Available online at www.sciencedirect.com

ScienceDirect

Journal of Economic Theory 183 (2019) 384–416

JOURNAL OF Economic Theory

www.elsevier.com/locate/jet

On the consensus effect ☆

David Dillenberger [a,*], Collin Raymond [b]

[a] *Department of Economics, University of Pennsylvania, USA*
[b] *Department of Economics, Purdue University, USA*

Received 22 December 2017; final version received 22 June 2019; accepted 3 July 2019
Available online 10 July 2019

Abstract

Individuals often tend to conform to the choices of others in group decisions, compared to choices made in isolation. We show that this behavior — which we term the consensus effect — is equivalent to a well-known violation of expected utility, namely strict quasi-convexity of preferences, which is shared by many popular non-expected utility models. In contrast to the equilibrium outcome when individuals are expected utility maximizers, quasi-convexity of preferences imply that group decisions may fail to properly aggregate preferences and strictly Pareto-dominated equilibria may arise.

JEL classification: D71; D80; D81

Keywords: Aggregation of preferences; Choice shifts in groups; Consensus effect; Non-expected utility; Quasi-convex preferences

☆ This paper was previously titled "Group-Shift and the Consensus Effect" and "Mixture Aversion and the Consensus Effect". We thank Kfir Eliaz, Andrew Ellis, Andrei Gomberg, Paul Heidhues, David Huffman, Gilat Levy, Yusufcan Masatlioglu, Ronny Razin, Chloe Tergiman, Peter Wakker, Jingni Yang, and especially Debraj Ray for useful comments. An editor and two anonymous referees provided valuable comments that improved the paper significantly. We also thank Kelly Twombly of Amherst College for her research assistance. Part of this work was done while Dillenberger was visiting the Department of Economics at NYU; he is grateful to this institute for its hospitality.

* Corresponding author.
E-mail addresses: ddill@sas.upenn.edu (D. Dillenberger), collinbraymond@gmail.com (C. Raymond).

https://doi.org/10.1016/j.jet.2019.07.002

1. Introduction

Group decision-making is ubiquitous in social, economic, and political life. Empirical evidence suggests that individuals tend to make different choices depending on whether the outcome of interest is a result of their choice alone or also the choice of others in a group. In particular, the existing evidence largely supports the idea that these choice shifts in groups, which are prominent in a variety of contexts across fields, are predicted by the expected choice of the majority of individuals. The phenomena that have been documented include the bandwagon effect in political science (e.g., Goidel and Shields, 1994; Niemi and Bartels, 1984, and Bartels, 1988); risky and safe shifts studied by psychologists (e.g., Brown, 1986; Stoner, 1961, 1968; Nordhöy, 1962; and Pruitt, 1971); and severity and leniency shifts in legal studies (Schkade et al., 2000; Sunstein et al., 2002; Sunstein, 2005). As an influential early article in sociology by Granovetter (1978) summarized it, "collective outcomes can seem paradoxical — that is intuitively inconsistent with the intentions of the individuals who generate them."

Models of group decisions typically analyze either private-value or common-value settings. Because, as will be explained below, with expected utility preferences in a private-value setting we should not observe choice shifts, much of the literature exploring choice shifts has focused on the common-value setting. In this context, group decisions aggregate private information regarding the relative value of possible outcomes.[1] In contrast, in this paper we maintain a private-value setting, but relax the assumption of expected utility. In particular, we show that well-known violations of expected utility can explain these commonly observed choice shifts, even in settings without private information.[2] Thus, our paper joins a literature discussing how relaxations of the main assumptions of expected utility can have important implications for behavior in strategic situations, as in auctions (Karni and Safra, 1989; Neilson, 1994; Nakajima, 2011; Baisa, 2013; Eisenhuth, 2019), pricing by firms (Heidhues and Kőszegi, 2008, 2014; Carbajal and Ely, 2016, Rosato, 2016), and incentives schemes (Herweg et al., 2010, Carbajal and Ely, 2012).

To see why a violation of expected utility may generate choice shifts in groups, note that an individual choice in a group decision matters only when that individual is pivotal, that is, when his vote actually changes the outcome. However, from an ex-ante perspective, when choosing for which option to vote, an individual does not know whether or not he will be pivotal. Thus, his choice is not a choice between receiving Option 1 or Option 2 for sure, but rather between *lotteries* defined over these two options — where if the individual turned out to be pivotal his selected option will be implemented, and otherwise the probability of each alternative to win depends on the probability that the group chooses it conditional on him not being pivotal. Violations of the independence axiom of expected utility imply that an individual may prefer Option 1 to Option 2 in isolation, yet prefer the lottery induced in the group context by choosing Option 2 over the one induced by choosing Option 1, thus accounting for the aforementioned choice shift.

[1] This literature, typified by Feddersen and Pesendorfer (1997), focuses on the ability of group decisions to aggregate private information rather than preferences. In Section 4.1 we contrast our findings with theirs as well as the larger literature on information aggregation in groups.

[2] Many lab experiments control for private information. Empirically, Sunstein (2005), who reported shifts toward the majority option in the context of mock juries, found that "when a majority of individuals initially favored little punishment, the jury's verdict *[...]* was systematically lower than the median rating of individual members before they started to talk with one another." When referring to these findings, Eliaz et al. (2006) argued that those shifts cannot be explained by asymmetric information, as all juries arrive to the trial without any prior information and they all receive the exact same evidence presented to them. And while juries may differ in their interpretation of the evidence, the authors argue that any such asymmetry "will most likely wash out in the deliberation process."

In Section 2 we formally link violations of expected utility with the phenomenon of choice shifts in groups. In doing so, we provide a relationship between two types of non-standard behavior, one observed at the individual level and one at the group level. Our first result states that individuals have preferences that are strictly *quasi-convex* in probabilities if and only if they will systematically exhibit what we call a *consensus effect* — an individual who is indifferent between two options when choosing in isolation will actually strictly prefer to vote for the option that is sufficiently likely to be chosen by the group. As discussed, the consensus effect captures the stylized fact that in group contexts individuals want to exhibit preferences that match those of the group as a whole. Consistent with the predictions of our model, Agranov et al. (2017) find evidence that individuals are more likely to vote for an outcome if they perceive it as more likely to win. Quasi-convexity, on the other hand, is a well established preference pattern in decision making under risk, according to which individuals are averse toward randomization between equally good lotteries.[3] Popular models of preferences over lotteries which can exhibit quasi-convexity include rank-dependent utility (Quiggin, 1982, hereafter RDU), quadratic utility (Chew et al., 1991), and Kőszegi and Rabin (2007)'s choice acclimating personal equilibrium model of reference-dependence. Moreover, as observed by Machina (1984), quasi-convexity occurs if, as in common in many applications such as insurance purchasing, before the lottery is resolved the individual is allowed to take an action that determines his final utility. As long as the optimal decision is affected by a change in the probabilities, the induced maximum expected utility will be convex in the probabilities, meaning that even if the underlying preferences are expected utility, induced preferences over the 'optimal' lotteries will be quasi-convex.

To gain some intuition for the link between quasi-convexity in probability mixtures and the consensus effect, consider Kőszegi and Rabin (2007)'s model. Suppose an individual is indifferent between either knowing for sure Option 1 is chosen, or knowing for sure Option 2 is chosen. For this to be true, each option has some benefits and some drawbacks relative to the other. However, if the individual expects that Option 1 will be chosen, and ends up with Option 2, the relative drawbacks will loom larger than the relative benefits (because of loss aversion relative to the reference point, which is Option 1). Thus, if a voter thinks Option 1 will often be chosen when they are not pivotal, they strictly prefer to ensure that it will also be chosen when they are pivotal, in order to align outcomes with expectations. We formalize this intuition in Sections 2.2 and 2.3, where we point to a deep connection between the notions of reference dependence, loss aversion, and the consensus effect.

To expand the applicability of our results, we further demonstrate how they extend to models of (i) globally quasi-concave preferences, where the opposite group behavior is predicted; (ii) preferences with both quasi-convex and quasi-concave regions, on which they can be applied locally; and (iii) behavior that may not be captured by maximizing a single preference relation.

In an earlier paper on choice shifts in groups, Eliaz et al. (2006, hereafter ERR) used the same model of group decision making but focused on group choices between particular pairs of options, safe and risky, where the former is a degenerate lottery that gives a certain outcome with probability one. They confined their attention to RDU preferences and established an equivalence

3 Our proof shows that having quasi-convex preferences is equivalent to adopting a "threshold" rule towards the level of support that others will exhibit for any given option (i.e., the probability that any given option is chosen when a voter is not pivotal). When the level of support for an option exceeds the threshold, the individual will strictly prefer to choose it in a group situation. These thresholds have similar intuition to the reasons provided for similar consensus effects in other fields; for example, Granovetter (1978) specifically discusses the effect thresholds will have on aggregate versus individual behavior.

between specific types of choice shifts and Allais paradox, one of the most documented violation of expected utility at the individual level. Since choice shifts in groups are observed in experiments even when all lotteries involved are non-degenerate, our results suggest that the choice shifts discussed in ERR are actually manifestations of the consensus effect. In Section 2.4 we relate our results to theirs. We extend their results for RDU preferences, but, more importantly, also demonstrate why the link to Allais paradox is restricted to that specific class of preferences. In particular, the consensus effect is in general consistent not only with Allais-type behavior but also with the opposite pattern of choice and, similarly, Allais-type behavior does not rule out the anti-consensus effect.

In Section 3 we analyze what type of equilibrium behavior results from quasi-convex preferences in conjunction with strategic considerations. We describe a majority voting game as a collection of individuals, each of whom has one vote to cast in favor of option p or option q (no abstentions are allowed). After observing their own preferences (which are drawn i.i.d. from some known distribution), but no other information, individuals vote. Whichever option receives the majority of the votes is implemented.

Since individuals with quasi-convex preferences do not like to randomize, voting games take on the properties of *coordination games*. These individuals benefit from coordinating their votes with others because it reduces the "randomness" in the election. They typically face a tradeoff between having the option they prefer selected and reducing the uncertainty regarding the identity of the chosen outcome.

We prove the existence of an equilibrium and describe the main properties of any possible equilibrium. When individuals exhibit the consensus effect, group decisions may fail to aggregate preferences properly because voters are willing to coordinate on either option, rather than voting for the option they prefer in isolation. Thus, strictly Pareto-dominated equilibria may result. This willingness to coordinate implies that our model features non-uniqueness of equilibrium not due to randomization by indifferent types (as in the expected utility case) but rather because of weak-preference reversals. We discuss conditions under which we would expect to see such preference reversals and how they relate to whether the equilibrium is unique or not. We further show that some individuals *necessarily* exhibit strict preference reversal when the group becomes large.

In Section 4 we discuss how our model relates to, and can be distinguished from, alternative models in the literature on voting, including costly voting and common value settings. For example, individuals with quasi-convex preferences may be unwilling to pay the cost of voting even if they know they will be pivotal, but may be willing to do so when they have a smaller chance of being pivotal but their vote can help reduce the randomness of the election. The key conceptual difference from a common value setting is that there uninformed independent voters who want to choose the best candidate will tend to vote against — or to compensate for — a majority that is formed of partisans, while uninformed voters with quasi-convex preferences will tend to vote in accordance with this partisan majority.

