belief

	Threshold	= BR	< BR	> BR	#obs.
EXO0	$T = 21$	11 (6)	2 (0)	23 (6)	36
	$T = 39$	12 (7)	4 (1)	20 (2)	
	$T = 57$	16 (8)	3 (1)	27 (2)	
EXO50	$T = 21$	12 (8)	0 (0)	21 (0)	33
	$T = 39$	12 (10)	4 (1)	17 (0)	
	$T = 57$	15 (8)	2 (1)	16 (0)	
EXO100	$T = 21$	10 (7)	2 (1)	11 (1)	33
	$T = 39$	13 (11)	1 (0)	19 (0)	
	$T = 57$	11 (8)	1 (0)	21 (2)	

Note: The numbers in parentheses show the number of subjects with focal expectations (i.e., $\frac{2}{3}T$)

Figure 5

Own Contributions as a Function of Expected Contributions by others in EXO50

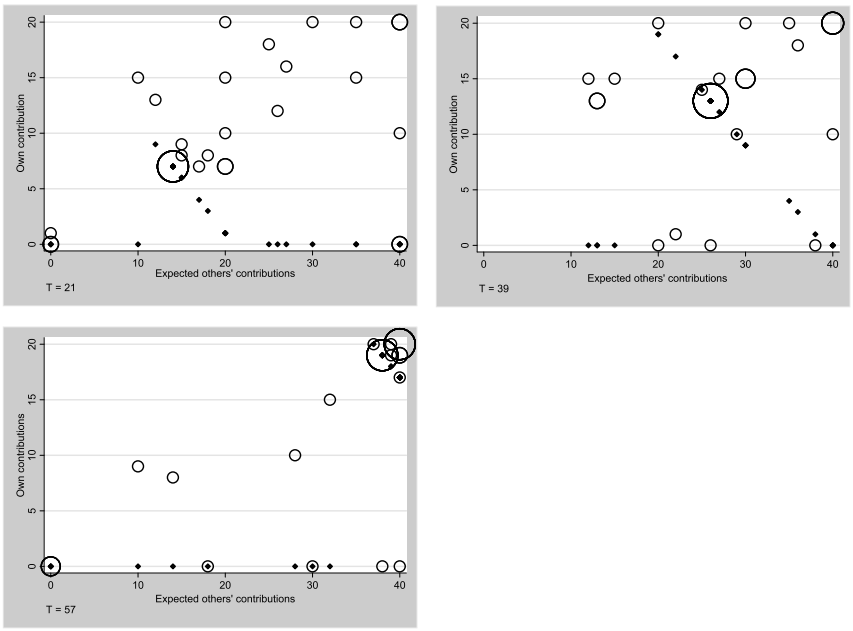

Note: Black diamonds indicate best response. Size of circles indicates the number of observations at a particular combination of contributions and beliefs

The observation that many subjects hold focal beliefs to which they choose a best response and that many subjects overcontribute given their beliefs explains why an increase in expected contributions translates into an increase in own contributions (note the strong effect of beliefs on contributions in regression (1) of table 3).

Figure 5 illustrates our discussion in the last paragraph for treatment EXO50.[14] The figure shows own contributions as a function of expected contributions by others for the low (upper left panel), intermediate (upper right), and ambitious (bottom panel) thresholds. The size of the circles indicates the frequency of observations for a particular combination of contributions and beliefs. The black diamonds show the theoretical best response. The figure serves to illustrate three regularities. First, contributions are positively associated with expected contributions. Second, as shown already in table 5,

14 The figures look quite similar for EXO0 and EXO100 (not reported here due to space constraints).

approximately one-third of the subjects choose the best response to their beliefs, and a large part of these beliefs are focal (i.e., beliefs are at $\frac{2}{3}T$). Finally, a rise in the threshold shifts up expected and actual contributions.

The third part of the answer is that, despite the positive effects of thresholds on expected contributions, coordination often fails. The most likely explanation is that expectations are imprecise and biased. The data shows a vast variation in the difference between expected and actual contributions: averaged over all thresholds and treatments, the standard deviation of this difference is 14.4 points. In addition, subjects tend to overestimate others' contributions. On average, the subjects expect others to contribute 2.1 points more than they actually do contribute ($p = 0.001$, two-sided Wilcoxon signed rank test). These figures do not come as a surprise, given the findings in the previous literature. With best-response functions that are kinked and decreasing for important ranges of expectations, coordination is difficult, especially in the one-shot game. In fact, Isaac et al. (1989) reported success rates similar to ours for the first period of the repeated game.

5.3. Results in END

This section reports the results for the endogenous treatments. We first ask which threshold is most popular. We then explore how voting affects choices and efficiency in the game.

Voting and aggregate outcomes: Table 6 shows the acceptance of the higher of the two thresholds in each pairwise vote by treatment. The numbers in parentheses show aggregate acceptance rates. In END0 and END50, the low threshold $T = 21$ is clearly the most popular, as it is the unique majority winner. In all pairwise comparisons, a majority of voters prefer $T_H = 21$ and $T_H = 39$ over $T_L = 0$, and prefer $T_L = 21$ over $T_H = 39$ and $T_H = 57$. This result suggests that the subjects anticipate the risk of miscoordination at

Table 6
Individual (Aggregate) Acceptance in Percent for T_H ($T_H > T_L$)

	END0	END50	END100
0 vs. 21	56% (58%)	56% (50%)	69% (58%)
0 vs. 39	53% (67%)	61% (58%)	72% (83%)
0 vs. 57	31% (25%)	39% (33%)	64% (75%)
21 vs. 39	44% (50%)	31% (42%)	64% (58%)
21 vs. 57	44% (42%)	31% (33%)	67% (75%)
39 vs. 57	36% (17%)	31% (25%)	56% (58%)

Table 7
OLS Regression Treatment END

Dep. Var.	(1) c_i	(2) π_i	(3) PG Surplus
T21	0.724	−1.883	0.021
	(0.802)	(1.184)	(0.047)
T39	1.962*	−3.897***	0.050
	(0.999)	(1.092)	(0.051)
T57	3.001**	−4.232***	0.025
	(1.412)	(1.459)	(0.072)
END50	−0.093	1.448	0.059
	(1.092)	(1.284)	(0.047)
END100	3.265***	6.649***	0.310***
	(1.002)	(1.330)	(0.049)
Belief	0.322***	–	–
	(0.042)		
Round	−0.376***	−0.369*	−0.024***
	(0.140)	(0.188)	(0.008)
Constant	2.032**	22.712***	0.334***
	(0.946)	(1.397)	(0.048)
Observations	324	324	324
R-squared	0.829	0.919	0.680
	$F(8, 107) = 191.69$***	$F(7, 107) = 391.25$***	$F(7, 107) = 77.87$***

Note: Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.1$

a high threshold. Indeed, when this risk is eliminated ($r = 1$), the majority winner is $T = 57$ in END100.

Despite the electoral support for positive thresholds, voting does not increase aggregate efficiency (see Dal Bo et al. (forthcoming) for efficiency-increasing effects of voting in a prisoners' dilemma game). Table 7 reports regressions for treatment END, i.e., when thresholds are endogenously chosen.[15] Column (1) reports results for determinants of contributions. These

15 The regression uses $3 \times 108 = 324$ data points. In principle, contributions may depend not only on the outcome of the vote, but also on the baseline option in the voting choice. To assure a constant default for the voting decisions, the regression uses only data for voting choices between $T_L = 0$ and $T_H > 0$. However, the overall picture remains the same if we also include the data for choices between $T_L > 0$ and $T_H > 0$.

results for END look very similar to those for EXO reported in column (1) of table 3 and are indeed not statistically different from those estimates.[16] As a consequence, there is no sizable difference between exogenous and endogenous thresholds on subjects' payoffs and the surplus from the public good (compare columns (2) and (3) across tables 7 and 3). Note, however, that the positive effect of an intermediate threshold on the surplus from the public good found for EXO disappears in END. As in EXO, the reason for a lack of positive effect of thresholds is that contributions often fall short of reaching the threshold (see table 4, lower panel). We now explore the effect of voting in more detail.

Signaling: As contributions are importantly driven by expected contributions, voting may make a difference if the voting outcome effectively signals others' contributions (see Tyran and Feld (2006) for evidence of signaling effects in voting on sanctions in public-good games). Evidence in support of a signaling effect is provided in regressions (1) and (2) in table 8. The regression uses the data that was elicited by the strategy method. In particular, it uses the decisions from the three pairwise votes between $T_L = 0$ and $T_H \in \{21, 39, 57\}$, where T_H has been accepted by the group. Therefore, variables T39 and T57 indicate the level of the threshold at $T = 39$ and $T = 57$ as compared to $T = 21$. Furthermore, to keep the aggregate voting outcome constant, the regression uses only the cases where the decision-maker plus at least one of the other subjects in the group have voted for $T_H > 0$.[17] The variable $\sum \text{vote}_j = 1$ indicates whether both of the others in the group have voted yes.

The coefficient for $\sum \text{vote}_j = 1$ in regression (1) shows that expected contributions strongly increase with others' acceptance. In particular, beliefs were on average about 3.8 points higher if the threshold was accepted with the approval of all others. The coefficient for the variable Belief in column (2) shows that contributions increase with expected contributions (and the effect is very similar to that in EXO; compare (1) in table 3), but contributions do not increase with the number of others' yes votes. Together, these estimates indicate that the expected support for a threshold provides is a signal for how much others are expected to contribute. Higher expectations, in turn, induce higher contributions.

16 To test, we estimated a single equation using the data both from EXO and END and employed a Chow test. None of the coefficients in column (1) is different between the treatments.

17 For the three pairwise votes between $T_L = 0$ and $T_H \in \{21, 39, 57\}$, we observe 145 cases in which T_H is accepted and the subject voted yes. Due to the strategy method, the subject takes two choices conditional on the number of yes-votes by others. Hence, the number of observations used in the regressions is 290.

Table 8
OLS Regression Treatment END: Signaling, Selection, and Commitment

Dep. Var.	(1) Belief	(2) c_i	(3) Belief	(4) c_i
T39	7.734***	−0.192	6.256***	0.223
	(1.253)	(0.620)	(1.167)	(0.703)
T57	15.324***	1.170	13.812***	0.977
	(1.751)	(1.131)	(1.251)	(1.084)
END50	1.444	−0.401	3.428**	−0.723
	(2.234)	(0.929)	(1.574)	(1.068)
END100	2.557	1.384*	2.425	1.877**
	(1.735)	(0.746)	(1.496)	(0.825)
Belief	–	0.383***	–	0.307***
		(0.044)		(0.048)
$\sum \text{vote}_j = 1$	3.841***	−0.293	–	–
	(0.715)	(0.373)		
$\text{vote}_i = 1$	–	–	2.157*	2.683***
			(1.112)	(0.746)
c_i PG	–	–	0.358***	0.126*
			(0.108)	(0.071)
Constant	12.124***	2.652**	14.353***	0.038
	(1.887)	(1.079)	(1.297)	(1.181)
Observations	290	290	324	324
R-squared	0.377	0.546	0.318	0.384
	$F(5, 71) = 36.33$***	$F(6, 71) = 79.54$***	$F(6, 107) = 36.40$***	$F(7, 107) = 49.33$***

Note: Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.1$

Selection and commitment: Voters who approve of a positive threshold may do so because they expect sufficiently high contributions from others. If so, voters would rationally select into a threshold regime according to their expectations. To evaluate this argument empirically, we regress subjects' beliefs on their voting behavior. Column (3) of table 8 shows the result. We use the data elicited from the strategy method analogously to columns (1) and (2) of table 8. However, to keep the aggregate voting outcome constant, we now only consider the cases where two of the other voters in the group

have voted for T_H.[18] Column (3) establishes a potential selection effect in that yes-voters expect higher contributions of others. However, the effect is significant only at the 10% level.

Selection into a threshold level may also happen due to heterogeneous social preferences. To obtain a measure of cooperativeness, we let the subjects play a one-shot standard linear public-good game prior to playing the threshold public-good game. The payoffs were determined according to $\pi_i = E - c_i + \alpha \sum_j c_j$, without mention of any threshold.[19] The subjects did not receive feedback on the outcome of this choice until the end of the entire experiment. The parameters and procedures were the same as the ones described above. Using the individual-level contributions from this game as a proxy for the subjects' cooperativeness, we find no correlation between this variable and voting.[20] We have also included this proxy for cooperativeness into regression (3) (see the variable c_iPG). The inclusion of this variable does not affect the impact of the own vote (vote$_i$).[21] To conclude: our results suggest at most modest selection effects that are not strong enough to improve the outcomes of endogenously chosen over exogenously imposed thresholds.

In addition to selecting into a threshold level, subjects may raise their contributions because they feel committed to their vote. In column (4) of table 8, we regress subjects' contributions on subjects' voting decision. The results show that yes-voters chose higher contributions than no-voters. Notice that this effect exists despite the fact that the regression allows for potential selection based on beliefs. These results indicate that subjects feel committed to their vote.

