

Who will vote quadratically? Voter turnout and votes cast under quadratic voting

Louis Kaplow^{1,2} · Scott Duke Kominers^{1,2}

Received: 31 August 2016 / Accepted: 30 September 2016 / Published online: 2 March 2017
© Springer Science+Business Media New York 2017

Abstract Who will vote quadratically in large- N elections under quadratic voting (QV)? First, who *will vote*? Although the core QV literature assumes that everyone votes, turnout is endogenous. Drawing on other work, we consider the representativeness of endogenously determined turnout under QV. Second, who will *vote quadratically*? Conditional on turning out, we examine reasons that, in large- N elections, the number of votes that an individual casts may deviate substantially from that under pure, rational QV equilibrium play. Because turnout itself is driven by other factors, the same determinants may influence how voters who do turn out choose the quantity of votes to cast. Independently, the number of votes actually cast may deviate dramatically from pure QV predictions because of the complex and refined nature of equilibrium play. Most plausibly, voting behavior and outcomes would be determined predominately by social and psychological forces, would exhibit few of the features emphasized in the analysis of hyper-rational equilibrium play, and would have consequential properties that require a different research agenda to bring into focus. Some of our analysis also has implications for voting behavior under other procedures, including one person, one vote.

Keywords Voting · Voter turnout · Paradox of voting · Quadratic voting · Pivotality · Elections

JEL Classification D71 · D72 · D82

✉ Louis Kaplow
meskridge@law.harvard.edu
Scott Duke Kominers
kominers@fas.harvard.edu

¹ Harvard University, Cambridge, MA, USA

² National Bureau of Economic Research, Cambridge, MA, USA

1 Introduction

Downs (1957) introduced the paradox of voting: many individuals vote in large- N elections despite the fact that the pivotality-based utility from voting is negligible whereas the cost of showing up is not. For example, if one's utility from switching the outcome in one's preferred direction is \$10,000 and the probability of being pivotal is 1 in 1,000,000, the utility from pivotality is literally 1¢, which is orders of magnitude smaller than even a modest time and hassle cost of, say, \$10.

Nevertheless, much of the literature on voter turnout and voting behavior given turnout examines hyper-rational models that focus on how a given voter's lone vote affects the outcome. Enos and Fowler (2014, p. 310) "reviewed all articles published on voter turnout since 1980 that appear in five leading political science journals. Of the 70 articles ... that addressed the causes of voter turnout, 41 made a clear appeal to the importance of pivotality or electoral competition." Other lines of work assume full turnout or focus on refined equilibrium calculations by hyper-rational voters who are motivated to turn out and decide how to vote based on pivotality.

Theoretical and empirical explorations have confirmed Downs's view that the probability of being pivotal in large- N elections is miniscule. This suggests that hyper-rational, conventionally self-interested voters are not voting in the first place. The many who do turn out presumably do so primarily for other social and psychological reasons, as suggested by Riker and Ordeshook (1968). Therefore, the turnout set is a selected sample largely determined by factors outside many of the equilibrium models of voter behavior. Moreover, it is a priori plausible that these qualitatively different motivations for turning out also influence how those individuals who do turn out cast their votes.

Against this background, a broader research program of ours (Kaplow and Kominers in preparation) seeks to build a theoretical framework that links voter turnout to voting outcomes in a manner that reflects a more empirically grounded perspective on the causes and process of turnout. The present article applies that overall approach to quadratic voting (QV). The core analytical papers on QV (Lalley and Weyl 2016a,b; Weyl 2015, 2017),¹ like much of the leading theoretical literature before them, mostly assume full turnout and that all voters are hyper-rational maximizers of their self-interests when determining how to vote—which, under QV, means not merely which way individuals will vote, but, importantly, how many votes each will cast (i.e., purchase).

We offer a preliminary conceptual investigation of two questions that have not been at the focus of the prior theoretical work on QV: *What determines who will turn out under QV in large- N elections?* And: *For those who do turn out, how many votes might they actually cast?* Specifically, we consider these questions in light of who turns out (selection), the motivations that led them to turn out (and how QV may affect those motivations), voters' difficulties in comprehending the hyper-rational QV equilibrium, and the refined nature of the QV equilibrium itself.

We regard our contribution as complementary to the core QV literature and to the theoretical voting literature more broadly. Our focus is positive. Much of our analysis suggests that voter turnout and voting outcomes under QV could diverge dramatically from what the pure theoretical models predict. As a normative matter (not our focus), this is not necessarily a criticism, both because some of our argument suggests that results under one person, one vote also depart from theoretical predictions and, as is familiar, outcomes

¹ See also Goeree and Zhang (2017).

under one person, one vote (and under other voting mechanisms) have important normative deficiencies, which is the motivation for considering QV in the first place. Interestingly—and ironically—it may well be that QV in practice will generate voter turnout and voting outcomes that are rather close to what prevails under one person, one vote. In any event, in assessing QV in isolation and comparing it to other methods, it is important to understand how QV may operate in practice.

Section 2 examines voter turnout and voting outcomes under the classical one person, one vote framework, drawing on our broader work on the subject. The purpose is to fix thinking about a number of issues, some familiar and some more novel, and to serve as a benchmark for comparison and a launching point for our examination of these questions under QV. First, we briefly recall the standard models of Downs (1957) and Riker and Ordeshook (1968). Second, we review some of the literature on pivotality in large- N elections, which suggests that, with turnout in the range actually observed, the probability that one's vote is pivotal is indeed miniscule. Last, we elaborate some of the implications under one person, one vote.

Section 3 presents our analysis of QV. We begin with basic QV mechanics, showing how the number of votes cast, the amount a voter spends casting those votes, and pivotality-based utility depend on the probability of being pivotal and on one's utility from influencing the outcome. Then we examine turnout under QV, applying Sect. 2's analysis and offering some specifics that are distinctive to QV. We also show how the results of Ledyard (1984), Myerson (2000), and others change dramatically under QV.

Most important, the rest of Sect. 3 addresses how many votes various individuals are likely to cast under QV, conditional on turning out. We first present a view ("View I") that takes hyper-rational QV equilibrium seriously and examines its implications. Both direct reflection and factors adduced previously lead us to advance an alternative view ("View II"), which presents a set of answers that are qualitatively and quantitatively different (by orders of magnitude) from those that obtain in the pure, hyper-rational QV equilibrium.

Importantly, we explain why the manner in which the number of votes cast varies across individuals with different preference intensities may not be well approximated by pure QV analytics. Instead, as mentioned, the results may be much closer to those under one person, one vote. Moreover, the conclusions we reach hold even if a significant fraction of voters—including in particular those with the most intense preferences—both turn out and vote as pure QV'ers, but taking into account that others are doing something else, perhaps as simple as casting one vote, paying one dollar. Indeed, if even a modest portion of the population behaves this way, the votes of remaining QV'ers will be drowned out. Forces entirely different from those analyzed in the core QV literature will then determine votes cast and voting outcomes.

We also examine affinity-group influences on voting. These may produce substantial and robust effects that are similar to those under collusion but stem from different motivations and hence are not undermined by forces that can render collusion unsustainable. We explain why understanding the social and psychological forces that determine affinity group behavior may be the most important consideration in predicting votes cast and voting outcomes under QV.

Finally, we extend our discussion of how individuals would choose votes and expenditures to the case of multi-contest elections, in which there are multiple races or referendums at a single election. Most theoretical work on voting, including the core QV literature, focuses on elections involving single choices, but under QV the presence of multiple choices at a single decision point could itself have important effects on voting outcomes.

Section 4 closes by emphasizing that the foregoing analysis calls for a new research agenda in order to advance our understanding of QV and other voting mechanisms (such as one person, one vote) in large- N elections.

2 Voter turnout and voting outcomes under one person, one vote

2.1 Standard model

Following Riker and Ordeshook's (1968) extension of Downs (1957), the standard model assumes that, under one person, one vote, an individual turns out if and only if the utility from voting, $PB - C + D$, is positive, where P is the probability that one's vote is pivotal, B is the gain one derives from swinging the election to one's preferred outcome, C is the cost of turning out (in time, hassle, and so forth), and D is one's utility from the act of voting per se, which is understood to be psychological or social in nature (avoiding guilt, doing one's civic duty, being able to tell others honestly that one voted, and the like). We will often use the notation U^{piv} to refer to the utility associated with pivotality, which, as stated, is PB under one person, one vote (but is qualitatively different under QV, as will be explained in Sect. 3.1).

Downs's paradox of voting considers the foregoing model without Riker and Ordeshook's addition of the D term. As illustrated in our introduction, if $P = 0.000001$ (one in a million) and $B = \$10,000$, then $U^{piv} = \$0.01$, so if C is as much as \$1, U^{piv} is swamped and hence no one would vote. Actually, as Ledyard (1984) emphasizes and we will elaborate below, if no one votes, P is large, so there may well exist an equilibrium with positive, albeit small turnout. However, if turnout is anything like that actually observed in large- N elections, P is quite small, so Downs's paradox stands—that is, unless one introduces a positive D term. The key message is that, when expected turnout is substantial, not only does one need D to explain turnout, but in addition U^{piv} does not itself play a substantial role in determining turnout and hence outcomes, contrary to the assumptions in much of the theoretical literature on voting.

