Rupert Sausgruber and Jean-Robert Tyran

Discriminatory Taxes are Unpopular. Even when they are Efficient and Distributionally Fair

Discriminatory Taxes are Unpopular

Even when they are Efficient and Distributionally Fair [+]

Rupert Sausgruber[*] and Jean-Robert Tyran[**]

December 2013

We explore the political acceptance of taxation in commodity markets. Participants in our experiment earn incomes by trading and must collectively choose one of two tax regimes to raise a given tax revenue. A "uniform tax" (UT) imposes the same tax rate on all markets and is fair in that it yields the same – but low – income to participants in all markets. The "discriminatory tax" (DT) imposes a higher burden on markets with inelastic demand and is therefore efficient but it is also unfair in that incomes are unequal across markets. We find that DT are unpopular, as predicted. Surprisingly, however, DT remain unpopular when they are both efficient and produce a fair (equal) distribution. We conclude that non-discrimination (equal treatment) is a salient fairness principle in taxation that shapes voting on commodity taxes above and beyond concerns for efficiency and equal distribution.

Keywords: taxation, behavioral public economics, voting, efficiency, fairness

JEL-Codes: C92, H21, D72

[+] We are grateful for financial support by the Austrian Science Fund (FWF) under Project No. S10307-G14

[*] rupert.sausgruber@wu.ac.at, Department of Economics, Vienna University of Economics and Business, Austria.

[**] jean-robert.tyran@univie.ac.at, Department of Economics, University of Vienna, Vienna, Austria and Department of Economics, University of Copenhagen, Denmark, and CEPR (London).

1 Introduction

This paper investigates the political acceptance of discriminatory commodity taxes. By discriminatory commodity taxation we mean a tax regime in which the consumption of some good *i* is taxed at a different rate than that of some other good *j*. We focus on discriminatory commodity (rather than income) taxation not only because it is not uncommon (e.g. excise taxes), but also because a case for the economic desirability of discriminatory commodity taxes is easy to make. The reason is that such taxes can be efficient. In the experimental laboratory, we vary the distributional fairness of discriminatory taxes to test when such taxes are popular with voters, as explained below.

To see that discriminatory taxes can be more efficient than uniform taxes, consider the simple case in which government wants to raise a particular revenue from commodity taxation, and suppose there are only two markets. Suppose that demand in one market is (own-price) elastic, i.e. demand falls strongly when the price of this good increases, while demand in the other market is inelastic (and suppose markets are independent, i.e. the cross-price elasticity is zero). It is easy to see that if the goal of the government is to minimize efficiency losses, government should tax the inelastic market more heavily (the so-called inverse-elasticity rule is a special case of the Ramsey (1927) rule).[1] This type of discriminatory taxation is thus efficient (in the sense that it minimizes the deadweight loss from taxation) but imposes a higher burden on consumers in the inelastic market.

A common argument forwarded to explain why such discriminatory taxes are unpopular despite being efficient is that they often induce unfair distributional outcomes.[2] To continue our example above, suppose that the poor spend most of their income on consuming the inelastic good (food, say), and the rich spend most of their income on the elastic good (a luxury good, say). The simple inverse-elasticity rule discussed above would now imply taxing necessities heavily and luxuries lightly. The tax would be regressive as it imposes a heavy burden on the poor and a light one on the rich. This outcome clearly flies in the face of distributional equity.

[1] The Ramsey rule essentially says that indirect taxes should be designed so as to cause an equi-proportionate reduction in the compensated demands for all commodities. Discriminatory taxes may not be optimal in more complicated settings e.g. when commodity consumption is complementary to labor supply (Atkinson and Stiglitz 1980), see e.g. Sørensen (2010) for a discussion.

[2] An extreme example of discriminatory taxation that is claimed to improve efficiency is to tax body height, i.e. a tall person of a given income should pay more in taxes than a short person of the same income, see Mankiw and Weinzierl (2010).

This paper investigates the political acceptance of discriminatory taxation in an experiment with consumers who participate in a democratic vote. Consumers obtain rents from purchasing units of a commodity in markets. They then vote on two tax regimes to raise a given tax revenue. The "discriminatory tax" (note that we use neutral labels in the experiment) is levied on the market with an inelastic demand only; the market with elastic demand remains tax exempt. The discriminatory tax (DT) is efficient but distributes the tax burden unequally between the subjects. In treatment LOW, it is the "poor" (participants who earn low incomes in the market) who bear the burden of the discriminatory tax. In this sense, DT is efficient but unfair. In contrast, the regime with a "uniform tax" (UT) is inefficient (it reduces total welfare more than the DT) but procedurally fair as it treats everyone in the same way (i.e. UT imposes the same formal tax rate on all participants) and distributionally fair as it yields the same ex-post incomes (i.e. UT eliminates inequality across markets and voters).

We find, unsurprisingly perhaps, that DT is unpopular. To test whether DT is unpopular because of the unfair ex-post distribution of income DT tends to generate, we run a treatment in which it is the "rich" (participants who earn high incomes in the market) who bear the burden of DT (we call this treatment HIGH). We find that DT remains unpopular even when it is both efficient and distributionally fair, and that voters at least initially are strongly biased toward UT. The reason seems to be that the "procedural fairness" of UT (treating everyone the same way) is salient to voters and this fairness aspect dominates concerns for efficiency or distributional fairness in the context of commodity taxation as studied here. With repetition, the effect of concerns for fairness is greatly reduced and most voters eventually vote in line with material self-interest.

Our paper contributes to a small experimental literature on voting to tax markets which can be thought of as a special case of the broader literature on endogenous policies and institutions. While experiments studying the properties of experimental markets are abundant (see Kagel 1995 for a survey) and experimental studies investigating voting and collective choice are common[3], there are only very few papers available which study voting on market parameters, in particular taxation (e.g. Cherry et al. 2012, Großer and Reuben 2013, Sausgruber and Tyran 2011). Note that this innovative combination of two institutions forces us to restrict ourselves to relatively simple market and voting institutions (see section 2). In a broader perspective, our paper also contributes to a literature on endogenous, i.e.

[3] Experimental studies on voting on other topics in political science have recently been so numerous that a new, especially dedicated journal, the Journal of Experimental Political Science, has been created to publish related research. A recent example of a voting study is e.g. Markussen et al. (2013).

democratically chosen, policies and institutions. In particular, we contribute to the literature that investigates how non-standard preferences shape voting outcomes (e.g. Höchtl et al. 2012, Feddersen et al. 2003) and market outcomes (e.g. Fehr and Falk 1999). Within this literature, considerable attention is devoted to studying the role of pro-social orientations (inequality-aversion, reciprocity etc.), only few papers have investigated the role of procedural fairness aspects such as non-discrimination, to the best of our knowledge (e.g. Bolton et al. 2005).

2 Experimental Design

Our design combines a market and a voting experiment. The subjects earn incomes by trading in a market and choose how to tax market transactions by majority vote. Section 2.1 explains the market institution as well as the procedures and rules for voting. Section 2.2 describes the sequence of events and the feedback subjects receive during the experiment. Section 2.3 explains the experimental treatments and discusses our predictions.

2.1 Market and Voting

All subjects participate as consumers in markets in which they can buy units of a hypothetical good from automated sellers. At the start of the experiment, the subjects are randomly assigned to one of two markets and remain in the same market throughout the entire experiment. Both markets are organized as a uniform-price sealed bid/offer auction (see Appendix II for instructions). In every period, consumers bid for units of the commodity they wish to buy.

Market outcomes are determined as follows: Bids are ordered from high to low, all bids exceeding the uniform offers are accepted, and the last accepted bid determines the price for all units traded. Thus, the market clears at a uniform price equal the lowest accepted bid. A buyer's payoff per traded unit is the difference between the induced value and the market price; the payoff is zero, if he does not trade. In our experiment, taxes in both regimes are transaction taxes, i.e. the tax only depends on the number of transactions, not on their value. A tax adds to sellers' cost. Since automated sellers bid their true (tax inclusive) cost, implementing a seller tax of t shifts the supply function up by t. At the end of a trading period, consumers learn the market price and their earnings for that period. A new trading period begins and every consumer starts to bid again.

Using this market institution has three important advantages for our purposes. First, the market is known from previous experiments to quickly converge to equilibrium predictions.[4] Second, trading in our market is simple and is easy to explain to participants, allowing participants to concentrate on the choice of the tax regime which is the main focus of our paper. Third, this market institution allows us to automate sellers which simplifies the analysis.[5]

Figure 1: Market parameters in treatment LOW

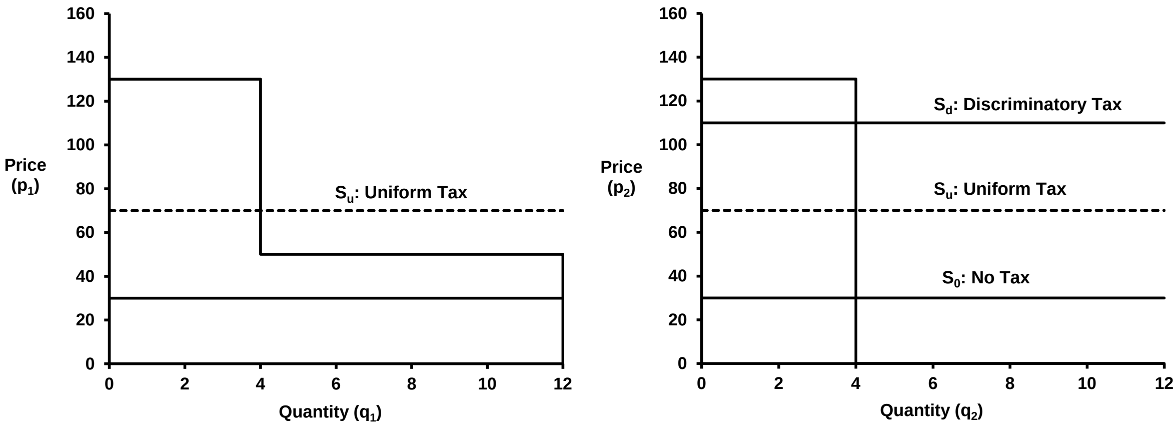

Market 1 with elastic demand *Market 2 with inelastic demand*

Figure 1 shows the parameters in the two markets that operate independently and simultaneously. Each market has four consumers and there are no spillovers between markets whatsoever.