Figure 6 illustrates these effects graphically. The left panel shows that yes-voters, i.e., those who themselves approve of the high threshold, expect slightly higher contributions from others than no-voters. However, the difference in expectations is small and insignificant for $T_H = 21$ ($p = 0.990$) and $T_H = 39$ ($p = 0.252$, according to a two-side Wilcoxon rank-sum test). The only significant effect occurs for $T_H = 57$ ($p = 0.012$). At least for a high threshold, this observation indicates that subjects' approval is partially due to fact that they expect high contributions from others. The right schedule of figure 6 shows the average own contribution conditional on the subject's

18 The number of observations consequently amounts to 324, i.e., 3 pairwise voting decisions × 108 subjects in END confined to the case that two other group members vote yes.

19 Subjects' earnings from playing this game were 2.08 euros in EXO and 2.14 euros in END. This money was paid out together with the other payments at the end of the experiment.

20 A linear probability model including a constant, the refund rate, and the threshold of the subjects' contributions to the public good does not explain any variation: the estimated parameter is 0.001 at $p = 0.684$.

21 This implies that the interaction is insignificant: using the specification of regression (3), the estimate of the interaction vote$_i$ × c_iPG is −0.115 at $p = 0.529$.

Figure 6
Average Expected (left) and own Contributions (right) by Vote

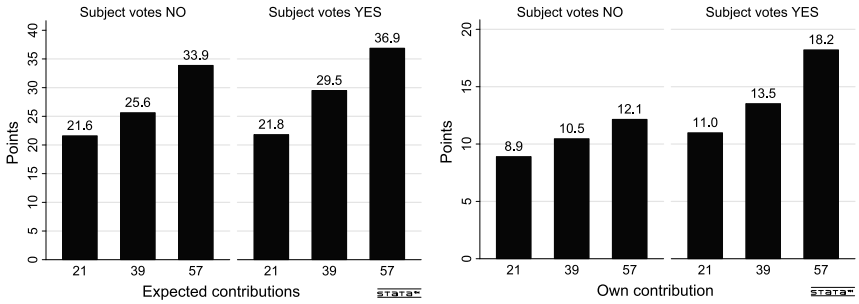

vote. It shows that yes-voters contribute more than no-voters for $T_H = 21$ ($p = 0.063$), $T_H = 39$ ($p = 0.042$), and $T_H = 57$ ($p = 0.001$). While both effects seem to play some role, the relatively large effect of the voting behavior on the own contributions suggests that the commitment effects are larger than any selection effects. Further support for this conjecture comes from analysis of best-response behavior. Presumably, subjects who choose a threshold based on their expectations are more likely to choose contributions in accordance with their best response. This link is broken if a subject raises own contributions because of a commitment with the own vote. To test, we calculate the difference between own contributions and the best response towards expected contributions from others. We find that no-voters deviate by less from their best response than yes-voters (2.76 versus 5.78 points, $p = 0.030$), thus providing further evidence for a commitment effect.

6. Discussion and Conclusion

We have studied the effect of introducing a threshold into a public-good game by voting. *A priori*, this is a promising approach when there is no predefined threshold for technological reasons and participants can be identified (and can vote), but there is no central authority with the power to enforce contributions or punish noncompliance. Examples range from the small (e.g., a fund-raising drive to start a community library only if a sufficient amount is raised to buy some minimal number of books) to the large (e.g., a voluntary agreement of nations to reach a minimum abatement of greenhouse gases).

We find that accepting a threshold has significant signaling and commitment effects, but these effects are insufficient to improve the efficiency of public-good provision in our design. In both endogenous and exogenous conditions, participants in our experiment are challenged to solve a difficult

coordination problem absent previous experience, absent opportunities for communicating or learning from mistakes within the game, and in a game that is rather complex (we use the strategy method) and context-free (we use neutral wording in the presentation of the situation to participants). These design aspects suggest that we provide a demanding test for the efficiency-improving effect of voting on thresholds. However, our design can also be considered to be favorable for coordination, since the players were symmetric, i.e., the endowments, costs of contributing, and benefits from the provision of the public good were the same for all participants, and the symmetry was common information. In naturally occurring contexts the number of players may be large and unknown, and information about important parameters of the game may be incomplete, which may further exacerbate coordination problems because contributions according to equal cost share are less focal (as suggested by Bagnoli and McKee 1991; see Rondeau et al. 1999 for a discussion). In addition, the thresholds to choose from were divisible by N in our design, which should have made equal contributions to reach the threshold focal. In most naturally occurring examples such as donations for a community library or greenhouse-gas abatement on a global level, players are not symmetric. For example, in the case of abatement, some countries are large (i.e., their emissions are large), some are poor (i.e., their opportunity cost of abatement is high), some are both, and some are neither. In this situation, coordination can be expected to be more difficult than in our experiment, because equal contributions are neither focal nor fair.

In summary, we provide the first study to show that introducing thresholds, either imposed or approved in a referendum, is no sure cure for the inefficient voluntary provision of public goods when no other, more intervention-intense mechanism is available. An interesting avenue for further research is to investigate the robustness of this result by adding more context (e.g., in a field experiment), allowing for learning (e.g., in a repeated laboratory experiment), making contributions sequentially, or adding opportunities for communication.

7. Appendix: Instructions for Treatment ''END''

(Instructions and screenshots are translated from German. Instructions for EXO are available from the authors on request.)

Welcome to the experiment. If you read the instructions carefully and follow the rules, you can earn money. Your earnings not only depend on your decisions, but also on those of the other participants. Your earnings will be paid to you in cash immediately after the experiment. During the experiment you earn points. These points will be converted to euros according to the following exchange rate: **10 Points = 80 cents (0.8 euros)**.

This experiment has seven phases. We now explain phase 0. For phases 1 to 6 you will get additional instructions at the end of phase 0.
In this experiment, the participants are randomly divided into groups of three, i.e., you are in a group with two other participants. The groups remain the same during the whole experiment, i.e., the groups have the same members in all phases.
At the beginning of each phase, every member is endowed with 20 points. Your task is to decide how many of your 20 points you want to contribute to a project and how many you want to keep for yourself. Likewise, every other group member has to decide how to use his endowment.
Your income comes from two sources: (1) the points you keep for yourself. (2) the income from the project. This income is calculated as follows: **Your income from the project** $= 0.5 \times$ **sum of the contributions of all group members to the project**.
The income of the other members is calculated in the same way, i.e., each group member gets the same income from the project. If, for example, the sum of contributions is 60 points, you and the other group members earn an income from the project of $0.5 \times 60 = 30$ points. If you and the other two group members together contribute only 10 points to the project, you and the other group members each get $0.5 \times 10 = 5$ points of income from the project.
For every point you keep for yourself, you earn an income of 1 point. If you add this point to the project instead, the sum of contributions to the project will increase by 1 point and your income from the project increases by $0.5 \times 1 = 0.5$ points. However, the income of all the other group members also increases by 0.5 points, so that the entire income from the project increases by $0.5 \times 3 = 1.5$ points. Thus, your contributions to the project generate income for the other members of your group. Conversely, you earn money for every point the other group members add to the project. For every point some other member of your group adds to the project, you get $0.5 \times 1 = 0.5$ points.
You make your decision on the computer. At the beginning of phase 0 the following decision monitor is displayed:
In the upper right corner you see how much time is left for you to decide. In the row in the middle you enter *your contribution* to the project. Please enter a number from 0 to 20 here. In the row below, we ask you to state your *expected contributions from others*. Please enter a number from 0 to 40 here. We ask you to think carefully about how much you expect the others to contribute. If your expectation matches the actual contributions from others, you will earn 10 points in addition to your other payments.
Once you have entered your decisions, we ask you to click on the gray button labeled **"Calculate"**. As a decision aid, the computer then calculates the outcomes given the numbers you (hypothetically) typed. In particular, the computer calculates the income for you and the others, given your contribution and assuming your estimated contributions from others are correct. You can change your input and push the "Calculate" button as many times as you like.

Example:

- Your contribution is 10 points. You expect the contributions by others to be 30 points. In this case your calculated income will be 30 points, and the expected average income of the others will be 25 points.
- Your contribution is 0 points. You expect the contributions of the others to be 30 points again. Your calculated income is now 35 points, and the expected average income of the others is 20 points.

You confirm your final decision by pushing the **OK button**. You will be informed about your income only at the end of the experiment, after phase 6 is completed.
If you have any question, please raise your hand, and an experimenter will come to you and offer help.

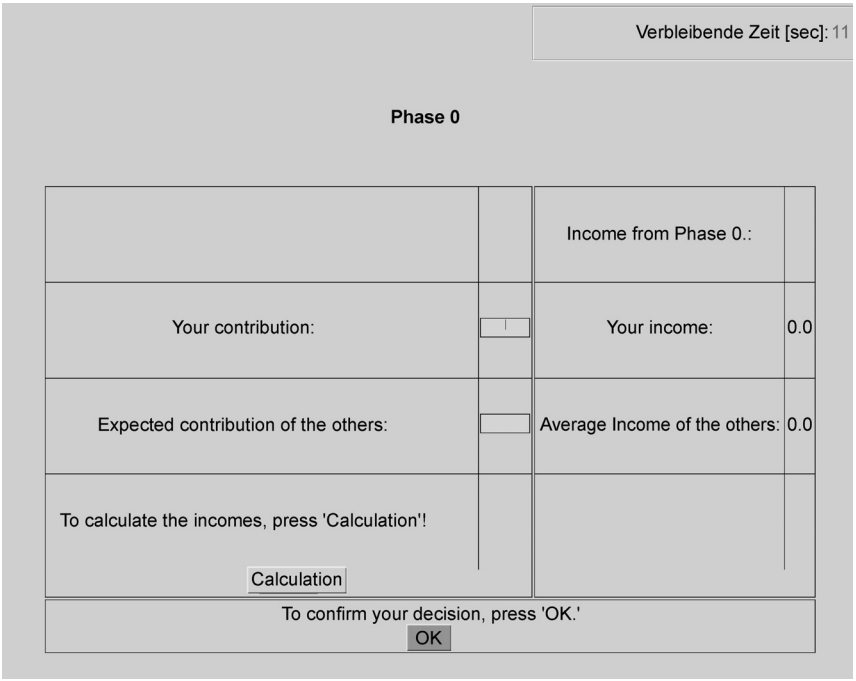

Instructions for Phases 1 to 6

As in phase 0, in phases 1 to 6 you decide how many of your 20 points you want to contribute to a project and how many you want to keep for yourself. As in phase 0, your income consists of the income from the project and the points that you keep for yourself.

Unlike in phase 0, however, in phases 1 to 6 you earn an income from the project only if the sum of the contributions of all group members exceeds a **minimal project size**. The minimal project size is a number. Specifically, this number will be 0, 21, 39, or 57 points, depending on the phase you are in.

1. What happens if the sum of contributions of all group members is **smaller than** the minimal project size? In this case you and the other group members **do not get an income from the project.** You will be refunded **50% [0%, 100%]** of your contributions to the project.
2. What happens if the sum of the contributions of all group members is **equal to or larger than** the minimal project size? In this case you and the other members of the group **get an income from the project**. You **will not** be refunded your contributions to the project. (Note: As in phase 0, your income from the project is 0.5 × sum of the contributions of all group members to the project.)

Example: Assume that the minimal project size was 39 points. The following table *(not shown)* illustrates some examples.

Example, row I: You contribute 13 points and the others contribute 26 point to the project. In this example the minimal project size is exactly reached. You earn an income of 19.5 (39 × 0.5) points from the project, plus 7 points, which you have kept for yourself. Thus your income is 26.5 points.

Example, row II: You contribute 12 points and the others contribute 26 points to the project. In this example the minimal project size is not reached. Thus you do not earn income from the project. **0% [50%, 100%]** of your contribution is refunded. This leads to an income of 14 [8, 20] points.
Example, row III: You contribute 20 points and the others contribute 26 points to the project. 7 points more than the minimal project size are contributed. You earn an income of 23 (46 × 0.5) points from the project and have not kept points for yourself. Thus your income is 23 points.
Example, row IV: You contribute 13 points and the others contribute 20 points to the project. The minimal project size is not reached. Note that while your contribution is as big as in the example in row I, your income just amounts to 13.5 [7, 20] points.

