

Does Vote Trading Improve Welfare ?*

Alessandra Casella[†] Antonin Macé[‡]

August 24, 2020

Abstract

Voters have strong incentives to increase their influence by trading votes, a practice indeed believed to be common. But is vote trading welfare-improving or welfare-decreasing? We review the theoretical literature and, when available, its related experimental tests. We begin with the analysis of logrolling – the exchange of votes for votes, considering both explicit vote exchanges and implicit vote trades engineered by bundling issues in a single bill. We then focus on vote markets, where votes can be traded against a numeraire. We cover competitive markets, strategic market games, decentralized bargaining, and more centralized mechanisms, such as quadratic voting, where votes can be bought at a quadratic cost. We conclude with procedures allowing voters to shift votes across decisions – to trade votes with oneself only – such as storable votes or a modified form of quadratic voting. We find that vote trading and vote markets are typically inefficient; more encouraging results are obtained by allowing voters to allocate votes across decisions.

Keywords: logrolling, vote trading, storable votes, quadratic voting, bundling, vote markets

*When citing this paper, please use the following: Casella A, Macé A. 2020. Does Vote Trading Improve Welfare ? *Annu. Rev. Econ.* 3: Submitted. DOI: 10.1146/annurev-economics-081720-114422.

[†]Department of Economics and Department of Political Science, Columbia University, NYC, USA; email: ac186@columbia.edu.

[‡]CNRS, Paris School of Economics and École Normale Supérieure, Paris, France; email: antonin.mace@ens.fr.

1 Introduction

Voting is a fundamental procedure for discovering collective preferences, codified in democratic political institutions, but in fact used, more or less formally, in all associations of people with an interest in consensual group action. The observation that voting lends itself to deals and trade is as common as the act of voting itself. Exchanging favors by supporting each other's favorite proposals or extracting immediate compensation in return for voting as a fellow group member desires are behaviors that do not surprise us. They may be described as unsavory but they are essential to the practice of politics:

“Logrolling is [...] the most characteristic legislative process... It is compromise, [...] in the practical form with which every legislator who gets results through government is acquainted. It is trading. It is the adjustment of interests. ..There never was a time in the history of the American Congress when legislation was conducted in any other way.”

Bentley (1908), *The Process of Government*, (pp. 370–371)

Given the prominence of vote trading in all groups' decision-making, it is very surprising how little we know and understand about it. Collecting rigorous data is difficult, but the theory itself remains underdeveloped to this day. The problem has attracted prominent scholars and some indeed have put vote trading at the center of politics, as Arthur Bentley advocated more than a century ago, but disagreements about both the mechanics and the effects of vote trading have never been fully resolved. Are trades of votes for votes different from trades of votes for more flexible compensation, like money or favors? Are trades most commonly decentralized, voter-to-voter, or are they coordinated by group leaders or others? Are they mostly one-to-one, or do they involve larger coalitions? Most fundamentally, are vote trades beneficial or welfare reducing?

It is not difficult to see why the answer to this last, crucial question is not a priori clear. Vote trades allow a voter to modulate her influence, increasing it on her highest priorities by acquiring additional votes, at the cost of reducing it on less important issues by give one's votes away. If voters differ on their priorities, the reallocation of voting power resulting from their trades should be beneficial. The difficulty, however, is that a vote trade affects all voters, not only those directly trading. The externalities will be both positive and negative, and there is no immediate reason why the net effect should be positive.

Research thrives on questions whose answer is not obvious. Why has vote trading remained relatively unexplored? One possible reason is that the difficulty here is deeper, and has to do with the very nature of the good traded – a vote. Imagine, as a fully abstract exercise, an egalitarian society where no-one is budget constrained and votes can be bought and sold on a market in exchange for a numeraire. The value of a vote is not in the consumption utility it provides; it is in the influence it grants its holder. But such influence does not depend only on the number of votes the voter holds; given the voting rule, it depends on the full distribution of votes among all other voters. As a result, standard economic concepts like competitive equilibrium or the core fail: in a market for votes, competitive equilibrium does not exist and the core is empty. The intuitive parallel between a vote and a tradable economic good is misleading.

As we will show in what follows, progress has been made in recent works, in part by defining alternative equilibrium concepts or focusing on the stability of the system, in part by proposing and studying alternative institutions. The overall message invites to be humble: neither logrolling (trading votes for votes) nor markets for votes (trading votes for a numeraire) are found to have properties that are both positive and robust: welfare gains relative to no-trade require strong and typically special assumptions. Other institutions, allowing voters to reallocate a given total budget of votes across multiple decisions without trading with others, seem more promising—both the theory and the existing experiments suggest gains, without requiring fine-tuning the design to the details of the environment. A realistic goal is a modest but reliable improvement over majority rule. It is with this objective in mind that practically implementable schemes, inspired by but crucially different from vote trading, are most likely to be successful.

We highlight results that we believe can be useful going forward but of course the survey is limited and neglects important related topics. First, because we consider the lack of a common theoretical approach the first challenge to understanding vote trading, we focus exclusively on theoretical contributions, and, when they exist, on their experimental tests, neglecting the empirical literature on vote trading.¹ Second, the survey does not discuss

¹In the absence of direct measures of vote trading, the empirical literature is particularly creative in drawing insights from the data that can be observed. Cohen & Malloy (2014) and Guerrero & Matter (2016) and the references cited there are good entry points for the interested reader.

the rich, related work on selling votes to candidates or lobbyists.²

The survey is organized in three parts. The first part, Section 2, discusses logrolling, or trading votes in exchange for other votes; the second part, Section 3, discusses trading votes for a numeraire; finally Section 4 discusses what voters can achieve when allowed to shift their own votes across different decisions, in a sense trading with themselves.

2 Trading votes for votes - logrolling

We begin our discussion of vote trading with its most intuitive form: the exchange of votes, or logrolling. Studies of vote trading multiplied in the 60's and 70's, originally triggered by Buchanan & Tullock (1962)'s *Calculus of Consent*. Buchanan and Tullock reacted to Arrow (1951)'s *Social Choice and Individual Values* by offering a different analysis of majority voting. Building on Condorcet's work, Arrow had highlighted that an alternative that is majority preferred to each of the others (the Condorcet winner) may not exist, and that in such a case majority voting generates cycles and indeterminacies. Buchanan and Tullock argued that voters in reality face not a single decision but multiple decisions, and disagreement on priorities leads them naturally to vote trading: "*The existence of a logrolling process is central to our general analysis of simple majority voting*" (p. 133). Noting that over multiple decisions, the Condorcet winner is defined over the combined result of the multiple decisions at stake, they stated that vote trading is not only common, it is the means through which voters overcome the limitations of majority voting. In articles that predate the book (Tullock, 1959, 1961) and in a single-authored appendix to *Calculus of Consent*,³ Tullock makes two statements, precise enough to be formally analyzed. First, if a Condorcet winner exists, vote trading will reach it: "*The problem which puzzled Condorcet, Carroll, Laplace, and Black was [...] finding a system of voting which would lead to [...] the genuine will of a majority of the group.[...] [L]ogrolling eliminates the basic problem.*" (p. 332–333). Second, if a Condorcet winner does not exist, vote trading will prevent the development of majority cycles: "*If logrolling is the norm, [...] then the problem of the cyclical majority*

²Classic references are Myerson (1993), Groseclose & Snyder Jr (1996), Dal Bo (2007), Dekel et al. (2008) and Dekel et al. (2009). See also Faliszewski & Rothe (2016) for a survey on vote-buying in the computer science literature.

³Tullock, "Theoretical Forerunners", Appendix 2 to the *Calculus of Consent*.

vanishes.” (p. 336).

These arguments, of great originality and influence, were supported by suggestive models that combined a striving for rigor with very effective story telling. The models however were too informal in their assumptions and definitions to provide definite support for the conjectures they proposed. It was the beginning of a very large literature that focused primarily on logrolling and alternatively took issue or supported Buchanan and Tullock’s positions. The number of articles that followed in a relatively short window of time was very large. We can refer here only to very few, which did deliver lasting results.⁴ We distinguish two forms of vote trading: explicit trades—explicit exchanges of votes on one issue for votes on a different issue, and implicit trades—the bundling of different issues in a single vote, such that either all issues pass or all fail.

2.1 Explicit logrolling

In a response to Coleman (1966), who was writing in support of Buchanan and Tullock’s position, Park (1967) made two important contributions. The goal was to analyze the claim that vote trading resolved majority cycles. Park’s first contribution was to formalize the question precisely, clarifying the meaning of stability. The model Park suggested remains very relevant and because clarity will help our discussion, we paraphrase it below.⁵

Consider a committee $\mathcal{C} = \{1, \dots, N\}$ of N (odd) voters who must approve or reject each of K independent binary proposals. Committee members have separable preferences represented by a profile of values, z , where z_i^k is the utility i experiences if k passes, positive if i is in favor of k and negative if i is opposed.⁶ It is convenient to describe the model in terms of cardinal intensities, but its logic depends exclusively on ordinal rankings over the full set of 2^K possible outcomes (taking into account all possible combinations of proposals passing or failing). Any change in numerical values that keep all ordinal rankings identical is irrelevant. Separability, on the other hand, is an important assumption because it restricts the set of admissible profiles of ordinal preferences over the set of outcomes. It will be maintained

⁴Among those we do not discuss, see for example Wilson (1969), Miller (1977a), Mueller et al. (1972) and Schwartz (1975).

⁵We borrow, with few modifications, from the description of the similar model in Casella & Palfrey (2019).

⁶The value of any proposal failing is normalized to 0.

throughout Section 2.⁷ Proposals are voted upon one-by-one, and each proposal k is decided through simple majority voting. Before voting takes place, committee members can trade votes. Call v_i^k the votes held by voter i over proposal k , $v_i = (v_i^1, \dots, v_i^K)$ the profile of votes held by i over all proposals, and $v = (v_1, \dots, v_i, \dots, v_N)$ a profile of vote holdings for all voters and proposals. At the initial vote allocation, denoted by v_0 , all voters hold 1 vote on each issue, or $v_{0i} = (1, \dots, 1)$ for all i . Let \mathcal{V} denote the set of feasible vote allocations.⁸

Given a feasible vote allocation v , when voting takes place on proposal k each voter has a dominant strategy to cast all votes in favor of the proposal if her proposal's value is positive ($z_i^k > 0$), and against the proposal if her proposal's value is negative ($z_i^k < 0$). We indicate by $\mathbf{P}(v)$ the *outcome* of the vote if voting occurs at allocation v , i.e. the set of proposals that receive at least $(N+1)/2$ favorable votes and therefore pass. Finally, we define $u_i(v)$ as the utility of voter i if voting occurs at v : $u_i(v) = \sum_{k \in \mathbf{P}(v)} z_i^k$.

The focus is on the existence and properties of vote allocations that hold no incentives for trading. Define:

Definition 1 *A coalition of voters $S \subseteq \mathcal{C}$, is said to **block** v if there exists a feasible vote allocation $\hat{v} \in \mathcal{V}$ such that $\hat{v}_i \neq v_i$ and $u_i(\hat{v}) > u_i(v)$ for all $i \in S$, and $\hat{v}_j = v_j$ for all $j \notin S$.*

Definition 2 *A vote allocation $v \in \mathcal{V}$ is **stable** if there exists no coalition of voters who can block v .*

Park's question is the relationship between a vote allocation that delivers the Condorcet winner, if it exists, and stability. Here he made an important observation, which we consider Parks' second contribution. Because it is instrumental to a number of subsequent results, we report it here as a lemma:

Lemma 1 (Condorcet Winner).⁹ *If the Condorcet winner exists, it can only be $\mathbf{P}(v_0)$, the outcome corresponding to the pre-trade vote allocation.*

⁷Non-separable preferences on a combinatorial domain are an important topic in voting, as reviewed by Lang & Xia (2016). Logrolling with non-separable preferences has been investigated in Miller (1977b), Schwartz (1981) and more recently in computer science by Grandi et al. (2019). In that paper, voters hold lexicographic preferences and can make binding side-payments promises to other voters. The authors show that these side-payments help sustain efficient outcomes.

