

Trading and Voting

David K. Musto and Bilge Yilmaz

University of Pennsylvania

Complete financial markets transform the political choice between candidates with different redistribution policies. If redistribution policies do not affect aggregate wealth, then financial trade implies that wealth considerations have no effect on voting and so do not affect who wins. However, an election in which one candidate would redistribute results in redistribution, and redistribution is the same whether or not he wins. Furthermore, he proposes, and if elected carries out, more redistribution than he prefers. If redistribution policies do affect aggregate wealth, then everybody expects more wealth if the candidate with the higher aggregate-wealth policy wins.

Elections assign the right to design tax policies, so to consumers they represent uncertainty over future wealth. An election pitting a candidate who proposes to redistribute wealth against another who does not indicates one future state of the world in which the wealthy lose wealth to the poor and another future state in which they do not. One implication of this uncertainty pertains to the election's outcome: other things equal, a voter chooses the candidate delivering him more wealth, so the interaction of the tax policies with the wealth distribution decides who wins. The literature has explored this implication extensively (see, e.g., Myerson 1993; Persson and Tabellini 1994; Lizzeri and Persico 2000), including multiperiod models focusing on accumulating debt (see, e.g., Aghion and Bolton 1990; Lizzeri 1999).

We are concerned here with a different implication, which can have

We thank Domenico Cuoco, Simon Gervais, Bob Inman, Antonio Merlo, Nicola Persico, Tom Rietz, Nick Souleles, participants in the University of Pennsylvania's Political Economy seminar, two anonymous referees, and Fernando Alvarez, the editor, for helpful advice and comments.

[*Journal of Political Economy*, 2003, vol. 111, no. 5]
© 2003 by The University of Chicago. All rights reserved. 0022-3808/2003/11105-0008\$10.00

strong implications for the first: consumers would respond to the wealth uncertainty, as they do to other financial risks, in their demand for financial instruments. If consumers can share the wealth risk by trading election-contingent securities (e.g., paying one if the redistributionist wins, zero if he loses) in a frictionless market, the outcome of the political process is quite different on all the important dimensions. The probability that the redistributionist wins is different, the amount of redistribution is different, and the timing of the redistribution is different too.

The intuition for this result is that the potential for redistribution creates offsetting risks. The dollars that the wealthy lose to redistribution match the dollars the poor gain, so the enthusiasm of the wealthy for buying insurance against the redistributionist's winning resembles the enthusiasm of the poor for selling it. If voters can share this risk with the state-contingent security, the equilibrium is full insurance; all consumers equalize wealth across the two states. This delivers both the result that wealth redistribution occurs before the election, rather than during or after, and also the result that wealth considerations do not determine who wins the election. With wealth the same whether or not the redistributionist wins, voters refer to their nonwealth preferences when making their choices. In other words, an externality of frictionless financial risk sharing is that ideological, and not pocketbook, concerns decide who governs.

The rest of the paper is organized in four sections. Section I covers the relevant background. Section II presents the model, Section III solves and interprets the model, and Section IV summarizes and presents a conclusion.

I. Background

Consumers can vote themselves other people's wealth. A candidate can communicate that he would redistribute if elected, and a majority of votes makes it happen. This would seem to have serious implications for the distribution of wealth and the incentive to accumulate wealth in the first place. When wealth is concentrated in a few voters, a candidate who favors redistribution would intuitively have an easy win over a candidate who does not, and wealth would accordingly even out.

Consider a two-candidate race. Two candidates communicate how they would govern, including how they would redistribute; then the vote occurs and then the governing, including the redistribution. When the candidates' governing policies are endogenous, that is, they communicate whatever policies they want to communicate, not just the policies that match their principles, the race becomes strategically interesting. If candidates just care about winning and not redistribution per se, they

will espouse redistribution anyway to buy votes (see, e.g., Myerson 1993). Similarly, political parties with ideological motives are nonetheless motivated to buy votes with transfers (Dixit and Londregan 1996). Throughout the extensive literature on this subject (see Drazen [2000] for a review), candidates' redistribution policies are viewed as key to electoral outcomes, being the major—or only—determinant of voters' preference orderings, and the candidates' policy choices are analyzed from this perspective.