2.2 What we know about pivotality in large- N elections

An early theoretical exploration of the probability of pivotality in large- N elections is offered by Chamberlain and Rothschild (1981). Examining a model in which the vote is between two alternatives, voters' preferences are drawn from a binomial distribution, and the pertinent probability for that distribution is itself drawn from another distribution, they derive that the probability of an individual being pivotal is on the order of $1/N$ in large- N elections.²

Gelman et al. (2012) offer an empirical analysis of the 2008 presidential election and find that the probability of a vote being pivotal (taking into account complexities due to the Electoral College) was approximately 1 in 60,000,000 (the total turnout was roughly 130 million). In a number of states, the probability was less than 1 in 1,000,000,000; at the other end of the spectrum, in a handful of states, the probability was around 1 in 10,000,000. Note further that the pivotality probability varied by orders of magnitude

² Specifically, in the limit as N goes to infinity, the probability of being pivotal is $1/2N$ times the magnitude of the (assumed to be continuous) density of the distribution of true probabilities, evaluated at 0.5.

across the states, whereas the turnout differentials undoubtedly did not reflect more than a sliver of this difference.³

Mulligan and Hunger (2003) assembled a substantial historical database of House elections and state legislative races, both of which have vastly smaller electorates than the nation as a whole. The pivotality probability was roughly 1 in 89,000 for the former and 1 in 15,000 for the latter.⁴

This general conclusion about the exceedingly low probabilities of being pivotal in large- N elections is stated aggressively by Enos and Fowler (2014, p. 310):

This model [$PB - C + D$] provides the basis for much of the current understanding of the decision to vote. However, P is infinitesimal for any large election..., so changes to P should have little effect on turnout. As Schwartz (1987) aptly points out, “Saying that closeness increases the probability of being pivotal is like saying that tall men are more likely than short men to bump their heads on the moon.” Gerber, Green and Larimer (2008), echoing Schwartz, note that “Because the probability of casting a decisive vote in an election is typically infinitesimal, the calculus of voting boils down to the relative weight of C and D .”

2.3 Voter turnout and voting outcomes

Motivated by the foregoing considerations, Kaplow and Kominers (in preparation) examine a turnout model driven entirely by C and D . Specifically, we use a continuum of voters, so that U^{piv} is literally 0, leaving only individuals’ costs and per se psychic benefits of turning out. To perform the analysis, we consider various properties of the joint distribution of C , D , and voters’ preferences. We here present four results (which we will apply later in our analysis of QV) in an oversimplified form for ease of exposition.

First, when preferences are distributed independently of C and D , the level of turnout is irrelevant. More precisely, when we vary the distributions of C and D , the distribution of voters by preferences is the same (as long as turnout remains positive). Hence, the voting outcome will be the same as well (as long as the voting mechanism exhibits scale invariance, which is a property of one person, one vote and many other mechanisms, including QV in an approximate sense).⁵

Second, when we relax the independence assumption, anything is possible. That is, for any given, unconditional distribution of preferences, it is possible for the preferences among those voters who do turn out to have any (other) distribution. (For example, we may have the opposite distribution, wherein the density of preferences of a given intensity for

³ Empirical work does suggest that closeness influences turnout to some extent (Geys 2006), but even those modest effects (relative to orders-of-magnitude differences in the pivotality probability) may be partly explained by the fact that closer races generate greater publicity, higher campaign expenditures, and more aggressive get-out-the-vote campaigns (Cox and Munger 1989). And surveys suggesting that voters who turn out have higher estimates of P may reflect cognitive dissonance rather than actual behavior-inducing core beliefs (Enos and Fowler 2014).

⁴ For the House races, the mean and median vote totals are roughly 110,000, and for the state legislative elections, the median and mean are roughly 14,000 and 24,000. Their results are in an important respect the most convincing because they impute the likelihood of a tie (or one vote from a tie) from the actual distribution of the vote differential in all of the very close contests that occurred (in a sample of roughly 17,000 House elections and 40,000 state contests). Hence, they rely less on the sort of parametric assumptions that are required for most other empirical or theoretical exercises.

⁵ As suggested in the core QV literature, the approximation statement is advanced in the limit and is conjectured to be true as a close approximation in large- N elections.

any choice becomes the density of preferences of that intensity in the turnout set in favor of the opposite choice.) Hence, any voting outcome that is in the range of the voting mechanism is possible, without regard to underlying preferences.

Third, simple correlations between C , D , and preferences generate predictable skews in the turnout set and thus in the voting outcome. Fourth, as one increases the noisiness of this correlation, there is (speaking quite loosely now) a tendency in the limit for turnout to approach 50% and for the distribution of preferences in the turnout set to become representative of the unconditional distribution of preferences in the population. Interestingly, higher turnout is not associated with more representative turnout in any general sense. In particular, when turnout starts above 50%, greater noise in the correlation tends to cause the representativeness of the turnout set to rise and the turnout level to fall.

Taken together, these results support a broader proposition that, if representativeness is the concern (a normative question that we do not address), one should pay more attention to differential selection than to turnout levels.

In our analysis of voter turnout and voting outcomes, the question of how individuals will cast their votes conditional on turning out is submerged. Specifically, individuals' preferences in our model refer to those preferences that individuals would express in their votes under a given mechanism. How these preferences relate to self-interest, narrowly construed, to the factors underlying C and D , or to the voting mechanism itself (and how it, in turn, may bear on C and D) is not considered. In Sect. 3.3's analysis of the question of how individuals will choose how many votes to cast under QV, we peek inside this black box to explore this dimension of the problem as well.

3 Voter turnout and voting outcomes under QV

This section begins by presenting basic QV mechanics, both for readers who may be unfamiliar with the mechanism and to develop some simple points not emphasized in the core QV literature regarding individuals' expenditures and the resulting pivotality-based utility, U^{div} , under QV. Next we examine voter turnout under QV. Whereas the core QV literature posits that turnout is 100%, we apply and extend the analysis described in Sect. 2.3, wherein turnout is endogenous, to examine some plausible similarities and differences between turnout under one person, one vote and under QV.

We devote most of our attention to the question of how individuals would actually vote under QV. Importantly, this includes not just the question of for whom or what they would vote but also how many votes they would cast (i.e., purchase). Here, as the Introduction mentions, we introduce two contrasting views. View I takes QV on its face and attempts to describe a world in which (most) voters behave in accord with its dictates—in particular, they derive and play QV-equilibrium strategies of the sort elaborated in the core QV literature. We contemplate how they might determine how much to vote, how much they would spend, and the utility they would derive.

View II, by contrast, imagines that a significant fraction of the electorate are ordinary people who turn out based on a $D - C$ calculus and vote in ways that may reflect their D -related motives for showing up in the first place—and, in any event, they do things like rounding up a provisional voting expenditure of a few micro-cents to perhaps a dollar. This depiction reflects that the vast majority of the population: (1) does not have training in statistics and game theory, (2) would be unlikely in any event to undertake the demanding analysis required to determine QV equilibrium strategies, (3) would not have shown up in

the first place if they had engaged in a purely rational self-interested calculus (because under QV, U^{piv} remains negligible in large- N elections), and (4) would be governed by various social and psychological considerations.

Moreover, as we will explain, once even a modest portion of actual voters no longer behaves in a hyper-rational QV fashion (that is, as if under View I), then the remaining pure-QV'ers will find it in their interest to cast a negligible quantity of votes. Therefore, almost all votes cast as well as voting outcomes will be essentially unaffected by the hyper-rational QV calculus; instead, social and psychological forces will dominate. In the simplest version, interestingly, actual voting behavior under QV may be vaguely like that under one person, one vote. However, a number of considerations may push voting in directions quite different from those under either pure QV or one person, one vote.

3.1 QV mechanics

Here we revisit the all-in utility from turning out under one person, one vote, which—as per the standard model of turnout—is $PB - C + D$. There, we had $U^{piv} = PB$. For quadratic voting, we find it convenient to continue to employ a closely analogous representation of $U^{piv} - C + D$, wherein we now reinterpret and rederive U^{piv} . First, individuals who turn out do not simply cast a single vote but rather choose the number of votes V to cast. Second, they must pay V^2 for the votes that they cast. We will include both of these features in the U^{piv} term.

Each vote an individual casts has a probability P of being pivotal, and if it is pivotal generates the utility benefit B , as before.⁶ The total utility from influencing the outcome is therefore PBV . The cost of casting V votes is V^2 . Therefore, from the first-order condition, an individual's optimal number of votes is $V^* = PB/2$. This number of votes generates an influence utility of $P^2B^2/2$ and requires an expenditure for casting V^* votes of $P^2B^2/4$. Thus, we have $U^{piv} = P^2B^2/2 - P^2B^2/4 = P^2B^2/4$ as well.

It is, of course, the central feature of QV that V^* rises linearly in B (the intensity of an individual's preferences) and also rises linearly with P . It will be useful below to have in mind as well the basic properties of the expenditure and U^{piv} (which have the same value, as just derived). First, observe that both are rising with the square of intensity, B . As we will see in Sect. 3.2, in a world in which U^{piv} did have an important turnout effect, this means that the turnout effect would be magnified for those with extreme preferences. Second, both of our values—expenditure and U^{piv} —are rising with the square of the probability, P , that one vote (here, one unit of vote cast rather than one voter) is pivotal. This proves to be quite important because P tends to be very small in large- N elections.

To illustrate the interaction of these effects, let us return to our simple, toy example in which $P = 0.000001$ (one in a million) and $B = \$10,000$. It is immediate that: $V^* = 0.005$, the optimal expenditure to cast this quantity of votes is \$0.000025 (one four-hundredth of a cent), and, likewise, $U^{piv} = 0.000025$.⁷

⁶ As Lalley and Weyl (2016b) emphasize, the statement in the text is an approximation for pure-QV behavior, one that becomes increasingly precise as N rises. As the reader will see, however, because we present the basic QV mechanics only to get a sense of the magnitudes involved, this subtlety is unimportant for our purposes.