Voting on the minimal project size

You and the other group members **choose the minimal project size** relevant to your decision. At the beginning of each phase you and the other two members vote between two alternatives A and B. Every alternative stands for a minimal project size. The alternative that receives the majority of votes (two or more) is realized and determines the minimal project size for the relevant phase.
You cast your vote on the computer. At the beginning of each phase the following decision screen is displayed:

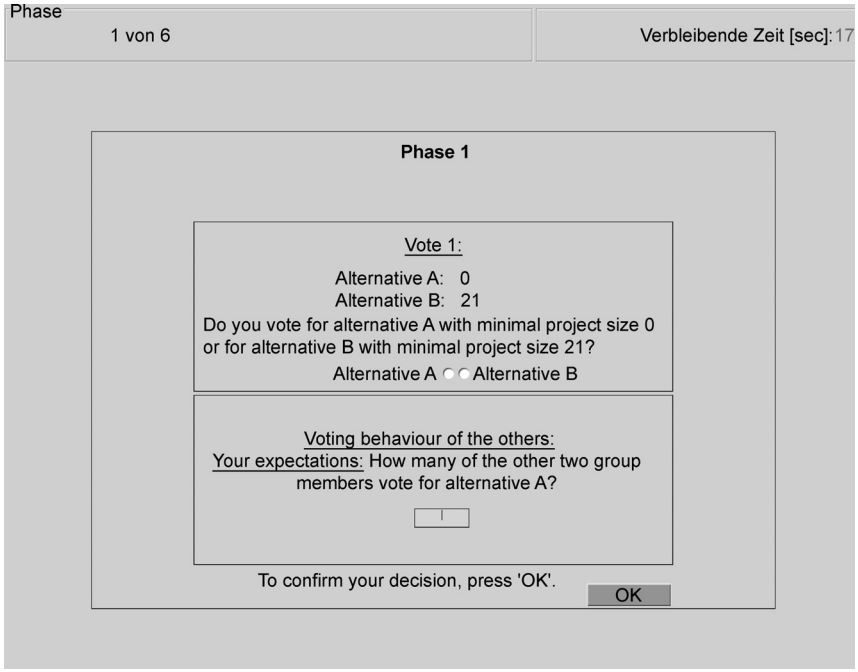

On the upper panel, you choose between alternatives A and B. In the example shown on the screen, alternative A is a minimal project size of 0 points, and alternative B is a minimal

project size of 21 points. That means that the vote is between a minimal project size of 0 and one of 21 points. Choose your preferred alternative.
In the panel below we ask you to state your **expectation** of the voting behavior of the other two group members. Please enter whether you expect 0, 1, or 2 other group members to vote for alternative A. Again, we ask you to think carefully about your decision. If you correctly predict the actual voting behavior of the others, you earn an additional 10 points. After the vote you choose your contribution on the following decision screen:

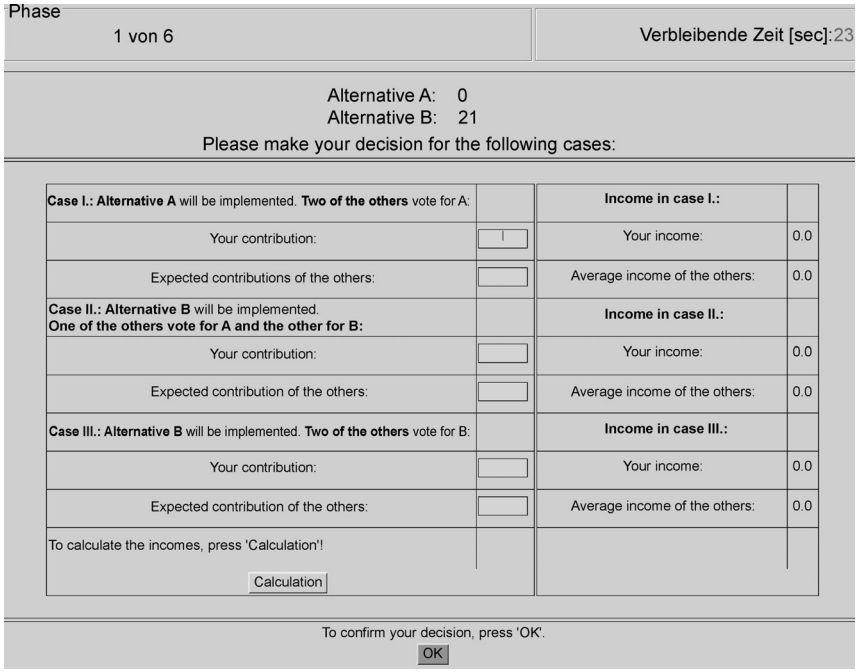

In this example the minimal project sizes are 0 for alternative A and 21 for alternative B. We ask you to take your decision for three cases:
Case 1: You choose given that both other members vote for alternative A, and hence – in this example – a minimal project size of 0 (alternative A) is implemented.
Case 2: You choose given that one of the other group members votes for alternative A and one for alternative B. Hence – in this example – a minimal project size of 0 is implemented if you voted for A, and a minimal project size of 21 is implemented if you voted for B. (Note that the screen shows the outcome assuming that you voted for alternative B.)
Case 3: You make your decision, given that both of the other group members vote for alternative B, and hence – in this example – a minimal project size of 21 (alternative B) is implemented.
Depending on the voting behavior of the others, only one of these cases will occur. Your income therefore is determined by the case that occurs. The other two cases are not relevant for your income. At the time of your decision you do not know which case will occur; we therefore ask you in your own interest to treat all three cases as if they actually occur.
For each of the three cases you take two decisions:

1. In the row "Your contribution" you enter how many points you want to contribute to the project. Please enter a number between 0 and 20 here.
2. In the row "Expected contribution of the others" we ask you to state your expectation of how much the other two group members will contribute to the group project. If your expectation is correct, you earn an additional income of 10 points.

Note: When you take your decision, please keep in mind the relevant minimal project size! Please further keep in mind that 50% [0%, 100%] of your contribution to the project will be rebated if the minimal project size was not reached.

As in phase 0, pressing the "Calculate" button will calculate your income and the average income of the others, based on your contributions and given that your expected contributions of others are correct. To confirm your final decision, press the "OK" button. Then, the next phase will start, which is essentially the same as phase 1 but with different alternatives A and B.

You will learn the outcome of the voting and your income only at the end of phase 6. Before that you will not receive any information about the contributions or voting outcomes for a particular phase.

The experiment will start soon. If you now have any questions, please raise your hand and wait until an experimenter comes by.

References

Allison, S. T., McQueen, L. R., and Schaerfl, L. M. (1992), Social Decision Making Processes and the Equal Partitionment of Shared Resources, Journal of Experimental Social Psychology 28, 23–42.

Bagnoli, M. and Lipman, B. (1989), Provision of Public Goods: Fully Implementing the Core through Private Contributions, Review of Economic Studies 56, 583–601.

Bagnoli, M. and McKee, M. (1991), Voluntary Contribution Games: Efficient Private Provision of Public Goods, Economic Inquiry 29, 351–366.

Cadsby, C. and Maynes, E. (1999), Voluntary Provision of Threshold Public Goods with Continuous Contributions: Experimental Evidence, Journal of Public Economics 71, 53–73.

Coats, J., Grosskopf, B., and Gronberg, T. J. (2009), Simultaneous versus Sequential Public Good Provision and the Role of Refunds: An Experimental Study, Journal of Public Economics 93, 326–335.

Croson, R. and Marks, M. (2000), Step Returns in Threshold Public Goods: A Meta- and Experimental Analysis, Experimental Economics 2, 239–259.

Dal Bo, P., Foster, A., and Putterman, L. (forthcoming), Institutions and Behavior: Experimental Evidence on the Effects of Democracy, American Economic Review.

Falk, A. and Heckman, J. J. (2009), Lab Experiments Are a Major Source of Knowledge in the Social Sciences, Science 326, 535–538.

Fischbacher, U. (2007), Z-Tree: Zurich Toolbox for Readymade Economic Experiments, Experimental Economics 10, 171–178.

Gerber, A. and Wichardt, P. C. (2008), Providing Public Goods in the Absence of Strong Enforcement Institutions, University of Zurich, unpublished paper.

Isaac, M., Schmidtz, D., and Walker, J. (1989), The Assurance Problem in a Laboratory Market, Public Choice 62, 217–236.

List, J. A. and Lucking-Reiley, D. (2002), The Effects of Seed Money and Refunds on Charitable Giving: Experimental Evidence from a University Capital Campaign, Journal of Political Economy 110, 215–233.

List, J. A. and Rondeau, D. (2003), The Impact of Challenge Gifts on Charitable Giving: An Experimental Investigation, Economics Letters 79, 153–159.

Marks, M. and Croson, R. (1998), Alternative Rebate Rules in the Provision of a Threshold Public Good: An Experimental Investigation, Journal of Public Economics 67, 195–220.

Rondeau, D. and List, J. A. (2008), Matching and Challenge Gifts to Charity: Evidence from Laboratory and Natural Field Experiments, IZA discussion paper No. 3278, January 2008.

Rondeau, D., Poe, G. L., and Schulze W. D. (1999), Voluntary Revelation of the Demand for Public Goods Using a Provision Point Mechanism, Journal of Public Economics 24, 131–155.

Rondeau, D., Poe, G. L., and Schulze W. D. (2005), VCM or PPM? A Comparison of the Performance of Two Voluntary Public Goods Mechanisms, Journal of Public Economics 89, 1581–1592.

Spencer, M. A., Swallow, S. K., Shogren, J. F., and List, J. A. (2008), Rebate Rules in Threshold Public Good Provision, NBER Working Paper No. 14559, December 2008.

Suleiman, R., Budescu, D. V., and Rapoport, A. (2001), Provision of Step-Level Public Goods with Uncertain Provision Threshold and Continuous Contribution, Group Decision and Negotiation 10, 253–274.

Tyran, J.-R. and Feld, L. (2006), Achieving Compliance when Legal Sanctions are Non-Deterrent, Scandinavian Journal of Economics 108, 135–156.

Jean-Robert Tyran
University of Copenhagen
Department of Economics
Building 26
DK-1353 Copenhagen K
Denmark
jean-robert.tyran@econ.ku.dk

Behavioral and Experimental Economics Do Inform Public Policy

Arno Riedl*

Received 22 December 2009; in revised form 28 February 2010; accepted 28 February 2010

Experimental and behavioral economics are well-established branches of economic science. This essay presents and discusses some results and behavioral regularities from these fields, which are of potential and actual importance for public policy. After a brief introduction to what experimental and behavioral economics are, some important behavioral regularities – presentation and framing effects, prosocial behavior, and reciprocity – are introduced, and it is reported how they interact with prominent trading institutions, taxation, and social and individual well-being. Throughout, some implications for public policy are discussed.

Keywords: public economics, experimental economics, behavioral economics, public policy

JEL classification: C 90, D 00, H 00

After consulting my advisory board of experimental and behavioral economists, I am confident that the reframing proposed in the new public policy program will increase subjective well-being by 34 percent and prosocial orientation by 27 percent at almost no cost.

Taken from Amir et al. (2005) (rephrased by the author)

1. Introduction

Despite the largely unpredicted recent financial crises and accompanied economic downturn, most, if not all, recent public policy choices still rely on the traditional economic concept of rational economic man and woman, the

* This paper is based on my keynote lecture "Sociality and Institutions" presented at the workshop on Behavioral Public Economics at the Venice Summer Institute 2009 and on my inaugural lecture "Facts and Fiction in Public Policy: How Behavioral and Experimental Economics Can Inform Public Policy" at Maastricht University. I owe thanks to the participants in the Workshop on Behavioral Public Economics at the Venice Summer Institute 2009 and to the organizers of that workshop for their hospitality. I also thank an anonymous referee and the editor for their helpful and valuable comments on an earlier version of this paper, which was circulated under the title "Behavioral and Experimental Economics Can Inform Public Policy: Some Thoughts."

ISSN 0015-2218
doi: 10.1628/001522110X503389

homo economicus. Interesting examples in this respect are recent reforms in the Netherlands concerning the health insurance market and the markets for the supply of gas and electricity. In all cases important principles of the reforms were (i) more competition between suppliers, and (ii) more choice possibilities for consumers. The main arguments in favor of such reforms are based on the received wisdom among economists that increased competition and an enhanced choice set for consumers ultimately increase (consumer) welfare.

To a large extent the presumed positive outcomes rely on the assumption that consumers will exercise their power to choose. For the energy market reform in the Netherlands, however, the expectations about the exercise of consumer power have been largely disappointed, at least if one measures this power by the fraction of consumers switching supplier; for only a negligible fraction of consumers switched to an alternative electricity or gas provider. After the reform of the health insurance sector, though, consumers did seem to take up their power to choose. It was reported that with the introduction of the new health insurance system, in total 21% of consumers have changed their insurance company (de Jong et al., 2008). However, even this switching behavior looks less impressive if one takes into account that a majority of those individuals who switched their insurance company did this within a so-called collective agreement where the employer (or another collective) and not the individual chose the insurance company. In all, only about 9% of all switchers decided to do so on an individual basis. Furthermore, in 2008 the percentage of people switching their health insurance provider decreased to a mere 4% (NIVEL, 2009).

However, any well-trained economist will argue that (non)switching behavior *per se* does not yet mean that consumers did not make the correct choices. Indeed, revealed-preference theory states that those who did not switch simply reveal that they had already chosen their utility-maximizing insurance package and/or energy supplier and hence had no reason to switch.

But let us be a little bit skeptical and ask if there is any way to assess if consumers indeed made good choices. Unfortunately, there is no study (at least that I am aware of) that investigates this question for the mentioned recent reforms in the Netherlands in a systematic way. Yet, an evaluation of a similar reform of the Swedish social security system in 2000 may help us to get some clues about how such reforms may work out and whether consumers indeed make the good choices traditional economic theory assumes. Cronqvist and Thaler (2004)[1] investigated consumer choice behavior after the introduction of the new system. In this system participants are allowed to choose their own portfolios, but there is also one default fund that is selected

1 See also Thaler and Sunstein (2008, ch. 9).

automatically for those who do not actively choose. The authors asked if, compared to the default, active choices are better choices. They conclude that "it would be hard to make the case on an *ex ante* basis that the actively selected portfolios were better than the default fund" (Cronqvist and Thaler, 2004, p. 427).[2] In addition, in those first three years, and indeed up to 2007, the actively chosen funds also did worse in terms of returns (see Thaler and Sunstein, 2008, p. 427). Another interesting observation is that since the Swedish government reduced the campaigns advertising active choice, most people (90%) have opted for the default fund and almost nobody has made any changes to the chosen portfolio or switched the chosen fund.[3]

This study strongly suggests that consumers do not (always) make the wise choices traditional economic models assume. However, there are too many unobservables (e.g., risk preferences, self-selection effects) that may influence behavior, and one may well beware of drawing too strong conclusions from this field evidence. Therefore, in the remainder of this contribution I shall present "clean" evidence that standard assumptions of economic models are indeed often violated and argue that neglecting the observed nonstandard behavioral regularities will lead to wrong predictions and worse public policy than necessary.