Market 1 (left panel) has elastic demand. Each of the i = 1…4 consumers in this market has redemption values for three units of the commodity. The values are $\nu_{i1} = 130$, $\nu_{i2} =$

---

4 The properties of uniform-price sealed-bid markets are well established in the literature; see Smith et al. (1982) or more recently Sausgruber and Tyran (2011). This literature finds rapid convergence to competitive equilibrium in uniform-price sealed-bid markets if the following holds: (i) privacy, i.e. each agent knows only his own valuation (or cost); (ii) exchange follows the rules described in the main text; (iii) stationary replication, i.e. aggregate market supply and demand is stable for a number of periods; and (iv) deep markets, i.e. there are at least four buyers and as many sellers. Consumers have no incentive to bid their true value in such markets. Yet, convergence to equilibrium is remarkably strong because the consumers have an incentive to bid the price down to marginal cost. In addition, when there is excess supply, human sellers would have an incentive to compete.

5 Automating sellers has a number of advantages. First, such sellers trade actively even when their equilibrium rent is zero. Second, their bidding behavior is perfectly predictable to buyers because they are commonly known to trade according to pre-determined rules (they ask their costs). We can thus avoid erratic bidding which might occur when participants are assigned the role of sellers. Third, we prevent potentially complicated and thus distracting comparisons of relative income between buyers and sellers.

50, and v_{i3} = 50 points. Consumers receive the same redemption values in every period of the entire experiment. A consumer earns the difference between these values and the market price per unit traded in the market. When there is no tax in this market, the supply curve (S_0) is horizontal at 30 points, i.e., the cost per unit is $c = 30$. Demand (D) intersects supply (S_0) at a price of $p_1(t_0) = 30$ points and a quantity of $q_1(t_0) = 12$ units. In equilibrium, where supply intersects demand, total consumer rent is $\Pi_1(t_0) = 560$ points and every individual consumer earns a rent of $\pi_1(t_0) = 140$ points. Consumer rents in points are converted into money at a rate of points 100 = €0.3 and paid out in cash to the subjects at the end of the experiment.

Market 2 (right panel) has inelastic demand. Again, there are four consumers $j = 1 \ldots 4$ in the market. Each of them has only one redemption value at $v_j = 130$ points (in all periods). Pre-tax supply (S_0) intersects demand at a price of $p_2(t_0) = 30$ points and a quantity of $q_2(t_0) =$ 4 units. In market 2, the total rent absent taxation is $\Pi_2(t_0) = 400$ and every consumer earns $\pi_2(t_0) = 100$ points, i.e. about 30% less than consumers in market 1.

Participants start trading in untaxed markets for 5 periods. After period 5, the experimenter halts trading and informs subjects that a tax revenue of given size has to be raised in each period for the remainder of the experiment by taxing market transactions. Subjects are told that they have the choice between two tax regimes (see Appendix III for the voting proposal).[6]

The "uniform tax" regime (UT) imposes the same tax $t_u = 40$ to transactions in both markets.[7] The tax increases the unit cost by 40 points and shifts the supply schedule from S_0 to S_u in both markets (see dashed line S_u in both panels of Figure 1). We chose the parameters such that the UT is distributionally fair, i.e. induces equal market incomes in both markets in equilibrium. In particular, after-tax equilibrium prices and quantities are the same in both markets: $p_1(t_u) = p_2(t_u) = 70$ and $q_1(t_u) = q_1(t_u) = 4$. The total tax revenue per period is $T(t_u) =$ 320. Total consumer rents are therefore equal in both markets: $\Pi_1(t_u) = \Pi_2(t_u) = 240$ and all consumers receive the same rent in equilibrium $\pi_{i1}(t_u) = \pi_{j2}(t_u) = 60$. UT is thus a "fair" tax in the sense that it eliminates the inequality that prevailed without tax (pre-tax equilibrium rents per capita in the two markets are 140 and 100, respectively). However, UT is an inefficient tax in the sense that it reduces the sum of payoffs.[8] In fact, the sum of the after-tax consumer

6 In our experiment, tax receipts are wasted. For experiments that use tax receipts for providing public goods or redistribution see, for instance, Güth and Mackscheidt (1984) and Güth at al. (2005).

7 The instructions generally used neutral labels. For example, we call the taxes in the two regimes "alternative 1" and "alternative 2" rather than "uniform" and "discriminatory" as in the main text.

8 This notion of efficiency is sometimes called Kaldor-Hicks efficiency.

rents and the tax revenue ($\Pi_1(t_u) + \Pi_2(t_u) + T(t_u) = 800$) is smaller than the sum of consumer rents before tax ($\Pi_1(t_0) + \Pi_2(t_0) = 960$). This deadweight-loss ($\mathrm{DWL}(t_u) = 960 - 800 = 160$) arises because buyers in the elastic market 1 are forced to reduce their consumption as they no longer can profitably trade their low-valuation units at the after-tax gross price.

The "discriminatory tax" regime (DT) imposes a tax $t_{d2} = 80$ on transactions in the inelastic market 2 and leaves transactions in the elastic market 1 untaxed ($t_{d1} = 0$). The after-tax supply (S_d) in market 2 shifts vertically to $c + t_d = 110$ (see solid line S_d in the right panel of Figure 1). The after-tax equilibrium price is $p_2(t_d) = 110$ and consumer rents in this market shrink to $\Pi_2(t_d) = 80$. Because market 1 remains untaxed (and because markets are independent), outcomes remain at the initial levels in market 1, i.e., $p_1(t_0) = 30$ and $q_1(t_0) = 12$. Total revenue from imposing the discriminatory tax on market 2 is $T(t_d) = 320$, which is exactly the same as under the uniform tax. DT differs from UT both with respect to equality of incomes and with respect to efficiency. While UT eliminates income inequality (compared to the untaxed state), DT exacerbates it (per capita equilibrium rents are 140 and 20 in market 1 and 2, respectively). But DT is more efficient than UT. DT avoids distortions of market outcomes, i.e., the tax does not reduce total welfare: $\mathrm{DWL}(t_d) = 0$ (because $\Pi_1(t_{d1}) + \Pi_2(t_d) + T(t_d) = \Pi_1(t_0) + \Pi_2(t_0) = 960$).

Voting is organized such that the four consumers in market 1 and the four consumers in market 2 simultaneously vote between the two alternatives. Voting is compulsory (no abstentions) and costless. Simple majority decides. In case of a tie (4:4 votes) the uniform tax serves as the default, i.e., the discriminatory tax DT is imposed only if it wins a majority of votes (i.e., 5 or more).[9] The instructions tell the subjects that the tax will add to the cost of the sellers. We explicitly inform the subjects that a seller will accept a bid only of it equals or exceeds the seller's cost plus the tax which applies in the market. With this feature of the design, the subjects can easily anticipate that the tax adds to the pre-tax market price.

2.2 Sequence and feedback

Figure 2 illustrates the sequence of phases in our design. An experimental session has 10 main phases with taxation which follow the phase 0 without taxation. Each main phase consist of a vote and a subsequent trading in the market.

---

9 We chose not to break ties randomly to obtain a sharper statistical comparison of acceptance rates across treatments by removing noise.

In phase 0, the consumers learn the total number of transactions in their market and the market price after each trading period. They also receive information about their period earnings and the number of their individual transactions.

Figure 2: Sequence of decisions in the experiment

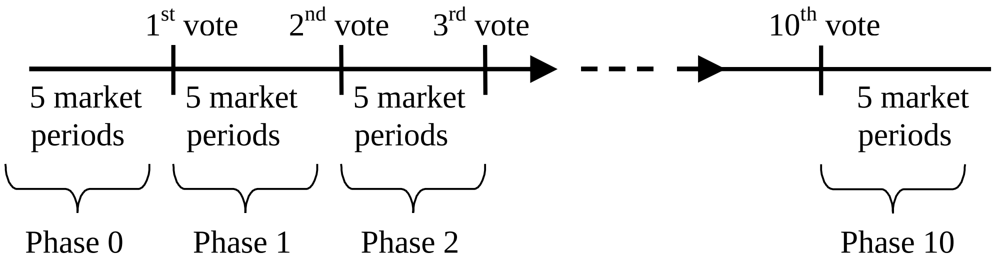

In the main phases (phase 1 and later), participants first vote on the two tax regimes. As feedback from voting, subjects learn the total number of votes for each alternative; they do not learn whether the votes for a particular alternative come from their own or the other market. The outcome of the vote is binding for trading in the next 5 market periods. This procedure is repeated 10 times, i.e. every subject votes 10 times.