⁷ As stated, this is a toy example; to do an actual one, we would need to specify the distribution of types and uncertainty over that distribution, state what our “large” N was, work out the actual equilibrium or equilibria, and then—for our later adjustments—build a new model to incorporate those phenomena. Any particular pivotality probabilities would be highly sensitive to those choices, and as of this writing there were no convenient examples to take from the core QV papers. One might, for example, use a larger

These figures are very small. By contrast, recall that, under one person, one vote, we had $U^{piv} = \$0.01$ (one cent). Small as that was (and dwarfed by even a modest C), we now have a value for U^{piv} that is much smaller. Moreover—and foreshadowing much that will follow—we have the spectacle of an individual, say, with a C of \$10, who nevertheless votes because D is \$15, arriving at the voting booth, feeling especially civic-minded, pulling out a credit card, and spending ... \$0.000025.

As per the QV logic, extremists spend much more. Consider an individual facing the same P but having $B = \$100,000$. Because expenditures and U^{piv} rise with the square of B , both values are two orders of magnitude larger: \$0.0025 (a quarter of a cent), but still much less than under one person, one vote, where for this example $U^{piv} = \$0.10$.

Against the foregoing background as well as the analysis from Sect. 2, we now examine voter turnout and the number of votes likely to be cast under QV.

3.2 Voter turnout under QV

Individuals turn out if their all-in utility from voting, $U^{piv} - C + D$, is positive. We have just seen that, compared to one person, one vote, which we analyzed in Sect. 2, U^{piv} falls by orders of magnitude in our toy example, even for extremists. Nevertheless, since U^{piv} was mere rounding error for most (if not all) voters to begin with, this large relative drop in itself is unlikely to be consequential.

The discussion thus far, however, is incomplete because our example simply posits a common P , whereas the equilibrium P will not generally be the same under QV as it is under one person, one vote.⁸ If View I prevails for (nearly) the entire population, the number of votes may well be lower and P will in any case differ (in a manner that depends on subtleties of the equilibrium). Nevertheless, the analysis in Sect. 3.3—supportive of View II—will reinforce the suggestion that a large number of votes likely will be cast (purchased) in large- N elections under QV even if most individuals are pure QV-ers, so P will plausibly be quite small regardless.

If this were the entire story, the four results from Kaplow and Kominers (in preparation) about voter turnout and voting outcomes presented in Sect. 2.3, as applied to one person, one vote, would carry over to QV. First, as long as C and D (however they may differ from under one person, one vote, as elaborated below) are distributed independently of preferences, turnout is irrelevant: as long as there is positive turnout, the results will be the same. The subtle difference under QV, as explained in the core literature, is that many of QV's properties have only been identified in the limit. If turnout, say, were only 50 or 25%, rather than the stipulated 100%, the effective N would be smaller and therefore the applicability of the limiting results would be qualified accordingly. Additionally, when U^{piv}

Footnote 7 continued

P (perhaps with one or two fewer zeros to the right decimal point) or a smaller B (perhaps with one or two fewer zeros); making both adjustments would be offsetting. Moreover, as will become apparent in our analysis that follows, whether the optimal expenditure for a typical voter is \$0.000025, \$0.025, or \$0.00000025 is of little significance (and, as Sect. 3.3.1 elaborates, once a modest fraction of the population just pays \$1 to cast 1 vote, the pivotality probability will become very low in any event).

⁸ Keep in mind that the probability P as well as the benefit B enter U^{piv} linearly under one person, one vote but each enters quadratically under QV, so for U^{piv} to rise under QV, an individual's B would need to exceed $1/P$ (by at least a factor of 4, recalling the full formula under QV). In large- N elections, as we elaborate below, this does not seem plausible for many (or any) individuals. In that event, note also that the magnitude of the derivative of U^{piv} with respect to B (intensity) is likely to be smaller under QV than under one person, one vote.

is not literally 0 but merely small, preferences distributed independently of both C and D will result in relatively more (but as per the above, not much more) turnout for those with more intense preferences. For large- N elections, these qualifications should be unimportant.

Second, anything can happen under QV as well (again subject to the sort of qualification just noted). If View I prevails, wherein most votes cast are those of pure-QV'ers, the outcome for a given turnout set is different from that under one person, one vote. Specifically, the pure QV outcome is the utilitarian outcome *for the individuals in that turnout set*, which, if unrepresentative, will not match that for the population. For example, if most individuals who turn out are those who favor the anti-utilitarian result, pure QV behavior in that turnout set will yield the anti-utilitarian outcome. Furthermore, the analysis showing that, regardless of the distribution of preferences in the population, one can get any turnout set (that is, any distribution of preferences on the turnout set) under one person, one vote holds under QV as well. Kaplow and Kominers (in preparation) prove this for the case in which $U^{piv} = 0$, which is likely to be approximately true under QV. But to whatever extent this does not hold, one can simply consider a conditional distribution of C such that those with preferences generating a higher U^{piv} also have a correspondingly greater C , which makes it possible to generate any target level of turnout, conditional on preferences, by the same sort of construction.

Also still applicable are our third and fourth results on how systematic correlation between C , D , and preferences can generate skewed turnout sets and hence skewed results, a skew that tends to be attenuated by greater noise in the correlation. As we now move to our discussion of particular reasons that QV may influence turnout, keep in mind that what matters for outcomes is *systematic bias* in turnout effects. To compare QV to one person, one vote, we need to compare the biases under the two mechanisms, along with the differences in how the mechanisms operate on a given turnout set. The foregoing analysis is presented under View I, wherein most votes are cast by pure, hyper-rational QV'ers. Note, however, that such QV equilibria are delicate in various respects, and once turnout is taken to be endogenous and, moreover, influenced by the sorts of forces considered below, the nature of equilibrium behavior and its properties remain to be identified.

Against this baseline in which nothing much changes with respect to generation of the turnout set, let us now consider (or really, speculate about) some reasons that turnout decisions may differ nontrivially under QV. First, many individuals might regard C to be higher under QV because of the mechanism's greater complexity and concomitant calculation costs. To be clear, we do not have in mind the difficulty of taking a derivative and solving the basic problem to determine, from a given P and B , how many votes to cast.⁹ Rather, in the spirit of View I, we are contemplating that individuals take QV seriously and attempt to think about the determinants of the equilibrium level of P . (Not to actually solve for the equilibrium, even roughly, which is unimaginable—a fact that can be verified by the reader's examination of the appendix in Lalley and Weyl (2016a)¹⁰—although key results

⁹ Of course, for most citizens, it is hardly obvious that this is simple, although one can imagine providing tables or enabling entries of tentative votes or expenditures that would allow individuals to develop some sense of the tradeoffs. Even so, the resulting opacity may contribute further to the phenomenon that we elaborate in the text.

¹⁰ We further note that this mind-boggling complexity arises even with a number of simplifying assumptions, including common knowledge, various approximations, and exogenous turnout. Even then, various results as of this writing have not yet been proved. See also Weyl (2017): "However, [Lalley and Weyl] also show that proving properties of the QV equilibria fully formally involves a rich statistical machinery; even in the simple environment they study the proof of their main theorem involves forty pages of dense

in the core QV papers suppose that individuals behave this way and, moreover, assume that all others are behaving likewise.)

Our suggestion is that individuals who really attempt to wrap their minds around this—even a little—will suffer significant costs due to mental effort and anxiety. (“Marge: Did you read the new voting instructions? It’s a math test! Actually, they said that it’s ‘Game Theory,’ whatever that is.”) We call this phenomenon *QV headaches*.¹¹ If turnout is indeed driven by C and D , and if C accordingly rises, then turnout will fall. Moreover, recall that D is the social or psychological utility from voting, perhaps from a sense that one is performing a civic duty. If that duty is now headache-inducing, the act of voting may lose some of its luster. In addition, foreshadowing the more behavioral discussion that follows, we note that the anxiety dimension may create dissonance; in particular, the act of performing one’s civic duty while being aware that one does not understand what one is doing might itself discourage turnout.¹²

If there would be a significant reduction in turnout, however, the immediate implication is unclear. Turnout levels per se do not influence outcomes. What matters instead is whether QV might skew turnout in some systematic fashion. For example, might more informed and thoughtful individuals get more headaches and suffer more anxiety from contemplating QV equilibrium behavior than do individuals who do not think very hard about how to vote? Or might individuals with more extreme preferences attempt to undertake greater effort and, in the process, suffer greater psychological costs? Or might the emphasis on calculations and gaming primarily dissuade individuals who find those sorts of tasks distasteful? Or, as we explore below, might these considerations—which push individuals away from undertaking a pure QV calculus—give a further boost to various alternative behaviors that we examine under View II, including affinity group behavior?

Consider next how the QV mechanism may affect turnout through its influence on D , which necessarily drives turnout when C is nontrivial relative to U^{piv} for almost everyone. We understand D to consist substantially of nonpecuniary utility associated with the social or moral character of the voting experience. Because of the subjective and social nature of D , it is plausible to suppose that it is endogenous to the voting mechanism, a proposition that has some empirical support.¹³

The most straightforward channel by which QV may influence D is that individuals must now pay to vote. A familiar phenomenon is that extrinsic motivations—usually taken to be financial rewards—sometimes drive out intrinsic motivations. See, for example, Ariely et al. (2009), Frey and Jegen (2001), and Gneezy et al. (2011). To the extent this

Footnote 10 continued

mathematics. It therefore seems impractical to study robustness of their conclusions to different modeling assumptions in the highly formalistic manner they pursue.”

¹¹ Whether to avoid headaches or because virtually no one has the capacity to figure this out, individuals may satiate, relying on guesstimates from votes cast in prior elections or indications from the latest polls. If so, many of the subtle properties explored in the core QV literature could become inoperative. Note, for example, that in the generic (mean not equal to 0) case examined in Lalley and Weyl (2016b), extremists are critical to the equilibrium in the case they emphasize, and that offsetting collusion (Weyl 2017) requires counter-behavior—but polls are unlikely to have a large enough sample to measure either extremists or colluders; nor would such individuals respond truthfully given that revelation would trigger countermeasures by other voters.