Other approaches to conformity typically add an (additive) exogenous conformity benefits term to an otherwise standard model. One way to view our contribution is to take violations of expected utility in decision making under risk as descriptively valid and analyze — without tying our hands to any specific functional form — to what extent non-expected utility models can generate new predictions in the context of group decisions. Thus, our analysis provides a non-expected utility foundation for group choice anomalies. In particular, as we have discussed above, since the expected choice of the group serves as a reference point when the individual is deciding how to make his own choice, our intuitions are closely related to some of the recent models of expectation-based reference dependence (as in Kőszegi and Rabin, 2007).

2. The consensus effect and quasi-convex preferences

2.1. Model

Our aim is to link an individual's private ranking of objects with his ranking of these same objects in a group context. We assume that any individual has preferences over simple lotteries. Formally, let X be the set of outcomes (which is assumed to be a compact metric space) and denote by Δ the set of lotteries with finite support over X. We identify an individual with his complete, transitive, and continuous preference relation $\succcurlyeq$ over Δ, which is represented by some monotone function $V : \Delta \to \mathbb{R}$.[4] Throughout the paper we denote by x, y, z generic elements of X and by p, q, r generic elements of Δ.

In describing group decision problems, we extend the model suggested by ERR (see Section 2.4). There is a group of N individuals. We identify a group decision problem as perceived by any individual i with a quadruple (p, q, α, β), consisting of two lotteries $p, q \in \Delta$ and two scalars $\alpha \in (0, 1)$ and $\beta \in [0, 1]$; α is the probability that individual i's decision is pivotal in choosing between p and q, and β is the probability that the group chooses p conditional on i *not* being pivotal.[5] For now, both α and β are exogenous and fixed; accordingly, we can interpret the choice from any such quadruple as determining an individual's best-response function. In Section 3 they will be derived as part of the equilibrium analysis. Note that the alternatives we consider are lotteries. For example, in a voting context we associate a candidate with a lottery over policies.[6]

Remark. Since we require preferences to be monotone, either p or q must be non-degenerate. However, our results would go through even if we allow alternatives to be final (i.e., degenerate) outcomes, at the cost that preferences would violate first-order stochastic dominance. We focus on monotone preferences in order to — as will be clearer in Sections 2.2 and 2.4 — more cleanly relate our results to the results of ERR and to standard models of non-expected utility preferences that explicitly impose monotonicity.

If, in the group context, the individual votes for q, the effective lottery he faces is the convex combination of p and q, given by:[7,8]

$$q^* = \alpha q + (1-\alpha)\left(\beta p + (1-\beta) q\right) = [\alpha + (1-\alpha)(1-\beta)]\, q + (1-\alpha)\beta p.$$

[4] Monotonicity means that $V(p) \geq V(q)$ whenever p first-order stochastically dominates q; the stochastic dominance order is with respect to the induced relation $\overline{\succcurlyeq}$ on X, defined by $x \overline{\succcurlyeq} y \iff \delta_x \succcurlyeq \delta_y$, where for any $z \in X$, δ_z is the Dirac measure at z.

[5] We omit the index i till Section 3, where we explicitly study strategic interactions between members of the group.

[6] This framework incorporates not just the typical election framework (where the threshold for adopting a policy may be majority rule or some other number), but also other types of group decision-making that feature some ex-ante uncertainty; for example, random serial dictatorship, where it is not clear at the time of choice who will be the dictator.

[7] For $p, q \in \Delta$ and $\lambda \in (0, 1)$, $\lambda p + (1-\lambda) q \in \Delta$ yields any $x \in X$ with probability $\lambda p(x) + (1-\lambda) q(x)$.

[8] We assume the reduction of compound lotteries axiom to only analyze single-stage distributions. This assumption is plausible in our framework, where voters face two kinds of uncertainty: whether they are pivotal, and what the outcome of the vote was. Both types of uncertainty are resolved in a standard voting environments at the same time, thus ruling out typical motives for not reducing compound lotteries (e.g., a preference for early resolution of information).

And if the individual votes for p, the effective lottery he faces is:

$$p^* = \alpha p + (1-\alpha)\,(\beta p + (1-\beta)q) = (1-\alpha)(1-\beta)q + [\alpha + (1-\alpha)\beta]p.$$

A choice shift is thus the joint statement of $p \sim q$ but $q^* \succ p^*$ or $q^* \prec p^*$.[9,10]

Our definition of the *consensus effect* below suggests a specific type of choice shift, whereby an individual tends to draw towards what others would do in the absence of him being pivotal. In particular, it captures the idea that if other members of the group are likely enough to choose p when the individual is not pivotal, then the individual himself will prefer to choose p as well.

Definition 1. The individual exhibits a consensus effect at (p, q, α, β^*) if $p \sim q$ and $\beta > \beta^*$ (resp. $\beta < \beta^*$) implies that $p^* \succ q^*$ (resp., $p^* \prec q^*$).

The individual exhibits the consensus effect if for all p, q, α with $p \sim q$, there exists β^* such that he exhibits the consensus effect in (p, q, α, β^*).

Anti-consensus effect at (p, q, α, β^*) and general anti-consensus effect are similarly defined. The threshold value β^* in Definition 1 is determined by preferences and is not necessarily equals 0.5. We describe the case where $\beta^* = 0.5$ for all p, q, and α as the *simple majority effect*. So if initially indifferent, the individual simply chooses the option he believes the group is most likely to choose when he is not pivotal. Proposition 2 of Section 2.2 characterizes the class of preferences that are consistent with the simple majority effect.

Since both p^* and q^* are convex combinations of p and q, if $\succcurlyeq$ satisfies the following *betweenness* property, $p \sim q$ implies $\gamma p + (1-\gamma)q \sim q$, then the individual will never display any choice shift in group. This property is weaker than the standard independence axiom,[11] which suggests that to accommodate such shifts, one needs to go beyond expected utility (or, more generally, beyond the betweenness class of preferences, suggested by Chew, 1983 and Dekel, 1986). To this aim, we consider the following two properties.

Definition 2. The preference relation $\succcurlyeq$ is strictly quasi-convex if for all $p, q \in \Delta$, with $p \neq q$, and $\lambda \in (0, 1)$,

$$p \sim q \Rightarrow \lambda p + (1-\lambda)\,q \prec p$$

and is strictly quasi-concave if

$$p \sim q \Rightarrow \lambda p + (1-\lambda)\,q \succ p.$$

Quasi-convexity implies aversion towards randomization between equally good lotteries; whereas quasi-concavity implies affinity to such randomization. (Betweenness preferences satisfy both weak quasi-convexity and weak quasi-concavity.)[12]

[9] Both p^* and q^* are functions of the group decision-problem, but for simplicity we will suppress the notation depicting this dependence.

[10] The consensus effect is defined where $p \sim q$. By continuity, the choice patterns that we study when the options are indifferent will persist even when one option is strictly preferred to the other.

[11] According to the independence axiom, $p \succcurlyeq q$ if and only if for any $r \in \Delta$ and $\gamma \in [0, 1]$, $\gamma p + (1-\gamma)r \succcurlyeq \gamma q + (1-\gamma)r$.

[12] The experimental evidence on quasi-convexity versus quasi-concavity is mixed. While it is a stylized empirical finding that betweenness is often violated, most of the experimental literature that documents violations of linear indifference

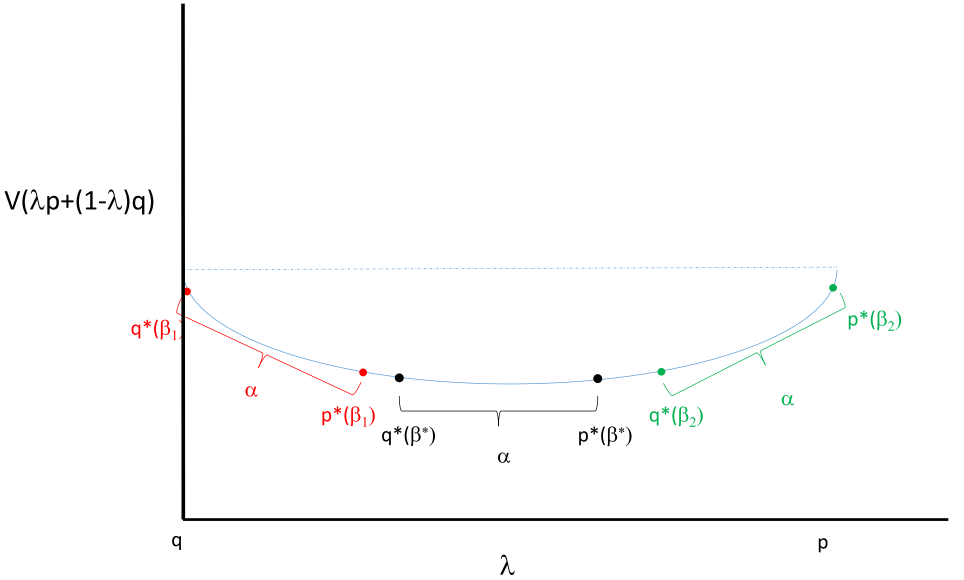

Fig. 1. $V(\lambda q + (1-\lambda)p)$ for $\lambda \in [0, 1]$.

Our main result links violations of expected utility at the individual level with a specific pattern of choices in group situations.

Proposition 1. *The preference relation $\succcurlyeq$ is strictly quasi-convex (resp., strictly quasi-concave) if and only if the individual exhibits the consensus (resp., anti-consensus) effect.*

All proofs are in the Appendix. To see the intuition behind Proposition 1, observe that p^* is always closer to p and q^* is always closer to q, with $p^* - q^* = \alpha(p - q)$. In addition, preferences are quasi-convex if and only if they are *single-troughed* between p and q. The proof is established by noting that an increase in β moves both p^* and q^* closer to p. Fig. 1 plots the function $V(\lambda q + (1-\lambda)p)$ for $\lambda \in [0, 1]$. The lotteries p^* and q^* are depicted for three different values of β ($\beta_1 < \beta^*$, β^*, and $\beta_2 > \beta^*$). Given a fixed α, for $\beta >$(resp., $<$)β^* the two lotteries get closer to p (resp., q), while the distance between them remains intact. More generally, even when the individual is not initially indifferent between p and q, his ultimate choice reflects a tradeoff between his preferred outcome and the desire to avoid randomization across lotteries (i.e., he wants to choose extreme lotteries that are as close to p or q as possible). If the latter effect is strong enough, he may reverse his preferences in a group context.