⁸ $v \in \mathcal{V} \iff \sum_i v_i^k = \sum_i v_{0i}^k = N$ for all k and $v_i^k \geq 0$ for all i, k .

⁹The lemma is implied by Park's results. It was later made precise and derived independently by Kadane (1972).

The logic is simple. Take any outcome such that $m \in \{1, \dots, K\}$ issues are decided against the majority, while the remaining $K - m$ follow the majority's preference. Consider any issue k decided against the majority. Because preferences are separable, resolving that issue in favor of the majority while keeping all other decisions unchanged must be majority preferred to the initial outcome. But since this is true for any $m \in \{1, \dots, K\}$, it follows that the only possible Condorcet winner is outcome $\mathbf{P}(v_0)$, where all issues are decided as the majority prefers. The reasoning is very intuitive but also very useful. Because $\mathbf{P}(v_0)$ is the pre-vote trade outcome, introducing vote trading cannot help to reach the Condorcet winner.¹⁰

Park proceeds to argue that only the vote allocation yielding the Condorcet winner is stable because any other is, by definition, blocked by a majority of voters. But the reasoning here depends on the timing and enforcement of the trading contracts. These are important details that Park's analysis left unspecified.

Defining precisely the process through which vote trades take place was one important contribution of Riker & Brams (1973)'s classic paper. Riker and Brams follow a suggestion by Mueller (1967) and suppose that vote trades are fully enforceable—each trade is a physical exchange of ballots, with each ballot specialized by proposal. Trading can then be analyzed as a dynamic process of successive vote exchanges. Riker and Brams make two further assumptions: trades are between pairs of voters only, and traders are myopic—that is, traders evaluate the payoff from the trade as if voting took place immediately after the trade (an assumption incorporated in the definition of stability formulated above).¹¹ Vote trading is thus defined as a sequence of pair-wise trades yielding myopic strict gains to both traders. Casella & Palfrey (2019) formalize Riker and Brams' process by representing it as an algorithm:

A family of vote-trading algorithms: *Start from any vote allocation v_0 . If there is no pair-wise strictly welfare-improving trade, stop. If there is one such trade, execute it. If there are multiple, choose one according to rule*

¹⁰Of course it could be that vote trading takes the vote allocation *back* to its pre-trade realization.

¹¹Especially when combined with non-enforceable trades, the effort to consider forward-looking behavior was a source of confusion in the early literature. It is often unclear whether trades are not taking place because forward-looking voters anticipate that they would be undone, or they do take place but are irrelevant because they are undone in the future. Note that this muddles the equilibrium concept.

R. Continue in this fashion until no further improving trade exists.

Rule *R* specifies how the algorithm selects among multiple possible trades and distinguishes individual algorithms in the family.

Having phrased vote trades so precisely, Riker and Brams identified one of their crucial properties:

Lemma 2 (Pivotal vote). *In any trade executed under an algorithm in the family, both sides of the trade must include a pivotal vote.*

The requirement that the trade be strictly improving for both voters makes the observation immediate: it must be that the trade changes the outcome on at least two proposals, and thus at least two pivotal votes, one initially held by each voter, must be changing hands. It is an essential addition to Buchanan and Tullock’s initial insight: vote trades cannot reflect preferences alone but are necessarily constrained by the “power” of the votes an individual holds, and the correct measure of power is the vote’s pivotality.¹²

Armed with these precise definitions, Riker and Brams study specific examples. They proposed the scenario reported in Table 1 below, with $K = 6$ and $N = 3$. Each column is a proposal (X, Y, W, Z, T , and V) and each row a voter (1, 2, and 3). Each cell $\{k, i\}$ reports z_i^k , the value attached by voter i to proposal k passing.

Table 1: Voter preferences in Riker and Brams’ example

	X	Y	W	Z	T	V
1	3	3	2	-4	-4	2
2	2	-4	-4	2	3	3
3	-4	2	3	3	2	-4

Absent trade, each proposal passes and $u_i(v_0) = 2$ for each voter. However, any two voters have a strict payoff-improving trade. Consider the following sequence of trades (corresponding to a specific *R* in the definition of the algorithm): 2 gives her X vote to 3, in exchange for 3’s Y vote (note that both votes are pivotal). After the trade, both X and Y would fail; $u_i(v_1) = 4$

¹²Coleman (1966) already noted that the value associated with a proposal had to reflect the distribution of power over it, but did not take the explicit step to pivotality.

for $i = 2, 3$, and $u_1(v_1) = -4$. But further payoff-improving trades are possible: voter 1 gives a W vote to 2, in exchange for a Z vote, leading both proposals to fail: $u_2(v_2) = 6$, $u_i(v_2) = -2$ for $i = 1, 3$. Finally, voters 1 and 3 trade votes on proposals T and V , reaching vote allocation v_3 such that $\mathbf{P}(v_3) = \emptyset$ —all proposals fail—and $u_i(v_3) = 0$ for each voter. Note that at v_3 trade stops: there is no further pair-wise improving trade.

Although each individual trade is profitable for the voters involved, the final outcome is Pareto inferior to the pre-trade outcome: everyone is worse-off. Riker and Brams' article was extremely influential because it provided a transparent counter-example to Buchanan and Tullock's optimistic view of vote trading. Because vote trades cause externalities to those individuals who are not part of the trades, there can be no presumption that trade will improve welfare. It remains one of few undisputed conclusions about vote trading, and we phrase it as a formal result:

Result (Lack of Pareto optimality). *Individually rational pair-wise vote-trading can deliver Pareto inferior outcomes.*

Buchanan and Tullock's conjectures about logrolling and the Condorcet winner remained open.¹³ Does vote trading always converge to the Condorcet winner when it exists? And when it does not exist, does vote trading converge to a stable allocation? Does trading converge when exchanges are not pair-wise but occur within coalitions of arbitrary size? Both Riker and Brams and Ferejohn (1974) conjectured that allowing coalitions imperils stability. Ferejohn concluded that even when exchanges are limited to pair-wise trades, there is no logical connection between the Condorcet winner and the stability of a vote allocation.¹⁴ Ferejohn (1974) is an unpublished but highly cited note, striving to create some order in the existing literature. Written after the large production of works on vote trading in the 60's and early 70's, the note ends with sobering words: “[W]e really know very little theoretically about vote trading. We cannot be sure about when it will occur, or how often, or what sort of bargains will be made. We don't know if it has any desirable normative or efficiency properties.” (Ferejohn, 1974, p. 25).

The frustration these words reflected effectively stopped theoretical research on trades of votes for votes until very recently. Casella & Palfrey (2019) goes back to the conceptual tools proposed by the early authors and

¹³Note that there is no Condorcet winner in Riker and Brams' example.

¹⁴Contradicting Park, as well as Bernholz (1973, 1974), Koehler (1975), Kadane (1972), Schwartz (1977, 1981).

closes the loop between the definition of stability proposed by Park and the dynamic process of trade suggested by Riker and Brams. The first question to be addressed concerns the stability of the system: as in Riker and Brams' example in Table 1, a trade of votes changes the vote allocation and typically creates the possibility of new trades, which themselves then generate incentives for further trading, etc.. Could the result be an infinite cycle?

In fact, the system has a built-in bias towards stability. It is easy to see it in the example: note that each trade requires a voter to give away a vote on a proposal to which the voter attributes an absolute value of 2 and acquire a vote on a proposal with absolute value of 4. Because trades must be strictly welfare-improving, the vote that is traded away must be pivotal on a proposal whose resolution matters less to the trader than the resolution of the proposal whose pivotal vote is instead acquired. Thus in each trade a voter must accumulate votes on higher (absolute) value proposals. But that means that for any given realizations of values, a voter can only be part of a finite number of trades, and since the number of voters is finite, the total number of trades must be finite.

As stated, this simple intuition is not robust to straightforward generalizations of the allowed trades—it is not robust to coalitional trades, or to voters adding to their trades non-pivotal votes¹⁵. Yet, with the appropriate qualifications, the intuition can be extended. In the general case, it is not true that all sequences of trades must converge to a stable vote allocations, but it is true that there always is one sequence that does. We can phrase the conclusion precisely:

Result (Stability of vote trading). *Extend the vote-trading algorithms to blocking coalitions of arbitrary size. For any number of voters N , any number of proposals K , any (separable) preferences z , there is always a rule R (hence a sequence of trades) that converges to a stable vote allocation. Such allocation must deliver a Pareto optimal outcome.*

The result contradicts the earlier belief that coalitional trades must induce instability, and answers one of Buchanan and Tullock's claims: vote trading—as formalized here—does not always eliminate the possibility of cycles, but there is always a sequence of trades that does converge to a vote allocation from which no further trade takes place and such that majority delivers an

¹⁵Such trades do not change outcomes and thus have no effects on utility. As a result, voters could give away votes on high value proposals in exchange for votes on low value ones.

unambiguous outcome. Moreover, as no restrictions are imposed on trading coalitions, any Pareto-inferior outcome would be blocked by the coalition of the whole. Thus any outcome reached from a stable vote allocation must be Pareto optimal: vote trading can always reach a Pareto optimal outcome.¹⁶

However, as Ferejohn had seen, the result has little to do with the existence or not of a Condorcet winner. The indeterminacy that is ruled out concerns cycles of never-ending trades, not majoritarian cycles over multiple alternatives, and with myopic traders and enforceable trades there is no logical connection between the two types of instability.

The lack of such a connection is stressed by the second of Casella and Palfrey's results. In partial rebuttal of Buchanan and Tullock's first conjecture, vote trading need not reach the Condorcet winner when one exists. It does if $K = 2$ or $N = 3$, but need not do so in general: the Condorcet winner can exist and yet vote trading can deliver a different outcome. Yet, the rebuttal is only partial because one aspect of the question is left open. Suppose a Condorcet winner exists. Is it the case that it can always be implemented by a stable vote allocations reached by vote trading? Within the precise assumptions of this model, the answer is not known, but we will return to the question below.

Following the simple logic of the vote trading algorithms allowed some progress in a complex problem. But does such a logic have any predictive power? In a companion paper, Casella & Palfrey (2018b) study the question in a lab experiment. Although controlled lab experiments should presumably be privileged tools for studying a question of difficult empirical documentation, experiments on vote trading are in fact very few. The pioneering study is McKelvey & Ordeshook (1980) which focused primarily on alternative cooperative solution concepts, anticipating in spirit some of the experiments on vote trading with social preferences that will be discussed below. Casella & Palfrey (2018a) study instead the dynamics of vote trading, implementing in

¹⁶Vote allocation v_3 in Riker and Brams' example above is stable relative to pair-wise trading (and Pareto inferior to no-trade), but is not stable relative to unconstrained coalitional trading because the coalition of the whole can trade back to the initial allocation and induce a strict gain for all its members. This full sequence of trades is thus a non-converging cycle. On the other hand, suppose that after trading on (X, Y) , voters 2 and 3 were to trade on (W, V) . The final vote allocation is now stable because there is no additional trade, either pair-wise or involving all three voters, that can benefit all parties to the trade. Casella & Palfrey (2019) shows that restricting coalitions to be pair-wise guarantees that *all* sequences of trades must converge to a stable vote allocation (but allows for the possibility that the final outcome be Pareto inferior).

the lab a scenario that mimics the model discussed here. The experimental data suggest that stability is indeed useful in thinking about convergence of vote allocations: more than two thirds of all vote allocations at which trading stops are stable. In addition, the incentives driving trades seem to mirror the theory well: averaging over participants assigned the same preferences, all those in voter roles predicted to acquire votes do so, and all those in voter roles predicted to give away votes do so, over all treatments and over all proposals. On the other hand, the precise dynamics of the vote trades observed in the lab are less well described by the model: although myopic strictly improving trades are observed and, according to statistical tests, appear intentional, a large fraction of trades involves non-pivotal votes and has no effect on payoffs. Participants tend to accumulate votes on high value proposals, even when the votes acquired do not change outcomes. It is tempting to interpret such behavior as prudential behavior, and thus as an invitation to question the assumption of myopia adopted in the theoretical approach.