¹² Consider results indicating that individuals tend to be passive when action would require making complex, difficult-to-comprehend decisions. See, for example, Iyengar and Lepper (2000) and Tversky and Shafir (1992).

¹³ For example, making it easier to vote remotely, such as by mail, so as to reduce C , should raise turnout on that account, but it may instead lower turnout by reducing D as well (Funk 2010).

occurs, turnout would fall, particularly in light of the fact that the extrinsic rewards themselves are negligible, as explained. Once again, looking ahead to voting outcomes, the question of interest would be whether any such turnout suppression would be nonrandom.

When we look more closely at QV, however, the analysis differs critically from familiar crowd-out considerations because here individuals are not being paid to vote but instead must pay to vote.¹⁴ The resulting monetary disincentive is negligible and already incorporated in U^{piv} , which in aggregate is plausibly quite small. Rather, any effect is likely to be attributable to the notion that being required to pay, even nominal amounts, to perform one's civic duty might undermine that sense of duty. Individuals may also be bothered by the fact that others are able to buy more votes.

Another possibility relates to the idea advanced by Shampanier, Mazar and Ariely (2007) that, even without regard to intrinsic motivations and social preferences, a price of zero is special: individuals, contrary to neoclassical assumptions, treat even small positive prices as significantly more of a disincentive than indicated by the cost itself. To the extent that this effect is nontrivial with ordinary goods, it might be significant for voting, and all the more so in combination with the foregoing consideration of the social and moral character of D .

Taken together, it seems possible that moving to QV may result in substantial changes—in particular, reductions in D . In the absence of empirical evidence, we will briefly speculate on ways that this effect may be nonrandom. One might suspect that the effect would be most pronounced among altruists. The main reservation is that, as per Downs, purely selfish individuals (for whom we will suppose that $D = 0$) are not voting to begin with. If everyone's D fell by the same amount or percentage, those no longer voting would be individuals who had nontrivially positive D s, but typically lower D s than those who continue to vote. Therefore, it may be that the smaller turnout set would exhibit, on average, even more of a pro-social orientation. Or perhaps the opposite would obtain, if the very high- D types are turned off and the more moderate- D types are influenced less. Recall in any event that pro-social voting is not obviously advantageous; the normative basis for QV is that it will induce the revelation of individuals' own preferences in settings where a planner and fellow citizens do not know which policy or candidate is actually superior. By contrast, because of the failure of one person, one vote to reflect intensity, pro-social turnout may improve its performance.

Before leaving the topic of QV and turnout, it is interesting to consider briefly how the results originated by Ledyard (1984) and extended by Myerson (2000) and Krishna and Morgan (2015), among others, change when one moves from one person, one vote to QV.¹⁵ The central model posits that individuals vote if their utility, given by $U^{piv} - C$, is

¹⁴ It is also true that under QV, at least as presented in the core literature, it is imagined that the proceeds will be rebated. The rebate does not, however, depend on one's own payment, and in any event it is not obvious that these events will be linked in individuals' minds. One could limit the rebates to those who actually turn out, in which case we would be back to the prior point about being paid to vote. Another possibility is that individuals could be given some allowance—perhaps an estimate of their rebate—out of which they could pay for votes, keeping the balance or topping it off, at their discretion. Anticipating the discussion in Sect. 3.3.2, such a scheme may result in the allowance becoming the focal point for expenditures, generating an outcome under QV that is much like that under one person, one vote.

¹⁵ The core QV literature notes some kindred spirit between QV and the literature on costly voting under one person, one vote. We see this duality as follows: in Ledyard (1984), a tiny fraction of the population casts a single vote, and turnout is proportional to intensity; under QV (with costless turnout), everyone casts a tiny number of votes, and the number of votes cast by each individual is proportional to intensity. On another matter, in Ledyard's paper that launched this line of work, candidates endogenously choose positions, and equilibria can have the form that literally no one votes because both candidates choose identical

positive. (Note, importantly, that this is the Downs (1957) model rather than the Riker and Ordeshook (1968) model that we have been employing.) Ledyard sought to characterize the equilibrium of the Downs model, noting that if few or none turn out, the pivotality probability P would be nontrivial, and this would induce some turnout.

This line of literature assumes that preferences and individuals' C s are distributed independently and, moreover, that the density of the distribution of C is positive on $[0,1]$ (with no point mass at 0). The intuition and results can be seen by recognizing that, when one examines neighborhoods of 0, the density will be approximately uniform. In a large but finite population (individuals with given preferences have their C s drawn from the aforementioned distribution), the number who turn out for any given intensity is proportional to that intensity under one person, one vote. This can be seen, roughly, by positing an equilibrium P (which, as N goes to infinity, we can take to be small) and noting that the fraction of voters who will have their randomly chosen C fall below PB will be proportional to B (in the limit). This turnout selection means that one person, one vote produces the utilitarian outcome because endogenous turnout generates a number of votes that is proportional to intensity. Endogenous turnout thus yields quadratic outcomes.

What about QV? Here, the number of votes cast will rise not linearly with intensity, as under pure QV with a stipulated 100% turnout, but rather with the *cube* of intensity, generating a tremendous overweighting by intensity. The reason is that the selection logic just presented means that turnout under QV would be proportional to the square of intensity. Recall that, under one person, one vote, $U^{piv} = PB$, which is why turnout rises linearly with B . Under QV, as explained in Sect. 3.1, we have $U^{piv} = P^2B^2/4$, which is rising with the square of B . Moreover, under QV, each person turning out does not cast a single vote but rather casts $V^* = PB/2$ votes, which rises in proportion to B . Combining these two points, the number of votes cast, for a given intensity, rises with the cube of intensity.

Summarizing the analysis of this model, under one person, one vote, endogenous turnout solves the failure to reflect intensity, whereas under QV, the absence of a problem when there is stipulated 100% turnout becomes a serious problem when turnout is endogenous. Ultimately, we regard this observation to be interesting but not practically important because, as we have argued from the beginning, Riker and Ordeshook's model of turnout based on $U^{piv} = C + D$ seems much more compelling.¹⁶ In the models building on Ledyard, which include C but not D , a vanishingly small fraction of the population turns out in large- N elections, contrary to what we observe and to what we would expect to prevail under QV.¹⁷ Indeed, as we suggested, it seems most likely that under QV, like under one person, one vote, U^{piv} would have at most a negligible effect on turnout, in

Footnote 15 continued

positions, at the electorate's preferred point (in the sense discussed in the text to follow). A candidate's deviation would induce turnout with the sorts of properties discussed in the text.

¹⁶ To be clear, Ledyard and the subsequent papers assume no mass point in the distribution of C at 0 and that C cannot be negative. We also find it interesting that, in a sense, the QV literature "needs" D to play a significant role in order to have a chance of plausibly functioning in the intended manner in large- N elections. This further motivates our attention in this subsection and the next to how the QV mechanism may interact with D .

¹⁷ Note that a result similar to that in Ledyard could be generated if one instead introduced D and assumed it to be proportional to preference intensity (perhaps for expressive reasons). If preferences, and hence D , were distributed independently of C , the result under one person, one vote would involve significant turnout (if the proportionality factor was sufficiently large) with that turnout being crudely proportional to intensity. (It would not be strictly proportional, and perhaps not even that close, because the argument that the density of the distribution of C is uniform in the neighborhood of 0 would no longer apply.) One person, one vote

which case the preceding analysis in this subsection is most relevant to understanding turnout under QV.

3.3 Voting outcomes under QV

Voting outcomes under QV are, importantly, determined not only by who turns out but also by how many votes those individuals cast (i.e., purchase). To begin, examine what we have seen thus far. A typical voter who does turn out incurs, say, a cost, C , of \$10 but shows up because her D is \$15, and then she spends less than a hundredth of a cent. “Hi Jill! Did you vote today?” “Of course!” “Just wondering, how much did you spend?” “25 millionths of a dollar, naturally.”

3.3.1 Two views of votes cast under QV¹⁸

This subsection and those that follow develop what we have referred to as View II, wherein most votes actually cast do not follow hyper-rational pure-QV behavior (View I). Our reasoning in favor of View II reflects two considerations: the complexity and fineness under pure QV pushes voters away from View I behavior, and social and psychological considerations simultaneously pull individuals toward other, qualitatively different behaviors. The strength of our belief in View II reflects the power of both the push and the pull, as well as their mutually reinforcing influences on individual voters. It is bolstered by the fact that, as we will explain, once even a moderate portion of individuals are understood to behave in the ways that we will describe (such as spending \$1 rather than \$0.000025), the number of votes cast by remaining hyper-rational pure QV’ers would plummet and their influence would be drowned out in any event.¹⁹

Regarding the push away from hyper-rational pure-QV behavior in determining the number of votes to cast (i.e., purchase), recall our discussion of QV headaches. We had suggested that the requisite effort and resulting anxiety may raise C (immediately, or in anticipation of one’s experience in the voting booth) and thereby suppress turnout. Here we are concerned instead with how many votes an individual will cast if that hurdle is overcome. Under one person, one vote, the only relevance of P is to the turnout decision, and we suggested that the effect was negligible. Once in the voting booth, one simply selects one’s preferred outcome.

Under QV however—and continuing to take View I seriously for the moment—individuals would need to know P in order to know how many votes to cast (i.e., purchase). The concern is not merely with the subtleties of approximations that are the focus of much of the analytics in the core QV literature. Instead it is with how a voter is to determine, even within an order of magnitude, what P is. Remember that, under QV, $V^* = PB/2$. Large (orders of magnitude) differences in what a voter thinks about P lead to dramatic differences in the votes one is supposed to cast. And, in order to guesstimate P , the

Footnote 17 continued

would (again) lead roughly to the utilitarian outcome, whereas QV would generate an outcome that roughly reflected the square of intensity.