---

curves (e.g., Coombs and Huang, 1976) found deviations in both directions, that is, either preference for or aversion to randomization. Camerer and Ho (1994) find support for a mixed pattern with quasi-convexity over gains and quasi-concavity over losses. A concrete example of the behavioral distinction between quasi-concave and quasi-convex risk preferences is the probabilistic insurance problem of Kahneman and Tversky (1979). They showed that, in contrast with experimental evidence, any risk averse expected utility maximizer must prefer probabilistic insurance to regular insurance. Sarver (2018) pointed out that this result readily extends to the case of quasi-concave preferences. In contrast, quasi-convex preferences can accommodate aversion to probabilistic insurance.

In Section 2.2 we demonstrate the applicability of Proposition 1 in the context of several well-known models, focusing on the case of quasi-convex preferences. We do so because stylized facts, as well as strong intuition, suggest that the consensus effect is much more prominent than the opposite effect. The results naturally extend, modulo standard reversal, to quasi-concavity.

In Section 2.3 we discuss two extensions of our primary result. First, although Proposition 1 is framed in terms of a global preference restriction, we show that the intuition also applies locally if preferences include both quasi- concave and quasi-convex regions. Second, we show how our results extend to environments where we only observe choices, which may not be rationalized by maximizing a preference relation. In particular, we show that a previously discussed notion of reference dependence defined on choice correspondences leads to the consensus effect.

2.2. Examples

We now discuss the implications of Proposition 1 for some popular non-expected utility models. In all these examples, preferences are defined over monetary lotteries, that is, the underlying set of outcomes is an interval $X \subset \mathbb{R}$.

Rank-Dependent Utility (RDU): Order the prizes $x_1 < x_2 < \ldots < x_n$. The functional form for RDU is:

$$V_{RDU}(p) = u(x_1) + \textstyle\sum_{i=2}^{n} g\left(\sum_{j \geq i} p\left(x_j\right)\right)\left[u(x_i) - u(x_{i-1})\right] \tag{1}$$

where the weighting function $g : [0, 1] \to [0, 1]$ is bijective and strictly increasing. If $g(l) = l$ then RDU reduces to expected utility.[13]

RDU preferences are quasi-convex if and only if the weighting function is convex (see Wakker, 1994). Convexity of the weighting function — which is also a necessary condition for risk-aversion within RDU — is typically interpreted as a type of pessimism: improving the ranking position of an outcome decreases its decision weight. This suggests the following corollary.

Corollary 1. *Suppose preferences are RDU. Then the individual is pessimistic (g is strictly convex) if and only if he exhibits the consensus effect.*

The consensus effect as in Definition 1 is weak, in the sense that it does not determine how likely it has to be that the group chooses p in the absence of the individual being pivotal. However, if we put more structure on preferences we can have stronger results. This motivates introducing the class of quadratic preferences.

Quadratic Utility: A utility functional is quadratic in probabilities if it can be expressed in the form

$$V_{\mathbb{Q}}(p) = \sum_{x}\sum_{y} \phi(x, y)\; p(x) p(y)$$

where $\phi : X \times X \to \mathbb{R}$ is symmetric. The quadratic functional form was introduced in Machina (1982) and further developed in Chew et al. (1991, 1994).

[13] Wakker (2010) offers an extensive treatment of RDU preferences under risk.

The following result shows that quadratic preferences are the *only* preferences that generate the simple majority effect ($\beta^* = 0.5$ for all p, q, and α). In other words, for these preferences the threshold value for the consensus effect coincides with a simple majority rule.

Proposition 2. *Preferences are strictly quasi-convex and can be represented by a quadratic functional if and only if the individual exhibits the simple majority effect.*

The proof of Proposition 2 relies on the observation that quadratic preferences imply mixture symmetry (Chew et al., 1991). The preference relation $\succcurlyeq$ satisfies mixture symmetry if for all $p, q \in \Delta$ and $\lambda \in [0, 1]$,

$$p \sim q \Rightarrow \lambda p + (1 - \lambda) q \sim \lambda q + (1 - \lambda) p.$$

Suppose $p \sim q$. Using mixture symmetry, we find (unique) two lotteries $\widehat{q}$ and $\widehat{p}$ such that $q^* \sim \widehat{q}$ and $p^* \sim \widehat{p}$. If $\beta < 0.5$, then p^* is a convex combination of q^* and $\widehat{q}$ and thus, by strict quasi-convexity, $q^* \succ p^*$. Similarly, if $\beta < 0.5$, then q^* is a convex combination of p^* and $\widehat{p}$ and thus $p^* \succ q^*$. If $\beta = 0.5$ then $p^* = \widehat{q}$ (and $q^* = \widehat{p}$).

The class of quadratic preferences is relatively large. Nevertheless, when confining attention to special cases, Proposition 2 allows us to link some known behavioral biases with the strong form of consensus effect, as reflected in the simple majority effect. The next popular functional form is a vivid example.

PERSONAL EQUILIBRIUM (KŐSZEGI AND RABIN, 2007): Kőszegi and Rabin (2007) develop several related notions of reference-dependent choice: Preferred Personal Equilibrium (PPE) and Choice Acclimating Personal Equilibrium (CPE), where the latter was independently introduced in Delquié and Cillo (2006).[14] Because the choices generated by PPE can violate the weak axiom of revealed preference, they cannot be represented as maximizing a single preference relation, and thus we consider PPE in Section 2.3 as an extension.

Under CPE, the value of a lottery p is

$$V_{\mathbb{CPE}}(p) = \underbrace{\sum_x u(x) p(x)}_{\text{consumption utility}} + \underbrace{\sum_x \sum_y \mu(u(x) - u(y))\ p(x) p(y)}_{\text{gain-loss utility}}$$

where u is an increasing utility function over final wealth and

$$\mu(z) = \begin{cases} z & \text{if } z \geq 0 \\ \kappa z & \text{if } z < 0 \end{cases}$$

is a gain-loss function with $0 \leq \kappa \leq 2$ denoting the coefficient of loss aversion. Loss aversion occurs when $\kappa \geq 1$. Masatlioglu and Raymond (2016) show that these preferences are the intersection of RDU and quadratic utility, and that they are quasi-convex if and only if $\kappa \geq 1$. Combining this observation with Proposition 2 above, yields the following result.

Corollary 2. *Suppose preferences have a representation $V_{\mathbb{CPE}}$. Then the individual is loss averse if and only if he exhibits the simple majority effect.*

---

[14] These preferences have been widely used in the behavioral literature (Heidhues and Kőszegi, 2008; Sydnor, 2010; Herweg et al., 2010; Abeler et al., 2011; Ericson and Fuster, 2011; Gill and Prowse, 2012; and Barseghyan et al., 2013.)

Corollary 2 links a notion of expectations-based reference dependence in individual choice with a similar notion (the consensus effect) in group choice. If the group is more likely to choose p than q when an individual is not pivotal, then this expected choice would naturally serve as a reference point when the individual is deciding how to make his own choice (which will only matter in the case where he is pivotal). This mirrors the underlying intuition often provided for a preference for conformity — it is a type of external (i.e., based on the actions of others) reference point.

2.3. Extensions

We now extend our results in two directions. These extensions expand the applicability of our results to additional models of economic behavior.

LOCAL ANALYSIS: As we have pointed out in the previous section, while we focus our discussion on preferences that are globally quasi-convex, the result of Proposition 1 holds also locally, whenever both options (as well as all convex combinations between them) lie in a region where all indifference curves have the same curvature. We illustrate this using a well known specification of RDU, in which the weighting function g has an "inverse S"-shape; it overweights small probabilities and underweights large probabilities.

Definition 3. Suppose $g : [0, 1] \to [0, 1]$ is continuous, differentiable, and strictly increasing weighting function. Then it is inverse-S shaped if there exists a $\bar{p} \in (0, 1)$ such that g is concave on the range $[0, \bar{p}]$ and convex on $[\bar{p}, 1]$.

Fix RDU preferences. We say that two lotteries p and q are *similarly good* (resp., *similarly bad*) with respect to these preferences if they both place high enough probability on the same best (resp., worst) outcome in their mutual support. Formally, for $i = 1, ..., n$, denote by x_{r_i} the rank-ordered outcomes in lottery r, ordered from best to worst. Lotteries p and q are similarly good if $x_{p_1} = x_{q_1}$ and $\min\{p(x_{p_1}), q(x_{q_1})\} > 1 - \bar{p}$; and they are similarly bad if $x_{p_n} = x_{q_n}$ and $\min\{p(x_{p_n}), q(x_{q_n})\} > \bar{p}$.

Proposition 3. *Suppose preferences are RDU with an inverse-S shaped weighting function. If both options are similarly good, then the individual exhibits the consensus effect. And if both options are similarly bad, then the individual exhibits the anti-consensus effect.*

CHOICE CORRESPONDENCES: There are models of choice that cannot be rationalized by the maximization of a single preference relation. Some of these models feature reference dependence in choice. We now formally extend the link between reference dependence and the consensus effect, we alluded to in Corollary 2, to such settings.

Let c be a choice correspondence on Δ.[15] Freeman (2019) introduces a property he calls Strong Reference Bias: If $p \in c(\{p, q\})$, then $p \in c(\{p, \lambda p + (1 - \lambda)q\})$ for all $\lambda \in (0, 1)$. Intuitively, the higher the exogenous probability the individual will end up with a certain option (p in this case), the stronger that option serves as a reference point and thus its desirability increases. This condition, as he explains, is consistent with behavior in experiments about reference depen-

[15] That is, $c : 2^{\Delta} \setminus \emptyset \rightrightarrows \Delta$ such that $c(A) \subseteq A$ for all $A \subseteq \Delta$.

dence, such as the one reported in Ericson and Fuster (2011). We consider the strict version of this property.

Definition 4. The choice correspondence c satisfies strict reference bias, if $p \in c(\{p, q\})$ implies $p = c(\{p, \lambda p + (1-\lambda)q\})$ for all $\lambda \in (0, 1)$.

Defining the consensus effect in the domain of choice correspondences (rather than preferences) requires to replace $p \sim q$ with $p, q \in c(\{p, q\})$ and $p^* \succ q^*$ with $p^* = c(\{p^*, q^*\})$.

Definition 5. The individual exhibits the consensus effect for choices, if for any triple $\alpha \in (0, 1)$ and $p, q \in \Delta$ with $\{p, q\} = c(\{p, q\})$, there exists a $\beta^* \in [0, 1]$ such that $\beta > \beta^*$ implies that $\{p^*\} = c(\{p^*, q^*\})$, and $\beta < \beta^*$ implies that $\{q^*\} = c(\{p^*, q^*\})$.

Proposition 4. *If c satisfies strict reference bias then the individual exhibits the consensus effect for choices.*

Note that unlike Proposition 1, here we only have a one-side implication rather than equivalence. Intuitively, the consensus effect only restricts behavior fixing α, but violations of strict reference bias can only occur when α is allowed to vary. If c were rationalized by a complete and transitive preference relation $\succcurlyeq$, then transitivity would allow us to link implications across different values of α. Indeed, in this case Proposition 4, in concatenation with Proposition 1, implies that c satisfies strict reference bias if and only if $\succcurlyeq$ satisfies strict quasi-convexity.