Indeed, returning to the theory, it is natural to ask whether vote trading would have better properties if voters were forward-looking. Introducing forward-looking behavior in a model of vote trading is not trivial, because any individual future trade has the potential to influence the profitability of any current trade and affect the whole dynamic path. One possible strategy is to borrow from the literature on farsightedness in cooperative game theory.¹⁷ Casella & Palfrey (2018a) propose extending the notion of stability by exploiting the concept of *farsighted core* (*F-core*):

Definition 3 *A vote allocation v is farsightedly-dominated (F -dominated) by $v' \neq v$ if there exists a chain of coalitional trades that leads from v to v' and such that all parties to any trade strictly prefer v' to the vote allocation at which they trade. A vote allocation v belongs to the F -core if there is no v' that F -dominates v .*

Coalitional stability as defined earlier corresponds to the myopic core. The extension to the F -core is then quite natural. Note the change in the blocking requirement: under myopia, all members of a trading coalition must

¹⁷See Chwe (1994), Mauleon et al. (2011), Ray & Vohra (2015), Dutta & Vohra (2017) and the references therein. However, none of these models can be adapted directly to the problem, both because of the externalities involved, and because the opportunities for trade depend on the vote allocation.

prefer the outcome corresponding to the vote allocation reached via their trade; under farsightedness, this needs not hold, but all must prefer the final vote allocation reached through the full sequence of trades.

Given the initial vote allocation v_0 , vote trading is said to converge to farsighted stability if either v_0 is itself in the F-core, and no farsighted vote trade takes place, or there exists a farsighted chain of coalitional trades that leads from v_0 to some allocation in the F-core. Farsightedness is extremely demanding, and improbable in real applications.¹⁸ From a theoretical point of view, however, we would expect it to lead to better outcomes. In fact the opposite is true. First, examples show that the earlier stability result is no longer guaranteed: there exist scenarios such that neither v_0 nor any vote allocation that F-dominates v_0 are in the F-core. While a sequence of trades that leads from the initial vote allocation to a myopically stable allocation always exists, no trading path to a farsightedly stable allocation may exist. Second, with farsightedness, vote trading becomes inconsistent with delivering the Condorcet winner. We know that the Condorcet winner can only be the outcome corresponding to v_0 , $\mathbf{P}(v_0)$; but with farsightedness, vote trading cannot lead to any allocation v such that $\mathbf{P}(v) = \mathbf{P}(v_0)$, because v cannot F-dominate v_0 if $\mathbf{P}(v) = \mathbf{P}(v_0)$. Farsightedness thus provides a negative answer to the question we left open earlier: the Condorcet winner may exist and yet it may be that no vote allocation delivering it is in the set of farsightedly stable allocations reachable from v_0 . These conclusions depend on the specific notion of farsightedness, and other concepts should be proposed and studied. As a first pass, though, the results are not encouraging.

On the whole, then, the message of the theoretical literature as it stands now is mixed at best: stability may be reached, although even this becomes doubtful with farsightedness, but, contra Buchanan and Tullock, the problem of delivering an outcome that is a sincere mirror of the will of the majority is not solved by vote trading.

2.2 Implicit logrolling and bundling

Besides explicit vote-for-vote trading, vote exchanges over multiple issues may occur implicitly, as recognized originally by Buchanan & Tullock (1962).¹⁹

¹⁸The notion of farsightedness described here is soundly rejected by the experimental data in Casella & Palfrey (2018b).

¹⁹See also Tullock (1970) and Haefele (1971) for early discussions. Tullock proposes a model where multiple policy decisions are consolidated in a single bill, and voters' pref-

To fix ideas, consider a simple environment with N voters and $K = N$ proposals, denoted by k_1, \dots, k_N , where each proposal k_i is strongly favored by voter i . Voter i receives positive utility $z_i^{k_i} = z^H > 0$ if proposal k_i is approved, but negative utility $z_i^{k_j} = z^L < 0$ from any other proposal passing. In this environment, each individual proposal would be rejected by a majority. However, if $z^H > (n - 1)|z^L|$, all voters would be better-off if all were approved.²⁰

The Pareto-superior outcome where all proposals are approved can arise from repeated interactions among the voters, as is usually the case in legislatures or international organizations. In the model of Carrubba & Volden (2000), each period of a repeated game features K sequential votes, with the full set of proposals being decided, proposal-by-proposal. As usual in these games, the superior outcome can be reached on the equilibrium path if voters are sufficiently patient, when the myopic incentive to vote against another voter's issue is dominated by the threat of having one's own issues turned down in the subsequent periods.²¹

Implicit vote exchanges could also be the result of social preferences. Fischbacher & Schudy (2014, 2020) study whether norms of trust and reciprocity can sustain deliberate votes against one own's interest, if issues are voted upon sequentially and votes are transparently recorded. Reciprocating voters approve other voters' issues if these voters have approved the reciprocating voters' issues in the past. Trusting voters approve other voters' issues, anticipating that these voters will reciprocate. In the lab, cooperation among voters is partial, in the sense that issues voted upon earlier (in the sequence of issues) are approved at a higher rate than later ones. In turn, this pattern can be strategically exploited by voters or agenda setters.²²

erences can be described spatially. Haefele studies the interesting case when implicit logrolling over multiple issues takes place through the election of a representative.

²⁰Kadane (1972) considers the opposite but analogous environment where each proposal is approved by a majority, but has a sufficiently strong opponent to yield a net loss. He observes that in such a case, it is easier to pass proposals when voted upon one-by one, rather than as a joint bundle.

²¹To be precise, in the paper, the number of proposals voted upon at each period corresponds to the size of a minimal winning coalition, but this does not change the nature of the argument.

²²In Fischbacher & Schudy (2014), if given the choice, voters favoring the first issues prefer to vote sequentially issue-by-issue rather than over the whole bundle. In Fischbacher & Schudy (2020), agenda setters choose to put their preferred issues at the beginning of the sequence of issues voted upon.

More surprisingly, in rich dynamic scenarios with multiple votes, some voters may offer their votes without receiving anything in exchange. Hortala-Vallve (2011) considers a model with three farsighted legislators voting on two issues, so that four outcomes (each issue passing or failing) are possible. In a dynamic game, each period's status quo is an outcome, which is confronted to another randomly chosen outcome, and next period's status quo is decided by legislators' vote. In the leading example, the first issue is favored by voter 1 only, the second issue is favored by voter 2 only, while voter 3 disapproves the second issue more than the first.²³ The key result consists of an equilibrium where the first issue is approved but not the second one, although both voters 2 and 3 would be myopically inclined to reject the first issue. The result relies on the seemingly generous behavior of voter 3, who approves of the first issue without apparently receiving anything in exchange, but in fact preventing the occurrence of the outcome where both issues are approved.

A more direct form of implicit vote trading is achieved by bundling, when an agenda setter proposes a bundle of proposals to be voted upon in one single vote. Either all proposals in the bundle fail, or all pass. In the simple environment considered above, any bundle of at least $(n + 1)/2$ proposals, each favored by a different voter, would be approved at a majority. Bundling can then function as a form of vote trading. It can, however, also serve other purposes. Imagine an agenda setter, adding to such bundle one additional proposal disliked by the voters, who all incur a disutility $-\varepsilon$ in case of approval. If ε is small enough, the joint bundle continues to be approved by a majority. The additional proposal would not have been approved after any explicit vote exchange; rather, it is imposed on the voters by the agenda setter.²⁴

This example raises the question, addressed by Câmara & Eguia (2017), of how to bundle issues optimally, from the point of view of an agenda setter intent on passing a maximal number of proposals against the preferences of a majority of the voters. Consider again the example with three voters and three issues, with each issue k_i , for $i \in \{1, 2, 3\}$, only favored by voter i , and with $z^H = 3$ and $z^L = -2$. The agenda setter has complete information on voters' preferences. If issues are voted upon separately, each is turned down by a majority of two against one. If all issues are bundled together,

²³Note that the example departs from the environment described above because of the asymmetry in voter 3's preferences.

²⁴This form of agenda power is notably prevalent in corporate governance, as discussed in Bebchuk & Kamar (2009) and Cox et al. (2015).

each voter incurs a loss of -1 in case of approval, so that the bundle too is rejected by all. Instead, the agenda setter can profitably bundle together issues 1 and 2 together, while leaving issue 3 separate. A majority consisting of voters 1 and 2 then approves the bundle of issues 1 and 2, while turning down issue 3.

In fact, Câmara & Eguia (2017) show that the agenda setter can do even better if allowed to simultaneously bundle *and slice* issues. For instance, issue 1 could be sliced into two issues, each one favored by voter 1 only, and with utilities $z^H = 3\alpha$ and $z^L = -2\alpha$ for the first slice and $z^H = 3(1 - \alpha)$ and $z^L = -2(1 - \alpha)$ for the second slice, for some $\alpha \in (0, 1)$. In the example above, the agenda setter can have all three issues approved by appropriately slicing and bundling. For example, she can propose as a first bill a bundle of one-half of issues 1 and 2 (approved by voters 1 and 2); as a second bill a bundle of one-half of issues 1 and 3 (approved by voters 1 and 3), and as a third bill a bundle of one-half of issues 2 and 3 (approved by voter 2 and 3). The agenda setter succeeds in having all issues passed, while all voters are worse-off than if voting took place issue-by-issue (with no logrolling), as in the earlier example of Riker & Brams (1973). The approach leverages results from Bayesian persuasion (Alonso & Câmara, 2016). It is particularly ingenious and at the same time opens new, realistic questions about the actual composition of bills in legislatures. We emphasize it here:

Result (Optimal bundling). *A strategic agenda setter able to slice and bundle issues may achieve voting outcomes that make all voters strictly worse-off.*

In the paper, Câmara and Eguia characterize the optimal slicing and bundling policy by the agenda setter for a broad set of preferences and show that the optimal agenda always makes a majority of voters weakly worse-off.²⁵

²⁵The general model relies on the following assumptions: (i) the bundle of all issues would be rejected; (ii) at least one bill can pass, and (iii) voters have strict preferences. The result follows because if a majority of voters were strictly better-off, the agenda setter could profitably modify the agenda by adding some issue slices to at least one bill. Adapting Alonso & Câmara (2016), one obtains that a majority of voters is strictly worse-off if any optimal agenda contains at least two bills that are expected to pass.

3 Buying and selling votes for a numeraire

As suggested, for example, by Coleman (1966), exchanges of votes could be theorized in a manner quite different from logrolling, as resulting from purchases and sales via a numeraire. The analysis starts from ignoring ethical and philosophical objections—the skewed results that would follow the inequality of budgets, and the deeply disturbing explicit link between political and economic voice. Even in an abstract world of full equality, there are reasons to be concerned. The problem lies in the nature of votes: the utility that can be derived from a vote depends not on holding the vote itself, as is the case for a private good, but on the influence the vote can confer on the final decision. Thus, for any given decision rule, the value of the vote depends on the distribution of the votes among all other voters. It is a particular form of externality that threatens not only the welfare properties of the market but its very functioning.

In the discussion that follows, there are N (odd) voters, a single binary decision ($K = 1$), and preferences are described in terms of cardinal intensities. Welfare is measured in terms of utilitarian efficiency, and thus relies on the cardinal representation. We again use the notation $z = (z_i)_{1 \leq i \leq N}$ to indicate the profile of preferences, and we denote by $x_i = |z_i|$ the *value* or *intensity* of voter i . Absent any trade, the outcome preferred by the majority would be implemented. Instead, we suppose that votes can be bought and sold for a numeraire, a private good that functions as means of exchange. Voters are risk-neutral, and a voter i receiving a net transfer t_i in the market derives utility $u_i = t_i + z_i$ if the proposal passes and $u_i = t_i$ otherwise.