¹⁸ At the Quadratic Voting and the Public Good conference, Nicolaus Tideman directed us to Margolis (1982b), which advances, in a related context, themes similar to some of those developed in this subsection.

¹⁹ We further observe that if, contrary to the preceding analysis, U^{piv} did have a nontrivial effect on turnout, the analysis here—wherein, with different voting behavior, U^{piv} would differ as well—implies that turnout would be influenced accordingly. As we will see, the effect may then be to suppress the turnout of pure QV’ers relative to those who cast votes in accord with View II, further reinforcing the central claim here.

individual has to understand the character of the QV equilibrium, calibrate it, and solve. We note that Weyl (2017) relies on QV's subtle features in developing arguments for why QV is robust to collusion. From Lalley and Weyl (2016b), in particular, we learn that there are two qualitatively different types of equilibria, depending on whether or not the mean of the distribution equals 0. Each of these equilibria is remarkably complicated and subtle even for trained researchers to comprehend; they depend on common knowledge of unobserved phenomena; and they are potentially sensitive to various factors that are ruled out by assumptions, many employed for purposes of analytic tractability.²⁰ Finally, in our simple example—which merely posited a P of one in a million—recall that the expenditure on voting by an individual with $B = \$10,000$ is \$0.000025.

One might wonder whether this challenge to View I behavior can be addressed by supposing that individuals would make modest, unbiased errors. Suppose, for example, that individuals might spend plus or minus 10% of the correct amount. This gives errors of \$0.0000025, with those on the high side spending \$0.0000275 and those undershooting spending \$0.0000225. Perhaps instead some would round \$0.000025 to \$1 and others to \$0. Then, if a million voters turned out, unbiased results would require that 5000 cast one vote and 995,000 cast 0 votes (or just stay home). If instead 10% rounded up to \$1 rather than down to \$0, we would have 20 times too many votes. Brief reflection suggests that the introduction of some unbiased noise cannot come close to rationalizing approximately pure-QV behavior in large- N elections with human voters.²¹

Recognizing the strong push away from anything like hyper-rational, pure-QV determination of the number of votes that individuals actually would cast does not, in itself, tell us what individuals are likely to do instead. The simplest (and most plausible) thing that many individuals might do is round their votes and/or expenditures to natural focal points. For example, as already suggested, individuals whose optimal expenditures are minuscule might round them up to \$1, which purchases 1 vote.²² Note further that the pure-QV expenditure for an individual with extreme preferences of $B = \$100,000$ amounts to \$0.0025, buying 10 times the number of votes and spending 100 times as much money as our more moderate individual with $B = \$10,000$. If both rounded up to \$1, 1 vote (which even for our extremist involves a 400-fold overpayment), their large difference in intensity vanishes; then, the number of votes cast and hence the voting outcome is precisely the same as under one person, one vote. (In the next subsection, we consider further possibilities regarding voting behavior under View II, including ones under which voters with more intense preferences purchase more votes.)

²⁰ As but one illustration, in the generic (mean not equal to 0) case (for the subcase emphasized by Lalley and Weyl (2016b)), for almost all voters the optimal number of votes to cast depends on their appreciation of the tail properties of the (common knowledge) probability distribution of preferences and, in particular, the probability and behavior of the (limiting) most extreme types with whom one disagrees.

²¹ In addition, introducing additional phenomena such as noise requires corresponding supplemental assumptions (Is there common knowledge about the noise? Do I view my own, say, high number of implied votes as informative about the underlying distribution of preferences?) and further complicates the already daunting problem a voter faces in solving the QV equilibrium.

²² Whether it is votes or expenditures that are rounded to natural focal points could matter greatly. The exception is the focal point of 1, which corresponds to the same number of votes and dollar expenditure. If one spends \$10 instead, the number of votes purchased is ~ 3.2 . If instead this person casts a round 10 votes, there is an expenditure of \$100. Therefore, if social norms, affinity groups, or other considerations lead to one or more focal points, it may matter a great deal whether they are in terms of votes or dollars. The ultimate outcome under QV would differ accordingly, but (echoing another of our themes) only if votes cast and resulting expenditures bear a nonrandom relationship to preferences.

Consider as well the feedback effects of such behavior. Suppose that 1,000,000 voters turn out under QV. Assume further that each of a random (with respect to underlying preferences) 25% of voters simply spends \$1 to purchase 1 vote. Then, even if the remaining 75% of the electorate were hyper-rational QV'ers—who in turn were aware that 250,000 random voters each simply casts a single vote—they would appreciate that P thereby must be quite small with so many total votes cast. Moderates would accordingly find it optimal to make tiny expenditures to cast slivers of a vote.²³ Even those with extreme preferences ($B = \$100,000$ in our example) would cast just fractions of a vote, being drowned out by the 250,000 (mostly moderates) who are each spending \$1 to cast an entire vote.²⁴

We can see that it does not take a very large share of the population to act under View II (as we suspect most individuals would) in order to overwhelm the influence of pure-QV'ers, even if, contrary to the earlier suggestions, pure QV'ers make up a large share of the population. View II voters cast far too many votes, thereby inducing View I voters to cast correspondingly fewer.²⁵ As one imagines the fraction of the population that behaves in accord with View II, rather than View I, rising gradually from 0 toward 1, the votes cast and voting outcomes in a QV election would fairly soon come to reflect a View II world. Of course, if almost everyone behaves in accord with View II in any event, this analytically powerful point is unnecessary for the conclusion that the QV world would be characterized by View II.

3.3.2 Elaboration of View II

Because we believe that View II largely would determine votes cast and voting outcomes under QV, it is important to examine—really, speculate—further on how individuals might determine votes cast and expenditures under QV. This question is the focus of the

²³ This point is admittedly a conjecture because we have not undertaken the highly complex exercise of attempting to solve for the relevant QV equilibrium behavior. Our suggestion is on target if votes cast by pure QV'ers would be similar in magnitude (on average, rather than just for extremists) to those under one person, one vote. (Using the empirical estimates reported in Sect. 2.2 yields pivotality probabilities between $2.17/N$ and $0.93/N$. For convenience, let us use $2/N$ for P in the formula $V^* = PB/2$, so we have $V^* = B/N$. Given our posited N of 250,000—which, conservatively, ignores all the votes cast by the QV'ers themselves—our moderates with $B = \$10,000$ would each cast 0.04 votes and our extremists would cast 0.4 votes. If all were moderates, the 750,000 QV'ers would cast 30,000 total votes; if 10% were extremists, QV'ers' vote total would still be under 60,000.) We note further that our reading of the pertinent formulas and conditions in Lalley and Weyl (2016b) suggests that the number of votes cast under QV would be much smaller than in elections under one person, one vote, raising the possibility that the foregoing guesstimates understate our point. Suppose, for example, that 60% of voters favor one outcome moderately (they have \$10,000 valuations) and 40% favor the other extremely (they have \$100,000 valuations). Using Table 1 of Chamberlain and Rothschild (1981), 100,000 random voters each casting 1 vote implies that P is less than 10^{-35} . Using the formula $V^* = PB/2$, it is clearly an equilibrium for the 900,000 pure QV'ers to collectively cast a total of fractional votes that is orders of magnitude smaller than the one vote that would be cast by a single individual who spends \$1.

²⁴ On one hand, this example may be conservative because our assumption that moderates and extremists benefit, respectively, \$10,000 and \$100,000 from their preferred outcomes is generous to casting large numbers of votes under pure QV. On the other hand, many elections, such as local contests, have notably smaller electorates, which boost the probability of pivotality and thus the incentive to purchase votes—although if the stakes scale downward in similar proportion, the ultimate effect would be much the same as in our example.

²⁵ If one were to take a weighted average of the two groups where the weights reflect votes cast (which is what matters for outcomes), we have the View II group substantially overweighted and the QV'ers underweighted.

remainder of this article and, as will become evident, requires a different research agenda to answer.

If almost everyone spent \$1 to cast 1 vote, the results would be essentially the same as under one person, one vote. Consider whether, instead, individuals' behavior might exhibit some QV-like properties. We suspect that it would. Interestingly, with regard to votes cast and voting outcomes, the bulk of any such effects would come from View II voters rather than View I voters.²⁶ As mentioned, the latter—any remaining hyper-rational pure QV'ers—cast slivers of votes that are drowned out, even if, within that group, the thickness of the slivers varies perfectly with intensity.²⁷ On the other hand, the normal people whom we had imagined might simply round to \$1, purchasing 1 vote each, may behave otherwise. One possibility is that they would follow different rounding conventions: perhaps some would cast 2 votes (paying \$4), or just spend \$10, purchasing approximately 3.2 votes. In that event, individuals attracted to multiple focal points may be more likely to choose higher (lower) ones if they have more (less) intense preferences. As a result, the total number of votes cast and thus voting outcomes would lean away from the result under one person, one vote toward the intensity-weighted outcome that would arise under pure QV.²⁸

Our suggestion of simple focal points naturally leaps to mind when one recognizes that pure-QV behavior implies optimal expenditures like \$0.000025. More broadly, individuals have difficulty with even simple calculations, they do not process small (here, very small) probabilities in accord with the predictions of expected utility theory, and complexity induces satisficing (both rationally and beyond such considerations). For these familiar reasons, we suspect that readers were not surprised by our above suggestion that many, and perhaps nearly all, individuals are more likely to spend \$1—or perhaps \$5, or \$10, or maybe \$0.25—than anything like \$0.000025 (even if an electronic system made such payments possible).