The plan of the rest of paper is as follows. First, I shall briefly discuss the scope of experimental and behavioral economics. Then I shall present important examples of violations of standard behavioral assumptions, based on questionnaire studies and laboratory experiments. Thereafter, I shall link up these observations with questions concerning public economics and public policy. Finally, I shall present some ideas about interesting and important further research directions.

2. The Coherence of Experimental and Behavioral Economics

2.1. Experimental Economics

The principle of science, the definition almost, is the following: The test of all knowledge is experiment. Experiment is the sole judge of scientific "truth."
Feynman (1964, chapter 1)

What precisely is experimental economics? As the name suggests, it is the branch of economics that uses experiments to investigate human behavior

2 In particular, in comparison with the default fund, actively chosen portfolios contained a higher equity exposure and much more local concentration (e.g., almost 50% of the equities are from Swedish firms), required more active management, and had higher fees.

3 When the system was introduced in 2000, two-thirds of participants actively selected a portfolio on their own. Interestingly, the proportion of people actively choosing their own portfolio decreased to 17.6% in 2001 and to only 8% in 2006.

in economic decision situations. Experimental economics is a method that brings real people to the laboratory (or the field), where they make real choices with which they earn (or lose) real money (Plott, 1982). An important feature of the method of economic experimentation, which distinguishes it from traditional empirical economic research, is that experiments allow the researcher to tightly control the environment in which people make choices. The controllable components of this environment comprise technologies, initial endowments, action spaces, timing of actions, accessible information, context, and – to some limited, but important, extent – also preferences. Varying these elements in a controlled way allows *ceteris paribus* inference and the isolation of true causes of change in human behavior to an extent unattainable by other methods of investigation. Additionally, laboratory experiments can be replicated by other researchers under the same or different conditions, thereby assessing the robustness of obtained results.

As economists, however, we know that there is no such thing as a free lunch. An often raised concern about the experimental method is the presumed lack of external validity of the obtained results. This is indeed a serious concern, in particular, when one aims at using the experimental method for informing public policy. I shall therefore come back to this issue at the end of the paper.

There are two hard and fast principles that experimental economists subscribe to and that also differentiate economic experiments from most experiments in psychology and marketing. Firstly, in economic experiments the monetary earnings subjects receive depend in a transparent way on the choices they make. The reason for the application of this principle is that it is one story to merely tell what one would do in a particular situation, but another story to actually take a particular action if it is linked with monetary consequences.[4] Secondly, deception of subjects is effectively banned – the main reason being that once deception is used it is likely that the news of it will leak out. Subsequently the knowledge of being deceived will spread through the subject population, which seriously undermines an important advantage of experiments, namely having control over the information and knowledge subjects have concerning the economic situation they are in. For a discussion of the effects and costs of using deception in experimental research see, e.g., Ortmann and Hertwig (2002) and Jamison et al. (2008).

2.2. Behavioral Economics

But also needed is imagination to create from these hints the great generalizations to guess at the wonderful, simple, but very strange patterns beneath them

4 For evidence that monetary incentives indeed make a difference, see, e.g., Camerer and Hogarth (1999) and Forsythe et al. (1994).

all, and then to experiment to check again whether we have made the right guess.
Feynman (1964, chapter 1)

Behavioral economics is not a synonym for experimental economics, although they share some common ground. Initially behavioral economics strongly relied on empirical evidence generated in psychological and economic experiments. Nowadays, behavioral economics is defined more broadly as an approach incorporating evidence – not necessarily experimental – from psychology and other disciplines to explore the limits of existing models of behavior and create new, parsimonious ones that can explain actual behavior in a better way than current models are able to.

Importantly, behavioral economics does not abandon the disciplining strict formality that distinguishes traditional theoretical economic modeling from "softer" approaches in some other social sciences. It also does not try to fit a new model for each new behavioral anomaly or regularity, but rather seeks parsimonious models and themes that can be applied to many different domains. Behavioral economics is also not another subdiscipline, next to labor economics, public economics, and the like, but understands itself as a modeling approach that should be applicable to a wide range of economic questions, the ultimate aim being "generating theoretical insights, making better predictions..., and suggesting better policy" (Camerer and Loewenstein, 2004, p. 3). In particular, the last of these is also theme of this contribution.

3. Examples of Behavioral Regularities

3.1. Presentation and Framing Effects

Traditional economic reasoning is usually silent about possible effects of the presentation or framing of a decision situation. For instance, from the viewpoint of revealed-preference theory it simply does not matter whether one has to make choices in sequence or simultaneously. Rational economic man and woman will always choose according to their true preferences. These preferences are assumed to be well-behaved, coherent, and invariant with respect to superficial variations in the way a choice problem is presented. However, one might start to wonder then why grocery and other stores often price their products at 9.90, 19.99, and the like – just a little bit below a round number. Is this just coincidence – does it happen that the true marginal costs of all these products are exactly these prices? Probably not.

Let me pose the problem more concretely by discussing an example adopted from Simonson (1990), who was among the first to demonstrate with the help of an experiment that it may greatly matter whether consumers have

to choose from an array of products simultaneously or sequentially. Suppose one is entering a grocery store today and this store offers a choice one snack out of six different brands of snacks for free.[5] The same will happen one week from today and two weeks from today. Hence, in each subsequent week one is free to choose the most preferred snack for free. Now, consider another grocery store that also offers free snacks for the next three weeks, again one per week. However, in this store the rule for choosing the free snack is slightly different. Here one has to choose today which brand of snack one would like to receive today, which one to receive in one week, and which one to receive in two weeks. Thus, the only difference is that in the second store one chooses today for today and the next two times, whereas in the first store one decides in each week on the spot.

When facing such offers, for a homo economicus, doing the right thing is a pretty straightforward decision. Just choose the most preferred brand. Note that this does not mean that one is going to choose the same brand for each of the three weeks, because one might like variety in snacks. Importantly, however, whether the decision has to be made simultaneously or sequentially should not make a difference. Consequently, it should not matter if one is confronted with the choice sequentially three weeks in a row as in the first store, or if one has to choose at once for all three weeks as in the second store. One might not choose the same snack in each week, but the variety of snacks one chooses should not differ under the two conditions. Now, the question asked was whether real consumers act in this way. At least, students in a laboratory study (Simonson, 1990) did not do so. In the sequential choice (grocery store 1) only 9% chose a different snack in each week, whereas in the simultaneous choice for sequential consumption (grocery store 2), this was the case for 64% of participants. These are by no means small differences, and even a skeptic should be ready to admit that these results are hard to reconcile with the assumption of stable and/or coherent preferences. In a follow-up study Simonson and Winer (1992) corroborated the laboratory findings in the field by using scanner data on actual yogurt purchases in a grocery store. They found that the variety of flavors chosen significantly increases with the number of purchases per occasion. The observed choices in the laboratory and the store strongly indicate that revealed preferences systematically depend on the way the choice set is presented, a dependence not accounted for in traditional economic models.[6]

5 If one does not like snacks at all, one can imagine brands of other products, e.g., beer or yogurt.

6 These observations are consistent with the concept of *choice bracketing* and its specific consequence of *taste change* (Read et al., 1999). The former refers to the fact that the way people make decisions, narrowly or broadly, affects their choices. The latter refers, specifically, to the effect that the choice people make today can change their tastes and, hence,

A second prominent example, which probably casts even more doubt on the assumption of coherent and stable preferences, is the famous study, known as the *Asian disease* problem, by Tversky and Kahneman (1981). Tversky and Kahneman conducted questionnaire studies with students (at Stanford University and the University of British Columbia) where they asked them to indicate their preference concerning different programs proposed to combat an unusual Asian disease. To study potential framing and presentation effects, the problem was presented in two economically equivalent but presentationally different formulations. Figure 1 reproduces the original text.

Figure 1
The Asian Disease Problem of Tversky and Kahneman (1981)

Imagine that the U.S. is preparing for the outbreak of an unusual Asian disease, which is expected to kill 600 people. Two alternative programs to combat the disease have been proposed. Assume that the exact scientific estimate of the consequences of the programs are as follows:

Problem 1: [$N = 152$]

If Program **A** is adopted, 200 people will be saved.	If Program **B** is adopted, there is 1/3 probability that 600 people will be saved, and 2/3 probability that no people will be saved.

Which of the two programs would you favor?

Problem 2: [$N = 155$]

If Program **A** (C) is adopted, 400 people will die.	If Program **B** (D) is adopted, there is 1/3 probability that nobody will die, and 2/3 probability that 600 people will die.

Which of the two programs would you favor?

Problem 1 presents the decision situation in a *positive* frame by emphasizing that lives can be *saved*. Problem 2, in contrast, presents the very same options in a *negative* frame by emphasizing that some people will have to *die*. Obviously, Programs A in Problems 1 and 2 are identical, since in both cases 200 people will be saved and 400 people will die for sure. The same is true

influence their choices in the future. The emergent property of *diversification bias* seems to be a robust phenomenon and is replicated in several studies (see Read et al., 1999, p. 178).

for Programs B, where in both problems 200 people will be saved and 400 people will have to die, in expectations. Hence, whatever people like more, there should be no significant difference in revealed preferences between the two problems.

Figure 2 depicts the frequency of actual choices. When the subjects are confronted with Problem 1 (positive frame), an overwhelming majority of 72% opt for program A, which saves 200 people for sure, whereas only 28% opt for the risky program where 200 people are saved only in expectations. When they are confronted with Problem 2, a dramatic shift in revealed preferences occurs. Now, only a minority of 22% go for the sure outcome of 400 dead people, but 78% are ready to accept the risky choice where 400 people die only in expectations. It should be obvious that such a strong framing effect, effectively inducing revealed-preference reversal, is hard to square with the assumption of coherent and stable preferences.

Figure 2
Actual Choices in Asian Disease Problems (Tversky and Kahneman, 1981)

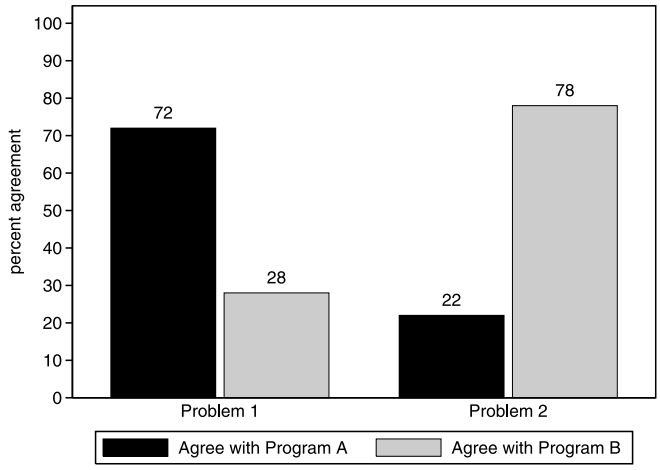

The offered options in Tversky and Kahneman's study are both rather undesirable, in the sense that one has to choose between sacrificing more people and sacrificing less people. Unfortunately, this makes them representative of many decisions made in the public domain. Consider, for instance, public investment decisions, especially, investment in infrastructure concerning safety. The decision not to invest in more secure highways or railway infrastructure means to effectively decide to accept deaths that otherwise could have been avoided. Similarly, not investing in research for an influenza vaccine means

risking avoidable deaths in case of an outbreak of influenza. Many more examples could be given, from airport security to school crossing guards ("lollipop men").[7]

...in situations where self-interest and ethical values with wide verbal allegiance are in conflict. Much of the time, most of the time in fact, the self-interest theory...will win.
Stigler (1980, p. 176)

3.2. Morality and Reciprocity

Although the neoclassical concept of utility is broad and flexible (critics might even say tautological and hence without content) and, therefore, in principle not restricted to narrow selfish preference orderings, most work and almost all applications in (public) economics assume that people are narrow-minded selfish material-wealth maximizers. One might argue that, as long as models based on the assumption of narrow self-interest describe behavior of real people sufficiently well and, hence, make correct predictions that can be used to evaluate and validate public policy, there seems to be no reason to abandon such models. This is precisely the argument endorsed by Milton Friedman, who argues that theories should be judged by the accuracy of their predictions but not by the accuracy of their assumptions Camerer (2005). However, if not only the basic assumptions are counterintuitive (and empirically proven to be wrong) but also the models' predictions are incorrect or at least misleading, then this approach is in deep trouble.

In the following I shall describe two prominent simple games where models based on narrow material self-interest turned out to do a bad job in predicting actual behavior. These examples will show that fairness considerations and, in particular, an inclination towards reciprocal behavior are important constituents of human behavior. I shall distinguish between *negative reciprocity* and *positive reciprocity*. Negative reciprocity describes the tendency to respond to an unkind act with an unkind act, whereas positive reciprocity describes a kind response to a kind course of action.[8] The following examples will make these differences clear.