Now consider the race from the consumers' point of view. The candidates partition the future into two possible states, one for each outcome, and each consumer expects more wealth in one state than in the other. So the election creates uncertainty over future wealth, and there is an extensive literature on that subject, too. The standard analysis is that consumers have strictly concave utility for wealth, so they are all risk-averse at any wealth level; thus they would prefer to hedge uncertainty over future wealth by trading financial securities. The uncertainty caused by a potential redistribution is well suited to such trade, for two reasons. First, if wealth is simply redistributed, rather than created or destroyed, then net redistribution is zero in each state. This suggests that the demand for securities that hedge against a candidate's winning matches the supply. Second, the contingency that consumers want to hedge is easily verifiable and therefore contractible.

In a complete and perfect capital market, consumers enjoy frictionless and unlimited access to a security paying one if a candidate wins and zero if he loses. In actuality, consumers enjoy at least some access. Securities contingent on the major U.S. elections trade on the Iowa Electronic Market, where there are no commissions but low position limits.¹ Securities sensitive to, but not defined by, the major U.S. elections trade without explicit position limits on the major exchanges. Examples include municipal bonds.² Since financial market access is intuitively weak for one side of this market, poor people,³ the result we find for complete and perfect markets can be viewed as cautionary, a prediction of the economy with easier access. And since some consumers might wish to manipulate such a market should it exist, our results best represent an economy in which manipulation is uneconomical or is outlawed.⁴

In summary, the existing literature on elections has not allowed for

¹ Traders can spend up to \$500 (see, e.g., Feder 2002).

² See "Presidential Race Induces Creation of Index Strategies" (2000) for more examples.

³ One could, however, view municipal bonds as sales to rich people made on their behalf.

⁴ One source of manipulation is insider trading. Political insiders could not only buy or sell on their information but also manipulate the market with strategic (potentially false) announcements. The goal could be either to make money or to influence voters through the market price. Another source is wealthy consumers; if campaign contributions have sufficient impact on candidates' chances (and thereby the market price), the rich could find them more economical than insurance.

consumers' adaptive response to the uncertainty over wealth that potential redistribution represents, and the response indicated by the literature on financial securities is to hedge by trading. To see that this response is potentially crucial, consider a situation in which consumers trade away all the risk, which is clearly possible since aggregate risk is zero. How would they vote? In anticipation of that, how would they trade in the first place? And what does this imply for the redistribution of wealth? The next two sections answer these questions with a simple but general model.

II. Model

A. The Setup

There are two dates, time 0 and time 1. There are two candidates, L and R , who announce at time 0 what they would do if elected at time 1. There are N consumers who can trade at time 0 and can vote and consume at time 1. What the consumers can trade is contracts that pay one if L wins and zero otherwise. They can buy or sell any amount of this contract. A candidate can propose a redistributive wealth tax that occurs immediately upon election. The tax applies to postcontract wealth: a consumer first pays or gets what his contract position dictates, and then his resulting wealth is redistributed.

Candidate R communicates that he would govern with ideology R and impose no tax, and L communicates that he would govern with ideology L and impose a redistributive tax of τ (the method of communication is not modeled here; we take as given that voters learn that the candidates would enact these policies). The time t wealth of consumer c is w_c^t , and the utility of consumer c over time 1 (postelection, postcontract, and postdistribution) wealth and ideology I (i.e., the ideology of the winner) is

$$u_c(w_c^1) + v_c(I),$$

where $u_c' > 0$, $u_c'' < 0$, and $v_c(R) \neq v_c(L)$ for all c .

The redistributive wealth tax collects τ of each consumer's wealth and distributes $1/N$ of the receipts to each consumer. To allow a potential effect of the tax on real activity, we let δ_c represent a potentially nonzero, pretax value change for consumer c . So if L wins, the wealth of consumer c changes by δ_c and the total wealth in the economy changes by the sum of δ_c across c . We denote the total wealth in the economy if L wins and if R wins as W^L and W^R , respectively, and we let \bar{w} stand for per capita wealth if L wins, that is, $W^L/N = \bar{w}$. So the net redistribution to a consumer with postcontract wealth w (which includes δ) is $\tau(\bar{w} - w)$.