At the end of each period in the main phases, subjects learn the market price and the total quantity traded in their market. They receive information about the number of own trades and own profit. At the end of a main phase, subjects in addition learn how the regime affected other participants in their own market, in the other market, and in both markets as a whole. In particular, the subjects are reminded of the tax that applies in a given market and they are informed about the corresponding average per-capita earnings over the last 5 periods in their own market, in the other market. In addition, subjects learn the per-capita tax payments and earnings averaged over both markets over the last 5 periods. Provided that there is some variation in aggregate voting outcomes, subjects can use this information to compare the welfare cost of taxation. To make this information more salient and to facilitate comparisons across tax regimes, subjects are required to copy the information on the feedback screen on a separate form at the end of each phase. We note however, that buyers ex ante have incomplete information about the distribution of the tax burden between the markets. The information conditions in our experiment are common to market experiments and mimic the fact that participants possess only private information about market parameters.

Consequently, buyers have to anticipate the effects of taxation with uncertainty, but they can learn to the extent that they gain informative experience.

2.3 Treatments

The design has two treatments. Treatment LOW implements the parameters as described in the previous section (see Figure 1). Treatment HIGH is the same as LOW but increases the redemption values of the four consumers in market 2 from $v_{j(\mathrm{LOW})} = 130$ to $v_{j(\mathrm{HIGH})} = 250$. Because this increase does not change equilibrium prices and quantities in either market, this increase simply implies that the consumers in market 2 have higher incomes in equilibrium in HIGH than in LOW (hence the treatment labels).

The purpose of the treatment variation is to test whether the reason for the (expected) low popularity of efficient commodity taxes is to be found in the distributional inequality it typically involves. If so, we should see that the discriminatory tax DT is more popular in HIGH than in LOW. The reason is that DT is both efficient and “fair” (in the sense of burdening “rich” voters more and eliminating income inequality across all participants) in HIGH but is efficient and “unfair” (in the sense of increasing income inequality) in LOW.

Table 1 shows the payoff consequences of the tax regimes in the two treatments. The numbers in the cells are the equilibrium rent in points per consumer and period. In section 3 we show that the markets in our experiment quickly converge to the equilibrium predictions. The numbers in Table 1 therefore closely approximate the per-period earnings of subjects in our experiment.

In treatment LOW, the uniform tax equalizes after-tax earnings at a level of 60 points, see column (2) of the table. In contrast, consumers earn 140 points in market 1 and 20 points in market 2 under the discriminatory tax, see column (3) of the table. Notice again that the uniform tax is inefficient: the sum of rents is smaller under UT than DT.

In treatment HIGH, under the uniform tax consumers earn 60 points in market 1 and 180 points in market 2 (see column 5). Now it is the discriminatory tax that equalizes the rents between the markets at 140 points (see column (6) of Table 1). Note that, except for the higher rents of consumers in market 2, treatments LOW and HIGH are identical. In particular, the welfare loss from uniform as compared to discriminatory taxation is held constant at 160 points, or 50% of the tax revenue, across treatments.

Table 1: Equilibrium consumer rents (in points, per period) by market, treatment, and tax regime (UT = Uniform Tax, DT = Discriminatory Tax)

	LOW			HIGH		
	(1) No tax	(2) UT	(3) DT	(4) No tax	(5) UT	(6) DT
	$t_1 = t_2 = 0$	$t_{u1} = t_{u2} = 40$	$t_{d1} = 0, t_{d2} = 80$	$t_1 = t_2 = 0$	$t_{u1} = t_{u2} = 40$	$t_{d1} = 0, t_{d2} = 80$
Market 1 4 consumers (elastic demand)	140	60	140	140	60	140
Market 2 4 consumers (inelastic demand)	100	60	20	220	180	140
∑ Consumer rents	960	480	640	1440	960	1120
Tax revenue	0	320	320	0	320	320
Deadweight Loss (DWL)	0	160	0	0	160	0
Total Rents	960	800	960	1440	1280	1440

2.4 Discussion of the Design

According to standard economics, the discriminatory tax is (narrowly) rejected in both treatments. Assuming that all voters are rational and self-interested and vote sincerely, the votes will be tied because material interests of the voters conflict across markets. Consumers in market 1 earn less under uniform taxation (UT) than with DT whereas the opposite is true for consumers in market 2 (see Table 1). All voters in market 1 are thus predicted to vote for DT, and none vote for DT (but all for UT) in market 2.[10] A tie therefore occurs at 4:4 votes, and given our voting rule, the uniform tax UT will be implemented as the default. The information conditions of the experiment are sufficient to predict voting of rational and self-interested subjects.[11] The same predictions apply for both treatments.

[10] For self-interested subjects, it is weakly dominant to vote for the alternative that offers the highest income. However, voters may not be pivotal and equilibria exist in which voters do not vote for their preferred alternative. We therefore assume that each voter faces a positive probability of being decisive and votes to maximize expected payoff. This assumption is reasonable because in our experiment the electorate is small and the probability of being pivotal is high.

[11] To obtain this prediction, we need to assume that buyers expect with positive probability market prices to increase in response to higher taxes. This assumption is very likely met in our experiment for a number of

The predictions discussed above are derived using the assumption of rational and self-interested agents who do not harbor concerns for fairness or efficiency in and of itself. But it is easy to think of non-standard, social preferences that may tip these predictions toward an efficient outcome (for a recent survey of experimental results regarding social preferences, see Cooper and Kagel 2013). Below, we discuss aversion to inequality, a taste for efficiency, and a preference for "fair treatment" (avoiding active discrimination) as additional motivations that may influence the popularity of the tax regimes.

First, support for discriminatory taxation can be rationalized by assuming a concern for equal distribution. To illustrate, suppose that subjects are averse to inequality in the distribution of after-tax market incomes (e.g. Fehr and Schmidt 1999, Bolton and Ockenfels 2000). In LOW, equilibrium payoffs in the two markets are the same under UT and different under DT. In HIGH, the opposite is true. Therefore, if subjects have a plain taste for equality, one may expect higher support for the discriminatory tax in HIGH than LOW. The prediction based on inequality aversion is subtle, though. To see this, note that consumers in market 1 who care for equality would have to give up 80 points for themselves to avoid a loss of 40 points for consumers in market 2. Such highly costly choices are empirically implausible and it is unlikely that the consumers in market 1 support the uniform tax because of inequality aversion (see e.g. Blanco et al. 2011).[12] Inequality-aversion motives are more plausible to make a difference for voting of consumers in market 2. For them, the material self-interest overlaps with a preference for equality in LOW, whereas the two motives pull in opposite directions in HIGH. Hence, inequality aversion on behalf of voters in market 2 implies that the discriminatory tax receives more support in HIGH than in LOW. This argument may matter for aggregate voting outcome because the standard model predicts a close defeat of the discriminatory tax. Under these conditions, a few fairness-minded voters from market 2 suffice to tip the balance (see Tyran and Sausgruber 2005 and Höchtl et al. 2012).

Second, a taste for efficiency in the sense of maximizing total market rents is an obvious candidate to rationalize support for DT. For example, a pivotal voter in market 2 who

---

reasons. In particular, buyers learn from experience with trading in phase 0 that bids are rejected if they fall below the sellers' unit cost, and buyers are told that the tax adds to this cost. Because the unit cost is constant, buyers typically experience a loss in earnings in response to a reduction of quantity traded in the market. Furthermore, the proposed taxes are large; they are particularly large enough to outsize even unusually large deviations from competitive equilibrium predictions.

12 In Fehr and Schmidt (1999), the parameter β_i captures aversion to "favorable inequality" (i.e. to being ahead). The parameter reflects the weight a subjects puts on the difference between own and another subject's payoff. To rationalize a choice as that for voters in market 1 in favor of the uniform tax, β_i would need to large (see Tyran and Sausgruber 2005). Blanco et al. (2011) find that such extreme inequality aversion is exceedingly rare.

supports the discriminatory tax loses 40 points, but another voter in market 1 gains more (80 points) which means that total market rents go up. In our design, a taste for efficiency would therefore bias choices in favor of the discriminatory tax.[13] Such a taste can be modeled in various ways. For example, Charness and Rabin (2002) assume that a subject's utility does not only depend on its own payoff, but also on the minimum of payoffs in a reference group, as well as on the sum of payoffs in this group. A concern for efficiency has received some empirical support in experimental studies, e.g. by Engelmann and Strobel (2004) in simple distributional games and by Kerschbamer et al. (2013) in markets for credence goods.[14]

Finally, there is a growing number of studies in economics suggesting that people not only care about outcomes of their interactions but also about the fairness of procedures (see, e.g., Babcock et al. 1995, Bolton et al. 2005, Chlaß et al. 2009, Krawczyk 2011, and Güth et al. 2013).[15] In our design, subjects are randomly assigned to markets and to different redemption values which implies random differences in equilibrium rents. One might therefore consider the pre-tax allocation as procedurally fair because everyone has the same chance to be member of market 1 or 2. Moreover, many voters may find it difficult to judge the fairness of outcomes in early votes, when information is still incomplete. In this case, voting for the discriminatory tax would discriminate against the consumers in market 2 and a taste for procedural fairness could bias voting in favor of the uniform tax.[16]

3 Results

In total, 192 undergraduate student subjects from the University of Innsbruck participated in our experiment, and about two thirds of them had some basic background in economics. We ran 12 pairs of markets in treatment LOW and 12 pairs in treatment HIGH. Each pair consisted of 8 subjects who were randomly assigned to market 1 or market 2. The experiment was split into 12 sessions with 16 subjects each. A typical session lasted for about 90 minutes.

---

[13] Note that the bias would exist irrespective of the treatments because the treatments do not change the marginal incentives to choose one alternative over the other.