3.1 Coordinated vote markets

One route to overcome the difficulties posed by vote markets is to assume that transactions are explicitly coordinated by a third party. In Philipson & Snyder (1996), a floor manager sets prices, receives buying and selling offers from the voters, and concludes trades. The identity of minority and majority voters is common knowledge, so that, if trade occurs, members of the minority buy while members of the majority sell. Philipson and Snyder assume that the floor manager enacts trades only if all majority voters offer to sell their votes. The veto granted to all majority voters prevent some beneficial

trades to occur,²⁶ but ensures that at any equilibrium where vote trading occurs, the trade is Pareto-improving. A similar approach, again concluding in support of a vote market, is taken by Koford (1982), where trade is coordinated by two benevolent party leaders who cooperate in maximizing their members' surplus. In both studies, the strong assumptions highlight that guaranteeing welfare gains from vote markets requires unusual restrictions. What happens when such restrictions are dropped? A natural starting point is the competitive market, the foundation of the economic theory of private exchange.

3.2 Competitive equilibrium

Consider then a competitive market for votes: a market where all transactions take place at the posted market price, all traders take the price as given, and each voter is free to approach and trade with whomever he chooses.

To see why such a market for votes can be problematic, consider the example described by Casella et al. (2012). There are two alternatives, A and B , and three voters, Ann, Alex, and Bob, each with one vote, choosing between A and B by majority voting. Ann and Alex both prefer A to B and derive utilities of 12 and 10, respectively, from A and 0 from B ; Bob, with more extreme preferences, derives utility of 30 from B and of 0 from A . Thus the majority prefers A but utilitarian efficiency favors B . It is natural to suppose that if Bob could buy Ann and Alex' votes for a price, the problem would be solved. For example, if Bob bought both votes for a price of 13, the outcome would be a Pareto improvement over no trade. But Bob will never demand two votes because one would be redundant: at any price below 30, Bob demands one vote. At any price above 12, both Ann and Alex offer to sell, and there is excess supply. At any price below 10, no-one sells. Consider then any price between 10 and 12: Alex is happy to sell, and her vote is pivotal. But if Alex sells, Ann's vote is worthless, and thus Ann is also willing to sell at any positive price: again there is excess supply. Because demands are interdependent, there is no price at which the market clears.

²⁶Philipson and Snyder further assume that if trade occurs, the revenue collected from the minority is divided equally among majority members. It follows that trade may occur at equilibrium if and only if $\sum_{i \in m} x_i > M \max_{i \in M} x_i$ (Proposition 1, p.251). Thus, if $z = (-4, -4, 1, 1, 5)$, so that $m = 2$ and $M = 3$, it is optimal for the minority to prevail, but equilibrium trades cannot achieve this outcome.

The non-existence of equilibrium in a decentralized market for votes is a general and robust problem, identified by the early literature (Coleman, 1967, Ferejohn, 1974, Shubik & Van der Heyden, 1978), and reiterated in more recent studies (Weiss, 1988, Piketty, 1994, Philipson & Snyder, 1996, Kultti & Salonen, 2005).²⁷ We can state it succinctly. Suppose, in line with the example above, that each voter’s preferred alternative (the sign of each z_i) is known, while intensities ($x_i = |z_i|$) may or may not be known. Voters are divided into a majority of M voters, who share the same sign of z_i , and a remaining minority of $m = N - M < M$ voters, who share the opposite sign.²⁸ Call p the market price of a vote, and $u_i(d, p)$ the utility of voter i for a net demand vector d and a price p (where $d_i = -1$ if i sells his vote).²⁹ Then:

Result (Non-existence of competitive equilibrium). *If $N > 2$, for any z such that $M \geq m > 0$, there is no price p and demand vector $d^* = (d_i^*)$ such that*

$$(i) \quad d_i^* \in \arg \max_{d_i} u_i(d_i, d_{-i}^*, p) \quad \text{and} \quad (ii) \quad \sum_{i=1}^N d_i^* = 0.$$

The logic is identical to the example above and we do not repeat the proof here. Because the problem comes from the discontinuity in the value of a vote at pivotality, intuition suggests that equilibrium can be recovered by inserting randomness and “smoothing” the problem. Thus a natural suggestion is to drop the assumption that the direction of preferences is known. However, introducing uncertainty over preferences will not help: there cannot be a purchase of two votes because one would be redundant, and there cannot be a purchase of a single vote because the person who is not trading would hold a useless vote, and thus would offer to sell it for any price. The problem has not changed.

Casella et al. (2012) re-establish the existence of competitive equilibrium by allowing traders to express mixed demands—i.e. to randomize between

²⁷Weiss (1988) studies coalitional agreements over the choice of a tax rate for financing a public good. He found that the core of the game is generically empty, a result with similar flavor to the non-existence of equilibrium in competitive vote markets.

²⁸We sometimes abuse notation and write M or m for the sets of majority and minority voters.

²⁹Note that individual demand depends not only on the price, but also on the demands of others, contrary to private good markets. We will return to this point below.

buying and selling.³⁰ The equilibrium price clears the market in expectation: the price must be such that expected supply equals expected demand. Ex post market clearing is achieved via a rationing rule. The equilibrium concept, called *ex ante equilibrium* by the authors, imposes the discipline of market balance while allowing individual randomizations: market clearing is violated ex post, but the violations must be non systematic and unexpected.³¹

Definition 4 *Given a rationing rule R , an ex ante competitive equilibrium is a vector of probabilistic net demands δ^* and a price p^* such that:*

$$(i) \quad \delta_i^* \in \arg \max_{\delta_i} u_i(\delta_i, \delta_{-i}^*, p^*, R) \quad \text{and} \quad (ii) \quad \mathbb{E}\left[\sum_{i=1}^N \delta_i^*\right] = 0.$$

In the model proposed in the paper, both intensity and direction of preferences are private information but each voter is equally likely to be in favor of either alternative. The authors construct an ex ante equilibrium for an arbitrary finite number of voters, arbitrary intensities, and two plausible rationing rules.³² In equilibrium, all voters with the exception of the two with highest values offer to sell their votes; the voter with highest value demands $(N - 1)/2$ votes, i.e. a majority of votes and full decision power; the voter with second highest value randomizes between selling her vote and demanding $(N - 1)/2$ votes. The price is always such as to keep the second highest value voter at indifference between the two actions, and her mixing probability guarantees expected market clearing.³³

The equilibrium has two important features. First, in equilibrium the market comes to approximate an all-pay auction for power: the only individuals bidding to purchase are the two with highest values, with the second highest at her indifference point. The result mirrors votes' special characteristics: utility derives not from owning votes, but from owning power. Thinking

³⁰Formally, a probabilistic net demand δ_i is a probability distribution on the set of all possible demands: $\{d_i \in \mathbb{Z} \mid d_i \geq -1\}$.

³¹Kultti & Salonen (2005) also propose allowing for probabilistic demands, but do not require balance between expected demand and expected supply.

³²Under rationing-by-voter, each voter with outstanding demand is considered in random order, and the order is satisfied only if it can be satisfied in full. Under rationing-by-vote, each vote offered for sale is randomly assigned to any voter with unsatisfied demand.

³³If the two highest values are very close, the highest value voter as well mixes between selling and demanding a majority. The range of values for which this occurs is always very small and disappears as N becomes large.

of votes as economic goods exchanged on a market is thus misleading, and the equilibrium forces us to recognize and clarify the confusion. Second, in any equilibrium with trade, there is always one individual holding a majority of votes: the market delivers dictatorship, a single voter controlling the decision. The market does redistribute votes towards those who value them most, but the redistribution is extreme. As a result, the welfare properties of the market are less promising. In particular, when the number of voters is large, and individual values are uncorrelated with the direction of preferences, as they are in this model, the following result must hold. Call W_N^{MR} ex ante expected utility under majority rule, and W_N^{VM} ex ante expected utility with the vote market, both evaluated before individual valuations are realized. Denote by $F_N(x)$ the joint probability distribution of the vector of valuations.³⁴ Then:

Result (Welfare in a market for votes). *Consider a sequence of vote markets, indexed by the size of the electorate N . For any sequence of distribution functions $\{F_N\}$, there exists a finite \bar{N} such that if $N > \bar{N}$, $W_N^{MR} > W_N^{VM}$.*

Under the assumptions of the model, if the number of voters is large enough, a market for votes is counter-productive.

But how robust is the conclusion? One particular concern is the strong symmetry of the scenario: any voter is equally likely to favor either alternative. Potential welfare gains from a vote market are usually discussed when there is a well-identified minority which majority voting consigns to the losing side of the decision, the opposite M and m groups discussed earlier, or in the model of Philipson & Snyder (1996). Does ex ante equilibrium exist in such a case? Casella & Turban (2014) find that the answer is positive. The article characterizes an ex ante equilibrium where the highest value voter on each of the two sides randomizes between selling her own vote and demanding a majority; all others sell their vote. The model has a number of unexpected implications. First, to counter selling by “weak” allies, in equilibrium votes are actively purchased from voters on the buyer’s own side. Second, in general, market exchanges result in supermajorities, even though the equilibrium price is positive. Third, belonging to the minority makes the active buyer more aggressive than the buyer belonging to the majority:

³⁴Valuations need not be i.i.d. When they are not, the result requires a strictly positive lower bound of the valuations’ support.

the probability of demanding votes can be higher at lower intensity of preferences. As long as there is trade, the side controlling the final outcome is simply the side whose highest value individual sees her demand satisfied, regardless of the relative size of the two groups. The relatively aggressive demand of the highest value minority voter translates into a relatively high frequency of minority victories. The result can be made precise. Call $\theta_N(m)$ the ex ante expected frequency of minority victories with market trading, and $\theta_N^*(m) = \mathbb{P}(\sum_{i \in m} x_i > \sum_{i \in M} x_i)$ the corresponding frequency under ex ante efficiency. Assume that intensities x_i are independently drawn from the same distribution F . Then

Result (Minority victories in a market for votes). *For any N , for any $M \geq m > 0$, and any F , $\theta_N(m) > \theta_N^*(m)$: the minority is expected to win strictly more frequently than is ex ante efficient*

The result becomes even stronger when the market is large. Suppose $m = \lfloor \alpha N \rfloor$ for any N , with $\alpha \in (0, 1/2)$, and consider a sequence of markets indexed by N . Then $\lim_{N \rightarrow \infty} \theta_N(m) = 1/2$ for any α and any F : in the limit, the minority is expected to win half of the time, *regardless of its share of the electorate*.³⁵ Not too surprisingly, the outcome is undesirable: the superiority of simple majority voting to the market of votes, stated earlier for the symmetric case, is confirmed in the presence of a known minority.

The notion of ex ante equilibrium in a market for votes finds experimental support both in Casella et al. (2012) and in Casella et al. (2014). The experiment described in the latter paper studies a minority of size 2, and a majority of size 3, trading votes either via a decentralized market or through one-to-one bargaining by the two group leaders.³⁶ When the leaders bargain, the majority leader “owns” the decision, and Myerson & Satterthwaite (1983)’s prediction applies: the theory predicts too little trade, or too few minority victories, relative to the first best, but welfare gains relative to majority rule. When trade is decentralized, on the other hand, ex ante equilibrium predicts too many minority victories and, with the experimental parameters, an efficiency loss relative to majority voting. All predictions are strongly supported by the experimental data.

³⁵In fact, the result is stronger: the expected frequency of minority victories converges to $1/2$ for all value realizations, except on a set with zero probability ($\theta_N(m) \xrightarrow{a.s.} 1/2$).

³⁶Contrary to Koford (1982), where leaders cooperate, here leaders bargain non-cooperatively.

3.3 Strategic market games

In an *ex ante* competitive equilibrium, agents' net demand depends not only on the price but also on others' demands. The approach is used in the analysis of competitive equilibria with externalities (Arrow & Hahn, 1971, pp.132-36) but moves away from the traditional competitive model, where the price contains all relevant information, and introduces aspects of a game. Xeferis & Ziros (2017, 2018) follow this alternative route and study vote markets as strategic market games (Shapley & Shubik, 1977): traders do not take the price as given, but strategically approach one another with offers to trade for money.

A second important assumption addresses the problem of equilibrium existence: the political system is described as “power-sharing” in the sense that the outcome of a vote between two options depends linearly on the number of votes received by each. More precisely, if preferences over a proposal are described by z , and the proposal receives a fraction $\alpha \in (0, 1)$ of the votes, then voter i receiving a transfer t_i derives utility $u_i = t_i + \alpha z_i$. Technically, the assumption removes the discontinuity of the outcome at $\alpha = 1/2$, and thus the discontinuity in the value of a vote at pivotality.