All consumers know the total wealth in the economy for both out-




FIG. 1

comes. Therefore, each consumer can calculate his wealth under both policies in period 1. Let $\mathbf{w}^L = (w_1^L, \dots, w_N^L)$ and $\mathbf{w}^R = (w_1^R, \dots, w_N^R)$ stand for period 1 wealth distribution. Similarly, $\mathbf{v} = (v_1, \dots, v_N)$ denotes the ideological preferences. At period 0, there is uncertainty about the collection of all consumers' ideology, \mathbf{v} , wealth distribution (and therefore period 1 wealth distribution, $(\mathbf{w}^L, \mathbf{w}^R)$), and identity of consumers. This uncertainty is represented by a finite set S of states. Consumers share a common prior about the true state of the world. Let ρ stand for this probability measure on S . Each consumer has private information in the sense that he knows at least his own ideological preference and wealth. This private information is described by $H_c: S \rightarrow 2^S$, a partitional information function. (Given the true state, $s \in S$, consumer c knows that the true state could be any element of $H_c(s)$.) In addition, there is residual uncertainty over events that will affect the election outcome such as turnout, revelations about candidates' private lives, wars, recounts, and so on. Let $f(d, p, s)$ be the joint probability distribution describing this residual uncertainty after trading at price p , where $d = 1$ if candidate L wins and zero otherwise. Therefore, we can rule out any trivial setting: At period 0, no consumer can be sure about the outcome of the election independent of the amount of information he has, that is, $0 < f(d = 1|p, s) < 1$ for all p and s . For further use, let π stand for the value of this conditional probability when $\mathbf{w}^L = \mathbf{w}^R$. Finally, we assume that no consumer is negligible in determining the election outcome: for all $c \in \{1, \dots, N\}$ and for all $s \in S$ and p , we have

$$f(d = 1|H_{-c}(s), H_c(s), p) \neq f(d = 1|H_{-c}(s), p),$$

where

$$H_{-c}(s) \subseteq \{H_1(s), \dots, H_{c-1}(s), H_{c+1}(s), \dots, H_N(s)\}.$$

We can summarize the model with the following chronology, illustrated in figure 1.

1. Before date 0, consumers learn that L will impose ideology L and tax τ if elected, and R will impose ideology R .

2. At date 0, consumers trade election-contingent securities; consumer c buys x_c contracts. After trading, there is residual uncertainty over the election's outcome.
3. At date 1, (a) consumers vote in the election, and the residual uncertainty is resolved. (b) If L won, consumer c receives x_c from his contracts and δ_c from the effect of the tax on real activity. (c) The winner's ideology is implemented, and if L won, then consumer c pays a tax of τ times his current wealth and receives τ times the average wealth \bar{w} .

B. Discussion of Modeling Choices

It would be simpler to solve a model with atomistic, and therefore price-taking, consumers.⁵ But while that approximation is acceptable in some settings, it is inappropriate here because it ruins the incentive to vote. That is, if a consumer's trading does not affect prices, then his voting should not affect the election. So we allow traders to move the market, in that the market price aggregates traders' information, though we do not explicitly model the trading mechanism. We solve for the equilibrium by first positing the existence of a price p^* at which there is no further trade and then solving for p^* . There could in practice be an incentive for candidates or their supporters to manipulate the market price since it can be taken as a de facto poll result (see, e.g., "Iowa Business School's Presidential Futures Market Still Too Close to Call," CNN Transcript 00110602V62, November 6, 2000). We are abstracting from that incentive in our analysis.

One potential concern with trading before voting is that the trading could reveal exactly who wins, so that the price goes to zero or one. To keep the focus on nontrivial trading outcomes, we assume that there is sufficient uncertainty over events intervening between trading and voting that the election's outcome cannot be predicted exactly with information known at trading time. So we do not endogenize the timing of the trading relative to the arrival of election-relevant news, but it is intuitive that consumers would want to take their position in election-contingent securities *before* a given source of uncertainty taps out, not after.