7 A nice and rather harmless example where the management of a semipublic enterprise seems to take framing effects into account is the following: recently the Dutch railways have stopped using the word "delay" when announcing that a train will arrive late at the railway station. Instead, it is announced that "the train will arrive in a few minutes."

8 In a sense positive and negative reciprocity may be viewed as just two sides of the same medal, since a non-unkind act is obviously a kind act. However, different emotions may be involved (e.g., anger versus joy) with negative and positive reciprocity, which is likely to make the responses psychologically and physiologically different. Additionally, reference points of fairness are important for the judgment of kind and unkind behavior.

3.3. Negative Reciprocity – the (Mini) Ultimatum Game

The ultimatum game by Güth et al. (1982) can be interpreted as a negotiation or bargaining situation that is stripped down to its most important constituents. It is a situation involving two people where one individual can make a take-it-or-leave-it offer to the other individual, who can, indeed, take it or leave it.[9] Figure 3 depicts – for the sake of the argument – an even further boiled-down version of the originally investigated ultimatum game (adapted from Falk et al., 2003). There one player, say Peter, has received 10 euros, which he has to split between himself and, say Rita, in a take-it-or-leave-it way. If Rita accepts the offer, both receive money according to the proposal. If she rejects it, then neither Peter nor Rita receives any money. For simplicity, Peter is given only two possible ways to split the money. He can make a rather unkind offer, "I take 8 and you get 2," leaving most of the money for himself, or he can decide to be kind and propose to split the money evenly, "I take 5 and you get 5," Rita, faced with one of these offers, has to decide whether to accept the offer or to turn it down. Traditional economics assuming narrow selfishness tells us that, because more money is better than less money, Rita will accept any offer. In terms of the figure it means that Rita will be kind after a kind proposal (5,5) and will also be kind after an unkind proposal (8,2).

Figure 3
The Mini Ultimatum Game

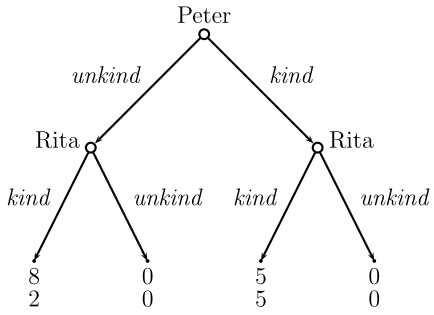

The empirical facts, however, deviate significantly from this prediction. A typical (qualitatively and quantitatively representative) result generated in many experiments is shown in Figure 4. It shows the percentage of ac-

9 Note that such (or similar) situations are not uncommon in everyday life. For instance, it is akin to shopping in Western supermarkets or shops, where one is usually not negotiating the price of the product but rather takes it or leaves it (on the shelves).

Figure 4
Acceptance and Rejection Rates in Mini Ultimatum Game

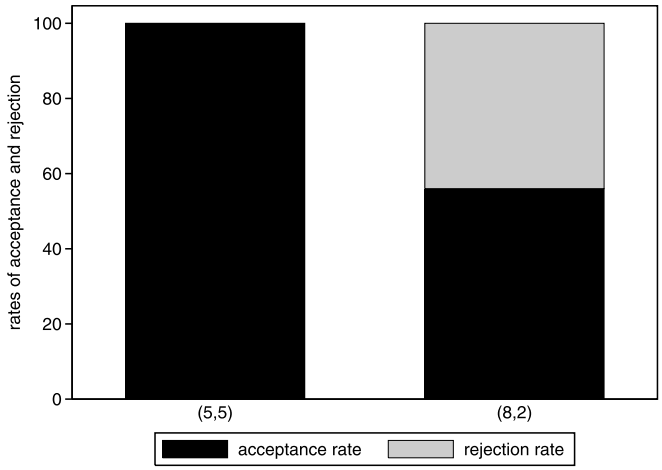

Data source: Falk et al. (2003)

cepted and rejected offers for both possible proposals. As one would expect, the kind proposal of (5,5) is never followed by the unkind response of a rejection. The situation looks quite different, however, if an unkind proposal of (8,2) is put on the table. Such a proposal is often followed by an unkind response, namely rejection. In the reported experiment this happens in more than 40% of the cases. It is important to see that the unkind response is costly, leaving both players without any monetary gain. This is precisely what makes it incompatible with traditional economic reasoning assuming narrow selfishness.[10]

3.4. Positive Reciprocity – the Gift-Exchange Game

Probably the best-known example of the existence of a predisposition towards positive reciprocity stems from a game termed the gift-exchange game.[11] In economics, the basic idea behind this game dates back (at least) to Akerlof (1982), who argued that gift exchange is an important constituent

10 Actually, the rejection rates for offers of only about 20% of the whole pie are usually higher than the 40% reported here. A likely reason is that Falk et al. (2003) applied the so-called strategy method, where subjects have to decide upon acceptance and rejection before they know the actual choice. That is, they make their decision in a cold (emotional) state, whereas responses to actual offers are made in hot (emotional) states (see Loewenstein, 1999).

11 A game very similar in nature is the so-called trust game (Berg et al., 1995).

of labor contracts that are genuinely incomplete. The *incompleteness* of the contract refers to the fact that the effort exerted by an employee is often not verifiable, because it cannot be observed by the employer and/or cannot be enforced by a third party. Akerlof's theoretical model, however, relied on assumptions about economic behavior of employees that are at odds with the assumption of narrow selfishness – namely, that employees respond to higher wages positively, in the sense that higher wages make them exert higher and more costly effort.

In Fehr et al. (1993)[12] this idea is put to a test in the experimental laboratory. In fact, the implemented situation more generally represents any kind of patron–client or principal–agent relationship where contracts cannot be (perfectly) enforced. More concretely, consider the following situation (based on Riedl and Tyran, 2005): A number of people are divided into a set of *employers* ("buyers") and a (larger) set of *employees* ("sellers"). The rules of the game are as follows. If an employer hires an employee who provides effort e and receives a wage w, then the employer's earnings π are $30 + 10e - w$. That is, the employer earns a lump sum of 30 plus 10 times the effort e exerted by the employee minus the wage w paid. The wage is between 1 and 100 ($1 \leq w \leq 100$), and the effort between 1 and 10 ($1 \leq e \leq 10$). The earnings u of the employee are then the wage w he receives minus a *cost of effort* $c(e) = e$ plus a lump-sum payment of 4 ($u = w - e + 4$). The sequence of actions is as follows. The employer first offers a *contract* specifying a wage w. When the contract is signed, the employee receives the wage *before* exerting any effort. Only after the wage is paid out does the employee decide on his effort, and he is completely free to choose any level of effort. Importantly, each employer–employee interaction is anonymous and essentially one-shot. That is, there is no possibility for reputation building or retaliation.

What will be the outcome of such an interaction? Consider first the employee who has received a wage and has now to decide on the effort level. Under the assumption of narrow selfishness, any employee will choose the effort level with the lowest cost, $e = 1$, no matter how high or low the wage received is. In effect, choosing a higher effort level only decreases his total earnings. A rational and selfish employer will perfectly anticipate this behavior and hence offer the lowest acceptable wage, $w = 5$, which is still accepted.[13]

However, in the experiment this predicted outcome is very rarely observed. Moreover, gift exchange is observed as a very strong behavioral regularity. Figure 5 illustrates this. The figure plots, on the vertical axis, the (average) effort chosen by the employees against the average wage offered

12 See also Fehr et al. (1997, 1998).

13 An employee without a contract receives a payment of 4.

Figure 5
Gift Exchange in an Experimental Labor Market

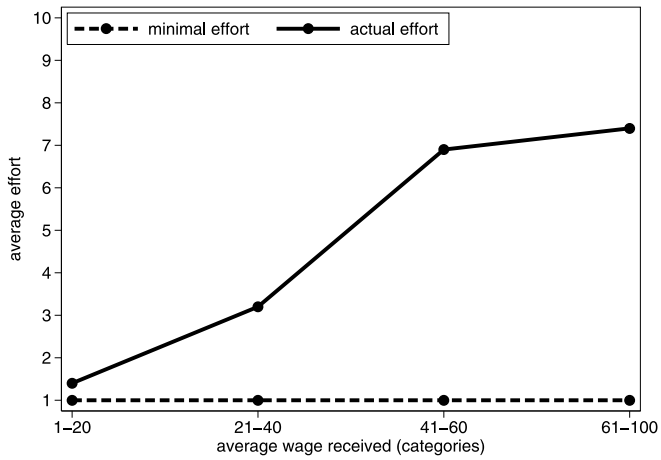

Source: Riedl and Tyran (2005)

by the employers on the horizontal axis. Under the assumption of narrow self-interest, employees should always choose the lowest, implying no positive (or negative) relationship between effort and wages. This is reflected by the straight line running through $e = 1$. The actual data, however, show a strikingly different pattern: effort is clearly and strongly increasing in the wage received. In the figure this is reflected by the increasing line, which connects average effort levels for wages smaller than 20, between 21 and 40, between 41 and 60, and larger than 61. Hence, in conclusion, this and many other studies (for a recent review see Fehr et al., 2009) clearly show that people respond positively reciprocally. An important side effect of the observed gift exchange is that it increases efficiency (in terms of surplus maximization) and decreases inequality in earnings, in comparison with the benchmark outcome predicted under traditional behavioral assumptions.

4. Behavioral Regularities, Public Economics, and Public Policy

One might wonder what all of this has to do with public economics and public policy. In the following I shall argue that the described behavioral regularities can indeed be of eminent importance for public economics and policy.

4.1. On the Interaction between Trading Institutions, Morality, and Tax Shifting

An interesting and important example where economic institutions and inclinations toward reciprocity interact in a nontrivial way is the case of *tax liability side equivalence*. Tax liability side equivalence is a basic tenet in public economics. It states that the statutory incidence (that is, who is legally responsible to pay a tax) is irrelevant to the economic incidence (that is, who actually bears the tax burden). In the words of one of the authorities in public economics, Richard A. Musgrave,

> *it is a matter of indifference whether a general tax on transactions is assessed on the seller's or on the buyer's side of the market.*
> Musgrave (1959, p. 351)

Importantly, under traditional economic assumptions this holds true independently of the trading environment (monopoly, oligopoly, competition, or bargaining), provided that prices can in principle adjust freely. Interestingly enough, however, much of the public debate about tax burden (and subsidy benefit) in the media and the political arena is concerned with statutory instead of economic incidence. This raises the question whether the public reasoning or the professional economic reasoning is incorrect. For public policy the answer to this question is obviously important, because it determines which groups of the society are actually going to carry the burden of a tax, and hence what the distributional and allocational consequences of the tax are. Neglecting the behaviorally true tax burden and relying (only) on normative prescriptions that are based on incorrect behavioral assumptions may, therefore, have very undesirable political and economical consequences.

To uncover the behaviorally true tax burden, ideally one would like to shift the statutory tax burden from one side of the market (e.g., buyers) to the other side of the market (e.g., sellers), *leaving everything else equal*. Naturally, such situations do not occur regularly in the field. Fortunately, laboratory experiments are an ideal method to do precisely this. The following briefly reports on three sets of experiments testing tax liability side equivalence under three important economic institutions: competitive markets with complete contracts, bargaining, and gift-exchange markets.

Three studies (Borck et al., 2002; Kachelmeier et al., 1994; Ruffle, 2005) experimentally investigate tax (and subsidy) incidence equivalence in competitive markets under various trading mechanisms. The results of these experiments can be summarized simply: tax liability side equivalence holds in competitive experimental markets independently of the trading institution. Traditional theoretical and behavioral tax incidence coincide. Kerschbamer and Kirchsteiger (2000) give an interesting twist to these results by study-

ing tax liability side equivalence in a simple bargaining environment. They modify the ultimatum game (similar to the one described above) so that in case of acceptance of a proposal one of the negotiators has to pay a tax from her gross earnings. In one treatment the proposer has to pay the tax. In a second treatment the statutory tax burden is placed upon the responder. If tax liability side equivalence holds, the distribution of earnings should be the same under both tax treatments. However, if statutory tax incidence implies also a moral obligation to actually bear part of the tax burden, then tax liability side equivalence breaks down. It is well known that in bargaining situations reciprocity and distributional fairness crucially influence behavior and economic outcomes. Hence, one might expect that tax liability side equivalence is especially vulnerable in such situations. Indeed, Kerschbamer and Kirchsteiger (2000) find that the side of the transaction that is legally obliged to pay the tax also bears a disproportionate part of the economic tax burden. The observed incompleteness of tax shifting is also economically significant. When the tax liability was taken from the responder and placed upon the proposer, the offered net (after tax) income for responders *increased* by up to 24%. In contrast, standard tax liability side equivalence predicts no change in net income at all. Therefore, this is a clear case where trading institution and moral behavior interact in a way traditional theory fails to predict.