A more nuanced sort of behavior may arise owing to individuals' appreciation that QV enables the expression of preference intensity. Accordingly, they might adjust their vote total or expenditure level more smoothly as a function of preference intensity. This need not reflect anything close to a pure-QV calculus, but rather may involve all manner of mental shortcuts. They may not solve the first-order condition, compute pivotality probabilities in light of subtle equilibrium properties, and so forth. Instead, they might simply spend more when they care more. It is a matter of speculation what functional forms this relationship might take for different voters.

For concreteness, consider the possibility that the expression of preference intensity would more likely be manifest (in individuals' oversimplified mental processes) in terms of dollars rather than votes. For example, one who cares ten times more might spend \$100 rather than \$10 (whereas, under QV, a B that is ten times larger results in ten times more

²⁶ Another irony is that the presence of a substantial number of View II voters may ease the guesstimation of P by pure QV'ers: because View II voters cast so many votes (in absolute numbers and as a fraction of the total), understanding their simpler behavior may enable one to approximate the equilibrium P —recognizing also that, as explained, the pure QV'ers are casting a very small fraction of the total votes, making it much less important to understand the subtleties of their behavior. That is, analysis of the behavior of View II voters will be first-order in the thinking of View I voters—much more important than their attempting to understand each other.

²⁷ Even this trivial effect is fragile and could point in the opposite direction if those with more intense preferences were more inclined to think hard about how much to spend rather than merely spending \$1 for 1 vote; then, extremists would be casting fractional votes, compared to moderates who cast a full vote.

²⁸ We interpret (extrapolate, since they are small N) experimental papers, such as Goeree and Zhang (2017), as suggestive of this sort of tendency under QV in the domain of large- N elections.

votes and thus a hundred times the expenditure—\$1000 in this example).²⁹ In that case, intensity would directly be reflected by the square root in terms of the votes generated (whereas under QV, of course, votes cast are linear in intensity).

The foregoing suggestions, however, are only the beginning. We must also consider the social, psychological, and moral nature of voting behavior that we have noted previously with regard to D , which drives turnout in the first place. Specifically, determinants of D may well bear on how many votes one would purchase under QV, and the QV mechanism's possible feedbacks on D (noted in Sect. 3.2 when considering turnout) may also influence the number of votes that individuals cast.

Recall the dialogue in which Jill responds that she spent \$0.000025, an entirely respectable View I magnitude in our example. But only asocial math geeks (like the authors and perhaps some of our readers) would answer this way or fail to respond with shock at such a response by someone else. If, for example, individuals turn out to vote because they wish (as a matter of pride or due to social expectations) to be able to say that they voted, they would be concerned with others' reactions, which in turn would reflect social norms about not just whether one should vote but, under QV, how much to vote. Exceedingly low (in our instance, fanciful) responses would be embarrassing, suggesting that the voter was not really serious. Or perhaps such would be a way to express discontent: exit polls and subsequent election coverage might point to significant numbers of "one cent votes" as an indication of the number of disgruntled voters, just as they point to low turnout rates today. Or voters might just stay home.

But voting and spending too much might send an unwanted message as well: Is a high-spending voter saying that she is more important than everyone else? Or, relate, is she flaunting her wealth in a setting where such behavior seems particularly inappropriate? Individuals can lie, of course, as some surely already lie about whether they vote. But if one is induced to vote in the first place because $D - C > 0$ (rather than just lying about whether one voted; see Harbaugh (1996) and DellaVigna et al. (2016)), one might be willing to spend a dollar, or a few, but not too many, in order to be able to answer truthfully in a way that generates the desired reaction.

Whether D reflects avoiding guilt, feeling civic pride, or a desire for social approval, what matters for QV is how this translates into an appropriate choice of V . Moreover, following our analysis when discussing turnout, what matters is how this translation interacts with individuals' preferences, including in particular their intensity. For example, if the applicable norm is egalitarian, the number of votes cast or expenditures may be independent of intensity, again generating a one person, one vote sort of outcome. But if the norm is that preference intensity should be expressed, we would have QV-like variation in the number of votes that individuals cast—driven not by the QV self-interested calculus (View I) but instead by social norms (View II).

A related question concerns the nature of the preferences that individuals hold and would express in casting their votes. The QV mechanism is designed to induce individuals to reveal their self-interested preferences, which the social planner cannot observe, in order to aggregate them and act accordingly.³⁰ This presents a bit of a conundrum when one

²⁹ Observe that individuals with greater wealth will place a greater dollar value on an outcome that delivers a common number of utils. For example, if preferences are additively separable and the utility of consumption is given by the log of consumption, individuals with ten times the wealth have ten times the willingness to pay for a benefit with common value (in utility). In such an instance, votes would, *ceteris paribus*, be proportional to wealth under pure-QV behavior, whereas in this example it is expenditures that are proportional to wealth.

³⁰ Individuals might instead vote based on their beliefs about the correct outcome from a social perspective. In that event, however, QV would not generate preference revelation but, instead, the revelation of uninformed individuals' guesstimates about others' preferences.

intersects decisions on turnout and the number of votes to cast. As we have explained, if an individual undertakes a self-interested calculus *ex ante*, she would appreciate that U^{piv} was essentially 0. With a positive C , such an individual would not turn out to vote—that is, unless she also has a sufficiently high D , which utility is social or moral in character. This creates a tension because one is voting in the first place as a consequence of having sublimated selfish concerns ($U^{piv} - C$) to social ones (D), but then one is spending money precisely to advance the selfish interests that cannot rationalize either showing up at all or spending more than, say, \$0.000025.³¹ Jill is thinking:

So I got out of bed early, drove to the polls, and waited in line to vote. Why? Because I'm not a selfish person; I'm a citizen, proudly performing my civic duty. I'm looking at those flags and smiling at the volunteers, basking in the glory of civic participation. Now, in the voting booth, at the peak of these emotions, I pull out my calculator, think selfishly as I compute how much I should spend, and charge that amount to my credit card.

Moreover, if Jill is an even slightly cognizant voter, she also appreciates that the likely impact of her vote, and her corresponding U^{piv} , is negligible. Hence, it is hardly obvious that preferences deriving from guilt, civic duty, and regard for social norms would best be satisfied by a voting expenditure calibrated to self-interest. Compare Goodin and Roberts (1975) and Margolis (1982a), who suggest that votes are guided by ethics or other social preferences rather than self-interest when U^{piv} is small.

Whatever combination of the foregoing hypotheses (and others) would actually guide decisions about votes and expenditures under QV, we should also contemplate how they might be calibrated. For example, if the norm is egalitarian, does that mean everyone spends \$1 for 1 vote, \$25 for 5 votes, or ... what? Will these norms be the same across groups? If they differ, will they be in terms of dollars or votes (keeping in mind that 10 times the dollars is only ~ 3.2 times as many votes)? And how might these amounts change over time? Might they cycle? (As groups—seeking to surpass their foes from the last contest—move from \$10 to \$20, and then higher and higher, might they at some point (\$50? \$100?) drop back to low levels, after which a new round of escalation would begin?) And so forth.

Taken together, the complexity and queasiness of hyper-rational, pure-QV voting, along with the miniscule expenditure that such voting involves, might push nearly everyone to View II behavior—and, as we explained, even if only a modest fraction behaves according to View II, the preferences of remaining QV'ers are drowned out. And how would View II individuals actually vote under QV? Familiar behavioral considerations favor focal-point voting (such as 1 vote for \$1) or other highly simplified rubrics, and the social nature of voting (which centrally explains why voters turn out in the first place) would likely exert a significant influence on which focal points or what expenditure levels individuals would choose. Taken together, the number of votes cast and thus voting outcomes under QV in large- N elections may be much closer to those under one person, one vote (and literally so if \$1, 1 vote is the focal point for everyone). However, the variety of social factors and how they may be influenced by the QV mechanism itself suggest other possibilities, making it

³¹ Another way to present this point is to begin with a simple model in which, under one person, one vote, many individuals turn out, induced by D . Suppose now that QV shifts the calculus of a significant portion of those individuals from a social one to a self-interested one. Although this subset would vote as QV'ers conditional on turning out, that same calculus would lead them not to turn out in the first place, leaving a smaller turnout set (a possibility suggested in Sect. 3.2) that is dominated by individuals who do not cast votes that reflect selfish preferences.

quite challenging to predict behavior under QV. Nevertheless, the foregoing analysis suggests that society is quite unlikely to behave according to the sort of equilibrium that would prevail under View I, wherein everyone is a pure QV'er, as in the theoretical models of the QV mechanism.

3.3.3 Affinity groups

We now consider another force that may significantly influence the number of votes cast by different individuals under QV. Affinity groups may have a qualitatively and quantitatively different influence on voting behavior and voting outcomes under QV from that under one person, one vote because they can affect not only turnout but also the number of votes that individuals casts. Moreover, as suggested in Sect. 3.3.2, the number of votes an individual casts can vary dramatically depending on whether she is a pure QV'er or adopts a different outlook. Regarding the latter, we suggested further that social norms may have a strong influence, but we did not say much about where such norms may come from for different groups of individuals.

An individual may feel connected to one or more groups of others, including groups of varying sizes that one does not know personally or interact with directly—from family to neighbors to coworkers to fellow congregants to members of the same faith, political parties, environmental organizations, gun owners, animal rights activists, Trekkies, or Cubs fans. Introspection, experience, and experiments teach us that such individuals often will be motivated to contribute to various sorts of public goods (contrary to narrow self-interest), as manifested by volunteerism, charitable donations, campaign contributions, and voting—the latter two behaviors being at the heart of our subject. To be sure, familiar obstacles such as the free-rider problem limit the extent of such activities. Yet, in the United States, individuals make charitable contributions of \$300 billion per year and in excess of 100 million turn out to vote in presidential elections. Moreover, a substantial portion of such activity—in terms of both magnitude and direction (which charity or candidate)—is probably influenced by an individual's felt connection with various affinity groups.