[14] Note that in our setup, maximin-preferences (according to which people also care about the poorest member in society) would result in the same deviation from opportunistic behavior as inequality aversion because all consumers within a market earn the same income.

[15] Considerable research on procedural justice and legitimacy is available in law and social psychology. See, for example, Tyler and Lind (2000). In the context of taxation, Wenzel (2002) finds that perceptions of procedural fairness also matter for tax compliance.

[16] Again, marginal incentives do not vary the treatments and a concern for procedural fairness can therefore not explain differential voting across treatments.

On average, a subject earned € 20, including a show-up fee of € 4. The experiment was programmed using the software z-Tree (Fischbacher 2007).

We present our results as follows. Section 3.1 essentially shows that markets converged very well to equilibrium predictions and that welfare consequences of the tax regimes are as theoretically expected. Section 3.2 essentially shows that the uniform tax (UT) was more popular than predicted but that voting converged to the standard theory prediction with repetition. Section 3.3 discusses learning from experience and section 3.4 provides a broad discussion of results.

3.1 Market Outcomes

Figure 3 shows transaction prices over time, averaged over all 24 markets of each type. When markets 1 remain untaxed (in phase 0 or because DT is accepted in the main phases) the predicted price is $p_1(t_0) = 30$ (see Figure 1). In markets with a uniform tax, the predicted prices are $p_1(t_u) = p_2(t_u) = 70$ for both markets 1 and 2. Finally, in markets 2 with the discriminatory tax implemented the price prediction is $p_2(t_d) = 110$. The figure reveals that realized transaction prices quickly converge (by design from above) to these predictions. Quantity traded also quickly converges to equilibrium. For example, in the final period of phase 0 buyers exploit between 93% and 95% of all opportunities to trade at a surplus in market 1 and market 2, respectively.

When realized prices and quantities are close to equilibrium, so are rents. In treatment LOW, the average earnings across all periods and subjects in market 1 are 129.2 points with the DT and 53.4 points with UT. The respective numbers for market 2 are 56.1 and 16.5 points. In HIGH, the average earnings in market 1 are 130.5 and 54.7 points with DT and UT, respectively. In market 2 the respective numbers are 171.1 and 131.1 points. These numbers are sufficiently close to the equilibrium predictions (see Table 1) such that our discussion of predictions for voting which was based on equilibrium considerations need not be reconsidered given the realized values. In particular, a materially self-interested voter would always prefer the discriminatory tax in market 1 and reject it in market 2, and this is true in both treatments.

Figure 3: Transaction prices (averaged over 24 markets per type)

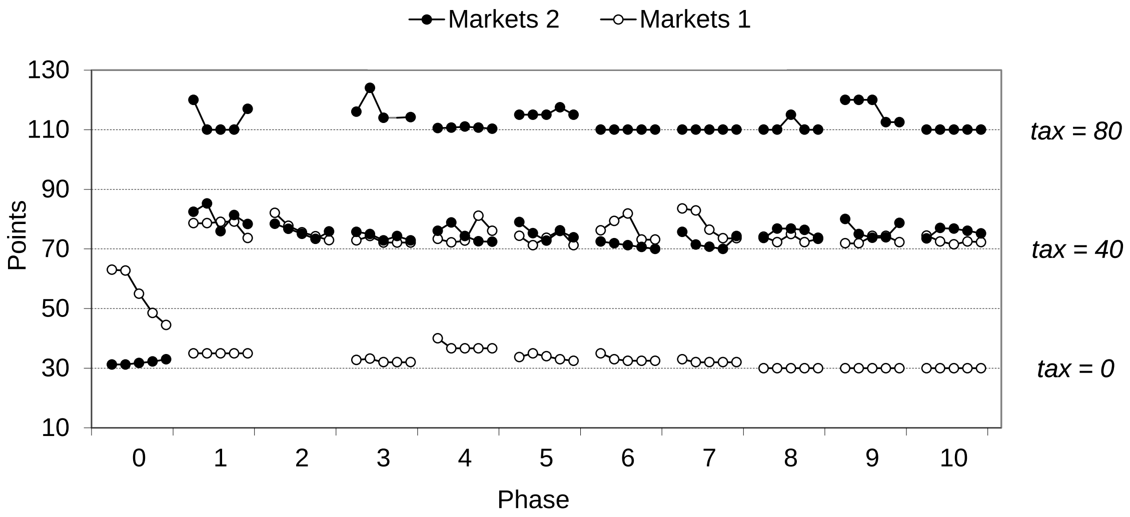

3.2 Voting Outcomes

Table 2 shows the share of individual votes for the two tax regimes by market in the first two columns of each panel. The third column shows aggregate acceptance rates, i.e. the share of referenda that lead to acceptance of the discriminatory tax (DT), averaged over all phases. Overall, two facts stand out. First, as predicted by standard assumptions, DT generally receives strong support from voters in market 1 and weak support from voters in market 2. Conversely, the uniform tax (UT) is more popular in market 2 than 1. Second, many votes are inconsistent with self-interested voting. This is particularly the case in market 1. For example, 31% of voters in market 1 in treatment LOW vote against their material self-interest by supporting UT (which yields an equilibrium payoff of 60 rather than 140 as with DT).

Our hypothesis is that making the efficient DT also distributionally fair and imposing the burden on the "rich" (in treatment HIGH) increases the popularity of DT. On average, support for this hypothesis is weak, however. The average share of individual votes in favor of DT (across both markets), is 45.1% (= 433/960) in LOW and 46.7% (= 448/960) in HIGH. These shares are not significantly different from each other (p = 0.434).[17] The discriminatory tax is rarely accepted by a majority. In fact, DT is approved by a majority in 13% (= 15/120) of all votes in LOW and in 20% (= 24/120) in HIGH. The effect on the aggregate is only

[17] Wilcoxon rank-sum (WRS) test using 12 independent observations per treatment (voting decision per market averaged over all periods).

weakly significant and only if based on a one-sided test (p = 0.058).[18] Moreover, the higher support for DT in HIGH (vs. LOW) is driven exclusively by voting in market 1. In this market, the average approval rate increases from 69% to 84% between LOW and HIGH (p = 0.014 WRS). In stark contrast, voters in market 2 vote less frequently for the discriminatory tax in HIGH than LOW (9% vs. 21%, p = 0.021, WRS).

Table 2: Approval rates (numbers in parentheses show equilibrium market rents)

	LOW			HIGH		
	Individual votes		Aggregate support	Individual votes		Aggregate support
	Uniform UT	Discrim. DT	Discrim. DT	Uniform UT	Discrim. DT	Discrim. DT
Market 1 4 consumers (elastic demand)	31% (60)	69% (140)	13%	16% (60)	84% (140)	20%
Market 2 4 consumers (inelastic demand)	79% (60)	21% (20)		91% (180)	9% (140)	

In our discussion in section 2.4 we have argued that inequality aversion is unlikely to matter for voting in market 1, as for them in LOW voting for the uniform tax (i.e., voting in line with inequality aversion) is very costly in own monetary terms. The question is then why the uniform tax nonetheless is popular for many voters in market 1.

Table 2 provides a first answer to this question. If we consider the votes which deviate from standard predictions in LOW, we see that voting is biased toward uniform taxation (compare 31% and 21% between the cells Market 1 & UT and Market 2 & DT). The same pattern applies in HIGH (compare 16% and 9% in the cells Market 1 & UT and Market 2 & DT, respectively). This result is surprising: it holds although UT is more costly in own monetary terms for voters in market 1 than is DT for voters in market 2; in addition, it holds irrespective of which tax regime leads to equal after-tax incomes. This observation implies that the uniform tax is appealing to many subjects irrespective of its effects on the level and

18 One-sided two-sample test of proportions.

distribution of after-tax incomes. Procedural fairness perceptions are a likely candidate to explain this pattern.

Figure 4 shows that there is a pronounced bias toward the uniform tax (UT) early in the experiment, but that this bias is almost eliminated by the end of the experiment. In our experiment, subjects voted 10 times on the same tax proposal. Figure 4 shows the individual and aggregate approval rates per phase, market and treatment (Appendix I shows the raw data of this figure). Recall that based on material self-interest, 100% of votes in market 1 and 0% in market 2 should be in favor of DT. The most striking observation from Figure 4 is the large discrepancy of this prediction for voters in market 1 in the early phases of the experiment.

In market 1 in LOW, only 38% of voters (= 18/48) approve of DT in the first vote. The support increases with repetition and the choices clearly converge toward the standard prediction. Despite this trend in LOW, 19% (= 9/48) of subjects vote against DT even in the last phase of the experiment. In HIGH, acceptance of DT in market 1 starts out higher in phase 1 at 52% (= 25/48) compared to LOW. From there it converges closer to the predicted value than in LOW, to 96% (= 46/48) in the last phase.