Perfectly competitive markets and pure (two person) bargaining situations are at the two extremes of actually existing economic exchange institutions. Although both are fairly common, they are probably not the most frequent trading institutions. Rather, a mixture of the two seems to be the most common one. Gift-exchange markets or markets with incomplete (or even no) contracts, as described above, incorporate both elements: competitive market interaction and bilateral bargaining. Riedl and Tyran (2005) investigate tax liability side equivalence in such markets. In one set of experiments buyers are legally obliged to pay a tax, and in another set of experiments sellers bear this obligation. If tax liability side equivalence holds, then there should be no difference in the outcomes of real variables between the two tax regimes. However, if moral obligations, as in the study of Kerschbamer and Kirchsteiger (2000), are important, then the side on which the tax is levied should also bear a larger part of the economic burden of the tax. In gift-exchange markets three important variables can be investigated: net prices, effort, and net earnings. Tax liability side equivalence may fail on each of these dimensions, with different consequences for the final distribution of income. Figure 6 depicts average net prices and exerted efforts in panel (a) and average net earnings in panel (b) for both regimes; taxes levied on buyers and taxes levied on sellers. From the figure it is obvious that there are no large differences between the two tax regimes on any of the three dimen-

Figure 6
Prices, Effort, and Earnings under the Two Tax Regimes

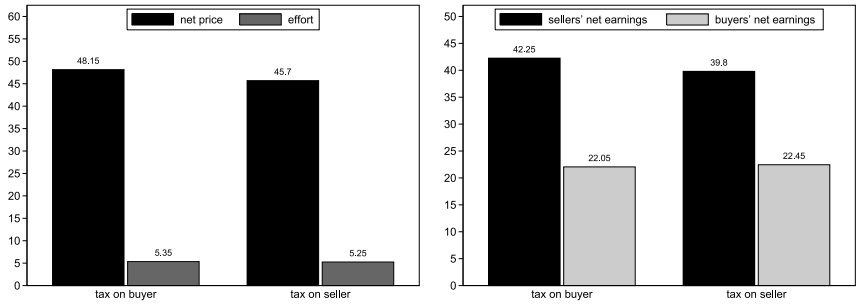

(a) Net prices and effort (b) Net earnings

Source: Riedl and Tyran (2005)

sions of comparison. Furthermore, the observed small differences are neither statistically nor economically significant. This is an important but also somehow puzzling result, given the observation that tax liability side equivalence breaks down in pure bargaining and not when it is coupled with a competitive market. It strongly suggests that the details of market interaction and perceptions about moral obligations to pay a tax interact in a nontrivial way, with real and economically significant results in terms of tax shifting.

One conclusion to be drawn from these studies is that researchers as well as policymakers should not solely rely on traditional economic theorizing when assessing the economic burden of a tax. Additionally, the institutional environment and its interaction with moral and reciprocal inclinations crucially affect the actual economic outcome. Where the precise borders for predictive accuracy of the standard economic model lie is still an open empirical question, though. If we want to understand under which circumstances tenets like tax liability side equivalence indeed hold or – more importantly – have to be modified, we need a research program that systematically evaluates such pieces of economic wisdom.

4.2. Identifiable Victims and Hidden Taxes

Real outcomes can be nontrivially influenced not only by the interaction between institutional design of trading institutions and behavior of economic agents, but also through pure presentation effects. In a 1968 article Thomas Schelling noticed that "[t]he life you save may be your own" and that "the death of a particular person evokes anxiety and sentiment, guilt and awe,...[but that]...most of this awesomeness disappears when we deal

with statistical death" (Schelling, 1968). This is probably the first account by an economist pointing towards how differently we perceive identifiable and statistical victims. Loewenstein et al. (2006) take up this issue and examine the public policy consequences of this human inclination to have stronger feelings towards an identifiable victim than towards a statistical victim. They argue that from a welfare economics point of view "people may be insufficiently sympathetic towards statistical victims." Mainly psychological research strongly supports the claim that individual concrete cases have a much more powerful motivational effect than statistical cases. This seems to be true even if the statistics are objectively more informative than the individual case. A typical example in this respect is that opinions about the abuse of welfare payments are shaped much more strongly by individual experience than by objective statistics.[14]

An important public policy implication of the identifiability effect is that for politicians hidden taxes tend to be much more popular than other taxes. For example, the value added (or sales) tax is for most consumers (including economists) simply part of the purchase price of a commodity, and hence has no identifiable victim. This concealment may make it politically easier to raise value added taxes than more direct taxes. How serious and economically important the misperception of hidden taxes is, is convincingly shown by Chetty et al. (2009). These authors conduct a field experiment in a grocery store where in one treatment the sales tax is made salient by explicitly showing it on the price tag whereas in another treatment only the tax-inclusive price is shown. The study finds a significant 8% decrease in purchases and sales revenues when the tax is made salient. This salience effect is corroborated with field empirical data using variations in taxes on beer among U.S. states. An important conclusion of this study is that it is crucial to distinguish between tax elasticities and price elasticities when thinking about tax policy – a distinction not necessary in traditional public economics.

Other prominent examples of hidden taxes are withholding income taxes, which make people think that the money transferred to the tax authority is not their own, and corporate income taxes, which make people think that shareholders pay the tax, although it is mostly the factor labor that actually carries the tax burden (McCaffery, 1994; McCaffery and Baron, 2006). In the mentioned examples the lack of identifiability makes the taxes themselves as well as an eventual increase of them much more acceptable than it would be the case for nonhidden taxes.[15]

14 For clean evidence from the laboratory (dictator game giving) as well as the field (housing for the needy), see Small and Loewenstein (2003).

15 The psychological appeal of hidden taxes is nicely summed up by the aphorism of Russell Long, one of the most powerful and influential tax legislators as chairman of the U.S.

Experimental and behavioral research has uncovered a number of other behavioral regularities that are particularly important for taxation economics, like misperception of the progessivity of the tax system, confusion of marginal and average income tax rates, and nonrationally high tax compliance rates (see, e.g., de Bartolome, 1995; Liebman and Zeckhauser, 2004; Slemrod, 2006). In his contribution to this issue Slemrod (2010) discusses some of these problems in more detail.[16]

4.3. Endogenous Preferences and Competition

One central assumption in economics is that people have "fixed lifetime preferences" (Bernheim and Rangel, 2007). In particular, this means preferences are assumed not to change across states of nature or institutional constraints. In this perspective, preferences are exogenously fixed and independent from the environment an individual is immersed in. In contrast to this traditional view, Bowles (1998) argues vividly in favor of *endogenous* preferences. He claims that our preferences are not well defined and stable, but rather are strongly dependent on the environment we have to deal with. However, all the evidence he puts forward in support of his claims is either indirect or open to alternative interpretations.

In a recent study, Brandts et al. (2008, 2009) experimentally test the direct influence of trading institutions on subjective well-being and (social) preferences. In a series of experiments subjects are divided into two groups. One group interacts in a competition-free environment, whereas subjects in the other group interact in a competitive environment. The hypothesis is that the experience of competition versus no competition *per se* leaves its traces in subjects' well-being (in the sense of Kahneman et al. (1997)'s "experienced utility") and in their social preferences (that is, their "social disposition towards others"). To test this hypothesis, methods from social psychology are combined with experimental economics. At the beginning of the experiment subjects are asked to perform a social value orientation test, which measures subjects' social preferences by letting them allocate real money between themselves and some anonymous other person. The same test is conducted also after subjects have experienced a competitive or a noncompetitive environment. Hence, social preferences are measured before and after the experience with a particular trading institution. If traditional economic reasoning were correct, then for the social orientation of participants it should not matter whether they have experienced a competitive or a noncompetitive

Senate Finance Committee: "Don't tax him, don't tax me, tax the man behind the tree" (quoted after Small and Loewenstein, 2003).

16 For earlier accounts of taxation and behavioral economics, see McCaffery and Slemrod (2006) and Kirchler (2007).

environment. This is not what is found. Firstly, there is a difference in the social orientation of subjects experiencing different institutions, and secondly, within the competitive institution the social orientation differs strongly between subjects who have experienced it in different roles (i.e., being on the long or the short side of the trade relationship).

Figure 7
Change in Social Preferences in Competitive and Noncompetitive Environment

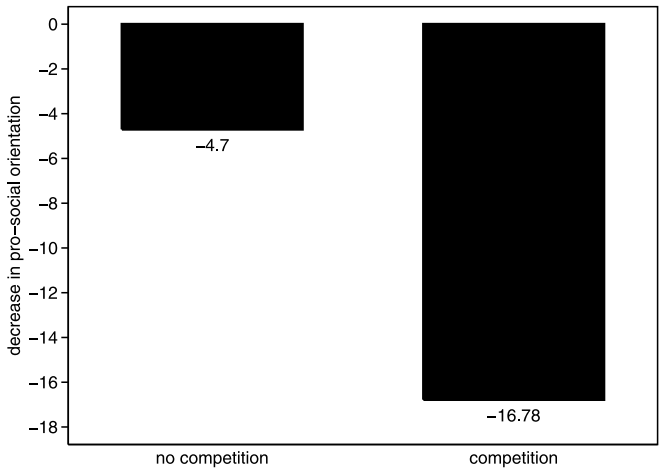

Source: Brandts et al. (2009, 2008)

Figure 7 gives a visual impression of the differences in social preferences. It shows the change in social orientations from before to after the experience with the competitive and the noncompetitive environment, respectively. [17] Interestingly, under both trading regimes prosocial orientation deteriorates. Importantly, however, compared to experience with the no-competition regime, the decrease in prosocial orientation is more than three times larger when subjects experience competition. This clearly indicates that preferences are indeed not – or at least only partly – exogenously given and are strongly shaped by subjects' institutional experience. An additional result of this study is that it is not only social orientation that deteriorates under competition, but that subjects also suffer a loss of experienced utility (i.e., subjective well-being). Similar findings are reported in a meta-study

17 For clarity of presentation, only the average of the strongest change in social preferences among subgroups across all investigated competitive environments is shown.

by Bowles and Polania Reyes (2009). The authors survey the data of 51 experiments and find that explicit economic incentives aimed at the narrowly materially selfish individual often have not only the intended effect of providing information and suggesting socially appropriate behavior, but also the unintended and unwanted effect of compromising intrinsic motivation and self-determination. An important conclusion drawn from reviewing these studies is that economic incentives and social preferences seem more often complements than substitutes.

In public and political debates, reforms aiming at more competition are often backed with the received wisdom of traditional economics that consumers will be better off in the presence of more competition. This argument refers to the notion of consumer surplus, which (in practice) measures welfare in purely material terms. This surplus may indeed increase with more competition and less regulation (although in the experiment discussed above even this is not the case). However, the ultimate aim of welfare economics and public policy should be the maximization of the citizens' well-being. Therefore, the findings of deteriorating social orientations and declining subjective well-being call into question the supposedly purely positive effects of more competition and less regulation. Unfortunately, there is as yet no generally accepted measure of subjective well-being developed, and much more research into the measurement and determinants of well-being is necessary.[18]

The evidence reported above also points to the important and not yet well understood interplay between heterogeneous social preferences and institution design and formation. There is mounting evidence that people differ quite substantially with respect to their social preferences (e.g., Andreoni and Miller, 2002; Engelmann and Strobel, 2004; Fisman et al., 2007; Bellemare et al., 2008), and there is also recent evidence that even subtle institutional differences may alter behavior substantially (see, e.g., Falk and Kosfeld, 2006; Reuben and Riedl, 2009). However, evidence on how these interact is only very recently emerging (Kosfeld et al., 2009).

4.4. Presentation Effects and Public Policy

Finally, to close the circle, let me give two examples where research into presentation and framing effects can very concretely inform public policy and thereby increase general well-being. The first one is taken from Amir et al. (2005) and highlights the importance of whether a decision task is presented simultaneously or sequentially. In many places police lineups are used to identify suspects of crime. In such a lineup eyewitnesses of crimes

18 Recently, some progress is reported concerning the measurement of subjective well-being. Kahneman and Krueger (2006) introduce a measure based on self-reports of peoples' emotional states (see also Brandts et al., 2008).

attempt to recognize one person in a group of suspects standing next to each other.

Evidence from psychological research casts serious doubts on this much-used method of identification. Lindsay and Wells (1985) designed a clever experiment and showed that the likelihood of false identification of an innocent subject is much higher under simultaneous than under sequential lineup. At the same time the frequencies of correct recognition of the guilty suspect did not differ significantly between the two conditions. This, together with subsequent research, clearly shows that in identification tasks – such as police lineups – the practice of presenting items or individuals one at a time leads to better (less biased) judgments than the practice of presenting them simultaneously. The straightforward and unambiguous policy recommendation is, therefore, to abandon simultaneous lineups in favor of sequential lineups. Actually, at least two U.S. states have adopted that alternative method.