Under either complete (Xeferis & Ziros, 2018) or incomplete (Xeferis & Ziros, 2017) information about preferences, an equilibrium exists where all voters are trading - on each side, voters with lower valuations sell their votes while voters with higher valuation buy votes. The welfare properties of this full-trade equilibrium depend on the informational environment. Under complete information, the results are surprisingly similar to Casella & Turban (2014). As in that model, in equilibrium there are only two buyers - the voter with the highest valuation on each side. The linearity of the outcome function translates in equilibrium vote shares that are asymptotically proportional to these two buyers valuations. As a consequence, welfare can be significantly far from the utilitarian optimum, and may be inferior to the situation without vote trading.

Xeferis & Ziros (2017) consider instead a model of incomplete information, where preferences (z_i) are drawn independently from the same distribution.³⁷ In stark contrast to Casella et al. (2012), they obtain a strong and positive welfare result:

³⁷The model further assumes that all voters interested in buying have to bid the same amount of the numeraire. The implication is that all buyers buy the same amount of votes.

Result (Efficiency of trade in a power-sharing system). *Under any symmetric equilibrium of the strategic market game with power sharing, all voters are weakly better off than under no trade.*

The finding reflects both the linearity of the outcome function and the symmetry of the model: from the point of view of any voter, not trading is always an available option and delivers in expectation the same utility as when trade is forbidden, because any favorable trade is as likely as any equivalent unfavorable trade. Thus in expectation trading exerts no externality on non-traders, and this is why trade only leads to Pareto-improvements, exactly as in a traditional exchange economy. The symmetry assumption is shared by Casella et al. (2012), but the importance of pivotality in this latter model leads to the opposite result.

Tsakas et al. (2020) test these predictions in the lab, both the positive welfare results under power-sharing systems and, in the same strategic market game setting, the negative results under majority rule. Both predictions are confirmed by the experimental data. However, the observed welfare gains under power-sharing are not as universal as the theory suggests, and some voters experience welfare losses in the long run.

3.4 Decentralized bargaining

Another possible and highly intuitive model for the study of vote trades studies agreements that result from decentralized bargaining. The model adopts the focus on strategic interactions of market games, but abandons the attention given to general equilibrium effects. A noticeable example of this approach is Iaryczower & Oliveros (2016). The article analyzes a complete information model of sequential, decentralized bargaining with three farsighted legislators voting on a binary decision. At each period, a pair of voters is randomly matched, allowing one voter (the proposer) to make a buy or sell offer of votes to the second voter.

The Markov perfect equilibria of the game reveal unexpected qualitative features. First, there are situations in which some legislators act as intermediaries, buying votes from one legislator only to resell them to a third one, even when the first and the third legislators would not trade directly. Second, the status of these political intermediaries is different from those of intermediaries in traditional exchange economies: a political middleman must have a stake in the final decision, and rather than extracting surplus

from both parties he interacts with, he is actually transferring some of his own surplus. The monetary expense the middleman incurs when acquiring votes is compensated in part by the monetary receipts when reselling them, but also, crucially, by the policy benefit he derives from passing the power to enact the policy to the third legislator.

3.5 Vote-buying mechanisms: Quadratic voting

Finally, another perspective, closer to mechanism design, is brought by a literature on vote-buying mechanisms.³⁸ The idea is that votes are not initially held, but rather acquired by the voters from the election organizer at a cost, before the voting stage. By shaping the cost function appropriately, the designer can influence how preferences are aggregated.

The most common and most promising cost function is quadratic: a voter i can acquire an amount $v_i \in \mathbb{R}_+$ at cost $c(v_i) = \kappa(v_i)^2$. Voters then vote on a binary decision, and each receives a rebate corresponding of the average of others' payments. The purchase decision trades off the benefit and the cost of acquiring votes. As long as the expected probability of shifting the decision by acquiring v_i votes is linearly increasing, at rate of *marginal pivotality* p , a voter maximizes $pv_i x_i - c(v_i)$, and with quadratic costs the first order condition yields $v_i = (\frac{p}{2\kappa})x_i$. Votes are then proportional to preference intensities, and one immediately obtains the following sufficient condition:

Result (Utilitarian efficiency of quadratic voting). *If the marginal pivotality p is constant and identical across voters and if the cost function is quadratic, the collective decision maximizes utilitarian welfare.*

The bulk of the work in these models consists in describing conditions under which the assumption of constant marginal pivotality is satisfied, or, if not, in showing that the exceptions are rare enough not to undo the welfare results. In Goeree & Zhang (2017), preferences z_i are independently drawn from a standard normal distribution. The authors show that in a large electorate there exists an approximate equilibrium such that everyone votes proportionally to one's own preference intensity, and marginal pivotality is approximately constant. To understand the result, suppose that all voters

³⁸See Tideman & Plassmann (2017) for an overview of the link between vote-buying mechanisms and the larger mechanism design literature. A first mechanism with a quadratic cost function, designed for the funding of public goods, is proposed by Hylland & Zeckhauser (1979).

but i buy an amount of votes proportional to their value, $v_j = \alpha_N x_j$, for some constant α_N .³⁹ Then, voter i chooses v_i so as to maximize his utility $u_i(v_i)$, which can be decomposed into three terms: the expected benefit from the collective decision, the cost of acquiring votes, and the rebate:

$$u_i(v_i) = -x_i \mathbf{1}_{\{z_i < 0\}} + x_i \mathbb{P}(v_i + \alpha_N \sum_{j \neq i} z_j > 0) - c(v_i) + \frac{1}{N-1} \sum_{j \neq i} c(v_j).$$

Since the sum of i.i.d. normally distributed variables also follows a Gaussian distribution, we obtain, denoting by Φ the c.d.f. of a standard normal distribution:

$$\mathbb{P}(v_i + \alpha_N \sum_{j \neq i} z_j > 0) = \mathbb{P}\left(\frac{\sum_{j \neq i} (-z_j)}{\sqrt{N-1}} < \frac{v_i}{\alpha_N \sqrt{N-1}}\right) = \Phi\left(\frac{v_i}{\alpha_N \sqrt{N-1}}\right).$$

As $\Phi\left(\frac{v_i}{\alpha_N \sqrt{N-1}}\right) \approx_{N \rightarrow \infty} \frac{1}{2} + \frac{v_i}{\alpha_N \sqrt{2\pi(N-1)}}$, it follows that

$$u_i(v_i) \approx \kappa' + \left(\frac{1}{\alpha_N \sqrt{2\pi(N-1)}}\right) v_i x_i - c(v_i)$$

for some constant κ' ,⁴⁰ so that the assumption of constant marginal pivotality is approximately satisfied. The first order condition yields $v_i = \left(\frac{1}{2\kappa\alpha_N\sqrt{2\pi(N-1)}}\right)x_i$. Thus, setting $\alpha_N = \left(\frac{1}{2\kappa\sqrt{2\pi(N-1)}}\right)^{\frac{1}{2}}$, the result is an approximate Bayes-Nash equilibrium for large electorates where voters vote in proportion to their values.⁴¹ Utilitarian efficiency then follows immediately, because the side with higher total value will hold more votes.

Beyond utilitarian efficiency, it is notable that, if the cost function scaling parameter κ does not depend on N , the number of votes bought by any voter vanishes as the electorate gets large.

Lalley & Weyl (2019) consider more general models, where preferences are independently drawn from an arbitrary bounded distribution, whose mean

³⁹We sketch here the argument of Goeree & Zhang (2017) for an arbitrary normalization of the cost function (the parameter κ), and thus introduce the proportionality parameter α_N .

⁴⁰We have $\kappa' = -x_i \mathbf{1}_{\{z_i < 0\}} + \frac{1}{N-1} \sum_{j \neq i} c(v_j) + \frac{1}{2}$.

⁴¹With $\alpha_N = \left(\frac{1}{2\kappa\sqrt{2\pi(N-1)}}\right)^{\frac{1}{2}}$, we obtain $\alpha_N = \frac{1}{2\kappa\alpha_N\sqrt{2\pi(N-1)}}$. Note that the approximation of the function Φ is valid whenever $\kappa \ll \sqrt{N}$ as this implies $\alpha_N \sqrt{N-1} \rightarrow_{N \rightarrow \infty} \infty$.

may be null or not. This latter case is more complicated as some voters with extreme preferences do not vote, even approximately, in proportion to their intensities. Yet, such voters must be sufficiently rare not to compromise the mechanism’s asymptotic efficiency.⁴² Eguia & Xefteris (2018) make the broader point that any welfare function of the form $\sum_{i \in N} \text{sign}(z_i) |z_i|^\rho$ can be implemented by a family of cost-functions including $c(v_i) = (v_i)^{1+1/\rho}$.

In a subsequent paper, Eguia et al. (2019) generalize the model further, by considering an election over multiple alternatives (as opposed to a Yes/No decision). Voters purchase votes at a quadratic cost, that they can use to vote in favor or against any alternative. In a setting of complete information, the optimality result remains.

Goeree & Zhang (2017) bring the mechanism to the lab and confirm that it leads to significant welfare gains relative to majority voting, mostly for voters with moderate intensities, benefiting from redistribution. Interestingly, moderates bid too high and extremists bid too low, suggesting that participants pay more attention to the cost of bidding than to the gain from pivotality. The experiment includes an endogenous choice of institution once subjects have gained enough experience: 90% of the groups prefer the vote-buying mechanism to majority rule.

Yet, the properties of quadratic voting in a real-world setting seem difficult to ascertain with precision, in particular for large elections. First, even if voters perceive their marginal pivotality to be constant, their estimates may well vary significantly across voters, which would threaten the utilitarian efficiency of the collective decision. Second, the quadratic cost creates transparent incentives for collusion in vote-buying, which may hinder the public trust in the system.⁴³ Third, calibrating the cost function correctly is not trivial. If the cost is high (for instance, if κ is constant and not too small), the equilibrium purchases of votes will be very low in a large population, so that actual purchases might be driven by considerations other than pivotality (Kaplow & Kominers, 2017). If the cost is low, (for instance, if κ is decreasing as $1/\sqrt{N}$), some voters will spend significant amounts of

⁴²These results are complemented by Chandar & Weyl (2019), who study exact quadratic vote-buying in finite populations, where optimality need not hold. The mechanism generally dominates standard majority voting, and when it does not, welfare losses remain limited.

⁴³Weyl (2017) establishes a number of precise and sophisticated conjectures on the robustness of quadratic voting to collusion, and concludes that the mechanism is less fragile than the VCG mechanism.

numeraire to purchase votes, which threatens the democratic legitimacy of the mechanism, and even its utilitarian efficiency once wealth inequalities are accounted for (Laurence & Sher, 2017).