Figure 4: Support for the discriminatory tax (DT)

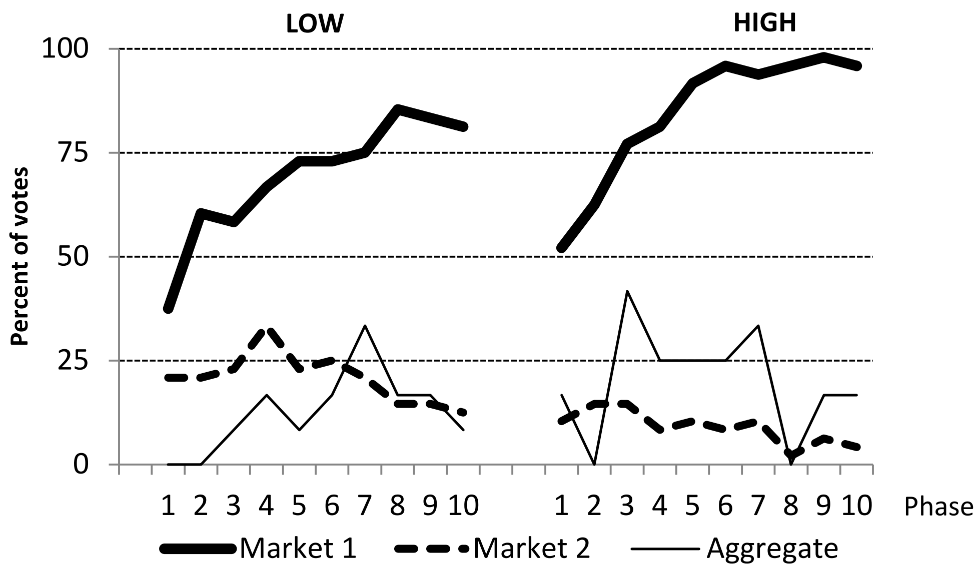

In market 2, voting behavior that is inconsistent with standard predictions could be explained by a concern for efficiency as well as errors. Note that inequality aversion is unlikely to drive non-egoistic voting on behalf of voters in market 2, as then approval rates would be higher in HIGH than in LOW.[19] Generally, voting is closer to the benchmark prediction in market 2 than in market 1.[20] Again, this observation indicates that voting is biased toward the uniform tax.[21]

Why does the bias toward the uniform tax decrease with repetition? One possible explanation is that subjects initially care more about the norm of treating everyone alike because, absent explicit information about market outcomes, the procedural fairness norm is more salient. As the experiment proceeds, subjects receive feedback that emphasizes the consequence of the tax on market outcomes, in particular, the level and distribution of incomes across the two markets. We believe that as this information becomes more salient, the behavioral relevance of procedural fairness norm erodes.[22]

3.3 Effects of Informative Experience

Our results show that voting is biased toward treating everyone in the same way, i.e. taxing uniformly, perhaps because voters think about taxing everyone equally as an aspect of procedural fairness. We have arrived at this conclusion because we find that voters deviate by more from the benchmark prediction in market 1 than in market 2 irrespective of the level and distribution of market rents across treatments. This bias towards uniform taxation is more pronounced in early phases than after subjects have gained some experience. In what follows, we therefore discuss how feedback and learning from experience may have shaped our

---

[19] We can think of two explanations for why voters in market 2 exhibit a lower acceptance of the discriminatory tax in HIGH than in LOW. First, there is some evidence that people who are rich relative to others contribute less to public goods (Buckley and Croson 2006) or donate less to charities (List 2011). These observations suggest that altruism may not be independent of status concerns. Second, the subjects in our experiment receive feedback about voting outcomes after each phase. Part of the dynamic we observe may therefore be driven by reciprocity. This explanation appears less plausible, though, because market-2 voters are less likely to vote for the discriminatory tax in HIGH than in LOW already in the very first phase of the experiment.

[20] The difference is significant ($p = 0.041$, according to a WRS-test using all 24 independent observations from both treatments i.e., using voting decisions per market averaged over all periods).

[21] Fehr et al. (2006) provide evidence that efficiency motives are in part driven by a subject pool effect according to which students in business and economics have a particularly strong taste for efficiency. In the light of their finding, the low and decreasing support of DT by voters in market 2 is all the more surprising. Our subject pool consists of two thirds students with basic training in economics, and one would therefore expect the taste for efficiency to be particularly pronounced in our experiment.

[22] Roth and Malouf (1979) find evidence suggestive of this interpretation. In their experiment, players bargained over lottery tickets that determined the probability with which each player wins a monetary reward. In the first condition the players only knew their own reward; in a second condition they had full information about their own and the opponent’s reward. Hence, the first condition highlighted procedural justice while the second condition put additional emphasis on allocation fairness. The result was that almost all players agreed on an equal division of lottery tickets in the first condition while the outcomes in the second condition shifted substantially toward an equal division of expected rewards.

results. However, experience is the result of (voting) choices, and because groups select into tax regimes, experience is endogenous and its effects cannot be interpreted as causal.

A first fact to note is that experience accumulates slowly and asymmetrically across groups. In phase one, less than 10 percent (2/24 markets) experience DT. By the middle of the experiment, half of the markets have seen both DT and UT, and the other half has only seen UT. In the final phase, three quarters (18/24) of the markets have experienced both UT and DT, but one quarter has never accepted DT. The subjects who have experienced both tax regimes can readily compare the distribution and welfare consequences of the two tax regimes and their choices can be taken as an informed expression of their truly preferred option. Yet, these subjects may systematically differ in relevant characteristics, such as tastes or cognitive skills, from those who have never experienced DT which means that simple comparisons will not reveal a truly causal effect.

Figure 5 presents the voting data over time broken down by experience. The time line shows renormalized phases such that zero is the time at which a group implemented DT for the first time. The grey bars (measured on the right scale) show the number of markets observed at a particular phase and time distance from that event. Left of zero we see the data from all markets and phases prior to experiencing DT. By construction, the markets shown in the left half of the diagrams only experienced the results of UT. The right half of the diagrams (i.e. the positive domain) shows the data from all markets and phases that had experienced DT previously.[23]

Figure 5 conveys two main messages. First, absent explicit information about the effects of DT vs. UT, voting is biased toward UT. This is particularly clear in LOW for voters in Market 1 (the black line in the left part of the upper diagram is far below 100%). The effect is somewhat weaker in HIGH, and there is also some deviation from self-interested voting in Market 2 (the dotted line in the upper diagram is above 0%). However, the net bias is clearly toward UT as can be seen by netting the vertical distance between the black line labeled "Market 1" and 100% and that between the dotted line labeled "Market 2" and 0%. Second, figure 5 suggests a strong and immediate effect of feedback on voting. Left of zero, when

[23] Here is an illustration of how to read the figure: At -4 on the horizontal axis, for example, we see the voting behavior (measured on the left scale labeled "Approval of DT in %") of all inexperienced markets four phases before they implemented DT for the first time. At -4, the height of grey bars (measured on the right scale labeled "Number of markets") is 9 in LOW and 5 in HIGH. This means that there were 9 markets in LOW and 5 markets in HIGH who were inexperienced in phase t, but implemented DT in phase t+4. By construction, the behavior of the remaining groups (3 groups in LOW and 7 in HIGH) who have at the same point in time implemented DT at least once in the past is indicated at value 6 (i.e., 10-4) on the horizontal axis.

voters are inexperienced, between 37.5% and 75% of voters in market 1 vote for DT and there is no clear trend. In sharp contrast, voting quickly converges to the money-maximizing choice when voters are experienced with both regimes, to the right of zero.

Figure 5: Experience and support for the discriminatory tax (DT) (phases renormalized, 0 = first phase of implementing DT)

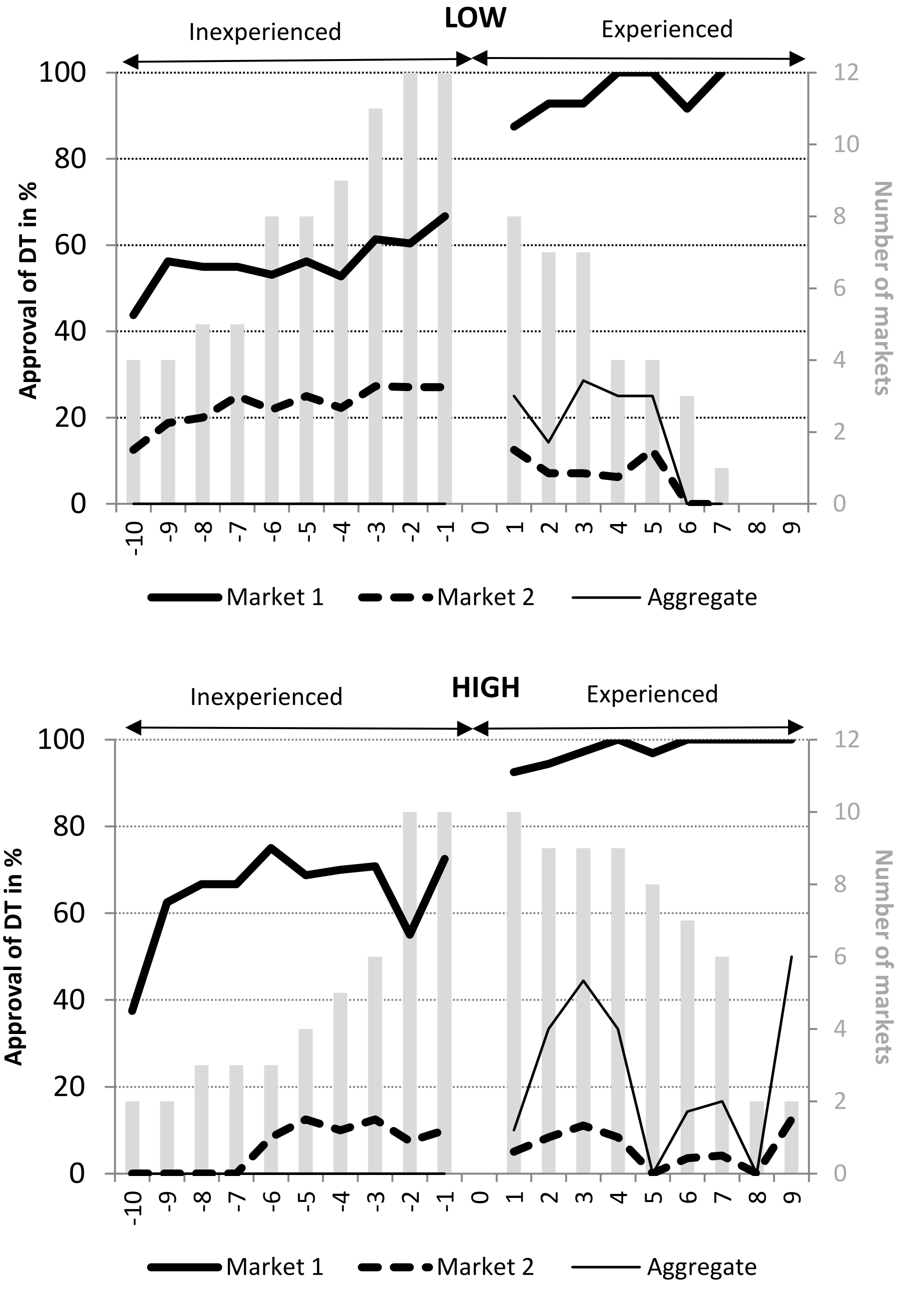

Note: Bars measure the number of markets on the scale to the right, while lines measure approval rates in the scale to the left.