Our analysis equates the state $\{L \text{ wins}\}$ to the state $\{L\text{'s policies are enacted}\}$. These states may not in practice be truly equivalent as a result of, for example, the competing agendas of other branches of government. In such cases the security is more accurately viewed as policy-

⁵ In an earlier version, we show that our main results hold under the assumption of a continuum of consumers/voters.

contingent rather than election-contingent. That is, it pays one if and only if L 's policies are enacted.

We model utility over wealth as separable from utility over ideology. This is the same approach taken by Dixit and Londregan (1996). It is not hard to think of campaign positions that relate to both wealth and ideology, such as federal funding of abortions or even redistribution itself. We are implicitly analyzing these positions as packages, combining wealth effects that affect consumers through u and therefore interact with other wealth effects such as security payoffs and ideological effects that are felt through v .

Finally, our representation of the effect of taxes on real activity, through δ_c , is reduced form and is not intended to be a rigorous analysis of that problem, which is analyzed extensively elsewhere (e.g., Meltzer and Richard 1981). The flat wealth tax is an approximation of federal taxes whose incidence generally increases with personal wealth. It is functionally equivalent to the linear income tax in Meltzer and Richard (1981), where voters start with no wealth.

III. Analysis

Let p^* be the equilibrium price of a contract. We first take it as given and then solve for it. If consumer c buys x_c contracts, then he pays $x_c p^*$ at time 0 and then gets x_c pretax in the state $\{L \text{ wins}\}$ and nothing in the state $\{R \text{ wins}\}$. The wealth he consumes is therefore

$$w_c^0 - x_c p^* + x_c + \delta_c + \tau[\bar{w} - (w_c^0 - x_c p^* + x_c + \delta_c)] = \\ \tau \bar{w} + (1 - \tau)[(1 - p^*)x_c + w_c^0 + \delta_c]$$

in $\{L \text{ wins}\}$ and $w_c^0 - x_c p^*$ in $\{R \text{ wins}\}$. To calculate c 's optimal contract position, we need the probability he puts to the outcome $\{L \text{ wins}\}$, which for the moment we call Π_c . With this notation, c 's problem is to choose the x_c that maximizes

$$\Pi_c [u_c(\tau \bar{w} + (1 - \tau)[(1 - p^*)x_c + w_c^0 + \delta_c]) + v_c(L)] \\ + (1 - \Pi_c) [u_c(w_c^0 - x_c p^*) + v_c(R)].$$

With w_c^L and w_c^R representing c 's terminal wealth in $\{L \text{ wins}\}$ and $\{R \text{ wins}\}$, respectively, the first-order condition can be written as

$$\Pi_c u'_c(w_c^L)(1 - \tau)(1 - p^*) = (1 - \Pi_c) u'_c(w_c^R) p^*. \quad (1)$$

If $\Pi_c \neq 0$ and $p^* \neq 1$, this can be rewritten as

$$\frac{u'_c(w_c^L)}{u'_c(w_c^R)} = \frac{(1 - \Pi_c) p^*}{\Pi_c (1 - \tau)(1 - p^*)}. \quad (2)$$

A. *Pure Redistribution*

We first focus on the pure redistribution case, that is, $\sum_c \delta_c = 0$.

PROPOSITION 1. In the unique equilibrium, all consumers equalize wealth across the possible electoral outcomes, and the outcome is determined solely by ideology rather than by the distribution of wealth.

Proof. The proof of uniqueness is presented in the Appendix. Here, we construct an informationally efficient rational expectations equilibrium (REE). In an informationally efficient REE, the equilibrium price is a sufficient statistic for all private information. Therefore, Π_c is the same for all c . Assume for the moment that $0 < \Pi_c < 1$ and $0 < p^* < 1$. The right-hand side of (2) is the same for all c , so all consumers equalize $u'_c(w_c^L)/u'_c(w_c^R)$ to the same number. This number must be one because if it were greater than one, then everybody would have more wealth in $\{R \text{ wins}\}$ than in $\{L \text{ wins}\}$, and this is not possible because aggregate wealth is the same in both states. Analogously, the number cannot be less than one. So it is one, implying $w_c^L = w_c^R$ for all c .