Vote-buying mechanisms are reminiscent of costly voting models, where voters incur a cost, possibly individualized, to participate in an election between two alternatives. The models have a long tradition and a straightforward intuition: suppose that each voter draws, independently and from the same distribution, a random cost linked to the act of voting. Participation will then be higher among people who attach higher value to the election: the redistribution of votes towards higher intensity voters occurs automatically, and improves utilitarian efficiency. Among the earliest contributions, Palfrey & Rosenthal (1983, 1985) made precise the strategic problem of a voter facing voting costs; in the context of political competition, Ledyard (1984) argued that costly voting leads candidates to locate at the utilitarian optimal policy, rather than at the median policy. The recent literature has focused on the stochastic properties of the distributions of the voting costs, of the preference intensities, and of the number of supporters on each side as determining the welfare properties of voluntary voting. A very useful entry into the current state of the literature is Krishna & Morgan (2015) who show that, in large binary elections, the equilibrium outcomes with heterogeneous voting costs are utilitarian efficient, when elections are held under majority rule.⁴⁴

4 Storing and distributing votes

Buchanan and Tullock’s discussion of majority voting started from the observation that committees are typically involved in multiple decisions, and thus committee members have the opportunity to trade votes over different issues. In fact, once we recognize that each voter votes over many decisions, trading off votes over one issue for extra votes over another does not require interpersonal trading. If the voting rule allows it, a voter can just as well “trade votes with oneself”, or shift the voter’s own voting weight across the different decisions. The logic is similar to logrolling, and as in the case of logrolling, the allocation of votes across decisions avoids the difficulties posed by mar-

⁴⁴Among other contributions, see Borgers (2004), Krasa & Polborn (2009) and Taylor & Yildirim (2010). Levine & Palfrey (2007) test several predictions of the theory in the lab.

kets for votes—in particular the differences in wealth among voters over which the rule designer has no control. The idea is proposed, in passing, in Mueller (1973), as an alternative to markets for votes. Mueller suggests endowing each voter with 1,000 points, to be divided over multiple binary decisions. Each decision is resolved according to the majority of total points cast on the decision. Mueller posits that each voter evaluates the impact of her votes on the probability of her favorite outcome prevailing through an arbitrary positive and concave function, equal across all decisions. The importance of the precise shape of such a function leads him to conclude against the scheme.⁴⁵

More recently, the same intuitive scheme was proposed by Casella (2005) (summarized and extended in Casella (2012)) under the name of *Storable Votes*, and, independently, by Hortala-Vallve (2012). In these works, the impact of cumulating votes on the probability of winning is captured precisely by the change in the probability of being pivotal. The term Storable Votes (SV) implies a temporal dimension—votes not cast on early elections can be ‘stored’ for future use. Time is relevant in a number of possible applications, but not necessary to the logic of the scheme. We thus present it in the same set-up used in Section 2 when discussing logrolling. A committee of N voters takes K binary decisions, and voters hold separable preferences summarized by the profile $z = \{z_i^k\}$. On each issue k , $x_i^k = |z_i^k|$ denotes voter i ’s preference intensity on that issue. We depart from the model of Section 2 however by assuming explicitly that preferences are private information, and by anchoring the welfare discussion on utilitarian efficiency, and thus on the cardinal representation of preferences. The stochastic properties of individual valuations z_i^k are central to the model. In the simplest case, individual preferences z_i^k are drawn, independently across individuals and across proposals from a probability distribution $F(z_i^k)$.⁴⁶ Each proposal is decided according to the majority of votes cast, with ties resolved randomly. As before, we indicate by \mathbf{P} the set of proposals that pass. The voter’s objective is to maximize total utility $\sum_{k \in \mathbf{P}} z_i^k$.

Under majority voting, each voter has one vote to cast on each proposal. Under SV, in the most intuitive formulation each voter is again endowed with

⁴⁵In these terms, the conclusion is actually due to Laine (1977) who reevaluates Mueller’s analysis.

⁴⁶To highlight the main intuition, we present the model under these assumptions. Casella (2012) discusses results when the distributions of values vary across proposals, or when either direction of preferences alone or both direction and intensities are correlated within subgroups.

K total votes, but is free to distribute them as desired over the K proposals. In general, on any one decision different voters will cast a different number of votes. And yet, all voters are treated equally and all have the same total number of votes.

Because each decision is binary, voting in the direction of sincere preferences is always weakly dominant. The interesting question is how to distribute the votes. The number of votes cast signals the intensity of the voter's preferences on the issue, and because the expense of such votes comes at the cost of votes on other issues, there is an intuitive sense in which the voting system induces incentive compatibility. The intuition is made precise in Jackson & Sonnenschein (2007): across problems as diverse as voting, bargaining, or the allocation of indivisible objects, linking decisions can work to overcome incentive constraints and achieve full efficiency in the absence of transfers.⁴⁷

Yet, outside the delicate mechanism specified by Jackson & Sonnenschein (2007), it is difficult to translate this intuition into designing a fully efficient voting rule because the decision of where to cast votes depends not only on the relative values attributed to winning the different issues, but also, as Mueller had remarked, on the probability that the vote makes a difference, i.e. on pivotality. Thus the number of votes a voter is willing to cast depends both on the voter's priorities and on the expectation of how many other votes are cast and in which direction. Beyond restricted scenarios, SV cannot make claims to full efficiency. It is instead designed to yield welfare improvements over majority voting, while remaining simple to implement.

SV changes the outcome of a decision, relative to simple majority rule, only if the minority casts on that decision more votes than the majority does. Minority victories become possible, but typically the minority will only win decisions that are both felt intensely by its members *and* weakly by the opposite larger group, the decisions that utilitarian criteria suggest the minority should win. SV thus protects minorities with intense preferences. However,

⁴⁷The result in Jackson and Sonnenschein is of great intellectual beauty but is not meant for practical implementation. It holds in the limit, as the number of linked decisions approaches infinity, as long as the values z_i^k are independent across both individuals and decisions, and as long as the distributions $F_i(z_i^k)$ are identical across all k . The planner requires agents to declare their values under the constraint that the distribution of each agent's declared values replicates the theoretical distribution. The mechanism is part of the literature on optimal mechanisms without transfers. Classic references are McAfee (1992) and Pendorfer (2000).

when the division between majority and minority remains constant across issues, a strong case can then be made for majority rule’s lack of legitimacy, regardless of intensity: voters who are consistently on the minority side are effectively disenfranchised by majority rule. SV is also appropriate, in fact possibly particularly so, in such situations.⁴⁸

In its protection of minorities, SV calls to mind existing semi-proportional voting systems. A natural parallel is Cumulative Voting (CV), actively used in the US in local and corporate elections (Sawyer & MacRae, 1962, Bhagat & Brickley, 1984, Guinier, 1995). Under CV, voters electing a commission of S members are each endowed with S votes and are able to cumulate as many of these votes as desired on as few candidates as desired. SV differs from CV because SV applies to multiple decisions, each of which has only two possible outcomes, while CV applies to a single election choosing multiple representatives from a larger set of candidates.⁴⁹

Casella (2012) reviews numerous works on SV, both theoretical and experimental, mostly applied to small size committees. Here we summarize the system’s properties when N , the number of voters, is large, and valuations are independent across voters and proposals, drawing on Casella & Gelman (2008). In this environment, it is generally optimal for a voter to cumulate all votes on the single decision the voter considers most important.⁵⁰ It is thus sufficient for the rule designer to endow each voter with a single “bonus vote” to be cast freely. To elicit some expression of voters’ preferences beyond their first priority, it is then desirable to complement the bonus vote with K “regular” votes—to be cast one on each proposal. The weight of the

⁴⁸Note that logrolling cannot help: a voter who is on the losing side of every issue has no valuable vote to give away. Casella et al. (2017) discuss, both theoretically and experimentally, how SV can grant a systematic minority a positive chance of prevailing occasionally. In their model there are no differences in cardinal intensities ($x = 1$ always) and occasional minority victories are defended as a goal in themselves. Lackner (2020) introduces an alternative class of mechanisms where a voter’s current votes are weighted according to the voter’s success in previous decisions. The mechanisms ensure a fair distribution of influence when voters vote sincerely.

⁴⁹CV has the explicit goal of giving representation to the minority. For example, it is easy to see that a cohesive minority that constitutes more than $1/(S+1)$ of the electorate and concentrates all its votes on a single candidate can *guarantee* itself one seat. CV is quite different from SV. It would correspond to SV only if each candidate could compete for one seat only, and each open seat were reserved for the winner between two specific candidates, one from the majority and one from the minority party.

⁵⁰See Casella & Gelman (2008), section 7.1.1. As shown by Eguia & Xefteris (2018) in the case of a centralized market for vote, linear costs lead to bunching.

bonus vote, relative to a regular vote, can be set to $B > 0$, where B is a choice variable for the designer. The modified design reflects SV's pragmatic inspiration and, if B is not large, moves the voting scheme towards majority voting. Overturning the majority's choice is possible but not easy.

With large N and independent values, realized preferences on each issue will converge to the theoretical distribution F . As N increases, the probability of approval of proposal k by any random voter converges to $\int_0^z dF(z)$ with probability 1, and thus is known ex ante, if the distribution F is pinned down. As a result, also known ex ante with probability approaching 1, are the majority voting outcomes and their efficiency properties. As pointed out by the literature, a better assumption is that the probability of approval for each proposal is stochastic (Good & Mayer, 1975, Margolis, 1977, Chamberlain & Rothschild, 1981, Gelman et al., 2002). Suppose that ex ante each voter i has probability ψ_k of being in favor of proposal k ($z_i^k > 0$), and $1 - \psi_k$ of being against ($z_i^k < 0$). The probability ψ_k is distributed according to a distribution H_ψ defined over the support $[0, 1]$, and each realized ψ_k , with $k = 1, \dots, K$, is an independent draw from H_ψ . To rule out any expected bias in favor or against the proposal, H_ψ is assumed to be symmetric around $1/2$. In addition, to rule out systematic expected biases in intensities that depend either on the direction of preferences or on specific proposals, conditional on being in favor or against, the distribution of intensities ($x_i^k = |z_i^k|$) is assumed to be identical both across individuals, for given proposal, and across proposals. It is denoted by $Q(x)$.

Casella and Gelman evaluate the welfare impact of SV, relative to majority voting. They construct the measure:

$$\omega_N = \frac{W_N^{SV} - W_N^R}{W_N^{MR} - W_N^R} \quad (1)$$

where W_N^{MR} is, as before, a voter's ex ante expected utility under majority voting, W_N^R is a floor, given by expected utility under random decision making (when any proposal passes with probability $1/2$), and W_N^{SV} is ex ante expected utility under SV. They conclude that for any weight of the bonus vote B , SV outperforms majority:

Result (Storable Votes in a large electorate). For all $B > 0$, for all distributions $H_\psi(\psi)$ and $Q(x)$, and for all $K > 1$, the measure ω_N converges and $\lim_{N \rightarrow \infty} \omega_N = \omega > 1$: SV increases expected welfare relative to majority voting.

In equilibrium, voters cast their bonus vote on the proposal to which they attach the highest intensity, and thus it is towards that proposal that they shift the probability of being on the winning side. As a result, expected utility increases. The assumptions on H_ψ and Q impose symmetry on the problem, an approach well-suited to a constitutional stage, preceding different decisions. When asymmetries are allowed, the result in general depends on the shape of the distributions, and thus the question is empirical, best studied in field experiments and surveys. We return to this point below.

We have argued earlier, with Casella and Gelman, that in a large electorate SV is best designed as a single additional bonus vote added to multiple regular votes because, if all votes can be moved freely across decisions, each voter would bunch them all on a single decision. As we learned from Eguia & Xefteris (2018), one alternative is to move away from the one-to-one transformation of votes shifted across decisions, towards convex costs. Quadratic Voting (QV) can be adapted to the problem on hand. With multiple binary decisions, QV could be implemented by paying for votes in an artificial currency, “voices”, which can be translated into votes at a quadratic cost. Casting v^k votes on proposal k requires spending $(v^k)^2$ voices on k . By building on the voters’ own rankings of the different decisions, QV becomes similar to SV, but with a quadratic cost. Relative to the original QV formulation, “QV with voices” has the advantage that the voters’ valuations are elicited from trading-off different decisions, as opposed to trading-off one public decision and private consumption, and thus no money is involved (Posner & Weyl, 2015, Lalley & Weyl, 2018).

There is no rigorous analysis of the equilibrium properties of QV in multiple elections. However, a simple model shows that in large electorates the outcome could approach efficiency if voters ignore the complex impact on pivotality of casting different numbers of votes. We know from Lalley & Weyl (2019) that the marginal impact of a vote in general is not constant, but the deviations may be too subtle for voters to take into account.