Treatment effects are weak if not entirely absent. In market 1, approval of inexperienced voters is slightly lower in LOW than HIGH (by 7.3 percentage points on average, p = 0.061); there is no treatment effect for experienced subjects though (2.4 points, p = 0.250, according to a regression-based t-test).[24] In market 2, outcomes do not differ between treatments, neither for inexperienced (average percentage point difference is -14.5, p = 0.433) nor experienced ones (-13.5, p = 0.160).[25] Finally, aggregate approval in markets with experience does not differ across treatments (averaged across phases past experience: 18.3% vs. 19.5%, p=0.462 WSR).

Our analysis of experience effects confirms our previous conjecture that subjects initially care more about the norm of treating everyone alike. As indicated by Figure 5, few markets accept DT already at early phases of the experiment (see the grey bars at the tails of the time line), and a sizeable number of markets remains without experience even for an extended period of time. As voting moves closely in line with predictions only once subjects have experienced the effect of DT, this result implies that voters who reject DT because they find it procedurally unfair have less opportunity to learn. This finding provides a first systematic reason for why a discriminatory tax may be generally unpopular.

Interestingly, neither procedural fairness nor any norm of outcome-based fairness seems to play a role once the subjects have received salient feedback on the consequences of the tax regimes. We observe that experienced subjects vote in line with self-interest in both treatments LOW and HIGH. As a consequence, the efficient tax gets rejected even when it is distributionally fair. One possible explanation for this surprising finding is that the cost of obeying a particular fairness norm is highly asymmetric between the markets, and voters are likely to learn that this is so. Experience, for example, has taught voters in market 1 that voting for uniform taxation hurts, as it imposes substantial costs on themselves. The disutility from bearing this cost is likely to swamp concerns for procedural fairness and equality. In contrast, voting for UT, i.e. voting in accordance with procedural fairness, is costless for subjects in market 2. For them, norms of fair treatment may unfold to an extent that overrides

[24] We use a linear probability model (with robust standard errors clustered in pairs of markets) and control for a time trend, a proxy for pivotality, a voter's past voting behavior and past profit (see Table 3 and the explanations given there for the variables). To test for potential treatment effects on experienced vs. inexperienced subjects, we include a treatment dummy and condition on whether subjects had DT implemented or not.

[25] In HIGH, voters in market 1 observe that they earn less than voters in market 2. Hence, inequality aversion may explain why voters are more likely to vote for DT in HIGH than in LOW even absent any experience with DT. As discussed in section 3.2, the results of voters in market 2 are difficult to reconcile with inequality aversion.

other norms based on fair outcomes. In any case, in our design only few voters with non-standard preferences in the guise of inequality aversion and a taste for efficiency would suffice to tip the vote against DT. If this does not happen in our experiment, it is difficult to believe that these types of preferences would increase the popular support of discriminatory taxes in situations where fair voters are much less likely to be decisive. This finding provides a second systematic reason for why a discriminatory tax may be generally unpopular.

3.4 Further Discussion of Results

We now explain that expressive voting may account for why some vote against their material self-interest (i.e. non-sincerely) and that expressive voting has an asymmetric effect across markets. This is so because of an asymmetry between voters in market 1 (voters1) and 2. In particular, because of the default rule we use, voters in market 2 are more likely to be pivotal than those in market 1.

To illustrate, suppose that voters in market 2 always stick to the standard prediction (they earn more with UT than DT). In this case, the DT is never implemented regardless of the choices by voters in market 1, i.e. voters in market 1 are never pivotal in this case. Conversely, suppose that the voters in market 1 always vote sincerely for DT (they earn more with DT because they remain untaxed). In this case, every voter in market 2 can tip the balance and induce acceptance of the discriminatory tax, i.e. voters in market 2 are likely to be pivotal. Voting against one's materially preferred choice is therefore more costly in expectation in market 2 than in market 1 (assuming that utility is linear in earnings, this holds as long as the pivot probabilities are at least twice as large in market 2 as in market 1).

The key idea of expressive voting theories is that there is a trade-off between expressing support (i.e. voting) for some "good cause" and the private (expected) cost of implementing the good cause. For example, the act of supporting environmentally friendly energy production may give voters some utility but if the policy is then implemented and energy is indeed produced in environmentally friendly ways, its cost (in the guise of higher electricity prices, say) may easily outweigh its benefits to this voter. But the cost of expressing support for the morally worthy cause depends on how likely it is that a particular voter affects the outcome, i.e. whether he is pivotal (see Brennan and Lomasky 1993). Voters would then be more likely to vote "morally" when they are unlikely to be decisive for the outcome, and the existence of sufficiently many of these voters may induce a "moralistic bias" (see Tyran 2004 and Feddersen et al. 2009 for experimental tests).

In the context of our experiment, both voting for DT (because of a concern for efficiency) in market 2 and voting for UT in market 1 (because of a concern for procedural fairness) could be thought of as a "moral" obligation. Note that the cost of acting "morally" in market 1 are twice as large as those in market 2 (40 vs. 80 points) but that because of high pivotality in market 2, the expected costs of voting morally are higher in market 2. Therefore, the theory of expressive voting predicts more insincere voting in market 1.

Table 3: Probit model (marginal effects): Dependent variable shows 1 for vote in favor of the discriminatory tax, and 0 else.

	Market 1	Market 2
Pivot $_{p-1}$	-0.049	-0.071**
	(0.050)	(0.036)
HIGH	0.059**	-0.259**
	(0.027)	(0.110)
Vote $_{p-1}$	0.415***	0.542***
	(0.062)	(0.068)
Profit $_{p-1}$	-0.001	0.002**
	(0.001)	(0.001)
Phase	0.005	-0.001
	(0.006)	(0.005)
Experience	0.155***	-0.047
	(0.033)	(0.039)
Observations	864	864
Pseudo R-squared	0.36	0.32
Wald	378.83	122.01

Notes: Robust standard errors (adjusted for clustering in pairs of markets) in parentheses.
* significant at 10%; ** significant at 5%; *** significant at 1%

Table 3 shows the determinants of the likelihood of voting for DT (more precisely, the marginal effects of a Probit regression) to confront this argument with the data. The variable Pivot_{p-1} indicates whether a particular voter has been pivotal in the vote of the previous phase.

We use this variable as a proxy for whether the voter expects to be decisive in the current phase. Specifically, the variable is set to 1 if there have been either 5 pro-votes including the own vote or 4 votes excluding the own vote in the previous phase, and set to 0 else. Hence, the variable indicates whether a voter would have changed the outcome, had he voted differently.[26] Counting the outcome of this variable per treatment reveals that the number of pivotal events is much lower in market 1 than 2 (19.5% = 187/960 cases in market 1 vs. 47.7% = 458/960 in market 2). Because voters in market 1 are less likely to be pivotal they have lower expected costs of voting sincerely. The difference is moderate, though, because the cost in terms of payoff foregone of tipping the outcome against one's material self-interest is twice as high in market 1 than 2, i.e. 80 vs. 40. Therefore, the expected cost is approx. 47.7% * 40 = 19.1 in market 2 which is about 22% larger than the expected cost in market 1 (19.5% * 80 = 15.6).

We find for market 2 that voters who are pivotal are significantly less likely (by 7.1%) than voters who are not pivotal to vote against their material self-interest (i.e. voters are less likely to vote for taxing themselves when it is likely to hurt). This is not the case in market 1. There we find no difference in voting across pivotal vs. non-pivotal voters.[27] Despite variable $Pivot_{p-1}$ being insignificant for market 1 and negative significant for market 2, the estimated parameters of variable $Pivot_{p-1}$ do not differ across markets (p = 0.543). It is therefore not likely that the difference in insincere voting across the markets explains our results.

Regarding the remaining variables, HIGH is a dummy variable taking the value 1 when the treatment is HIGH. Consistent with the earlier reported results, the first coefficient in the regression for Market 1 shows that voters are 5.9% more likely to vote for the discriminatory tax in HIGH than in LOW. This finding is consistent with the upward shift of the share of voters for DT in market 1 in Figure 4 (compare top line in HIGH vs. LOW) and suggests that a motive of fair distribution may have played some role in voting. Conversely, voters in Market 2 are about 25.9% less likely to vote for DT in HIGH than in LOW (see downward shift of the share voting for DT in Figure 4 when comparing LOW and HIGH).