In order to demonstrate the applicability of this result, we consider an example of perhaps the most notable form of reference dependence that cannot be captured by a preference relation: Kőszegi and Rabin, 2007 Preferred Personal Equilibrium (PPE). It relies on similar intuitions as CPE, but features a different solution concept.

To illustrate, we first define the utility of lottery p given that lottery $\hat{p}$ is taken as the reference point:

$$V_{KR}(p|\hat{p}) = \sum_x u(x)p(x) + \sum_x \sum_y \mu(u(x) - u(y))\ p(x)\hat{p}(y)$$

where u is an increasing utility function over final wealth and μ is the piecewise linear gain-loss function as defined for CPE.

Given a choice set S, lottery p is in the set of personal equilibrium of S (denoted $PE(S)$) if $V_{KR}(p|p) \geq V_{KR}(\hat{p}|p)$ for all $\hat{p} \in S$. Intuitively, p is a PE if when it is used as a reference point, there is no other available option that gives higher utility than it. Lottery p is a preferred personal equilibrium of S (denoted $PPE(S)$) if $p \in PE(S)$ and for all $\hat{p} \in PE(S)$, $V_{KR}(p|p) \geq V_{KR}(\hat{p}|\hat{p})$. Thus, PPE represents a refinement of PE, found by applying the CPE formula discussed previously to the set of personal equilibria.

Freeman (2019) shows that given a piecewise linear gain-loss function PPE satisfies Strong Reference Bias. Our next result shows that it also satisfies strict reference bias, and consequently the consensus effect.

Proposition 5. *If $c(S) = PPE(S)$ for all S, then c satisfies strict reference bias.*

Combining the last two results shows that the consensus effect is tightly linked to the reference dependent behavior of the kind developed by Kőszegi and Rabin (2007).

2.4. Risky shifts, cautious shifts, and Allais paradox

In this section we focus on group choices between particular pairs of options, s(afe) and r(isky), where s is a degenerate lottery, that is, a lottery that yields a certain prize $x \in X$ with probability 1, and r is some nondegenerate lottery. A group decision problem is then (r, s, α, β). In this context, we refer to *risky shift* (resp., *cautious shift*) as the joint statement $r \sim s$ and $r^* \succ s^*$ (resp., $r^* \prec s^*$), where

$$r^* = [\alpha + (1-\alpha)(1-\beta)]r + (1-\alpha)\beta s$$

and

$$s^* = (1-\alpha)(1-\beta)r + [\alpha + (1-\alpha)\beta]s.$$

These shifts are clearly a subset of the more general shifts discussed under the consensus effect. For a particular r, s, and α, there exists a β^* where an individual always exhibits a risky shift for $\beta \leq \beta^*$ and a cautious shift for $\beta \geq \beta^*$ if and only if the individual exhibits the consensus effect at (s, r, α, β^*).

ERR used this setting and focused on RDU preferences (see Section 2.2). Below we generalize their contribution within RDU, but also demonstrate that their main message *is not* necessarily valid for other types of non-expected utility preferences. Segal (1987) showed that within RDU, a convex distortion function g in equ. (1) implies (and is implied by) behavior that accommodates a version of Allais paradox — also known as the common consequence effect — which is one of the most prominent evidence against expected utility. Formally, fix any three prizes $x_3 > x_2 > x_1$ and denote by (p_1, p_2, p_3) the lottery that yields the prize x_i with probability p_i. The following definition formalizes this notion of the Allais paradox.[16]

Definition 6. An individual exhibits the Allais paradox if for every pair of lotteries $(1-\alpha, \alpha, 0)$ and $(1-\beta, 0, \beta)$ with $\alpha > \beta$, $(1-\alpha, \alpha, 0) \sim (1-\beta, 0, \beta)$ implies $(1-\alpha-\gamma, \alpha+\gamma, 0) \succ (1-\beta-\gamma, \gamma, \beta)$ for all $\gamma \in (0, 1-\alpha]$.

Theorem 1 in ERR states that within RDU, an individual exhibits the Allais paradox if and only if for any $r \sim s$ and $\alpha \in (0, 1)$ there exists $\beta^* \in (0, 1)$ such that he exhibits risky (resp., cautious) shift if $\beta < \beta^*$ (resp., $\beta > \beta^*$). ERR thus suggest an equivalence between a commonly known violation of expected utility and a robust phenomenon in the social psychology of groups when choosing between risky and safe options. Because Allais-type behavior is equivalent to the convexity of the weighting function and therefore to quasi-convexity of preferences, it is also the case that within RDU we have additional equivalences, as the following corollary summarizes.

Corollary 3. *Consider the rank dependent utility model (equ. (1)). The following statements are equivalent:*

1. *An individual exhibits the Allais paradox*

[16] In Allais' original questionnaire, $x_3 = 5M$; $x_2 = 1M$, and $x_1 = 0$. Subjects choose between $A = (0, 1, 0)$ and $B = (0.1, 0.89, 0.01)$, and also between $C = (0, 0.11, 0.89)$ and $D = (0.1, 0, 0.9)$. The typical pattern of choice is the pair (A, D). Definition 6 is more general than the original paradox proposed by Allais, since it puts behavioral restrictions also when no certain outcome is involved.

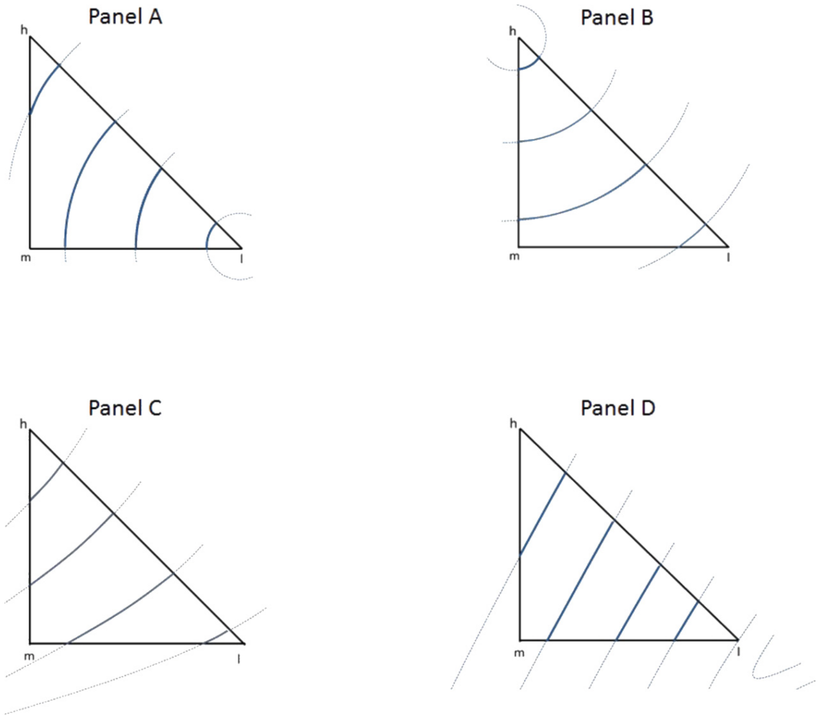

Fig. 2. Attitudes towards randomization and fanning properties of indifference curves.

2. *For all $r \sim s$ and α there exists β^* such that the individual exhibits the consensus effect at* (r, s, α, β^*)
3. *An individual's preferences satisfy quasi-convexity*
4. *An individual exhibits the consensus effect*

While these logical equivalences are quite strong (in the sense that they link specific behavior regarding r and s to arbitrary behavior for any p and q) and have an intuitive appeal (in that they link preferences for a risky versus safe option in the Allais questionnaire to similar preferences in group choice), they — as well as ERR's original results — are derived in the narrow context of RDU preferences. We will now argue that they are specific to that class and do not hold in general. In other words, empirical evidence that refutes RDU also challenges the aforementioned relationship between Allais-type behavior and consensus effects. The intuition, which we make more concrete in the two examples below and in Fig. 2, is the following: Definition 6 puts restrictions on how the slope of indifference curves change as we move *between them* in a specific direction in the probability triangle. Quasi-convexity and quasi-concavity, on the other hand, put restrictions on how the slope of a single indifference curve changes as we slide *along it*. In general, these two restrictions are independent.

To demonstrate this, first observe that the pattern of risky and cautious shifts discussed in ERR is implied by the consensus effect. Thus, in constructing our examples, we show that both quasi-convexity and quasi-concavity are consistent with both Allais-type behavior and with the opposite pattern of individual choice. We further note that any lottery p over fixed three outcomes $l < m < h$ can be represented as a point (p_l, p_h) in a two-dimensional unit simplex, where the probability of l (p_l) is on the x-axis and that of h (p_h) is on the y-axis. Showing that indifference curves become steeper, or fanning out, in the 'north-west' direction is sufficient for Allais-type behavior, while the opposite pattern, fanning in, is sufficient for anti-Allais-type behavior.

Fig. 2 shows all possible combinations of attitudes towards randomization and fanning properties of indifference curves. Panel (A) plots indifference curves of individuals who exhibit both the Allais paradox and the consensus effect. Panel (B) demonstrates preferences which exhibit both anti-Allais behavior and the anti-consensus effect (both sets of indifference curves can be generated by the functional form $V_{\mathbb{CPE}_M}$ (Section 2.2), which, as shown in Masatlioglu and Raymond (2016), also admits an RDU representation and thus falls under the results of ERR). In contrast, Panel (C) depicts preferences that exhibit the Allais paradox but the anti-consensus effect, while Panel (D) shows preferences which generate anti-Allais behavior but also the consensus effect.[17]

The following functional forms generate the types of behavior depicted in Panels (C) and (D).

Example 1 *(quasi-concavity with Allais-type behavior).* Consider the quadratic functional,

$$V(p) = \mathrm{E}[v(p)] \times \mathrm{E}[w(p)]$$

which is quasi-concave (since $\log V$ is concave).[18] For three outcomes, $l < m < h$, define v and w as follows: $v(l) = 1, v(m) = 2, v(h) = 4$; $w(l) = 2, w(m) = 3, w(h) = 4$. We show in the Appendix that the indifference curves of this functional fan out.

Example 2 *(quasi-convexity with anti-Allais-type behavior).* Consider again three fixed outcomes, $l < m < h$, and the utility functional be defined as

$$V(p_l, p_h) = -6p_l + p_l^2 + 7.82p_h - 3.2p_l p_h + 2.56p_h^2$$

We show in the Appendix that this functional represents quasi-convex preferences, but its indifference curves fan in.[19]

Examples 1 and 2 show that Allais-type behavior and risky and cautious shifts (and the consensus effect more generally) are not necessarily related outside RDU. In the Appendix we provide another example, which demonstrates that even the equivalence between risky and cautious shifts and quasi-convexity (and so the consensus effect) that Corollary 3 describes does not extend. While quasi-convexity is a sufficient condition for ERR's risky/cautious shifts, it is not necessary.