Consider the following scenario.⁵¹ As at the start of this section, individual valuations z_i^k are independently drawn from the known probability distribution $F(z_i^k)$. Each voter is endowed with a budget of voices y_i equal to 1 to allocate across proposals. On each proposal, voices are transformed

⁵¹The model, more transparent than Lalley & Weyl (2018) originates in a quick sketch due to Glen Weyl. Its discussion is borrowed from the online appendix to Casella & Sanchez (2019).

into a number of votes equal to the square root of the dedicated voices: if v_i^k denotes the votes cast on proposal k by voter i , and y_i^k the corresponding voices, then $v_i^k = \sqrt{y_i^k}$, or $\sum_{k=1}^K (v_i^k)^2 = \sum_{k=1}^K y_i^k = 1$. Both voices and votes are fully divisible, and voters vote sincerely. Each voter i faces the constrained maximization problem:

$$\max_{v_i} (-x_i^k \mathbb{1}_{\{z_i^k < 0\}} + \sum_{k=1}^K p_i^k(v_i^k) x_i^k) \quad \text{subject to} \quad \sum_{k=1}^K (v_i^k)^2 = 1,$$

where $p_i^k(v_i^k)$ is the probability that proposal k is decided as i prefers when i casts v_i^k votes. Recall that $x = |z|$ denotes absolute values, or intensities.

Assume:

$$\frac{\partial p_i^k(v_i^k)}{\partial v_i^k} = c_i \tag{2}$$

In words, the marginal impact of each vote is constant, across proposals and for any number of votes cast. Then for each proposal k , the first order condition yields: $v_i^k = \frac{c_i x_i^k}{2\lambda_i}$ where λ_i is the Lagrange multiplier linked to the budget constraint. Substituting the budget constraint $\sum_{k=1}^K (v_i^k)^2 = 1$, one obtains $\frac{c_i}{2\lambda_i} = \sqrt{\frac{1}{\sum_{k=1}^K (x_i^k)^2}}$ and thus:

$$v_i^k = \frac{1}{\sqrt{\sum_{k=1}^K (x_i^k)^2}} x_i^k. \tag{3}$$

The optimal number of votes cast on each proposal equals the voter's intensity, normalized by the Euclidean norm of the voter's intensities across all proposals. If such norms are similar across voters or if each individual's norm is used to normalize cardinal values in the welfare criterion, then utilitarian efficiency follows immediately. The number of votes cast in each proposal is proportional to the voter's intensity, and thus each proposal is won by the side with higher total intensities.

The assumption that voices and votes are divisible eliminates uninteresting integer constraints and recommends endowing voters with a large number of voices; the requirement that intensity norms be equal across voters suggests asking voters to distribute their voices over a large number of issues. In practical implementations, both ingredients could be confusing for voters. But the more substantive approximation is equation 2, the assumption of constant marginal impact of additional votes, which we know from Lalley &

Weyl (2019) will not hold in general even asymptotically at large N . Yet, here practical considerations may help: what matters are the voters' beliefs, and if voters miss the subtleties of pivotality and behave *as if* marginal effects were constant, then the desirable properties of QV may well be realized. In short:

Result (QV over multiple issues). *If voices and votes are divisible, K is large, and each voter i believes that $\partial p_i^k(v_i^k)/\partial v_i^k = c_i$ for all k and all v_i^k , then QV with voices yields utilitarian efficiency.*

Note that, even QV with voices provides opportunities for collusion. Consider two voters having the same preferred positions on two issues, but different ranking of intensities: the first voter finds the first issue more important and wishes to spend 4 voices on it but 0 voices on issue 2, while the opposite is true for the second voter. Then, the pair can profitably exchange by having each voter spending 1 voice per issue: the outcome is the same, but they both spare 2 voices to be spent on other issues. We see that the logic is very similar to logrolling, but, to our knowledge, the precise analysis of this interaction remains open.

At the end, beyond the theoretical exercises, the answers to the questions we care about are empirical: What shapes do the distributions of valuations take? How do voters think about pivotality? Are voters able to use effectively a voting rule like QV with voices? To begin addressing these questions outside the controlled environment of the lab, Casella & Sanchez (2019) ran an incentivized survey on Mechanical Turk, asking approximately 600 California residents to vote, using either SV or a simplified version of QV, on four propositions that were being prepared for popular vote in the 2016 California ballot.⁵² The first part of the survey elicited participants' direction of preferences and asked them to distribute 100 "importance points" among the four propositions. The points allocated could then be combined into a measure of utilitarian welfare reflecting voting outcomes.

⁵²In the SV sample, subjects were granted one extra vote, in addition to four regular votes (or $B = 1$). In the QV sample, subjects were asked to choose one of four different classes of votes: four blue votes, each of weight 1, three green vote, each of weight 1.2; two yellow votes, each of weight 1.5; or a single red vote of weight 2. Votes could not be cumulated, and thus a person choosing green/yellow/red votes could cast votes on only three/two/one propositions. Note that the simple four-class classification respects the convex cost of concentrating votes: the decline in total voting weight is increasing with concentration, and increasing at an increasing rate.

Because the two voting systems rely on the single budget constraint tying the different voting decisions, the survey data correspond to two data points only, one for SV and one for QV. To estimate the probability of minority victories, as well as the efficiency properties of the two schemes, the authors bootstrapped the data, generating two samples of 10,000 joint elections each (i.e. 10,000 instances of voting on all four propositions), one for SV and one for QV.

Both systems perform well. Both result in non-negligible minority victories: depending on details of the bootstrapping procedure, between 30% and 40% of simulations in both samples have at least one proposition won by the minority side. And while making minority victories possible, both systems consistently raise the measure of total welfare, appropriating at least 80% and as much as 99% of the available surplus. QV appears more sensitive to minority preferences than SV, inducing more minority victories and appropriating a larger share of the surplus. However, when the scope for improvement over majority voting is small, such sensitivity has a cost and induces frequent inefficient minority victories. SV is simpler, a strong argument in its favor, but its performance is less consistent across plausible subjects' behaviors.

5 Concluding remarks

Sixty years after Buchanan & Tullock (1962)'s seminal analysis, the optimism that characterizes their work on vote markets is difficult to share. Recent works have overcome technical difficulties: even though traditional equilibrium concepts generally fail in vote markets, researchers have found that these markets may generate stable outcomes – stable vote allocations under logrolling and ex ante competitive equilibrium, or strategic market game equilibrium when votes are exchanged against a numeraire. But the technical advances have helped to show that final outcomes may be far from desirable – the Condorcet winner need not be reached via logrolling, and the utilitarian efficiency of vote markets is generally inferior to simple majority rule.

Nevertheless, theoretical analyses of vote markets have been useful in advancing a research agenda aiming at improving collective decision rules. When one public decision is at stake, quadratic voting – making voters pay for their votes at a quadratic cost – can theoretically achieve utilitarian efficiency. Yet, the theoretical promise relies most blatantly (but not exclusively) on the assumption that inequalities in resources are not binding and evokes

what we believe is justified repugnance associated to trade-offs between political influence and private wealth.⁵³ We thus view as most encouraging the mechanisms that only allow voters to trade with themselves, when several decisions are to be voted upon. Both storable votes and a modified form of quadratic voting fulfill this role, while varying in the precise details of their implementation.

We highlight a few promising directions for further research. First, all the analyses we discuss rely on the assumption of private values. There is however a long tradition considering collective decisions as the aggregation of private information in a setting of common values, dating back to Condorcet (1785). This perspective opens new questions on vote trading: would trade channel votes to better informed individuals, thereby improving information aggregation?⁵⁴ How would the previous analyses evolve if voters' preferences have both private and common values components?⁵⁵

Second, most of the welfare results we presented refer either to the Condorcet winner or to utilitarian welfare. The welfare criterion must be linked to the broader objectives of the researcher, and more generally, of society. For example, if protection of the minority is a goal, identifying and enacting the Condorcet winner is the wrong normative criterion: when voters are divided into two groups of different sizes with systematically opposed preferences, the Condorcet winner simply selects the majority-preferred alternative on each issue. Utilitarian welfare can give some voice to minority interests, by weighting preference intensities or when utilities over multiple decisions are non-separable (Koriyama et al., 2013). Nevertheless, if the explicit goal is to grant groups more balanced chances of prevailing in policy choices when multiple decisions are at stake, a normative criterion derived directly from fairness objectives becomes the appropriate reference point. Such concerns are at the heart of recent works by computer scientists, on multi-winner elections (Aziz et al., 2017) and participatory budgeting (Freeman et al., 2019)

⁵³Walzer (1983)'s fundamental reminder that domains of influence should remain separate is very relevant here.

⁵⁴Focusing on corporate finance, a setting where the common values assumption fits well, Christoffersen et al. (2007) argues that vote trading does occur (with prices close to 0) and achieves vote reallocation from less informed to more informed traders. Esö et al. (2015) propose a first model where the information aggregation argument is rigorously identified.

⁵⁵In an unpublished manuscript, Piketty (1994) highlight the fragility of information aggregation under vote trading in the presence of extremist private value voters.

in particular. We believe that they will also be relevant for the analysis of vote trading.

Finally, beyond theoretical analyses and lab experiments, further work is needed to understand better how the mechanisms we favor might work in real-world settings, and in which environments they should be preferred to standard majority voting. The Mechanical Turk study on storable votes and quadratic voting with voices finds that both schemes outperform majority voting. Yet one important question is left open. In those data, majority voting delivers good results as well, appropriating between 72% and 95% of surplus, depending on the simulations. Do the welfare gains produced by either of the new schemes justify abandoning majority rule—a transparent, legitimate system voters understand well? It is a question in political philosophy and well beyond the scope of this survey, but a question that reformers of voting systems will need to address.

Disclosure statement

The authors are not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

Acknowledgments

We are grateful to the institutions that, throughout the years, have supported the works discussed in this survey, in particular the National Science Foundation, through its grants SES-00214013 and SES-0617934, the Agence Nationale de la Recherche, through its grants ANR-17-CE26-0003 and ANR-17-EURE-001, the Guggenheim Foundation, the Institute for Advanced Studies, the Straus Institute at the NYU Law School, and finally the Alliance Program at Columbia and Université Paris 1 that made our collaboration possible. We are also very grateful to the colleagues, friends, and students who have worked with us, in particular Tom Palfrey, whose collaboration shaped both specific papers and broad ideas, Jean François Laslier, who brought us together and whose teachings and moral compass helped to keep us straight, and Rafael Treibich for fruitful discussions and in particular comments on this survey which opened new perspectives.

References

- Alonso R, Câmara O. 2016. Persuading voters. *American Economic Review* 106:3590–3605
- Arrow K, Hahn F. 1971. General competitive analysis. San Francisco: Holder Day
- Arrow KJ. 1951. Social choice and individual values. New York: Wiley
- Aziz H, Brill M, Conitzer V, Elkind E, Freeman R, Walsh T. 2017. Justified representation in approval-based committee voting. *Social Choice and Welfare* 48:461–485
- Bebchuk LA, Kamar E. 2009. Bundling and entrenchment. *Harvard Law Review* 123:1549
- Bentley AF. 1908. The process of government: A study of social pressures. Chicago: University of Chicago Press
- Bernholz P. 1973. Logrolling, Arrow paradox and cyclical majorities. *Public Choice* 15:87–95
- Bernholz P. 1974. Logrolling, Arrow paradox and decision rules – a generalization. *Kyklos* 27:49–62
- Bhagat S, Brickley JA. 1984. Cumulative voting: The value of minority shareholder voting rights. *The Journal of Law and Economics* 27:339–365
- Borgers T. 2004. Costly voting. *American Economic Review* 94:57–66
- Buchanan JM, Tullock G. 1962. The calculus of consent. University of Michigan press Ann Arbor
- Câmara O, Eguia JX. 2017. Slicing and bundling. *The Journal of Politics* 79:1460–1464
- Carrubba CJ, Volden C. 2000. Coalitional politics and logrolling in legislative institutions. *American Journal of Political Science* 44:261–277
- Casella A. 2005. Storable votes. *Games and Economic Behavior* 51:391–419