The organization of organ donation is a another important domain where it has been shown that framing effects can crucially alter behavior and, in this case, make the difference between life and death. Johnson and Goldstein (2003) investigate and compare two different organ donation systems around the world. They distinguish between two so-called no-action defaults. No-action defaults are the choices implicitly imposed on individuals who do not take an active decision. In the case of organ donation the most widely used default decisions are *presumed consent* and *explicit consent.* Presumed consent means that people are assumed to be organ donors as long as they do not actively indicate otherwise. Explicit consent means that individuals have to actively register for being a donor; otherwise they are not. The authors investigate the effective consent rates for being a donor across the two defaults, experimentally as well as by cross-country comparisons. Traditional economic theory assumes that preferences are fixed and known to the economic agent, which implies that for effective consent rates it should not matter whether one or the other default option is taken. Figure 8 shows the facts. The four leftmost bars depict effective consent rates for four countries (Denmark, the Netherlands, the UK, and Germany) that apply explicit consent as the default. The bars to the right show effective consent rates for countries with presumed consent as default. The differences are striking. Whereas in the explicit-consent countries the effective consent rates are maximally 27.5% (the Netherlands), the minimum consent rate in countries with presumed consent is 85.9% (Sweden). These differences are surely too large to be explained by effort or transaction costs of actively opting in and opting out in the explicit-consent and presumed-consent countries, respectively. To exclude definitively this potential explanation, the authors conducted an experiment where effort and transaction costs were virtually zero. In the experiment the difference in effective consent rates between the two default

Figure 8
Effective Consent Rates in Countries with Explicit Consent (Four Leftmost Bars) and Presumed Consent (Seven Rightmost Bars)

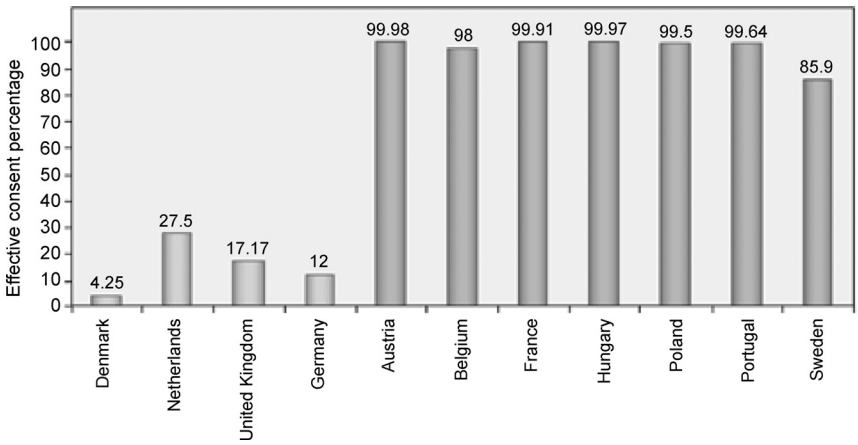

Source: Science 302, 1338 (2006)

options is slightly smaller than in the cross-country study, but still a long way from zero. Interestingly, a neutral framing without any default led to the same effective consent rate as the presumed-consent default. This allows the conclusion that the unbiased "true" preferences concerning organ donation are better elicited with presumed consent than with explicit consent. That some policymakers are – at least sometimes – aware of the fact that default options make a difference is nicely reflected in arguments brought forward during a recent discussion about reforming the donor registration system in the Netherlands. In 2005 a coalition of parties in the parliament actually did bring forward a motion that would have changed the explicit consent default into a presumed consent default. In the end, however, there was not enough political support for such a radical change of the donor registration system.[19] In any event, this example highlights how insights from behavioral and experimental research are related to important political decisions and

19 A clear case of a presentation and framing effect is also reported in an article in the Dutch newspaper *NRC Handelsblad* (September 2 & 3, 2006, pp. 41–42) about the Dutch immigration and naturalization service (IND). The newspaper reports that if employees of the IND *reject* an application for a temporary residence permit, they have to explain their decision in writing. For the case of hardship of asylum seekers, the IND employees have to explain the decision if they *accept* the application. The result is that cases of hardship are hardly ever positively assessed, whereas temporary residence permits are relatively easily issued.

can inform public policy more accurately than traditional economic reasoning would be able to.

5. Further Directions: the Field and the Brain

A common argument of skeptics against the use of laboratory experiments in general, and as a policy advice instrument in particular, is their supposed lack of external validity. This is indeed an important concern, because if regularities observed in the laboratory do not carry over to the field, any conclusions and public policy advice drawn from these experiments could be dangerously misleading. The potential problem of lack of external validity is not unique to economic (or psychological) experiments, however.[20] In physics, the feather and the stone, which fall with the same speed in vacuum but with different speeds in "real life," constitute a well-known illustrative example. As the air resistance in the terrestrial atmosphere affects the fall velocity of the two objects, in economic situations many factors one can control for in the laboratory but not in the real world can influence behavior and blur or even wipe out behavioral regularities observed in the laboratory. The experimental method offers a unique way to tackle this problem, which is adding pieces of real-life context to the dry laboratory environment *in a systematic way*. In this way one can trace if and how such pieces of reality alter behavior.[21]

Another important way to check for external validity is to replace the usual student subjects with subjects who are experienced with the decision situation at hand and/or are more representative than students. Such experiments have been conducted for a variety of decision situations (see, e.g., Fehr and List, 2004; Egas and Riedl, 2008, among many others). The general upshot from these experiments is that experts often do not make significantly different decisions from students in the same situation, although there are sometimes subtle and surprising differences. For instance, Alevy et al. (2007) investigate the behavior of financial market professionals regarding information cascades and find that "professionals are less Bayesian than students" (ibid., p. 161), but report only little evidence for differences in cascade formation. Haigh and List (2005) investigate the difference in myopic loss aversion (MLA) between students and professional traders and find that "traders exhibit behavior consistent with MLA to a *greater* extent than students" (ibid.,

20 Note that theoretical reasoning is confronted with exactly the same potential lack of external validity.

21 For a recent discussion of this and other pros and cons regarding the use of laboratory experiments in economics, see Falk and Heckman (2009).

p. 523; emphasis in original). Hence, sometimes behavioral regularities found with students are even amplified with nonstudent subjects.

A third way to test external validity is to conduct field experiments. Field experiments might be seen as an extreme combination of the two already described ways of adding 'realism' to 'experiments'. In field experiments researchers try to add the largest possible number of 'pieces of reality' and use nonstudent subjects in their "natural" environment. Harrison and List (2004) provide a typology of field experiments and define them, crudely speaking, as experiments where one is "recruiting subjects in the field rather than in the classroom, using field goods rather than induced valuations, and using field context rather than abstract terminology in instructions" (Harrison and List, 2004, pp. 1009–1010).

Without doubt, field experiments are an interesting and important development in economic research, but they also have their disadvantages. For instance, in comparison with the use of induced valuations, the use of field goods contributes to a loss of control regarding participants' true valuations of commodities, the use of special nonstudent subjects in their special circumstances does not necessarily not allow one to draw conclusions beyond the investigated group in the investigated circumstances, and field experiments usually do not allow for replications as precise as in the laboratory. Therefore, especially in research that is concerned with policy advice, field experiments are best viewed as an important complementary research method.

In the ideal case, an economic policy reform is evaluated with *all possible* scientific methods before a political decision is made: theoretically, experimentally in the lab and the field, and with traditional applied econometrics. It should be obvious that a thorough, scientifically sound examination of a policy reform that reduces the risk of implementing bad policies is much cheaper (at least in expectations) than the costs of an actually implemented bad policy. A rare example of a first – albeit incomplete – attempt at such a scientific approach to policy issues is the evaluation of the so-called Plan Van Elswijk in the Netherlands, which proposes a radical reform of the (Dutch) financing system for unemployment benefits. In the evaluation of this plan, simulation studies, laboratory experiments, and a small field experiment were conducted (see van Winden et al., 1999, 2000; Riedl and Winden van, 2001, 2007, 2008, and the references therein). These studies produced a rather clear picture of the likely benefits and disadvantages of the proposed reform.[22] An important lesson learned from these studies is that laboratory experiments can indeed provide valuable information even for such complex questions as unemploy-

22 In the end, despite the clear-cut results, policymakers chose to interpret them differently than most involved researchers, and to a large extent they ignored the scientifically obtained outcomes.

ment benefit reform plans. Thus, one may safely conclude that laboratory experiments, together with field experiments and theoretical reasoning, will play an important role in public policy advice in the future.[23]

A further stream of research – quite different, but potentially equally relevant, and in the long run probably even more important though more debated – is the recent combination of neuroscience with experimental and behavioral economics. This new research branch – *neuroeconomics* – uses knowledge about brain mechanisms to study the biological foundations of behavioral regularities observed in the laboratory and the field. Kevin McCabe, one of the pioneers in this new field of the behavioral sciences, defines it as follows:

Neuroeconomics is an interdisciplinary research program with the goal of building a biological model of decision making in economic environments. [More specifically, it] is the study of how the embodied brain interacts with its external environment to produce economic behavior. Research in this field will allow social scientists to better understand individuals' decision making, and consequently to better predict behavior.
McCabe (2003, p. 294)

At first sight this may sound abstract and remote from public policy issues. Indeed, in the above definition, the first part refers to the pure scientific element of neuroeconomics. The second part, however, reflects the potential of this approach for public policy making. Having *good models* of human behavior is crucial for making *good predictions* of human behavior in economic situations. Neuroeconomics has the potential to significantly contribute to this quest for better models of economic decision-making. When evaluating the potential of this new emerging field, one may want to recall the history of game theory and laboratory experiments in economics. In its beginnings, game theory was largely dismissed as being too academic and thought to be useful at most for war strategists, but surely not for the analysis of ordinary economic interactions. Nowadays game theory is used for policy advice on a large scale in many countries. (Recall, e.g., the commercial frequency auctions a few years ago, where game theorists played a crucial role in designing auction formats.) Later, experimental and behavioral economics was smiled at by many economists because it was the received wisdom in economics that "economists...cannot perform...controlled experiments" (Samuelson and Nordhaus, 1985), a view that changed quite a bit over the years, as the following statement by the very same authors seven years later testifies: experimental economics is an "exciting new development"

23 An interesting recent example where laboratory experiments informed politics is given in Jacob K. et al. (2005). For a survey of experiments for economic policy in the context of industry regulation, see Normann (2004) and Hinloopen and Normann (2009).

(Samuelson and Nordhaus, 1992). It is not unlikely that neuroeconomics awaits the same fate. Bernheim (2009) and Rustichini (2009) are recent critical appraisals of this new development in economics.

In any event, both directions – towards the experimental field and towards the brain – can be expected to be among the most lively areas of research in economics, and both are likely to produce results that will lead to better models, better predictions, better advice, and ultimately also – we may hope – better-informed public policies.

6. Conclusion

In the course of this contribution I have discussed a few areas in experimental and behavioral economics that are of importance for public economics and public policy. Naturally, many at least equally important issues were not even slightly touched upon – for instance, voluntary contributions to public goods (see, e.g., Gächter and Herrmann, 2009); time-inconsistent intertemporal decision-making, e.g., with respect to retirement decisions and pension systems (see, e.g., Frederik et al., 2002); decision-making under risk and uncertainty, e.g., with respect to health insurance and social security in general (see, e.g., Kahneman and Tversky, 1979; Wakker et al., 2007; Heinemann et al., 2008); or the role of moral property rights in policy reform (see, e.g., Gächter and Riedl, 2005) – to name only a few. For the future, also the fundamental issue of *welfare analysis beyond revealed preferences* is awaiting public economics research and, more generally, economic theory. First important steps in this direction have been made by Bernheim and Rangel (2009) and by Herings and Rohde (2006).

In recent years, the facts of behavioral regularities have been shown to be incompatible with the traditional wisdom of economic theory. This has led to the development of new theoretical approaches and models. It seems clear that for good public policy we need both good accounts of the behavioral facts and a theoretical knowledge that gives us the tools to deal with the upcoming challenges in an accurate way. This is all the more important in that politicians are increasingly more willing to listen to the advice of behavioral and experimental economists. In fact, the 2005 quotation that opens this paper has been overtaken by reality: *Time* magazine reports that behavioral economists are advising U.S. President Barack Obama (Grunwald, 2009).

References

Akerlof, G. A. (1982), Labor Contracts as Partial Gift Exchange, Quarterly Journal of Economics 97, 543–569.

Alevy, J. E., Haigh, M. S., and List, J. A. (2007), Information Cascades: Evidence from a Field Experiment with Financial Market Professionals, Journal of Finance 62, 151–180.

Amir, O., Ariely, D., Cooke, A., Dunning, D., Epley, N., Gneezy, U., Koszegi, B., Lichtenstein, D., Mazar, N., Mullainathan, S., Prelec, D., Shafir, E., and Silva, J. (2005), Psychology, Behavioral Economics, and Public Policy, Marketing Letters 16, 443–454.

Andreoni, J., and Miller, J. (2002), Giving According to GARP: An Experimental Test of the Consistency of Preferences for Altruism, Econometrica 70, 737–753.

Bellemare, C., Kröger, S., and van Soest, A. (2008), Measuring Inequity Aversion in a Heterogeneous Population Using Experimental Decisions and Subjective Probabilities, Econometrica 76, 815–839.

Berg, J., Dickhaut, J., and McCabe, K. (1995), Trust, Reciprocity, and Social History, Games and Economic Behavior 10, 122–145.

Bernheim, B. D. (2009), On the Potential of Neuroeconomics: A Critical (but hopeful) Appraisal, American Economic Journal: Microeconomics 1, 1–41.

Bernheim, B. D., and Rangel, A. (2007), Behavioral Public Economics: Welfare and Policy Analysis with Nonstandard Decision-Makers, in: Diamond, P. and Vartiainen, H. (Eds.), Behavioral Economics and Its Applications, Princeton University Press, Princeton, 7–84.

Bernheim, B. D., and Rangel, A. (2009), Beyond Revealed Preference: Choice-Theoretic Foundations for Behavioral Welfare Economics, Quarterly Journal of Economics 124, 51–104.