3. The consensus effect in equilibrium

Our analysis so far has been restricted to understanding the behavior of an individual who is facing a fixed, exogenous decision process. While, similar to ERR, our interpretation of the environment is of a group decision problem, the exact same analysis would apply also if the

[17] Graphically, the indifference curves in Panels (A) and (B) are sections of concentric circles, those in Panel (C) are sections of hyperbolas, and those in Panel (D) are sections of parabolas.

[18] In this example, $\phi(x, y) = \frac{v(x)w(y)+v(y)w(x)}{2}$.

[19] V is a quadratic functional, with $\phi(l,l) = -5$, $\phi(m,l) = -3$, $\phi(h,l) = 2.51$, $\phi(m,m) = 0$, $\phi(h,m) = 3.91$, and $\phi(h,h) = 10.38$.

environment reflects a situation where the individual gets to choose with some probability, and with the remaining probability a computer chooses for him. To explicitly captures the strategic interaction, in this section we extend our analysis to a full equilibrium setting, and in doing so refer to individuals as voters.

We will show that, in contrast to settings where voters are expected utility maximizers, quasi-convex preferences can lead to phenomena such as group polarization, preference reversals, and multiple equilibria. This is driven by the fact that quasi-convex preferences give the voting game properties of a *coordination game*.[20]

We describe a majority voting game as a collection of N individuals (where N is odd), each of whom receives one vote to cast in favor of either p or q (no abstentions are allowed).[21,22] Whichever option receives the majority of the votes is implemented. Throughout this section when we refer to equilibria we mean "voting equilibria", that is, Nash equilibria where voters do not use weakly dominated strategies.[23]

Voters come in three *types*: Those that prefer p to q (Type P), those that prefer q to p (Type Q), and those that are indifferent (Type I). Each individual is drawn at random from each of the three types with probabilities f_P, f_Q, and f_I, respectively, where $f_P + f_Q + f_I = 1$. We denote the vector of probabilities by F. Each individual observes his own type and votes for either option p or option q.

As a benchmark, we first review the set of equilibria that emerge if all voters have expected utility preferences.

Proposition 6. *An equilibrium always exists. Moreover, a set of strategies is an equilibrium if and only if*

1. *Type Ps vote for p*
2. *Type Qs vote for q*
3. *Any given i of Type I votes for p with probability $r_i \in [0, 1]$*

Observe that in this equilibrium people vote for the option they favor in individual choice, or arbitrarily randomize between outcomes they are indifferent between.

We now turn to voters with quasi-convex preferences. Types P and Q can now come in two different sub-types. We call them $P1$, $P2$ (and $Q1$, $Q2$). Types $P1$ and $Q1$ have monotone preferences between q and p. For example, $P1$ (resp., $Q1$) strictly prefers $\lambda p + (1 - \lambda)q$ to $\delta p + (1 - \delta)q$ if and only if $\lambda > \delta$ (resp., $\lambda < \delta$).[24]

---

[20] Chew and Konrad (1998) also study the effect of quasi-convex preferences in a (majority) voting setting and show how it implies the bandwagon effect, namely, the idea that if individuals believe others will vote for a certain option, they themselves are more likely to vote for that option as well. Chew and Konard did not provide an equilibrium analysis and their (purely decision theoretical, like in ERR) reasoning is confined to the case where the probability of being pivotal is arbitrary small, as in, for example, large elections.

[21] An alternative assumption, often taken in the voting literature, is that the number of voters is a random variable that has a Poisson distribution. Such an assumption will not change our results.

[22] Identical results will be obtained if voting is assumed instead to be voluntary but costless. Costless/required voting is a reasonable assumption in many settings, such as committee votes, where members are required to be present regardless of whether they choose to vote. We discuss costly voting in Section 4.

[23] While there are equilibria that involve coordination motives even when individuals are expected utility maximizers, they rely on voters using weakly dominated strategies and are ruled out by this assumption.

[24] Expected utility preferences are always monotonic between q and p.

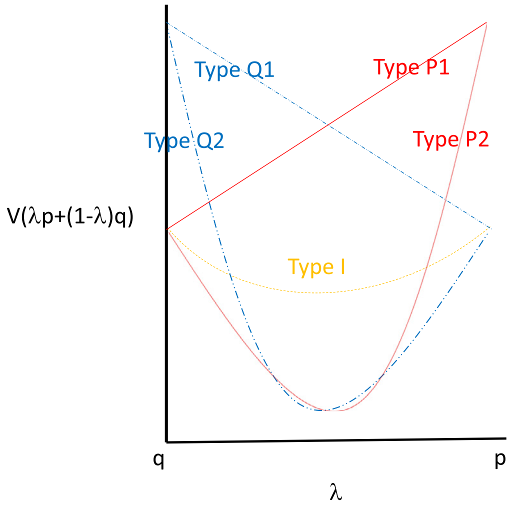

Fig. 3. Five types of preferences.

In contrast, $P2$'s preferences are non-monotonic between q and p. By strict quasi-convexity, $P2$'s preferences are single-troughed between p and q and there exists a unique λ^* such that $\lambda^* p + (1-\lambda^*)q \sim q$. Thus, for all $\lambda < \lambda^*$ we have $\lambda p + (1-\lambda)q \prec q$, which means that even though a $P2$ type prefers p to q, so long as p^* and q^* are both close enough to q he will prefer q^* to p^*. Similarly, for $Q2$, there exists a λ^* such that $\lambda^* p + (1-\lambda^*)q \sim p$, and $\lambda p + (1-\lambda)q \prec p$ for all $\lambda > \lambda^*$. We will refer to types $P1$ and $Q1$ as monotone types and to the others as non-monotone types. Fig. 3 illustrates the utility of each type over all convex combinations of q ($\lambda = 0$) and p ($\lambda = 1$). We assume that the distribution F has a full support, that is, it places strictly positive probabilities on all possible types.

We assume that individuals within each type have the same preferences, so that given a group problem, λ_i^* is the same for all i of type $P2$ (similarly for $Q2$ and I) and hence β_i^* is the same as well.[25]

A key property, which will be the main driving force behind many of the formal results below, is that with quasi-convex preferences, the majority voting game takes on aspects of a *coordination game* — non-monotone types experience benefits from coordinating their votes with others because it reduces the amount of "randomness" in the election, in the sense that it pushes p^* (resp., q^*) towards p (respectively, q).[26] In other words, if a non-monotone type expects that one of the options is very likely to be chosen, then voting for it is better than voting for the other one.

We turn now to studying some of the properties of the Nash equilibria of the voting game. First, we demonstrate that an equilibrium always exists. In particular, we prove the existence of

25 We focus on the situation where all individuals in each type have the same preferences for analytic convenience, although the results naturally extend to situations where they do not.

26 A key technical aside; as Crawford (1990) points out, games in which individuals have quasi-convex preferences may oftentimes admit no Nash equilibrium. He suggest a new notion "equilibrium in beliefs"which coincides with standard Nash equilibrium under expected utility, but also exists when players have quasi-convex preferences. We simply focus on Nash equilibrium, which, as we show in Proposition 7, always exists because of the benefits of coordination.

an "anonymous Nash equilibrium"that is, a Nash equilibrium in which each individual's strategy depends only on his preferences (i.e., his type) and not on his identity. Although the exact set of equilibria will depend on the distribution F, we will highlight some of the salient features that differ from the expected utility case.

Proposition 7. *An anonymous Nash equilibrium always exists. Moreover, in any equilibrium (not necessarily anonymous)*

1. *Generically,*[27] *all individuals strictly prefer to vote for one option or the other. Moreover, no individuals randomize*
2. *Type $P1$s vote for p*
3. *Type $Q1$s vote for q*

In contrast to Proposition 6, here no individual randomizes and, in fact, strictly dislikes randomizing. Thus, we will expect to observe choice shifts in the group — individuals who are indifferent between p and q in individual situations strictly prefer one or the other in a group setting. Proposition 7, however, does not specify whether the shift would be towards q or towards p.

In order to provide intuition for the actual pattern of voting that can be observed in equilibrium, we will analyze the best response function of a voter. We index the number of possible voting combinations by m. Consider voting pattern $\mathbb{V}^m$ and suppose individual i is of type Γ. Given this, observe that F and $\mathbb{V}^m$ generate a probability $\alpha(\mathbb{V}^m, F)$ of an individual being pivotal, and so a threshold probability $\beta^*(\mathbb{V}^m, \Gamma, F)$. Denote the set of types that vote for p (resp., q) given $\mathbb{V}^m$ as $\mathbb{P}(\mathbb{V}^m)$ (resp., $\mathbb{Q}(\mathbb{V}^m)$).[28]

The probability that p is chosen when i is not pivotal is:

$$\beta_{i,\mathbb{V}^m,F} = \frac{\sum_{k=\lceil \frac{N}{2}+2 \rceil}^{N} \binom{N}{k} (\sum_{\tau\in\mathbb{P}(\mathbb{V}^m)} f_\tau)^k (\sum_{\tau\in\mathbb{Q}(\mathbb{V}^m)} f_\tau)^{N-k}}{1-\sum_{k=\lceil \frac{N}{2}-1 \rceil}^{\lfloor \frac{N}{2}+1 \rfloor} \binom{n}{k} (\sum_{\tau\in\mathbb{P}(\mathbb{V}^m)} f_\tau)^k (\sum_{\tau\in\mathbb{Q}(\mathbb{V}^m)} f_\tau)^{N-k}}$$

Individual i's best response is to choose p if $\beta_{i,\mathbb{V}^m,F} > \beta^*(\mathbb{V}^m, \Gamma, F)$ and q if the inequality is reversed. Thus, a voting pattern is an equilibrium if it is the case that $\mathbb{V}^m$ generates $\beta_{i,\mathbb{V}^m,F}$ that are consistent with it.[29]

Because of the coordination nature of the majority voting game, we expect some of the non monotone types (but none of the monotone ones) to vote strategically in the group context, that is, to vote against the option they would prefer in isolation. Our next result groups together some sufficient conditions for this to happen, focusing on the case of sufficiently large elections.

[27] Generically here, as well as in Proposition 8, is in the set of distributions, using the standard weak* topology. That is, generic means open dense in the weak* topology.

[28] In defining $\beta_{i,\mathbb{V}^m,F}$, we assume that all individuals of the same type behave the same; a similar construction — albeit more complicated — can be performed without assuming anonymity.

[29] Quasi-convexity of preferences alone provides no restrictions on the ordering of the thresholds $\beta^*(\mathbb{V}^m, \Gamma, F)$ across different non-monotone types. Additional restrictions, such as that all preferences are in the quadratic class, do ensure that the thresholds are ordered in the "intuitive" fashion.