With wealth equalized across outcomes, a consumer prefers the outcome $\{L \text{ wins}\}$ to $\{R \text{ wins}\}$ if and only if $v_c(L) > v_c(R)$. So the probability of $\{L \text{ wins}\}$ is π , which by assumption is strictly between zero and one. This also implies that $0 < p^* < 1$ because if $p^* = 1$, everyone would be better off selling more contracts, and if $p^* = 0$, everyone would be better off buying more. Market clearance follows immediately from lemma 2. Q.E.D.

The equilibrium contract price is easily inferred.

LEMMA 1. The price per contract p^* is $(\pi - \pi\tau)/(1 - \pi\tau)$.

Proof. Set the right-hand side of (2) equal to one and solve for p^* . Q.E.D.

Note that p^* is always less than π for $\tau > 0$. We can also solve for the number of contracts purchased.

LEMMA 2. Consumer c buys

$$x_c = \left(\frac{1 - \pi\tau}{1 - \tau}\right) [\tau(w_c^0 - \bar{w}) - (1 - \tau)\delta_c]$$

contracts.

Proof. Set $w^L = w^R$, plug in the equilibrium value of p^* , and solve for x_c . Q.E.D.

This can also be written as

$$x_c = \left(\frac{\pi}{p^*}\right) [\tau(w_c^0 - \bar{w}) - (1 - \tau)\delta_c],$$

which leads to our next major result.

PROPOSITION 2. When consumers trade before voting, the wealth redistribution occurs before the election, is unrelated to the outcome,

and is the product of the probability that the redistributionist wins when votes depend only on ideology and the redistribution that would have occurred without trading if the redistributionist won.

Proof. In both states, wealth equals $w_c^0 - x_c p^*$, which is

$$w_c^0 - \left(\frac{\pi}{p^*}\right) [\tau(w_c^0 - \bar{w}) - (1 - \tau)\delta_c] p^*,$$

or $w_c^0 + \pi[\tau(\bar{w} - w_c^0) + (1 - \tau)\delta_c]$. So the wealth redistribution is $\pi[\tau(\bar{w} - w_c^0) + (1 - \tau)\delta_c]$ regardless of who wins, and this is π times the redistribution that would have occurred in $\{L \text{ wins}\}$ without trading. Q.E.D.

This is a big departure from the standard economic analysis of elections. When consumers can trade before voting, the wealth effect of a candidate's redistribution plan no longer affects his chances of winning, but it does affect the resulting redistribution whether or not he wins. The magnitude of the effect depends on his chances of winning, but his chances of winning depend solely on his ideological appeal.⁶ The wealth effect on the median wealth voter is not relevant. Consumers can trade wealth but not ideology across states, and this is what happens.

Because it affects the state probabilities, the trade in election-contingent securities is not simply Pareto-improving risk sharing. Poor people could view it as a coordination problem. The redistributionist might have been an almost sure thing if wealth distribution influenced voting, but not with wealth equalized, so the net expected redistribution to consumer c goes from close to $\tau(\bar{w} - w_c^0) + (1 - \tau)\delta_c$ to $\pi[\tau(\bar{w} - w_c^0) + (1 - \tau)\delta_c]$. This is an adverse development if $w_c^0 < \bar{w}$ (and δ_c is small), but the consumer is better off trading than not even though he would be best off if nobody traded. The poor would like to avoid this effect of trading by coordinating if they could find an incentive-compatible mechanism that implements their preferred no-trade outcome distribution. A constitutional amendment or referendum banning election-contingent trade may serve this purpose (though this vote would itself be susceptible to hedging).