[26] Remember that subjects were informed of the total number of votes for each alternative.

[27] This is not equivalent to say that there is no support of expressive voting. If we split variable $Pivot_{p-1}$ into the two possible pivotal events (3:5* or 4*:4 previous votes for UT: DT, respectively, where a * indicates the tally that includes the vote of the consumer considered) we find a close to significant effect of pivot4*:4 for voters in market 1 (-0.083, p = 0.103) and a significant effect of pivot3:5* for voters in market 2 (-0.102, p = 0.006). Hence, voters from both market 1 and 2 seem to react to whether their vote has been pivotal or not.

This effect is not consistent with a concern for fair distribution (voting for DT in HIGH produces an equal outcome).[28]

Finally, $Vote_{p-1}$ is the decision in the previous phase, $Profit_{p-1}$ is the subject's profit in the previous phase; we include these variables as controls. Phase captures a trend over repeated voting, and Experience indicates whether the market had implemented the discriminatory tax in one of the previous phases. These variables shed light on the dynamics of voting. As can be seen, voting is fairly persistent in both markets, but experience (and therefore feedback on the incidence of DT) matters for voters in market 1. In line with our discussion in section 3.3, these results suggest that in market 1 the support of DT increases with repetition mainly because of experience effects.

4 Conclusion

Discriminatory taxes are sometimes desirable from an efficiency perspective but frequently unpopular with voters. We have hypothesized that the acceptance of a discriminatory commodity tax depends on voters' fairness perceptions. To test this intuition we have proposed an experiment involving voting on taxing market transactions. In one treatment the discriminatory tax (DT) results in an unfair distribution of incomes, in the other it contributes to establish fair market outcomes.

Our results show that in early phases of the experiment, voting is biased toward taxing all voters alike (UT). For example, among voters who have a material interest in voting for DT (because they remain untaxed), more than 50% vote for UT in the first vote. The details of our results suggest that this pattern is driven by voters' concerns for procedural fairness. Many voters seem to dislike the idea of imposing a discriminatory tax on a subgroup of voters. We do, however, find that this tax appears to be more acceptable to (untaxed) voters when it hits consumers who are relatively rich. Moreover, we find that the procedural fairness norm erodes when voters are more experienced and receive salient feedback about the consequences of taxation on the distribution of market incomes. Our study therefore establishes the relevance of multiple fairness norms in voting on commodity taxation.

[28] See Footnote 19 for possible explanations of this effect.

We think our results are rather surprising and our study invites further investigations into the topic. In fact, the choice of our design was motivated by the hypothesis (which seems to capture a popular view among policy-oriented economists) that efficient discriminatory taxes are unpopular because they tend to be distributionally unfair, and that such taxes would meet with more support if they were distributionally fair and imposed the burden on the "rich" (in treatment HIGH). Yet, the results did not support the hypothesis. In retrospect, we think the reason why the results of our study were hard to anticipate is that there are various social concerns in additional to distributional fairness that might potentially matter and cause a bias away from the standard prediction, and some of these biases may operate in opposite ways. In particular, we study a setting in which standard theory predicts an inefficient outcome, and in which some types of social preferences have the potential to improve efficiency (e.g. inequality aversion, expressive voting or a concern for efficiency) while others (in particular a concern for procedural fairness) tends to favor the inefficient uniform tax. Given that our results suggest that referendum choices of inexperienced voters are importantly driven by a concern for procedural fairness, a useful next step would be to see how strong these concerns are in a setting in which standard theory predicts an efficient outcome and procedural fairness biases voting away from that outcome.

If our conclusion that procedural fairness matters in voting on such tax regimes were to prove robust under such alternative designs, our study has potentially important policy implications. Our results seem to suggest that the distribution of the tax burden is of second-order importance in the fairness perceptions of voters. Rather, commodity taxes are politically accepted if they formally tax everyone alike. The reason, we think, is to be found in the salience of the respective dimensions of what is fair. The distribution of market rents resulting from commodity taxation is typically hard to observe in the field (unlike the distribution of incomes from labor say). General value added and sales taxes are close to the ideal that they treat everyone alike. The perceived fairness of this property may explain the rising trend of these taxes in the evolution of tax systems (for a related argument regarding the role of bounded rationality in shaping the tax system, see, McCaffery 2000). On the other hand, general consumption taxes tend to be regressive. Because the income distribution is skewed to the right it is often said that the tax increases inequality at the cost of a majority.[29] We have

[29] This is certainly true in the short run if marginal propensities to consume fall with income. In this case, the "rich" bear a lower average tax burden than the poor. In the long run, the case is less clear since income must be consumed at some point (or passed on to the next generation through bequests).

seen in our study that procedural fairness norms may eventually give place to distributional fairness considerations in such a case.

References

Atkinson, A.B. and Stiglitz, J.E. (1980): *Lectures on Public Economics*. London: Mc-Graw-Hill.

Babcock, L., Lowenstein, G., Issacharoff, S. and Camerer, C.F. (1995): Biased Judgments of Fairness in Bargaining. *American Economic Review* 85(5): 1337-1343.

Blanco, M., Engelmann, D. and Normann, H.-T. (2011): Within-Subject Analysis of Other-Regarding Preferences. *Games and Economic Behavior* 72(2): 321-338.

Bolton, G. and Ockenfels, A. (2000): A Theory of Equity, Reciprocity, and Competition. *American Economic Review* 100(1): 166-193.

Bolton, G.E., Brandts, J., and Ockenfels, A. (2005): Fair Procedures – Evidence from Games Involving Lotteries. *Economic Journal* 115: 1054-1076.

Brennan, G. and Lomasky, L. (1993): *Democracy and Decision*. Cambridge. Cambridge University Press.

Buckley, E. and Croson, R. (2006): Income and Wealth Heterogeneity in the Voluntary Provision of Linear Public Goods. *Journal of Public Economics* 90: 935-955.

Charness, G. and Rabin, M. (2002): Understanding Social Preferences with Simple Tests. *Quarterly Journal of Economics* 117(3): 817-69.

Cherry, T.L., Kallbekken, S. and Kroll, S. (2012): The Acceptability of Efficiency-Enhancing Environmental Taxes, Subsidies and Regulation: An Experimental Investigation, *Environmental Science & Policy* 16: 90-96.

Chlaß, N., Güth, W. and Miettinen, T. (2009): Purely Procedural Preferences - Beyond Procedural Equity and Reciprocity. *Jena Economic Research Papers* No. 2009-069.

Cooper, D.J. and Kagel, J. (2013): Other Regarding Preferences: A Survey of Experimental Results. In J. Kagel and A. Roth (Eds.): *The Handbook of Experimental Economics*, Vol. 2, (forthcoming) Princeton: Princeton University Press.

Engelmann, D. and Strobel, M. (2002): Inequality Aversion, Efficiency, and Maximin Preferences in Simple Distribution Experiments. *American Economic Review* 94(4): 857-869.

Feddersen, T., Gailmard, S. and Sandroni, A. (2009): Moral Bias in Large Elections: Theory and Experimental Evidence. *American Political Science Review* 103: 175-192.

Fehr, E. and Falk, A. (1999). Wage Rigidity in a Competitive Incomplete Contract Market. *Journal of Political Economy* 107: 106-134.

Fehr, E. and Schmidt, K.M. (1999): A Theory of Fairness, Competition, and Cooperation. *Quarterly Journal of Economics* 114(3): 817-868.

Fehr, E., Naef, M. and Schmidt, K.M. (2006): Inequality Aversion, Effciency, and Maximin Preferences in Simple Distribution Experiments: Comment. *American Economic Review* 96(5): 1912-1917.

Fischbacher, U. (2007): z-Tree: Zurich Toolbox for Ready-made Economic experiments, *Experimental Economics* 10(2): 171-178.

Großer, J. and Reuben, E. (2013): Redistribution and Market Efficiency: An Experimental Study. *Journal of Public Economics* 101: 39-52.

Güth, W. and Mackscheidt, K. (1984): Die Erforschung der Steuermoral durch Experimente. University of Köln, mimeo.

Güth, W., Levati, M.V. and Ploner, M (2013): Does Procedural Fairness Crowd Out Other-Regarding Concerns? A Bidding Experiment. *Journal of Institutional and Theoretical Economics*. Forthcoming.

Güth, W., Levati, V.M. and Sausgruber, R. (2005): Tax Morale and (De-) Centralization. *Public Choice* 125(1): 171-188.

Höchtl, W., Sausgruber, R. and Tyran, J.-R. (2012): Inequality Aversion and Voting on Redistribution. *European Economic Review* 56(7): 1406-1421.

Kagel, J. H. (1995): Auctions: A Survey of Experimental Research. In J. H. Kagel and A. E. Roth (Eds.): *The Handbook of Experimental Economics*. Princeton, NJ: Princeton University Press, 501-586.

Kerschbamer, R., Sutter, M. and Dulleck, U. (2013): How Distributional Preferences Shape Incentives on (Experimental) Markets for Credence Goods. University of Innsbruck. Working Paper. Feb. 2013.

Krawczyk, M.W. (2011): A Model of Procedural and Distributive Fairness. *Theory and Decision* 70: 111-128.

List, J.A. (2011): The Market for Charitable Giving. *Journal of Economic Perspectives* 25 (2): 157-180.

Mankiw, N.G. and Weinzierl, M.C. (2010): The Optimal Taxation of Height: A Case Study of Utilitarian Income Redistribution. *American Economic Journal: Economic Policy* 2(1): 155-176.