Proposition 8. *There exists an N^*, such that for all $N > N^*$ the following statements are true:*

1. *There is a f^*_{P1} (resp., f^*_{Q1}) sufficiently close to 1, such that for all $f_{P1} \geq f^*_{P1}$ (resp., $f_{Q1} \geq f^*_{Q1}$) the unique equilibrium is for all non-monotone types to choose p (resp., q).*
2. *Generically in all equilibria, all non-monotone types take the same action.*
3. *Suppose p is the status quo option and that the threshold T needed to replace p with q increases from 50% in favor of q. There is a T^* sufficiently close to 1 such that for all $T \geq T^*$ the unique equilibrium is for all non-monotone types to choose p.*

Item (1) states that whenever there are enough voters that strongly favor one of the options (i.e., in a monotone fashion), it is the case that all non-monotone types vote for that option as well. The result generates an intuitive type of preference reversal — individuals coordinate on voting for an outcome strongly favored by many others.

Item (2) says that in large elections we should always expect to see preference reversals, meaning that large elections will almost surely fail to aggregate preferences. Intuitively, as N grows large, both the proportions of each type of voters and (since voters generically do not randomize) the proportions of votes for each option are known with almost certainty. This implies that p^* and q^* are arbitrarily close to either p or q, and so all individuals will prefer to vote for either one or the other.

Item (3) considers what happens as the voting rule shifts. Intuitively, as the threshold increases, the probability of q being chosen falls, and so non-monotone types become less likely to vote for it. Eventually, the unique equilibrium is for non-monotone types to votes for p.

Unlike the scenario in item (1) of Proposition 8, individuals may also coordinate on equilibria that are not necessarily strongly favored, as shown by the following proposition. This result highlights how benefits from coordination generate multiple equilibria.

Proposition 9. *For any given N, if there exists a small enough proportion of non-monotone types, then generically there is a unique equilibrium.*[30] *In contrast, for large enough N, if the proportion of non-monotone types is sufficiently close to 1, then there are always at least two equilibria.*

That is, if non-monotone types form a large enough proportion of the population, they can all vote the same to ensure an outcome gets elected with very high probability. Voting against the group leads to additional uncertainty, which reduces ex-ante utility. In other words, when there is a sufficient number of any non-monotone type, the benefits of coordination become so large that multiple equilibria must exist. This can have counter-intuitive effects on voting outcomes. For example, imagine that all individuals are of type $P2$ and hence, when choosing individually, will choose p. However, when choosing as a group they could not only coordinate on an equilibrium where everyone votes for p but also on one where everyone votes for q. The latter is clearly Pareto sub-optimal, but exists because of the benefits of coordination. Thus, we can observe preference reversal not just because an individual knows many other voters have "extreme" preferences, but also because an individual knows that many other voters have preferences where they would like to coordinate.

[30] Moreover, if preferences are quadratic and if f_{P1} and f_{Q1} are sufficiently close to one another, then $P2$ (resp., $Q2$) types all vote for p (resp., q).

4. Discussion

Our discussion of quasi-convex preferences has focused on preferences that are explicitly non-expected utility. However, as we mentioned in the introduction and was initially pointed out by Machina (1984), if an expected utility maximizer is allowed to take a payoff-relevant action before the lottery he faces is resolved, then his induced preferences over the ex-ante 'optimal' lotteries will be quasi-convex. To see this, suppose there are two individuals, facing two lotteries, p and q, between which they are both indifferent. There are three outcomes, and p is a binary lottery over the best and middle outcomes while q is a binary lottery over the best and worst outcomes. Both individuals are indifferent between p and q. The individuals vote as in our voting game. After voting, but before the chosen alternative is revealed, each individual can take one (and only one) of two 'insurance' action; a_1 or a_2. Action a_1 fully insures against the realization of the middle outcome, but not the low outcome, while a_2 insures against the realization of the low outcome, but not the middle outcome. Thus, even if the two individuals have expected utility preferences over lotteries, they have a strict incentive to coordinate their votes, because they would like to know which insurance action to take.

Because many applications focus on groups choosing between two options, we have also restricted our analysis to binary choices. Our predictions for the consensus effect, however, are readily extended. For example, suppose that the group must choose over Ω possible lotteries, denoted $p_1, ...p_\Omega$, and that an individual is indifferent between all of them. Then, the individual will strictly prefer to vote for option j as long as the probability that it will be chosen, when he is not pivotal, is sufficiently large.

4.1. Alternative approaches to group choice and consenses-type effects

Our results are related to the large literature on voting and the aggregation of preferences or information in elections.[31] The literature has made different assumptions regarding how individuals value outcomes and about the cost of voting. We have focused on the situation where voters have private values and have either compulsory or costless voting. As mentioned, with expected utility preferences and either compulsory or costless voting, all voters vote sincerely (i.e., individuals vote as part of the group in the same way they would choose in isolation) and all individuals vote, meaning that preferences are aggregated. By contrast, with quasi-convex preferences we find Pareto-dominated equilibria where preferences are not properly aggregated.

There are alternative models of group decisions, some of which can generate behavior that is consistent with our model. We now summarize these alternative specifications and compare their predictions to ours. In particular, voting models with common value components and exogenous conformity benefits can generate behavior akin to the consensus effect. We discuss in what ways our explanation can be distinguished from theirs.

COSTLY VOTING: Many papers in the literature on voting with private values assume that voting is costly, meaning that, as Ledyard et al. (1981, 1984) and Palfrey and Rosenthal (1983, 1985)

---

31 An important distinction is that while we assume that alternatives in the voting game are lotteries, most papers suppose they are final outcomes. Of course, this complicates thinking about our results in relation to the pre-existing literature; for example, in a common-value setting, private signals would then need to be about a particular outcome in the support of p or q. We nevertheless believe our assumption is natural in many instances; for example, if voters value candidates by what policies they will implement and there is a degree of uncertainty about what campaign promises candidates will actually follow through with.

point out, individuals need to compare the cost of voting to the benefit of voting, namely the chance of being pivotal. This implies that the proportion of votes cast for each side will now depend not only on the fraction of supporters for each option, but also on the cost and benefit distributions of both types of supporters (Taylor and Yildirim, 2010). However, as in models with compulsory voting and in contrast to the preference reversals we predict, conditional on voting, voters will still truthfully reveal their preferences over the options.

To build intuition, consider the case of an expected utility maximizer with utility function V_{EU}, who favors option p. In a private value setting, the only two factors affecting the decision of whether or not to vote is the utility gap between p and q and the probability of being pivotal, α; the benefit of voting is then $\alpha(V_{EU}(p) - V_{EU}(q)) = V_{EU}(p^*) - V_{EU}(q^*)$. Clearly, as the individual's probability of being pivotal increases, he requires a lower cost of voting to actually go to the poll. Note that the decision of whether to cast a vote (and whom to vote for) is independent of β.

For individuals with quasi-convex preferences (who favor p^*), the benefit of going to the poll is again $V_{EU}(p^*) - V_{EU}(q^*)$, which now depends not only on p, q, and α, but also on β. The dependency on β can generate some interesting dynamics in voting. For example, suppose that the individual is indifferent between p and q. Expected utility maximizers would never vote regardless of their pivotality. Similarly, voters with quasi-convex preferences will never pay the cost of voting when they know they are choosing along. But there is a threshold c^*, such that for all costs $c \leq c^*$ there are $\alpha, \beta \in (0, 1)$ under which the individual will pay the cost of voting.

For individuals with quadratic preferences, these situations take on a simple form. In particular, there are costs and options that generate similar enough utility so that although the individual will always abstain when the probability of being pivotal is high, for smaller values of α the individual will abstain only for intermediate value β, and instead vote for more extreme values of β. Thus, we *may* observe individuals who would not vote even if they knew that they are very likely to be pivotal ($\alpha \cong 1$), or if they believe the election is likely to be close ($\beta \cong 0.5$); but will vote when they are less likely to be pivotal, but think the election will be a blow-out (β is close enough to either 0 or 1).

COMMON VALUES: The other major assumption in the literature is that outcomes have a common-value component and voters receive private signals about it.[32] With compulsory voting, as Austen-Smith and Banks (1996) and Feddersen and Pesendorfer (1998) first noted, sincere voting is in fact not an equilibrium. Surprisingly, despite this, information is still aggregated in large elections, other than in knife-edge situations. The prediction is slightly different with voluntary, instead of compulsory, participation. Krishna and Morgan (2012) demonstrate that if participation is voluntary (either free or costly) then although some individuals may not vote, individuals who do vote will do so sincerely, and information is aggregated in large elections.

In a related environment, with common values and private signals, Sobel (2014) shows that without restricting the informational environment, any action is rationalizable. Roux and Sobel (2015) impose additional restrictions on the environment to identify when group decisions are more variable than individual decisions.

A key difference between the predictions of our model and the common-value literature is that in our model individuals may vote insincerely to avoid randomness, whereas with a common-

[32] Although we consider the two assumptions about values separately, Ghosal and Lockwood (2009), Feddersen and Pesendorfer (1999), Feddersen and Pesendorfer (1997) and Krishna and Morgan (2011) consider elections in the presence of both common-value and private-value components.

value component, individuals vote insincerely to help ensure the selected option is optimal given the (unknown) state. These two motivations can imply different behaviors in some circumstances. For example, adding partisan individuals who will always vote for p will push uninformed quasi-convex voters who want to match the state towards choosing p. However, as the example below illustrates, in a common-value setting individuals will instead want to more often vote against p; as Feddersen and Pesendorfer (1996) note "[uninformed independent agents] vote to compensate for the partisans". More generally, in a model with both features (quasi-convexity preferences and a common-value component) individuals may vote insincerely not only for strategic reasons but also for reasons related to their desire to reduce the randomness of the election. These mixed motivations for insincere voting will impede information aggregation. The following simple example highlights these issues.

Example 3. Suppose there are three voters in a majority rule election. Voter 1 is a partisan and will always vote for option p (as specified below). Voters 2 and 3 care about both what state will be realized and what alternative was chosen; in particular, they want to match the chosen alternative to the state. There are two equally likely states, s_p and s_q. Suppose there are three final outcomes $\bar{x} > x > \underline{x}$. The alternatives are two lotteries p and q. p (resp., q) gives x with probability ρ regardless of the state, $\bar{x}$ with probability $1-\rho$ if the state is s_p (resp., s_q), and $\underline{x}$ with probability $1-\rho$ if the state is s_q (resp., s_p). Moreover, Voter 2 receives a perfectly revealing private signals about the state prior to voting, while Voter 3 receives no signal at all.

If all voters have expected utility preferences, then consider a situation where Voter 2 always votes in accordance with his perfectly revealing signals. Voter 3 now wants to condition his vote on being pivotal, which happens only when the state is s_q (otherwise both of the other voters are voting for p). Thus, he should always cast his vote for q. It is easy to show that such behavior on the parts of Voters 2 and 3 constitute an equilibrium which aggregates information.

Now, to make the minimal deviation from the standard model, suppose only Voter 3 has quasi-convex preferences (everyone else still has expected utility preferences), that are non-monotone between p and q. One can easily construct preferences such that p^* is preferred to q^*. In this case there will be no equilibrium that aggregates information.

This intuition readily extends even when the number of voters becomes large (as in the results in the literature on information aggregation). In this case, information aggregation will still fail, even with many voters.