Under one person, one vote, affinity groups manifest themselves in terms of turnout (by influencing D) and expressed preferences (which candidate or other outcome will be selected, conditional on voting). The mechanisms include volunteerism, contributions, discussions, and statements of recognized leaders. One supposes that these influences have a moderate effect, perhaps by moving turnout rates a few percentage points here and there (on a base of, say, 50%)—generating some nonrepresentativeness of the turnout set—in addition to guiding how some voters cast their votes.

Under QV, the same phenomena would be present, but they now will operate on an additional margin: how many votes individuals cast (i.e., purchase). This, indeed, is the purpose of QV, but as we emphasize throughout, actual voting behavior is in fact much more driven by D than by U^{piv} —the simple and powerful point being that those with $D = 0$ are not in the turnout set to begin with.

The magnitude of plausible affinity group effects on the number of votes cast or the magnitude of expenditures individuals make (depending on which is focal) may well be large under QV and depends critically on whether View I or View II generally prevails. Suppose, for example, that some group of religious leaders tells their flock: “When you vote for the candidate who reflects our values, you surely should be spending at least as much as you put on the weekly collection plate.” And imagine that this is typically about \$10. (Note that average annual donations to religion by those who so give exceed \$1500, so

this is under 1% of that total.) Similarly, a celebrity, editorial writer, interest-group leader, blogger, candidate, or meddling economist might plant a similar impression. Against such possibilities, keep in mind our previous discussion to the effect that most individuals will feel at sea attempting to determine their V^* under pure QV, will be inclined to grab some focal point, and will most likely be guided by social norms in choosing some round number of votes to cast or dollars to spend. That is, we have both a push (for most, a shove) away from pure QV behavior as well as a pull—here, the pull of one or more affinity groups—toward something else.

If View I largely prevailed, wherein V^* might, as in our earlier example, be something like 0.005, then the ~ 3.2 votes cast by each individual induced by some group to spend \$10 would have an enormous influence. The battle among even a handful of such affinity groups—or a single Pied Piper—with followers collectively consisting of a few percent of the electorate would determine the outcome of the election, drowning out a large mass of pure QV'ers, pretty much regardless of the intensity of their preferences. Of course, who identifies with one or more affinity groups and the targets they successfully implant as social norms for their members will likely reflect preference intensity, but these will also reflect many other confounding effects. In addition, as explained in Sect. 3.3.1, once such groups cast significant numbers of votes, the recalculated V^* of pure QV'ers will fall even further; even when an election is likely to be close, P is very small once the vote total is large.³² In that event, predicting outcomes in large- N elections under QV would depend almost entirely on the analysis of affinity group behavior.³³

Under View II, affinity groups would be quite important, although they are no longer the entire story. Here we imagine, as a baseline, that most individuals are doing something like spending \$1 for 1 vote (rather than \$0.000025 for 0.005 votes). If affinity groups influence a few percent of the population to instead spend \$10 for ~ 3.2 votes, they will have a notable impact in close contests, but not otherwise.

However, a significant fraction of the population may be influenced by affinity groups. If some broad group can influence many of its members to spend, say, \$2 instead of \$1, purchasing ~ 1.4 votes rather than 1 vote, the collective impact would be large. We have

³² In light of the discussion in the text, we do not regard the analysis of Weyl (2017) on the ineffectiveness of small-group collusion under QV pertinent to affinity group behavior, which we see as qualitatively different. First, it is easy to imagine thousands, even millions, spending \$10 rather than \$0.000025, for the reasons presented. Second, Weyl (2017) argues that deviation is likely unless a colluding group is very small, but it seems unlikely that an individual commanded, say, by a religious leader to spend a mere \$10 would defect in a heartbeat, given the psychological and social consequences of going against the group. (The individuals behaving as part of an affinity group are, recall, already making campaign contributions, volunteering, making contributions to the group itself, and, importantly, incurring the cost C to vote in the first place.) Third, unlike the analysis in Weyl (2017), growth in N need not dilute affinity groups' influence because the number and size of affinity groups will tend to increase with N as well. Finally, the large increase in the number of votes cast as a consequence of affinity groups will cause P to fall, reducing the V^* of the remaining QV'ers, and if they predict a skewed outcome (even if contrary to their perhaps intense preferences), their perceived P , and hence their V^* , will fall even further. In light of these latter points, a larger N further dilutes the influence of even a large portion of voters who are pure QV'ers without diluting the impact of affinity groups. Although there may exist assumptions on limited collusion that will generate equilibrating tendencies in a world where all but a few individuals are pure QV'ers, such a world would be sharply altered by affinity groups, along with other factors that we have noted.

³³ Once there is some such affinity-group behavior, the calculus for remaining QV'ers also changes qualitatively. Instead of contemplating subtle properties of limiting equilibrium behavior under particular assumptions about common knowledge and features of probability distributions over preferences, one would be studying social psychology and following certain individuals' Facebook and blog posts. One still would have to process this information in the form of having probability distributions over how many will follow which affinity group, in what manner, and so forth.

suggested that individuals probably would choose round-number focal points, but whether that would be 1 vote, 2 votes, or 5 votes—or instead \$1, \$5, or \$10—is hard to say a priori, but is likely to be influenced greatly by social norms. Here we are adding that affinity groups will now devote much of their electioneering to boosting the voting norm for their followers, while trying simultaneously not to drive up that of their opponents.³⁴ The battle among focal points across various affinity groups could be the most important determinant of election outcomes.

Accordingly, prediction of QV outcomes depends importantly on understanding affinity group behavior. As mentioned, this phenomenon is also central to understanding outcomes with one person, one vote. However, because under QV affinity groups also affect the number of votes that individuals cast, they can have a much more powerful influence. Indeed, this difference could be the most important impact of moving from one person, one vote to QV.

3.3.4 Multi-contest elections

Until now, we have supposed (as in much of the literature on voting, including the core QV papers) that in each election there is a single choice: one office or one referendum. In many political systems, however, it is common to consolidate multiple offices and referendums into a single election. These choices will vary tremendously in perceived benefits and in the probability of pivotality.

If one makes various simplifying assumptions, the combination of multiple decisions into a single election under one person, one vote may affect turnout (in our view, primarily by raising D) but would not have a direct influence on how individuals would vote conditional on turning out. Likewise, in a pure, hyper-rational QV world (View I), individuals who turn out would, as before, cast V^* votes, using the appropriate B and P , in each contest.³⁵ However, under View II, the analysis could change substantially.

Necessarily engaging in another layer of speculation—in light of the lack of evidence and consideration of a regime that may qualitatively influence turnout and votes cast even with single-contest elections—we suggest some possibilities. Individuals might, to a significant extent, treat a single election, despite the many contests on the ballot, as in one mental account. Perhaps they would allocate a budget—maybe \$10 or \$50—to the election as a whole. Or, if permitted to check their total expenditures before finalizing all of their votes, they might adjust their individual votes if the totals deviate from a gestalt budget. (For example, if they cast 10 votes each on the first three ballot items, with seventeen remaining, and then realize that their partial bill is \$300, they might cut back dramatically.) Another plausible behavior would be to spend \$0 on most items, particularly if the list is long; this effect would be especially pronounced for items later on the list if voters allocate votes from a budget in sequence and if they neither look ahead to later ballot items nor go back to adjust votes on early ones.

Note that this sort of behavior will generate some tendency to reflect intensity: *ceteris paribus*, more of one's budget would be allocated to more important items. Of course, it

³⁴ Will “Five (dollars) for Fran” beat “Two for Teresa” or lose to “Ten for Teresa”?

³⁵ Of course, there may well be interdependencies: P in one race may be influenced by P in another, and the B s may depend on outcomes of other races. The need to determine the equilibrium for each contest, with such equilibria being interdependent, massively complicates the pure-QV calculus in such settings, but these considerations are second-order for our purposes because, as suggested earlier, we believe that View II would prevail.

may also reflect guesstimates of pivotality in some loose sense: nearly everyone might spend \$0 on most of the contests that do not seem close, which would blunt the degree to which intensity is reflected.

The phenomena just described could turn out to be dangerous. Suppose that a few crazy proposals make it on the ballot. One possibility is that nearly everyone spends \$0 so as to avoid wasting their budgets, in which event some of the crazy propositions might actually pass (perhaps because of the behavior of some affinity group)—even if they would have had no chance under one person, one vote or other schemes or, indeed, under hyper-rational pure-QV behavior.³⁶ Another possibility is that such proposals would draw most of individuals' expenditures in order to ensure that none would pass, leaving little room for preference expression regarding the core questions on the ballot, perhaps generating peculiar outcomes if many otherwise have exhausted their budgets. Note further that these possibilities create substantial incentives for agenda manipulation (ballot cramming, in contrast to ballot stuffing).

Our brief suggestions here only begin the inquiry. They would also have to be intersected with the psychological and social considerations that we have elaborated previously, including affinity group behavior. This subsection merely suggests that multi-contest elections further complicate the challenge of predicting voters' likely behavior under QV.³⁷

4 Conclusion

Under QV—like under one person, one vote—pivotality-based utility, U^{piv} , tends to be negligible for most (perhaps all) individuals in large- N elections. Under one person, one vote, $U^{piv} = PB$, whereas under QV, we have $U^{piv} = P^2B^2/4$. In our simple, toy example in which $P = 0.000001$ (one in a million) and $B = \$10,000$, $U^{piv} = \$0.01$ under one person, one vote and $U^{piv} = \$0.000025$ under QV. An extremist with $B = \$100,000$ has values of \$0.10 and \$0.0025, respectively.