1. Strategic Policy Choices

So far we have not specified a set of preferences for the candidates, but rather taken their policy choices as given and analyzed the consumers' reaction. However, if candidates care about the enacted policies (both

⁶ Our results do not depend on a perfect hedge. If we instead assume that the contract pays off in the wrong state (i.e., R wins) with probability ϵ , we can show numerically with specific utility and distributional assumptions that the resulting equilibrium converges to the one solved here as $\epsilon \rightarrow 0$ (results available on request).

ideological and redistributive), they will strategize over their policy choices. Maximizing expected utility, each candidate must consider a policy's electability as well as its desirability. Let $U_L(\tau, \mathcal{I}, d)$ denote the preference of candidate L .⁷ We assume that there exists a redistribution $\hat{\tau}_L$ such that $U_L(\hat{\tau}_L, \mathcal{I}, d) > U_L(\tau', \mathcal{I}, d)$ for all $\tau' \neq \hat{\tau}_L$ and for all \mathcal{I} and d . The preference of candidate R is analogous, and to simplify, we assume that candidate R prefers no redistribution, that is, $\hat{\tau}_R = 0$. Prior to trading, each candidate announces a tax rate and an ideological policy.⁸ We further assume that candidates can commit to policies and negative redistribution is not possible, that is, $\tau \geq 0$. At the time candidates announce policies, they know that the election will depend solely on \mathbf{v} because of backward induction, and their utilities can increase with their odds of winning because of either the effect on net distribution (as argued in the previous proposition), the direct utility from winning, or the chances of implementing their ideologies. Therefore, each may prefer to announce more moderate ideological policies in order to increase the probability of winning and maximize expected utility. On the other hand, the choice of tax rate will have no effect on the outcome of the election. However, we have shown that the tax rate will affect the wealth distribution independently of who wins, so to cause the redistribution that tax rate $\hat{\tau}_L$ causes in the absence of trading, candidate L commits to a tax rate *higher* than $\hat{\tau}_L$.⁹

PROPOSITION 3. When candidates choose their policies strategically, candidate L chooses a tax rate that causes more gross redistribution than the amount of net redistribution that L prefers or expects.

The proof follows immediately from the previous proposition and the fact that $\pi < 1$.

B. Economic Efficiency

The propositions to this point all assume that $\sum_c \delta_c = 0$, that is, no aggregate real effect of the tax. We can allow for an aggregate real effect by relaxing this assumption. In this case we can no longer peg the right-hand side of (2) at a given number, but we *can* establish whether it is greater or less than one, which allows us to relate the aggregate effect to wealth preferences.

PROPOSITION 4. When consumers trade before voting and aggregate

⁷ Note that candidate L 's utility is not a function of δ_c ; thus he is indifferent over individual wealth effects and cares only about the aggregate redistribution. This simplifies our next proposition.

⁸ Implicitly, we assume that there exists a set of possible policies that contains R and L .

⁹ With no hedging possibilities, Roemer (1998) shows that a leftist candidate proposes a lower redistribution (possibly zero) when two candidates compete in a two-dimensional policy space.

wealth depends on who wins, all consumers expect after trading to have more wealth if the higher-wealth candidate wins.

Proof. Without loss of generality, let $\{R \text{ wins}\}$ be the higher-wealth state. In an informationally efficient REE, the right-hand side is the same value for all consumers. This value must be greater than one because if it were less than or equal to one, then all consumers would have more wealth (or the same, respectively) in $\{L \text{ wins}\}$ than they do in $\{R \text{ wins}\}$, both of which are not possible. With the right-hand side greater than one, all consumers have more wealth in $\{R \text{ wins}\}$. Q.E.D.

So the wealth effect is strictly in the direction of the higher-wealth candidate. The resulting effect on the probability that this candidate wins depends on the relative strength of ideological preferences. In the intuitive case in which aggregate wealth decreases as τ increases, L decreases his probability of winning (with ideologies held constant) as he increases τ . In this case, L chooses a higher tax rate in economies in which leftist ideology is widely popular: higher taxes destroy more wealth, so all consumers have less wealth; but L can afford this loss of voters' wealth as long as they have sufficiently high $v_c(L)$ relative to $v_c(R)$. On the other hand, if there is no ideological difference between the candidates, so that $v_c(R) = v_c(L)$ for all c , we get the stronger result that consumers trade to the point at which they all want the higher-wealth candidate to win.