Markussen, T., Putterman, L. and Tyran, J.-R. (2013): Self-Organization for Collective Action: An Experimental Study of Voting on Sanction Regimes. *Review of Economic Studies*. Forthcoming.

McCaffery, E.J. (2000): Cognitive Theory and Tax. In: C.R. Sunstein (Ed.): *Behavioral Law and Economics.* Cambridge: Cambridge University Press: 398-421.

Ramsey, F.P. (1927): A Contribution to the Theory of Taxation. *Economic Journal* 37: 47-61.

Roth, A.E. and Malouf, W.K. (1979): Game Theoretic Models and the Role of Information in Bargaining. *Psychological Review* 86: 574-594.

Sausgruber, R. and Tyran, J.-R. (2011): Are We Taxing Ourselves? How Deliberation and Experience Shape Voting on Taxes. *Journal of Public Economics* 95(1): 164-176.

Smith, V.L., Williams, A.W., Bratton, W.K. and Vannoni, M.G. (1982): Competitive Market Institutions: Double Auction versus Sealed Bid-Offer Auctions. *American Economic Review* 72(3): 58-77.

Sørensen, P.B. (2010): The Theory of Optimal Taxation: New Developments and Policy Relevance. *Nationaløkonomisk Tidsskrift* 148 (2010): 212-244.

Tyler, T.R. and Lind, E.A. (2000): Procedural Justice. In: J. Sanders and V.L. Hamilton (Eds.). *Handbook of Justice Research in Law*: 63-91. New York: Kluwer.

Tyran, J.-R. (2004): When Money and Morals Conflict. An Experimental Test of Expressive Voting. *Journal of Public Economics* 88(7): 1645-1664.

Tyran, J.-R. and Sausgruber, R. (2006): A Little Fairness May Induce a Lot of Redistribution in Democracy. *European Economic Review* 50(2): 469-485.

Wenzel, M. (2002): The Impact of Outcome Orientation and Justice Concerns on Tax Compliance: The Role of Taxpayers' Identity." *Journal of Applied Psychology* 87(4): 629-645.

Appendix I: Number of votes and number of paired markets accepting the discriminatory tax per treatment and phase

	LOW			HIGH		
Phase	**Market 1** # votes out of 48 for discriminatory tax	**Market 2** # votes out of 48 for discriminatory tax	**Aggregate** # acceptance out of 12	**Market 1** # votes out of 48 for discriminatory tax	**Market 2** # votes out of 48 for discriminatory tax	**Aggregate** # acceptance out of 12
1	18	10	0	25	5	2
2	29	10	0	30	7	0
3	28	11	1	37	7	5
4	32	16	2	39	4	3
5	35	11	1	44	5	3
6	35	12	2	46	4	3
7	36	10	4	45	5	4
8	41	7	2	46	1	0
9	40	7	2	47	3	2
10	39	6	1	46	2	2
Sum and %	*333/480 = 69%*	*100/480 = 21%*	*15/120 = 13%*	*405/480 = 84%*	*43/480 = 9%*	*24/120 = 20%*

Appendix II: Instructions for the auction (Market 1, translated from German)

General Instructions for Participants

You are now participating in an economics experiment which funded by the Austrian Science Fund. The purpose of the experiment is to analyze decision making in markets. You are paid Euro 4 for showing up on time. If you carefully read the instructions and follow the rules you can earn more. The €4 and all additional amounts of money earned during the experiment will be paid to you in cash immediately after the experiment. In the experiment you earn points. Points will be converted to Euros according to the following exchange rate: Points 100 = €0.3. During the experiment we ask you not to speak to other participants. If you have a question, please ask us. We will gladly answer your questions in private. It is very important that you follow this rule. Otherwise the results of the experiment have no value from a scientific perspective.

You are now participating in a market experiment. In the market, you can buy units of a hypothetical commodity. You earn money by trading. How much you earn depends on your and the decisions of others. The experiment consists of two practice periods followed by a number of trading periods. In the practice periods you do not yet earn money; you should take these periods seriously, though, since you will gain valuable experience for the paid trading periods.

Detailed Instructions for Buyers

In this experiment each participant is a buyer. You can buy units from automated sellers. The sellers will sell to you according to the rules described below. In your market there are 4 buyers who can buy units from sellers in each of the trading periods.

<u>How the market works</u>:

As a buyer you state a price at which you would buy a unit. We call this your “bid”. You can buy three units at most. You can submit a bid for each of these three units.

The sellers bear a cost for each unit they sell. The sellers sell their units to buyers whose bids are above the cost for a unit. Bids that are below the sellers’ cost for a unit are rejected.

Example: Assume we collect four bids in a market period. The highest bid is 130, the 2nd highest bid is 100, the 3rd highest bid is 70, and the 4th highest bid is 20.

	1.	2.	3.	4.
Bids	130	100	70*	20

Assume that the sellers’ cost for a unit is 25. In this case, the first 3 bids are higher than the cost for a unit of the sellers. Therefore, the 3 units sell to the buyers who have made these bids. The 4th bid at 20 is rejected. The buyer who submitted this bid does not buy the unit. <u>Please note</u>: Buyers do not know the bids of others, nor do they know the sellers’ cost for a unit.

How your earnings are computed: Every unit you buy obtains a <u>value</u> for you. Your will learn your value in the experiment. If you buy a unit you have to pay a <u>price</u>. Each buyer pays the same price per unit. This market price is equal to the last accepted bid in the order explained above. In our example, 3 units have been bought. The bid of the last accepted unit is 70 (marked with an asterisk *). The uniform market price is therefore 70. Please note that the first two units of the order are also bought at a price of 70.

Your earnings per unit bought are calculated as follows: <u>Earning</u> = <u>value</u> *minus* <u>price</u>. Note that if you buy a unit you will pay less than your bid unless your bid is exactly equal to the market price. In our example, suppose that you submitted the bid of 100 for a unit which has a value of 130. This bid is above the market price of 70. Since you buy this unit at a market price of 70 your earnings will be 130-70=60.

If you do not buy a unit, its value expires. On the other hand, you also do not pay a price. Your earnings are therefore zero.

Subjects' original instructions contained figures showing how the computer screen would look like in the experiment.

How the trading is presented on the computer screen? In each trading period a Decision Screen appears (Figure 1). At the end of each period an Outcome Screen appears (Figure 2). After 5 trading periods a History of Results appears (Figure 3).
In the top area of the Decision Screen on the left you see the number of the current trading period (here: 2) and the total amount of trading periods (here: 5). Each trading period ends after a time limit. The remaining time within a period is shown in the top area on the right (here: 19 Seconds).
In the first row you see your value for buying your first unit. In this example the buyer has a value of 130. The input field below your value serves to enter your bid. To enter a bid, click on the field labeled 'Your Bid' and type in a number. To submit that bid, click on the 'Submit' button. In the second and third row everything is repeated for your 2nd and 3rd unit.
Rules for bidding: An important rule for trading is to "trade at no loss". Therefore, you may not submit a bid above your value for a unit. In the example shown in Figure 1, this buyer's bid must not be above 130. If you violate this rule, a message box appears. The message disappears if you press the 'OK' button.
The Outcome Screen (Figure 2) appears at the end of a trading period. Here you see your value, your bid, and the price for each of your units. In the center area you see the total number of units traded in the market (Market Quantity) and the uniform price per unit. In this example, there were 4 units traded at a price of 70. The next line shows the number of units bought by you (here 1). This means that this buyer has bought one of three units. The last row of the table shows your earnings for the current period. In this example, the earnings from buying the unit are 60 (=value of 130 minus price of 70).
A History of Results (Figure 3) appears after 5 trading periods. Here, the results from each period are summarized. Next to the market quantity and the price, you see again the number of units bought by you and your earnings per period (the example of the screen contains entries only for one period; in the experiment the screen will contain the entries for all five periods). The row labeled "Your total earnings" show your earnings summed over the last 5 periods.

In the lower area of the screen you find information regarding your market and "another market", as well as the two markers together. As in your market, there are 4 buyers in the other market who buy units from their sellers in their market. The identities of the buyers in the other market remain the same during the entire course of the experiment. The column labeled "**tax"** informs you about the tax the sellers must pay in your and the other market (in the example of the screen, the tax is 0 in both markets). The column labeled "**tax per buyer"** informs you of the total tax payments divided by the number of buyers in each market. Finally, the column labeled "**earnings per buyer"** informs you of the average income of a buyer in your market, the average income of a buyer in the other market, and the income per buyer averaged over both markets and over the last 5 periods.

Appendix III: Voting proposal (Market 1, translated from German)

You now vote between two alternatives to raise a given amount of tax. In your market there are 4 buyers, including yourself. There are also 4 buyers in the other market. All 8 voters cast a vote. The vote is anonymous; this means that no other participant is informed about how you vote.

The two alternatives are:

Alternative 1:	**Alternative 2:**
Sellers in **the other market** pay a tax of **80** points for every unit they sell**.**	Sellers in **your and the other market** pay a tax of **40** points for every unit they sell.

Decision rule

Alternative 1 is implemented for the next 5 periods if a majority of buyers (this is 5 or more) votes for it. Otherwise, alternative 2 gets implemented for the next 5 periods.

Please note

Under alternative 1, because of the tax, the sellers' cost for a unit in the other market increases by 80 points. Therefore, in the other market bids will be accepted only if they exceed the sellers' cost augmented by 80 points.

Under alternative 2, because of the tax, the sellers' cost for a unit in your and the other market increases by 40 points. Therefore, in your and the other market bids will be accepted only if they exceed the sellers' cost augmented by 40 points.