Borck, R., Engelmann, D., Müller, W., and Normann, H.-T. (2002), Tax Liability Side Equivalence in Experimental Posted Offer Markets, Southern Economic Journal 68, 672–682.

Bowles, S. (1998), Endogenous Preferences: The Cultural Consequences of Markets and Other Economic Institutions, Journal of Economic Literature 36, 75–111.

Bowles, S., and Polania Reyes, S. (2009), Economic Incentives and Social Preferences: A Preference-Based Lucas Critique of Public Policy, CESifo Working Paper No. 2734.

Brandts, J., Riedl, A., and van Winden, F. (2008), On Competition and Well-Being. An Experimental Investigation into Rivalry, Social Disposition, and Subjective Well-Being, Maastricht University, mimeo.

Brandts, J., Riedl, A., and van Winden, F. (2009), Competitive Rivalry, Social Disposition, and Subjective Well-Being: An Experiment, Journal of Public Economics 93, 1158–1167.

Camerer, C., and Loewenstein, G. (2004), Behavioral Economics: Past, Present, Future, in: Camerer, C., Loewenstein, G., and Matthew, R. (Eds.), Advances in Behavioral Economics, Princeton University Press, Princeton and Oxford, 3–51.

Camerer, C. F. (2005), Behavioral Economics, Lecture at World Congress of the Econometric Society, 2005.

Camerer, C. F., and Hogharth, R. M. (1999), The Effects of Financial Incentives in Experiments: A Review and Capital-Labor-Production Framework, Journal of Risk and Uncertainty 19, 7–42.

Chetty, R., Looney, A., and Kroft, K. (2009), Salience and Taxation: Theory and Evidence, American Economic Review 99, 1145–1177.

Cronqvist, H., and Thaler, R. H. (2004), Design Choices in Privatized Social-Security Systems: Learning from the Swedish Experience, American Economic Review 94, 424–428.

de Bartolome, C. A. M. (1995), Which Tax Rate Do People Use: Average or Marginal?, Journal of Public Economics 56, 79–96.

de Jong, J. D., van den Brink-Muinen, A., and Groenewegen, P. P. (2008), The Dutch Health Insurance Reform: Switching Between Insurers, A Comparison Between the General Population and the Chronically Ill and Disabled, BMC Health Services Research 8, 1–9.

Egas, M., and Riedl, A. (2008), The Economics of Altruistic Punishment and the Maintenance of Cooperation, Proceedings of the Royal Society – B 275, 871–878.

Engelmann, D., and Strobel, M. (2004), Inequality Aversion, Efficiency, and Maximin Preferences in Simple Distribution Experiments, American Economic Review 94, 857–869.

Falk, A., Fehr, E., and Fischbacher, U. (2003), On the Nature of Fair Behavior, Economic Inquiry 41, 20–26.

Falk, A., and Heckman, J. J. (2009), Lab Experiments are a Major Source of Knowledge in the Social Sciences, Science 326, 535–538.

Falk, A., and Kosfeld, M. (2006), The Hidden Costs of Control, American Economic Review 96, 1611–1630.

Fehr, E., Gächter, S., and Kirchsteiger, G. (1997), Reciprocity As Contract Enforcement Device: Experimental Evidence, Econometrica 65, 833–860.

Fehr, E., Goette, F., and Zehnder, C. (2009), A Behavioral Account of the Labor Market: The Role of Fairness Concerns, Annual Review of Economics 1, 355–384.

Fehr, E., Kirchsteiger, G., and Riedl, A. (1993), Does Fairness Prevent Market Clearing? An Experimental Investigation, Quarterly Journal of Economics 108, 437–459.

Fehr, E., Kirchsteiger, G., and Riedl, A. (1998), Gift Exchange and Reciprocity in Competitive Experimental Markets, European Economic Review 42, 1–34.

Fehr, E., and List, J. A. (2004), The Hidden Costs and Returns of Incentives – Trust and Trustworthiness Among CEOs, Journal of the European Economic Association 2, 743–771.

Feynman, R. P. (1964), The Feynman Lectures on Physics, Addison-Wesley, Reading, MA.

Fisman, R., Kariv, S., and Markovits, D. (2007), Individual Preferences for Giving, American Economic Review 97, 1858–1876.

Forsythe, R., Horowitz, J. L., Savin, N., and Sefton, M. (1994), Fairness in Simple Bargaining Experiments, Games and Economic Behavior 6, 347–369.

Frederik, S., Loewenstein, G., and O'Donogue, T. (2002), Time Discounting and Time Preference: A Critical Review, Journal of Economic Literature 60, 351–401.

Gächter, S., and Herrmann, B. (2009), Reciprocity, Culture and Human Cooperation: Previous Insights and a New Cross-Cultural Experiment, Philosophical Transactions of the Royal Society B: Biological Sciences 364, 791–806.

Gächter, S., and Riedl, A. (2005), Moral Property Rights in Bargaining with Infeasible Claims, Management Science 51, 249–263.

Grunwald, M. (2009), How Obama Is Using the Science of Change, Time Magazine, April 2, 2009.

Güth, W., Schmittberger, R., and Schwarze, B. (1982), An Experimental Analysis of Ultimatum Bargaining, Journal of Economic Behavior and Organization 3, 367–388.

Haigh, M. S., and List, J. A. (2005), Do Professional Traders Exhibit Myopic Loss Aversion? An Experimental Analysis, Journal of Finance 60, 523–534.

Harrison, G., and List, J. A. (2004), Field Experiments, Journal of Economic Literature 62, 1013–1059.

Heinemann, F., Förg, M., Jonas, E., and Traut-Mattausch, E. (2008), Psychologische Restriktionen wirtschaftspolitischer Reformen, Perspektiven der Wirtschaftspolitik 9, 383–404.

Herings, P. J.-J., and Rohde, K. I. M. (2006), Time-Inconsistent Preferences in a General Equilibrium Model, Economic Theory 29, 591–619.

Hinloopen, J., and Normann, H.-T. (Eds.) (2009), Experiments and Competition Policy, Cambridge University Press, Cambridge.

Jacob, K. G., Offerman, T. J. S., and Schram, A. J. H. C. (2005), Using First-Price Auctions to Sell Heterogeneous Licenses, University of Amsterdam, Working Paper.

Jamison, J., Karlan, D., and Schlechter, L. (2008), To Deceive or not Deceive: The Effect of Deception on Behavior in Future Laboratory Experiments, Journal of Economic Behavior and Organization 68, 477–488.

Johnson, E. J., and Goldstein, D. (2003), Do Defaults Save Lives?, Science 302, 1338–1339.

Kachelmeier, S. J., Limberg, S. T., and Schadewald, M. S. (1994), Experimental Evidence of Market Reactions to New Consumption Taxes, Contemporary Accounting Research 10, 505–545.

Kahneman, D., and Krueger, A. B. (2006), Developments in the Measurement of Subjective Well-Being, Journal of Economic Perspectives 20, 3–24.

Kahneman, D., and Tversky, A. (1979), Prospect Theory: An Analysis of Decision Under Risk, Econometrica 47, 263–292.

Kahneman, D., Wakker, P. P., and Sarin, R. (1997), Back to Bentham? Explorations of Experienced Utility, Quarterly Journal of Economics 112, 375–405.

Kerschbamer, R., and Kirchsteiger, G. (2000), Theoretically Robust but Empirically Invalid? An Experimental Investigation into Tax Equivalence, Economic Theory 16, 719–734.

Kirchler, E. (2007), The Economic Psychology of Tax Behaviour, Cambridge University Press, Cambridge.

Kosfeld, M., Okada, A., and Riedl, A. (2009), Institution Formation in Public Goods Games, American Economic Review 99, 1335–1355.

Liebman, J. B., and Zeckhauser, R. J. (2004), Schmeduling. Technical Report, Harvard University, unpublished paper.

Lindsay, R. C. L., and Wells, G. L. (1985), Improving Eyewitness Identifications from Lineups: Simultaneous versus Sequential Lineup Presentation, Journal of Applied Psychology 70, 556–564.

Loewenstein, G. (1999), A Visceral Account of Addiction, in: Elster, J. and J, S. O. (Eds.), Getting Hooked: Rationality and Addiction, Cambridge University Press, Cambridge, 235–264.

Loewenstein, G., Small, D. A., and Strnad, J. (2006), Statistical, Identifiable, and Iconic Victims, in: McCaffery, E. J. and Slemrod, J. (Eds.), Behavioral Public Finance, Russell Sage Foundation Press, New York, 32–46.

McCabe, K. (2003), Neuroeconomics, in: Nadel, L. (Eds.), Encyclopedia of Cognitive Science, volume 3, Nature Publishing Group, Macmillan Publishers Ltd, 294–298.

McCaffery, E. J. (1994), Cognitive Theory and Tax, UCLA Law Review.

McCaffery, E. J., and Baron, J. (2006), Thinking about Tax, Psychology, Public Policy, and Law 12, 106–135.

McCaffery, E. J., and Slemrod, J. (Eds.) (2006), Behavioral Public Finance, Russell Sage Foundation, New York.

Musgrave, R. A. (1959), The Theory of Public Finance, McGraw-Hill, New York.

NIVEL (2009), NIVEL jaarboek 2008.

Normann, H.-T. (2004), Experiments for Economic Policy Making, Technical report, ENCORE, University of Amsterdam.

Ortmann, A., and Hertwig, R. (2002), The Costs of Deception: Evidence from Psychology, Experimental Economics 5, 111–131.

Plott, C. R. (1982), Industrial Organization Theory and Experimental Economics, Journal of Economic Literature 20, 1485–1527.

Read, D., Loewenstein, G., and Rabin, M. (1999), Choice Bracketing, Journal of Risk and Uncertainty 19, 171–197.

Reuben, E., and Riedl, A. (2009), Enforcement of Contribution Norms in Public Good Games with Heterogeneous Populations, Maastricht University, Discussion Paper.

Riedl, A., and Tyran, J.-R. (2005), Tax Liability Side Equivalence in Gift-Exchange Labor Markets, Journal of Public Economics 89, 2369–2382.

Riedl, A., and van Winden, F. (2001), Does the Wage Tax System Cause Budget Deficits?, Public Choice 109, 371–394.

Riedl, A., and van Winden, F. (2007), An Experimental Investigation of Wage Taxation and Unemployment in Closed and Open Economies. European Economic Review 51, 871–900.

Riedl, A., and van Winden, F. (2008), Input versus Output Taxation in an Experimental International Economy, Maastricht University, Working Paper.

Ruffle, B. J. (2005), Tax and Subsidy Incidence Equivalence Theories: Experimental Evidence from Competitive Markets, Journal of Public Economics 89, 1519–1542.

Rustichini, A. (2009), Is There a Method of Neuroeconomics?, American Economic Journal: Microeconomics 1, 48–59.

Samuelson, P. A., and Nordhaus, W. D. (1985), Economics, 12th edition, McGraw-Hill, New York.

Samuelson, P. A., and Nordhaus, W. D. (1992), Economics, 15th edition, McGraw-Hill, New York.

Schelling, T. C. (1968), The Life you Save May be Your Own, in: Chase, S. B. (Eds.), Problems in Public Expenditure Analysis, The Brookings Institute, Washington, DC, 127–162.

Simonson, I. (1990), The Effect of Purchase Quantity and Timing on Variety-Seeking Behavior, Journal of Marketing Research 27, 150–162.

Simonson, I., and Winer, R. S. (1992), The Influence of Purchase Quantity and Display Format on Consumer Preference for Variety, Journal of Consumer Research 19, 133–138.

Slemrod, J. (2006), The Role of Misconceptions in Support for Regressive Tax Reform, National Tax Journal 59, 57–75.

Slemrod, J. (2010), Old George Orwell Got it Backward: Some Thoughts on Behavioral Tax Economics, FinanzArchiv/Public Finance Analysis 66, this issue.

Small, D. A., and Loewenstein, G. (2003), Helping a Victim or Helping the Victim: Altruism and Identifiability, Journal of Risk and Uncertainty 26, 5–16.

Stigler, G. J. (1980), Economics or Ethics?, Tanner Lectures on Human Values.

Thaler, R. H., and Sunstein, C. R. (2008), Nudge: Improving Decisions abouth Health, Wealth, and Happiness, Yale University Press, New Haven and London.

Tversky, A., and Kahneman, D. (1981), The Framing of Decisions and the Psychology of Choice, Science 211, 453–458.

van Winden, F., Riedl, A., Wit, J., and van Dijk, F. (1999), Experimenteel economisch onderzoek naar het Plan Van Elswijk, Eindrapport, CREED, Universiteit van Amsterdam.

van Winden, F., Riedl, A., Wit, J., and van Dijk, F. (2000), Experiment: het Plan Van Elswijk, Economisch Statistische Berichten 10, 197–199.

Wakker, P. P., Timmermans, D. R. M., and Machielse, I. (2007), The Effects of Statistical Information on Risk and Ambiguity Attitudes, and on Rational Insurance Decisions, Management Science 53, 1770–1784.

Arno Riedl
Maastricht University
School of Economics and Business
Department of Economics (AE1)
P.O. Box 616
NL-6200 MD Maastricht
The Netherlands
a.riedl@maastrichtuniversity.nl