- Casella A. 2012. Storable votes: protecting the minority voice. Oxford University Press
- Casella A, Gelman A. 2008. A simple scheme to improve the efficiency of referenda. *Journal of Public Economics* 92:2240–2261
- Casella A, Laslier JF, Macé A. 2017. Democracy for polarized committees: The tale of Blotto’s lieutenants. *Games and Economic Behavior* 106:239–259
- Casella A, Llorente-Saguer A, Palfrey TR. 2012. Competitive equilibrium in markets for votes. *Journal of Political Economy* 120:593–658
- Casella A, Palfrey T. 2018a. Trading votes for votes: A dynamic theory. *Social Science Working Paper No. 1444, California Institute of Technology*
- Casella A, Palfrey T. 2019. Trading votes for votes. a dynamic theory. *Econometrica* 87:631–652
- Casella A, Palfrey T, Turban S. 2014. Vote trading with and without party leaders. *Journal of Public Economics* 112:115–128
- Casella A, Palfrey TR. 2018b. An experimental study of vote trading. *Social Science Working Paper No. 1445, California Institute of Technology*
- Casella A, Sanchez L. 2019. Storable votes and quadratic voting. an experiment on four California propositions. *NBER WP. No 25510*
- Casella A, Turban S. 2014. Democracy undone. systematic minority advantage in competitive vote markets. *Games and Economic Behavior* 88:47–70
- Chamberlain G, Rothschild M. 1981. A note on the probability of casting a decisive vote. *Journal of Economic Theory* 25:152–162
- Chandar B, Weyl EG. 2019. Quadratic voting in finite populations. *Available at SSRN 2571026*
- Christoffersen SE, Geczy CC, Musto DK, Reed AV. 2007. Vote trading and information aggregation. *The Journal of Finance* 62:2897–2929
- Chwe MSY. 1994. Farsighted coalitional stability. *Journal of Economic Theory* 63:299–325

- Cohen L, Malloy CJ. 2014. Friends in high places. *American Economic Journal: Economic Policy* 6:63–91
- Coleman JS. 1966. The possibility of a social welfare function. *The American Economic Review* 56:1105–1122
- Coleman JS. 1967. The possibility of a social welfare function: Reply. *The American Economic Review* 57:1311–1317
- Condorcet JANdC. 1785. Essai sur l’application de l’analyse à la probabilité des décisions rendues à la pluralité des voix. Paris, l’Imprimerie Royale
- Cox JD, Ferri F, Honigsberg C, Thomas RS. 2015. Quieting the shareholders’ voice: Empirical evidence of pervasive bundling in proxy solicitations. *Southern California Law Review* 89:1179–1242
- Dal Bo E. 2007. Bribing voters. *American Journal of Political Science* 51:789–803
- Dekel E, Jackson MO, Wolinsky A. 2008. Vote buying: General elections. *Journal of Political Economy* 116:351–380
- Dekel E, Jackson MO, Wolinsky A. 2009. Vote buying: Legislatures and lobbying. *Quantitative Journal of Political Science* 4:103–128
- Dutta B, Vohra R. 2017. Rational expectations and farsighted stability. *Theoretical Economics* 12:1191–1227
- Eguia JX, Immorlica N, Ligett K, Weyl G, Xefteris D. 2019. Quadratic voting with multiple alternatives. *Michigan State U. Dept. of Economics Working paper 2019-01*
- Eguia JX, Xefteris D. 2018. Implementation by vote-buying mechanisms. *Available at SSRN 3138614*
- Esö P, Hansen S, White L. 2015. A theory of vote-trading and information aggregation. Tech. rep.
- Faliszewski P, Rothe J. 2016. Control and bribery in voting. In *Handbook of Computational Social Choice*, eds. F Brandt, V Conitzer, U Endriss, J Lang, AD Procaccia, chap. 7. Cambridge University Press, 146–168

- Ferejohn JA. 1974. Sour notes on the theory of vote trading. *Social Science Working Paper No 41, California Institute of Technology*
- Fischbacher U, Schudy S. 2014. Reciprocity and resistance to comprehensive reform. *Public Choice* 160:411–428
- Fischbacher U, Schudy S. 2020. Agenda control and reciprocity in sequential voting decisions. *Economic Inquiry*
- Freeman R, Pennock DM, Peters D, Wortman Vaughan J. 2019. Truthful aggregation of budget proposals, In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pp. 751–752
- Gelman A, Katz JN, Tuerlinckx F. 2002. The mathematics and statistics of voting power. *Statistical Science* 17:420–435
- Goeree JK, Zhang J. 2017. One man, one bid. *Games and Economic Behavior* 101:151–171
- Good IJ, Mayer LS. 1975. Estimating the efficacy of a vote. *Behavioral Science* 20:25–33
- Grandi U, Grossi D, Turrini P. 2019. Negotiable votes. *Journal of Artificial Intelligence Research* 64:895–929
- Groseclose T, Snyder Jr JM. 1996. Buying supermajorities. *American Political Science Review* 90:303–315
- Guerrero OA, Matter U. 2016. Revealing the anatomy of vote trading. *Saïd Business School WP 26*
- Guinier L. 1995. Tyranny of the majority: Fundamental fairness in representative democracy. Free Press
- Haefele ET. 1971. A utility theory of representative government. *The American Economic Review* 61:350–367
- Hortala-Vallve R. 2011. Generous legislators? A description of vote trading agreements. *Quarterly Journal of Political Science* 6:179–196
- Hortala-Vallve R. 2012. Qualitative voting. *Journal of Theoretical Politics* 24:526–554

- Hylland A, Zeckhauser R. 1979. A mechanism for selecting public goods when preferences must be elicited. *Working paper, Kennedy School of Government, Harvard University*
- Iaryczower M, Oliveros S. 2016. Power brokers: Middlemen in legislative bargaining. *Journal of Economic Theory* 162:209–236
- Jackson MO, Sonnenschein HF. 2007. Overcoming incentive constraints by linking decisions. *Econometrica* 75:241–257
- Kadane JB. 1972. On division of the question. *Public Choice* 13:47–54
- Kaplow L, Kominers SD. 2017. Who will vote quadratically? Voter turnout and votes cast under quadratic voting. *Public Choice* 172:125–149
- Koehler DH. 1975. Vote trading and the voting paradox: A proof of logical equivalence. *The American Political Science Review* 69:954–960
- Koford KJ. 1982. Centralized vote-trading. *Public Choice* 39:245–268
- Koriyama Y, Laslier JF, Macé A, Treibich R. 2013. Optimal apportionment. *Journal of Political Economy* 121:584–608
- Krasa S, Polborn MK. 2009. Is mandatory voting better than voluntary voting? *Games and Economic Behavior* 66:275–291
- Krishna V, Morgan J. 2015. Majority rule and utilitarian welfare. *American Economic Journal: Microeconomics* 7:339–75
- Kultti K, Salonen H. 2005. Market for votes. *Homo Oeconomicus* 23:323–332
- Lackner M. 2020. Perpetual voting: Fairness in long-term decision making. *AAAI2020*
- Laine CR. 1977. Strategy in point voting: A note. *The Quarterly Journal of Economics* 91:505–507
- Lalley S, Weyl EG. 2019. Nash equilibria for quadratic voting. *Available at SSRN 2488763*
- Lalley SP, Weyl EG. 2018. Quadratic voting: How mechanism design can radicalize democracy, In *AEA Papers and Proceedings*, vol. 108, pp. 33–37

- Lang J, Xia L. 2016. Voting in combinatorial domains. In *Handbook of Computational Social Choice*, eds. F Brandt, V Conitzer, U Endriss, J Lang, AD Procaccia, chap. 9. Cambridge University Press, 197–222
- Laurence B, Sher I. 2017. Ethical considerations on quadratic voting. *Public Choice* 172:195–222
- Ledyard JO. 1984. The pure theory of large two-candidate elections. *Public Choice* 44:7–41
- Levine DK, Palfrey TR. 2007. The paradox of voter participation? A laboratory study. *American Political Science Review* 101:143–158
- Margolis H. 1977. Probability of a tie election. *Public Choice* 31:135–138
- Mauleon A, Vannetelbosch VJ, Vergote W. 2011. von Neumann–Morgenstern farsightedly stable sets in two-sided matching. *Theoretical Economics* 6:499–521
- McAfee RP. 1992. Amicable divorce: Dissolving a partnership with simple mechanisms. *Journal of Economic Theory* 56:266–293
- McKelvey RD, Ordeshook PC. 1980. Vote trading: An experimental study. *Public Choice* 35:151–184
- Miller NR. 1977a. Graph-theoretical approaches to the theory of voting. *American Journal of Political Science* 21:769–803
- Miller NR. 1977b. Logrolling, vote trading, and the paradox of voting: A game-theoretical overview. *Public Choice* 30:51–75
- Mueller DC. 1967. The possibility of a social welfare function: Comment. *The American Economic Review* 57:1304–1311
- Mueller DC. 1973. Constitutional democracy and social welfare. *The Quarterly Journal of Economics* 87:60–80
- Mueller DC, Philpotts GC, Vanek J. 1972. The social gains from exchanging votes: A simulation approach. *Public Choice* 13:55–79
- Myerson RB. 1993. Incentives to cultivate favored minorities under alternative electoral systems. *American Political Science Review* 87:856–869

- Myerson RB, Satterthwaite MA. 1983. Efficient mechanisms for bilateral trading. *Journal of Economic Theory* 29:265–281
- Palfrey TR, Rosenthal H. 1983. A strategic calculus of voting. *Public Choice* 41:7–53
- Palfrey TR, Rosenthal H. 1985. Voter participation and strategic uncertainty. *The American Political Science Review* 79:62–78
- Park RE. 1967. The possibility of a social welfare function: Comment. *The American Economic Review* 57:1300–1304
- Pesendorfer M. 2000. A study of collusion in first-price auctions. *The Review of Economic Studies* 67:381–411
- Philipson TJ, Snyder JM. 1996. Equilibrium and efficiency in an organized vote market. *Public Choice* 89:245–265
- Piketty T. 1994. Information aggregation through voting and vote-trading. *unpublished, available at: <http://www.jourdan.ens.fr/piketty/fichiers/public/Piketty1994c.pdf>*
- Posner EA, Weyl EG. 2015. Voting squared: Quadratic voting in democratic politics. *Vanderbilt Law Review* 68:441
- Ray D, Vohra R. 2015. The farsighted stable set. *Econometrica* 83:977–1011
- Riker WH, Brams SJ. 1973. The paradox of vote trading. *The American Political Science Review* 67:1235–1247
- Sawyer J, MacRae D. 1962. Game theory and cumulative voting in Illinois: 1902-1954. *The American Political Science Review* 56:936–946
- Schwartz T. 1975. Vote trading and pareto efficiency. *Public Choice* 24:101–109
- Schwartz T. 1977. Collective choice, separation of issues and vote trading. *American Political Science Review* 71:999–1010
- Schwartz T. 1981. The universal-instability theorem. *Public Choice* 37:487–501

- Shapley L, Shubik M. 1977. Trade using one commodity as a means of payment. *Journal of Political Economy* 85:937–968
- Shubik M, Van der Heyden L. 1978. Logrolling and budget allocation games. *International Journal of Game Theory* 7:151–162
- Taylor CR, Yildirim H. 2010. A unified analysis of rational voting with private values and group-specific costs. *Games and Economic Behavior* 70:457–471
- Tideman N, Plassmann F. 2017. Efficient collective decision-making, marginal cost pricing, and quadratic voting. *Public Choice* 172:45–73
- Tsakas N, Xeferis D, Ziros N. 2020. Vote trading in power-sharing systems: A laboratory investigation. *Available at SSRN 3261141*
- Tullock G. 1959. Problems of majority voting. *Journal of Political Economy* 67:571–579
- Tullock G. 1961. Problems of majority voting: reply to a traditionalist. *Journal of Political Economy* 69:200–203
- Tullock G. 1970. A simple algebraic logrolling model. *The American Economic Review* 60:419–426
- Walzer M. 1983. Spheres of justice: A defense of pluralism and equality. Basic Books: New York
- Weiss JH. 1988. Is vote-selling desirable? *Public Choice* 59:177–194
- Weyl EG. 2017. The robustness of quadratic voting. *Public Choice* 172:75–107
- Wilson R. 1969. An axiomatic model of logrolling. *The American Economic Review* 59:331–341
- Xeferis D, Ziros N. 2017. Strategic vote trading in power sharing systems. *American Economic Journal: Microeconomics* 9:76–94
- Xeferis D, Ziros N. 2018. Strategic vote trading under complete information. *Journal of Mathematical Economics* 78:52–58