In models such as those of Ledyard (1984) and Myerson (2000) that introduce only the cost of voting, C , U^{piv} does matter, only individuals with extremely low C s turn out in large- N elections, and the total number of votes cast is proportional to intensity under one person, one vote (yielding the utilitarian outcome)—but it is proportional to the cube of intensity under QV, vastly overweighting extremists. These models have negligible turnout, so we believe it more useful in order to understand both voter turnout and votes cast under QV to build on the Riker and Ordeshook (1968) model in which the all-in utility from voting takes the form $U^{piv} - C + D$.

Regarding turnout, as developed in our broader work (Kaplow and Kominers in preparation), all of the action is in C and D . And, as we have explained here, these results largely carry over to QV: When both C and D are uncorrelated with preferences, the level of turnout is irrelevant: as long turnout is positive, the turnout set is representative. If one admits any correlation between preferences, C , and D , then anything is possible; that is, the

³⁶ This last point is important: if various of the necessary assumptions hold and everyone solves the hyper-rational pure-QV equilibrium, the masses against would cast enough slivers of votes so that the few who favor the proposal would almost never find it optimal to purchase enough votes to prevail. But once many of those tiny slivers are rounded to \$0, as in the View II world, defeat of such proposals is hardly assured.

³⁷ We have also focused on votes cast given turnout, whereas multi-contest elections may also influence how voters think about turnout.

distribution of preferences in the turnout set can be anything, without regard to the distribution of preferences in the population. Turnout sets and voting outcomes depend on the specific relationship between preferences, C , and D . Systematic skews can arise in straightforward ways (for example, if individuals with pro-left preferences tend to have lower C s), but there is a tendency for noise in any such correlations to produce a more representative turnout set.

In light of our analysis, we consider a number of ways in which C and D may differ under QV. (Recall that we include the cost of purchasing the votes themselves as a component in U^{piv} , not in C , which we take to be the cost of turning out per se.) For example, the effort and anxiety associated with determining how many votes to cast under QV may raise C , and the requirement to purchase one's votes under QV may reduce D . We emphasize, however, that even though both effects may suppress turnout, what matters for the representativeness of the turnout set and hence for voting outcomes are possible systematic biases: relationships between these sorts of effects and preferences. We speculate on some possibilities, but largely leave the matter for further study.

We also give a number of reasons why votes cast under QV could differ dramatically from the predictions of the core QV literature. The incredible complexity of divining the equilibrium (which is necessary in order to calculate P and hence to choose how many votes to cast) combined with the refined nature of the implications (spending \$0.000025 in our example) would strongly push (nearly) everyone substantially away from hyper-rational pure-QV equilibrium behavior (which we call View I) in large- N elections. In synergy with this push, a number of forces pull individuals toward specific alternative forms of behavior, which we refer to as View II. Rounding and focal points would become central: if the true (but impossible to calculate) optimal expenditure was anything like \$0.000025, to purchase 0.005 votes, individuals might just spend \$1 (40,000 times too much) to purchase 1 vote (200 times too many). If most voters who turn out did so, the results would be much like those under one person, one vote.

We also explain that even if only a modest fraction of individuals behaved in such vaguely normal ways, P would fall greatly (from an already low level) and any remaining pure-QV'ers would find it optimal to further reduce their votes cast, leading them (including extremists) to be drowned out. That is, it does not take much ordinary behavior to render pure-QV'ers essentially irrelevant. In practice, we would expect essentially everyone to be in the former group in any event.

Additional factors may govern individuals' decisions about how many votes to cast. In light of the fact that D drives turnout to begin with (recalling that $U^{piv} - C$ is negative for almost everyone in large- N elections), we suggest that the determinants of D , which are social and psychological, would be first-order determinants of votes cast under QV. The manner in which these factors may skew voting, relative to pure QV, will again depend on the relationship to underlying preferences. In addition, we question whether voters under QV would indeed cast votes reflective of their narrow preferences—generating the sort of preference revelation that motivates QV—because such may be inconsistent with the reason that one is voting in the first place (again, D greatly dominates U^{piv}).

An important force influencing how many votes individuals would cast is the activity of affinity groups. If religious leaders, celebrities, or spokespeople for various causes are able to induce their followers to spend, say, \$10 rather than \$1, members of those groups will each cast ~ 3.2 votes rather than 1, more than tripling their influence on the outcome. Under both QV and one person, one vote, all manner of social influences can affect turnout somewhat, but under QV they can also affect votes cast given turnout, perhaps substantially.

For this reason and others (including the presence of multi-contest elections), understanding how social interactions relate to voting is likely to be extremely important in predicting voting behavior and voting outcomes under QV. We also explain that the subtle nature of pure-QV equilibrium play in large- N elections is unlikely to pertain to many (if any) voters, and even if it did, those voters' voices would be drowned out by the other forces we describe.

We consider psychological and social considerations—rather than hyper-rational analysis that determines one's U^{piv} —to likewise be central to understanding voter behavior under other voting mechanisms, notably, one person, one vote. Accordingly, a qualitatively different research agenda is required if we are to advance our understanding of voting behavior and voting outcomes under existing and proposed voting procedures.

Acknowledgements We thank Ryan Bubb, Glen Weyl, and participants at the conference on Quadratic Voting and the Public Good for comments; Jimin He, Andrea Lowe, Andrew Paik, and John Rady for research assistance; and National Science Foundation grants CCF-1216095 and SES-1459912, the John M. Olin Center for Law, Economics, and Business at Harvard University, and the Harvard Milton Fund for financial support. *Disclaimer:* Kaplow occasionally consults on antitrust cases, and his spouse is in the legal department of a financial services firm.

References

- Ariely, D., Bracha, A., & Meier, S. (2009). Doing good or doing well? Image motivation and monetary incentives in behaving prosocially. *American Economic Review*, *99*, 544–555.
- Chamberlain, G., & Rothschild, M. (1981). A note on the probability of casting a decisive vote. *Journal of Economic Theory*, *25*, 152–162.
- Cox, G. W., & Munger, M. C. (1989). Closeness, expenditures, and turnout in the 1982 U.S. House elections. *American Political Science Review*, *83*, 217–231.
- DellaVigna, S., List, J. A., Malmendier, U., & Rao, G. (2016). Voting to tell others. *Review of Economic Studies*, *84*, 143–181.
- Downs, A. (1957). *An economic theory of democracy*. New York: Harper and Row.
- Enos, R. D., & Fowler, A. (2014). Pivotality and turnout: Evidence from a field experiment in the aftermath of a tied election. *Political Science Research and Methods*, *2*, 309–319.
- Frey, B. S., & Jegen, R. (2001). Motivation crowding theory. *Journal of Economic Surveys*, *15*, 589–611.
- Funk, P. (2010). Social incentives and voter turnout: Evidence from the Swiss mail ballot system. *Journal of the European Economic Association*, *8*, 1077–1103.
- Gelman, A., Silver, N., & Edlin, A. (2012). What is the probability your vote will make a difference? *Economic Inquiry*, *50*, 321–326.
- Geys, B. (2006). Explaining voter turnout: A review of aggregate-level research. *Electoral Studies*, *25*, 637–663.
- Gneezy, U., Meier, S., & Rey-Biel, P. (2011). When and why incentives (don't) work to modify behavior. *Journal of Economic Perspectives*, *25*(4), 191–210.
- Goeree, J. K., & Zhang, J. (2017). One man, one bid. *Games and Economic Behavior*, *101*, 151–171.
- Goodin, R. E., & Roberts, K. W. S. (1975). The ethical voter. *American Political Science Review*, *69*, 926–928.
- Harbaugh, W. T. (1996). If people vote because they like to, then why do so many of them lie? *Public Choice*, *89*, 63–76.
- Iyengar, S., & Lepper, M. (2000). When choice is demotivating: Can one desire too much of a good thing? *Journal of Personality and Social Psychology*, *79*, 995–1006.
- Kaplow, L., & Kominers, S. D. (in preparation). On the representativeness of voter turnout.
- Krishna, V., & Morgan, J. (2015). Majority rule and utilitarian welfare. *American Economic Journal: Microeconomics*, *7*, 339–375.
- Lalley, S. P., & Weyl, E. G. (2016a). An online appendix to “quadratic voting” (unpublished).
- Lalley, S. P., & Weyl, E. G. (2016b). Quadratic voting (unpublished).
- Ledyard, J. O. (1984). The pure theory of large two-candidate elections. *Public Choice*, *44*, 7–41.
- Margolis, H. (1982a). *Selfishness, altruism, and rationality*. Chicago: University of Chicago Press.

- Margolis, H. (1982b). A thought experiment on demand-revealing mechanisms. *Public Choice*, 38, 87–91.
- Mulligan, C. B., & Hunger, C. G. (2003). The empirical frequency of a pivotal vote. *Public Choice*, 116, 31–54.
- Myerson, R. B. (2000). Large Poisson games. *Journal of Economic Theory*, 94, 7–45.
- Riker, W. H., & Ordeshook, P. C. (1968). A theory of the calculus of voting. *American Political Science Review*, 62, 25–42.
- Shampanier, K., Mazar, N., & Ariely, D. (2007). Zero as a special price: The true value of free products. *Marketing Science*, 26, 742–757.
- Tversky, A., & Shafir, E. (1992). Choice under conflict: The dynamics of deferred decision. *Psychological Science*, 3, 358–361.
- Weyl, E. G. (2015). The efficiency of quadratic voting in finite populations (unpublished).
- Weyl, E. G. (2017). The robustness of quadratic voting. *Public Choice*. doi:[10.1007/s11127-017-0405-4](https://doi.org/10.1007/s11127-017-0405-4).