COROLLARY 1. When consumers trade before voting and aggregate wealth depends on who wins but utility from ideology does not, then after trading, all consumers prefer that the higher-wealth candidate win.

The frictionless trading opportunity biases the election toward the outcome with greater real activity.

IV. Summary and Conclusion

An election creates wealth risk, and a securities market reallocates wealth risk. The wealth risk created by an election with redistribution at stake is well suited for trade in that demand naturally equals supply. If one candidate would redistribute (but not create or destroy) wealth with a linear tax and this trade is frictionlessly available, the result is a transformed election, with wealth considerations separated completely from voting decisions and redistribution separated completely from the election's outcome. These results constitute a baseline case for arguments that redistribution buys votes or that the amount of redistribution depends on the election's outcome. For those arguments to go through, there must be some departure from our assumptions, such as transactions costs, incomplete markets, or effects of taxation on aggregate wealth. When we introduce an effect on aggregate wealth, we find that all consumers, after trade, expect more wealth in the higher-wealth state.

Today's markets depart from the idealized trade in our analysis with their positive transactions costs, imperfect hedges, and position limits. So one perspective on our results is that they warn of consequences from eliminating these frictions. Trade in Arrow-Debreu state-contingent securities might seem obviously Pareto-improving, but when state probabilities are endogenous, this is no longer clear. Elections are just one, well-defined, example of this endogeneity; the point applies more generally.

Another perspective on the results is that consumers' financial exposures to an election have qualitatively different implications for the outcome and net effect of the election than their other exposures do. Financial exposure can be traded across states, and risk aversion encourages this trade. So elections determine wealth redistribution differently from the way they determine other policies at stake, raising the question as to whether they are equally efficient at resolving distributional and ideological disputes.

Appendix

In this Appendix, we shall first define the equilibrium concept and recall the definition of common knowledge. Second, we shall prove the uniqueness of equilibrium.

DEFINITION A1. Price $p^*(H_1(s), \dots, H_N(s))$ is an equilibrium if the following conditions hold:

1. At $t = 1$, each consumer votes to maximize his expected utility given his postelection wealth, w_c^1 , and ideology, $v_c(\cdot)$.
2. At $t = 0$, each consumer chooses his demand $x_c(p, H_c(s))$ to maximize expected utility assuming that the probability Π_c is given by $f(d = 1 | p = p^*, H_c(s))$, where $f(d = 1, p^*, H_1(s), \dots, H_N(s))$ is the true joint probability distribution given $p^*(H_1(s), \dots, H_N(s))$ and optimal voting at $t = 1$.
3. Given consumers' demands, $p^*(H_1(s), \dots, H_N(s))$ is market clearing for all realizations of consumers' private information.

DEFINITION A2. An event $F \subseteq S$ is *self-evident* between consumers c and c' if for all $s \in F$ we have $H_c(s)$ and $H_{c'}(s) \subseteq F$. An event $E \subseteq S$ is *common knowledge* between consumers c and c' in state s if there is a self-evident event F for which $s \in F \subseteq E$.

LEMMA A1. There exists a unique equilibrium.

Proof. We prove our claim in two steps. First, we show that the full information economy has a unique equilibrium. This result implies that there can be at most one informationally efficient REE. Second, we show that there are no partially revealing or nonrevealing REE.

In a full information economy, Π_c is the same for all consumers. Furthermore, Π_c is strictly between zero and one by the residual uncertainty assumption. Therefore, equation (2) must hold for every equilibrium of the full information economy. However, this implies the equilibrium characterized by proposition 1. This concludes the first step.