Exogenous Conformity Benefits: One explanation for group shifts is an explicit benefit of conformity or for being on the winning side (for example, Callander, 2007; Callander, 2008; Hung and Plott, 2001; Goeree and Yariv, 2015; and Moreno and Ramos-Sosa, 2017). Our model generates an endogenous benefit of conformity; individuals are willing to vote against what they would choose in isolation in order to reduce the uncertainty of the outcome, or in other words to conform to what they expect to happen. Their interest in doing so is not explicit, but rather depends on the distribution of types and expected number of voters. Models of exogenous conformity often generate diminishing willingness to vote as α falls. Although this can happen in our model, it is also possible that these values can increase (for a time) with α, for the same reasons discussed under costly voting.

More generally, the motivation for conformity in our model, i.e., reducing randomness, is distinct from potential other motivations, which often rely on a desire to feel socially integrated and so may depend on factors such as the observability of one's vote, or the extent to which the choice is being made by other voters (versus an objective randomization device).[33]

VOTING WITH SUBJECTIVE UNCERTAINTY: Ellis (2016) also relaxes the assumption of expected utility in a voting setting. He considers a common-value voting game with subjective uncertainty, where voters have max-min utility as in Gilboa and Schmeidler (1989). Because in a subjective environment max-min utility implies a preference for hedging, he shows that voters have a desire to randomize; i.e., they exhibit an anti-consensus effect.

Appendix

Before we prove the results in the main text, whenever we consider two arbitrary options p and q, we adopt the following normalization: Recall that for all values of $\alpha, \beta \in [0, 1]$, q^* and p^* are on the line segment connecting q and p in some multidimensional simplex. In order to simplify notation, we will rotate the probability simplex so that for any given p and q under consideration, this line segment runs from the origin through $e_1 = (1, 0, 0, 0...)$ and associate q with the origin. Moreover, we can now focus on the 1 dimensional case, and think of the line segment connecting 0 and 1 where we associate q with 0 and p with 1. We will thus associate a lottery $zp + (1 - z)q$ for $z \in [0, 1]$ with the point z. Note that since $p^* - q^* = \alpha(p - q) = \alpha$, we have that $p^* \geq q^*$ given our normalization.

Moreover, we fix representation of the preference relation $\succcurlyeq$ for each given type V_Γ, which can depend on the type Γ (we will frequently omit the dependence on Γ to simplify notation). For $z', z'' \in [0, 1]$, let $\gamma(z', z'') = V(z') - V(z'')$ measure the utility gap between z' and z''. Observe that γ depends on the exact representation V. However, we will be concerned with ordinal rather than cardinal properties of γ and V.

Lemma 1. *$\succcurlyeq$ satisfies strict quasi-convexity if and only if for all p and q such that $p \sim q$ there exists a $z^* \in (0, 1)$ such that V is strictly decreasing on $[0, z^*]$ and strictly increasing on $[z^*, 1]$.*

Proof of Lemma 1. First we show the *if* part. Observe that the assumption implies that $V(z) < V(p) = V(q)$ for all $z \in (0, 1)$. This implies quasi-convexity since it holds for arbitrary p and q such that $p \sim q$.

We now show the *only if* part. Suppose not. Then for some pair p and q such that $p \sim q$ there is no z^* with the properties as in the premise. This implies that there exists at least one interior local *maximum*, denoted $Z \in (0, 1)$. Then, by continuity, there exists a neighborhood $[\underline{z}, \bar{z}] \ni Z$ such that $V(\underline{z}) = V(\bar{z}) \leq V(Z)$, violating strict quasiconvexity. □

Lemma 2. *For all p and q such that $p \sim q$ there exists a $z^* \in (0, 1)$ such that V is strictly decreasing on $[0, z^*]$ and strictly increasing on $[z^*, 1]$, if and only if for all p and q such that $p \sim q$ and $\alpha \in (0, 1)$, there exists a pair $z', z'' \in [0, 1]$ with the following three properties:*

1. *$z' - z'' = \alpha$ and $\gamma(z', z'') = 0$.*
2. *For all $\widetilde{z}' > z'$, $\widetilde{z}'' > z''$, and $\widetilde{z}' > \widetilde{z}''$, $\gamma(\widetilde{z}', \widetilde{z}'') > 0$.*

[33] Note that the case of $\alpha = 0$, that is, that the individual is never pivotal, is ruled out when defining a group decision problem. In the case of $\alpha = 0$, indifference between p and q in the private setting implies indifference in the group setting as well, independently of β. This would be a way to distinguish our model from one that posits a utility from conformity.

3. *For all* $\widetilde{z}' < z'$, $\widetilde{z}'' < z''$, *and* $\widetilde{z}'' < \widetilde{z}'$, $\gamma(\widetilde{z}', \widetilde{z}'') < 0$.

Proof of Lemma 2. We prove the only if part first. To see that 1 is implied, first consider all pairs z', z'' such that $z' - z'' = \alpha$. Observe that both $\gamma(1, 1-\alpha) > 0$ and $\gamma(\alpha, 0) < 0$ hold by definition. By continuity there must be a point $z \in [\alpha, 1]$ such that $\gamma(z, z-\alpha) = 0$.

To see that 2 is implied, observe that since $\gamma(z', z'') = 0$, $z^* \in [z', z'']$ (if not, then the line $[0, 1]$ would have at least two local minima, a contradiction). There are two cases. If $\widetilde{z}'' > z^*$, then by Lemma 1 we have $\gamma(\widetilde{z}', \widetilde{z}'') > 0$. In contrast, if $\widetilde{z}'' < z^*$ then $V(\widetilde{z}'') < V(z'')$, and since $V(\widetilde{z}') > V(z')$, we have $V(\widetilde{z}') > V(\widetilde{z}'')$, or $\gamma(\widetilde{z}', \widetilde{z}'') > 0$. The proof that 3 is implied is exactly analogous.

To prove the if part, suppose it is not the case so that there is an interior local maximum in the interval, denoted $Z \in (0, 1)$. Then, by continuity, there exists a neighborhood $[\underline{z}, \bar{z}] \ni Z$ such that $V(\underline{z}) = V(\bar{z})$. Thus there exists an α' such that $\underline{z} - \bar{z} = \alpha'$. Observe that the pair $\underline{z}, \bar{z}$ satisfies condition 1, but not conditions 2 or 3. □

Proof of Proposition 1. By construction $p^* - q^* = \alpha$. Given that, Condition 1 implies that at β^* we have $p^* = z'$ and $q^* = z''$. By Conditions 2 and 3 of Lemma 2, $\beta > \beta^*$ (resp., $\beta < \beta^*$) implies that $\gamma(p^*, q^*) > 0$ (resp., < 0). Conversely, the pair p^*, q^* at β^* satisfies the properties of $z', z'' \in [0, 1]$ in Lemma 2. □

Proof of Corollary 1. Wakker (1994) shows that convexity of g is equivalent to quasi-convexity of preferences. The result follows from Proposition 1. □

Proof of Proposition 2. Chew et al. (1991) show that quadratic preferences imply mixture symmetry. The preference relation $\succcurlyeq$ satisfies mixture symmetry if for all $p, q \in \Delta$ and $\lambda \in [0, 1]$,

$$p \sim q \Rightarrow \lambda p + (1-\lambda) q \sim \lambda q + (1-\lambda) p$$

Suppose $q \sim p$. By mixture symmetry, we have

$$q^* = [\alpha + (1-\alpha)(1-\beta)] q + (1-\alpha) bp \sim (1-\alpha) bq + [\alpha + (1-\alpha)(1-\beta)] p \equiv \widehat{q}$$

If $\beta < 0.5$, $k = \frac{(1-\alpha)(1-2\beta)}{\alpha+(1-\alpha)(1-2\beta)} \in (0, 1)$ and we have $p^* = kq^* + (1-k)\widehat{q}$. By strict quasi-convexity $q^* \succ p^*$.

Moreover, by mixture symmetry we have

$$p^* = (1-\alpha)(1-\beta) q + [\alpha + (1-\alpha)\beta] p \sim [\alpha + (1-\alpha)\beta] q + (1-\alpha)(1-\beta) p \equiv \widehat{p}$$

If $\beta > 0.5$, $l = \frac{(1-\alpha)(2\beta-1)}{\alpha+(1-\alpha)(2\beta-1)} \in (0, 1)$ and we have $q^* = lp^* + (1-l)\widehat{p}$. By strict quasi-convexity $p^* \succ q^*$.

And if $\beta = 0.5$ and $q \sim p$ then, by mixture symmetry,

$$q^* \sim \widehat{q} = (1-\alpha)\beta q + [\alpha + (1-\alpha)(1-\beta)] p = (1-\alpha)(1-\beta) q + [\alpha + (1-\alpha)\beta] p = p^*$$

and hence $q \sim p \Rightarrow q^* \sim p^*$

To show that quadratic preferences are the only class for which the consensus effect is a majority effect, we first establish that strict quasi-convexity is implied. By way of contradiction, suppose preferences do not satisfy strict quasi-convexity everywhere. If preferences satisfy betweenness in some segment, then in that region the decision-maker is indifferent to convexification, and so will not exhibit the consensus effect. Similarly, if preferences satisfy strict quasi-concavity somewhere, then we observe an anti-consensus effect in that region. Either case violates our assumption, and so preferences must satisfy strict quasi-convexity everywhere.

Observe now that for $\beta = 0.5$, if we let $\gamma = \alpha + 0.5(1-\alpha) \in [0.5, 1]$ then

$$p^*_{|\beta=0.5} = \gamma p + (1-\gamma)q \text{ and } q^*_{|\beta=0.5} = \gamma q + (1-\gamma)p$$

We know that the individual exhibits the consensus effect at $(p, q, \alpha, .5)$, that is, $p^*_{|\beta=0.5} \sim q^*_{|\beta=0.5}$ for all γ. This immediately implies mixture symmetry. And since preferences are strictly quasi convex, Theorem 4 of Chew et al. (1991) implies that preferences are proper quadratic. □

Proof of Corollary 2. Masatlioglu and Raymond (2016) show that under $\mathbb{CPE}_M$, individuals are loss averse if and only preferences are strictly quasi-convex. Moreover, they show that if preferences can be represented with $V_{\mathbb{CPE}_M}$ then they also have a quadratic representation. The result follows. □

Proof of Corollary 3. The equivalence of 1, 2, and 3 is shown by ERR. The equivalence of 3 and 4 is Proposition 1. □

Proof of Proposition 3. We prove the case if both options are similarly bad. The proof if they are similarly good is analogous. Observe that for any outcome x_i, the strict de-cumulative distribution function under p, $\sum_{j>i} p(x_j)$, must be less than or equal to $\bar{p}$ (or equivalently, the weak de-cumulative distribution function is less than $\bar{p}$ for all outcomes other than the worst). The same must be true for q, and thus, also for any convex combination of p and q. It follows that all probabilities involved in the mixing operation are less than or equal to $\bar{p}$, and so mixing takes place only on the strictly concave portion of g. As Wakker (1994) shows, mixing in a concave portion of g implies quasi-concavity of preferences. From Proposition 1 we know this implies the anti-consensus effect. □

Proof of Proposition 4. We prove the contrapositive by contradiction: we suppose that both the consensus effect is violated and the axiom holds, and show it yields a contradiction.