Now we proceed with the second step that there are no partially revealing or nonrevealing REE. Suppose not. Then there is an equilibrium price p' that is

not informationally efficient. Equation (2) must hold for this equilibrium as well given the residual uncertainty assumption. If each consumer has equalized wealth across states, $\mathbf{w}^L = \mathbf{w}^R$, then $(1 - \Pi_c)p'/[\Pi_c(1 - \tau)(1 - p')] = 1$ for all c . This implies that $\Pi_c = p'/[(1 - \tau)(1 - p') + p']$ for all c . However, Π_c cannot be the same for all c given that at least one consumer has private information. Therefore, the only possibility left is $(1 - \Pi_c)p'/[\Pi_c(1 - \tau)(1 - p')] \neq 1$ for at least one consumer. Therefore, there exists at least one consumer, say c , who has not equalized wealth across states. Without loss of generality, assume that this consumer has more wealth if L wins. From the market clearance condition, there must exist another consumer, say c' , who has more wealth if R wins. Therefore,

$$\frac{u'_c(w_c^L)}{u'_c(w_c^R)} < \frac{u'_{c'}(w_{c'}^L)}{u'_{c'}(w_{c'}^R)}.$$

Consequently, from equation (2), we must have $\Pi_c > \Pi_{c'}$, that is, $f(d = 1|H_c(s)) > f(d = 1|H_{c'}(s))$. Market clearance also implies that this inequality is common knowledge among these two consumers. (Equivalently, we say that event

$$E = \{s \in S | f(d = 1|H_c(s)) > f(d = 1|H_{c'}(s))\}$$

is common knowledge.) In the rest of the proof we shall show that common knowledge of such a disagreement cannot occur in equilibrium. Given that the event $f(d = 1|H_c(s)) > f(d = 1|H_{c'}(s))$ is common knowledge, there must be an event $F \ni s$ that is a subset of E and is a union of members of the information partitions of both consumers, that is, $\bigcup_{s \in F} [H_c(s) \cup H_{c'}(s)] = F \subseteq E$. Given that $f(d = 1|H_c(s)) > f(d = 1|H_{c'}(s))$ is common knowledge for all $s \in F$, this inequality $f(d = 1|H_c(s)) > f(d = 1|H_{c'}(s))$ must hold for all $s \in F$. Therefore, we have

$$\sum_{s \in F} \rho(s) f(d = 1|H_c(s)) > \sum_{s \in F} \rho(s) f(d = 1|H_{c'}(s)).$$

But since F is a union of members of each consumer's information partition, both sides of this inequality are equal to $\rho(F) f(d = 1|F)$. However, this contradicts the inequality above and concludes the second part of the proof. Therefore, neither partially revealing nor nonrevealing equilibria can exist. Q.E.D.

References

- Aghion, Philippe, and Bolton, Patrick. "Government Domestic Debt and the Risk of Default: A Political-Economic Model of the Strategic Role of Debt." In *Public Debt Management: Theory and History*, edited by Rudiger Dornbusch and Mario Draghi. Cambridge: Cambridge Univ. Press, 1990.
- Dixit, Avinash, and Londregan, John. "The Determinants of Success of Special Interests in Redistributive Politics." *J. Politics* 58 (November 1996): 1132–55.
- Drazen, Allan. *Political Economy in Macroeconomics*. Princeton, N.J.: Princeton Univ. Press, 2000.
- Feder, Barnaby J. "To Learn What People Want, Trade 'Idea Stocks.'" *New York Times* (February 10, 2002), p. 4.
- Lizzeri, Alessandro. "Budget Deficits and Redistributive Politics." *Rev. Econ. Studies* 66 (October 1999): 909–28.
- Lizzeri, Alessandro, and Persico, Nicola. "The Drawbacks of Electoral Competition." Manuscript. Philadelphia: Univ. Pennsylvania, 2000.
- Meltzer, Allan H., and Richard, Scott F. "A Rational Theory of the Size of Government." *J.P.E.* 89 (October 1981): 914–27.

- Myerson, Roger B. "Incentives to Cultivate Favored Minorities under Alternative Electoral Systems." *American Polit. Sci. Rev.* 87 (December 1993): 856–69.
- Persson, Torsten, and Tabellini, Guido. "Is Inequality Harmful for Growth?" *A.E.R.* 84 (June 1994): 600–621.
- "Presidential Race Induces Creation of Index Strategies." *Wall Street J.* (September 25, 2000).
- Roemer, John E. "Why the Poor Do Not Expropriate the Rich: An Old Argument in New Garb." *J. Public Econ.* 70 (December 1998): 399–424.