Suppose the consensus effect is violated. Then, adopting our previous notation, we can find a $p, q \in \Delta$ such that

- p^* and q^* satisfy $p^* - q^* = \alpha$
- $\hat{p}^*$ and $\hat{q}^*$ satisfy $\hat{p}^* - \hat{q}^* = \alpha$
- $p^* > \hat{p}^*$
- $p^* \neq c(\{p^*, q^*\})$ but $\hat{p}^* = c(\{\hat{p}^*, \hat{q}^*\})$

There are two cases to consider:

1. $p^* > \hat{p}^* > q^* > \hat{q}^*$. Observe that in this case there exists a $\lambda \in (0, 1)$ such that $\hat{p}^* = \lambda p^* + (1-\lambda)q^*$. Therefore, by strict reference bias, $q^* = c(\{\hat{p}^*, q^*\})$.
   Similarly, there exists a $\lambda \in (0, 1)$ such that $q^* = \lambda\hat{p}^* + (1-\lambda)\hat{q}^*$. But by Strict Reference Bias, we now have $\hat{p}^* = c(\{\hat{p}^*, q^*\})$. This is a contradiction.
2. $p^* > q^* > \hat{p}^* > \hat{q}^*$.
   Note that either p^* or $\hat{q}^*$ must be chosen from $\{p^*, \hat{q}^*\}$. Suppose it is the former. Then there is a $\lambda \in (0, 1)$ such that $q^* = \lambda p^* + (1-\lambda)\hat{q}^*$, and so by strict reference bias it must be the

case that $p^* = c(\{p^*, q^*\})$, a contradiction. We can similarly obtain a contradiction for the latter as well. □

Proof of Proposition 5. There are two cases to consider:

1. Suppose that $p = c(\{p, q\})$. Then it must be the case that $V_{KR}(p|p) \geq V_{KR}(q|p)$ and either: (i) $V_{KR}(p|q) > V_{KR}(q|q)$ or (ii) $V_{KR}(q|q) \geq V_{KR}(p|q)$ but $V_{KR}(p|p) > V_{KR}(q|q)$. In either case, the weak inequalities in the proof Proposition 2 of Freeman (2019) become strict and so strict reference bias is immediately satisfied.
2. Suppose that both p and q are chosen from $c(\{p, q\})$. Then we know that $V_{KR}(p|p) \geq V_{KR}(q|p)$, $V_{KR}(q|q) \geq V_{KR}(p|q)$ and $V_{KR}(p|p) = V_{KR}(q|q)$. If either $V_{KR}(p|p) > V_{KR}(q|p)$ or $V_{KR}(q|q) > V_{KR}(p|q)$ then the weak inequalities in the proof of Proposition 2 of Freeman (2019) become strict and so strict reference bias is satisfied.
   We will show that one of the two strict inequalities must hold by way of contradiction. Suppose that both $V_{KR}(p|p) = V_{KR}(q|p)$ and $V_{KR}(q|q) = V_{KR}(p|q)$ and recall that $V_{KR}(p|p) = V_{KR}(q|q)$.
   Observe that V is linear in the probabilities of its first argument, fixing the second, and in the second argument, fixing the first. Then for any $p^* = \lambda p + (1-\lambda)q$ for $\lambda \in (0, 1)$:

$$\begin{aligned}V(p^*|p^*) &= V(\alpha p + (1-\alpha)\beta p + (1-\alpha)(1-\beta)q|\alpha p + (1-\alpha)\beta p \\ &\quad + (1-\alpha)(1-\beta)q) \\ &= (\alpha + (1-\alpha)\beta)V(p|\alpha p + (1-\alpha)\beta p + (1-\alpha)(1-\beta)q) \\ &\quad + (1-\alpha)(1-\beta)V(q|\alpha p + (1-\alpha)\beta p + (1-\alpha)(1-\beta)q) \\ &= (\alpha + (1-\alpha)\beta)^2 V(p|p) + (\alpha + (1-\alpha)\beta)(1-\alpha)(1-\beta)V(p|q) \\ &\quad + (1-\alpha)(1-\beta)(\alpha + (1-\alpha)\beta)V(q|p) + (1-\alpha)^2(1-\beta)^2 V(q|q) \\ &= V(p|p)\end{aligned}$$

   But, we know from the proof of CPE that it is strictly quasi-convex, which implies $V(p|p) > V(p^*|p^*)$. Therefore, either $V_{KR}(p|p) > V_{KR}(q|p)$ or $V_{KR}(q|q) > V_{KR}(p|q)$. □

Proof of Example 1. This utility functional does not exhibit Allais-type behavior. To see this, denote the probability of h by p_h and the probability of l by p_l. The utility of a lottery $(h, p_h; m, 1 - p_l - p_h, l, p_l)$ is then

$$\begin{aligned}&p_l^2[\phi(m,m) - 2\phi(m,l) + \phi(l,l)] \\ &+ p_l p_h[-2\phi(h,m) + 2\phi(h,l) + 2\phi(m,m) - 2\phi(m,l)] \\ &+ p_h^2[\phi(h,h) - 2\phi(h,m) + \phi(m,m)] \\ &+ p_l[-2\phi(m,m) + 2\phi(m,l)] \\ &+ p_h[2\phi(h,m) - 2\phi(m,m)] \\ &+ \phi(m,m)\end{aligned}$$

First, we will normalize the utility values. Chew, Epstein, and Segal (1991) show that ϕ is unique up to affine transformation. So we will set $\phi(m, m) = 0$ and $\phi(m, l) = \phi(l, m) = -1$ (recall that $\phi(m, m) \geq \phi(l, m)$ by monotonicity). The other relevant values of ϕ will be stated below.

Second, recall that Allais-type behavior is equivalent to indifferent curves fanning out in the probability simplex, where the value of p_l is on the horizontal axis and that of q_h on the vertical axis. Fanning out is equivalent to the slopes of the indifference curves becoming less steep moving horizontally in the simplex. The slope of the indifference curves is equal to

$$\mu(p_l, p_h) = -\frac{2p_l[2+\phi(l,l)] + p_h[-2\phi(h,m)+2\phi(h,l)+2]-2}{p_l[-2\phi(h,m)+2\phi(h,l)+0+2]+2p_h[\phi(h,h)-2\phi(h,m)]+[2\phi(h,m)]}$$

Taking the derivative $\frac{\partial\mu(p_l,p_h)}{\partial p_l}$ and observing that its denominator is always positive, we know that to determine its sign (which tells us whether we get fanning out or fanning in) we only need to consider its numerator.

First, we focus on fanning out along the $p_l - axis$, and so will set $p_h = 0$ after calculating $\frac{\partial\mu(p_l,p_h)}{\partial p_l}$. Note that the derivative of the numerator of $\mu(p_l, p_h)$ with respect to p_l is $2[2 + \phi(l,l)]$, while the derivative of the denominator of $\mu(p_l, p_h)$ with respect to p_l is $[-2\phi(h,m) + 2\phi(h,l) + 2]$. We also have that at $p_h = 0$, the numerator of $\mu(p_l, p_h)$ equals $2p_l[2+\phi(l,l)]-2$ and the denominator of $\mu(p_l, p_h)$ equals $p_l[-2\phi(h,m)+2\phi(h,l)+2]+[2\phi(h,m)]$. Therefore, the numerator of $\frac{\partial\mu(p_l,p_h)}{\partial p_l}$ equals $-4\phi(h,m) - 4\phi(l,l)\phi(h,m) - 4\phi(h,l) - 4$, meaning that we get fanning out horizontally along $q = 0$ if and only if

$$-\phi(h,m) - \phi(h,l) - 1 - \phi(l,l)\phi(h,m) < 0$$

Given our specified v and w functions, we can represent ϕ using a matrix

$$\begin{pmatrix} \phi(l,l) & \phi(l,m) & \phi(l,h) \\ \phi(l,m) & \phi(m,m) & \phi(m,h) \\ \phi(l,h) & \phi(m,h) & \phi(h,h) \end{pmatrix}$$

Substituting in our actual values (only for the lower triangle, because of the symmetry of ϕ) gives

$$\begin{pmatrix} 2 & \phi(l,m) & \phi(l,h) \\ 3.5 & 6 & \phi(m,h) \\ 6 & 10 & 16 \end{pmatrix}$$

To normalize $\phi(m,m) = 0$ and $\phi(m,l) = -1$, we subtract 6 from all payoffs and then divide by 2.5. This yields the ϕ matrix

$$\begin{pmatrix} -8/5 & \phi(l,m) & \phi(l,h) \\ -1 & 0 & \phi(m,h) \\ 0 & 8/5 & 4 \end{pmatrix}$$

We then have $-\phi(h,m) - \phi(h,l) - 1 - \phi(l,l)\phi(h,m) = -1/25 < 0$, so indifference curves are fanning out. This proves fanning out along the line $p_h = 0$.

In order to extend fanning out throughout the unit simplex, we use the notion of expansion paths, defined by Chew et al. (1991). We will use their definition, tailored to our example, which is as follows.

Given three outcomes $l < m < h$, consider the probability simplex (triangle) over those three outcomes, as described in the text (where p_h denotes the probability of h and p_l the probability of l). Suppose that indifference curves in this space are always differentiable inside the simplex, where, as above, $\mu(p_l, p_h)$ denotes the slope of the indifference curve passing through any given point (p_l, p_h). An expansion path collects the set of all points, the indifference curve through which have the same slope (that is, (p_l, p_h) and (p_l', p_h') are on the same expansion path if $\mu(p_l, p_h) = \mu\left(p_l', p_h'\right)$.)

Chew et al. (1991) show that for quadratic preferences which are not expected utility, expansion paths are linear (in the case of expected utility all points in the simplex are in the same expansion path). Moreover, they show that either[34]

- no two expansion paths intersect (in other words expansion paths are parallel); or
- all expansion paths intersect at a single point (i.e., if two expansion paths intersect at (p_l', p_h') then all expansion paths must intersect there), which may or may not be inside the unit simplex (i.e., it is possible that the point where they intersect has p_l and p_h values greater than 1 or less than 0)

We now turn to applying expansion paths to our example. In Example 1, the "reduced form" utility function over lotteries defined over the three outcomes (taking into account our normalized values) is:

$$V(p_l, p_h) = -2p_l + \frac{2p_l^2}{5} + \frac{16p_h}{5} - \frac{6p_l p_h}{5} + \frac{4p_h^2}{5}$$

Observe that $\left(\frac{-6}{5}\right)^2 - 4 \times \frac{2}{5} \times \frac{4}{5} = \frac{36}{25} - \frac{32}{25} = \frac{4}{5} > 0$, and so we know the indifference curves take the shape of *hyperbolas*, and thus all expansion paths intersect at a single point.[35] To find this point of intersection, we simply need to find the critical point of the utility function.[36] The first order conditions demonstrate that this is at $p_l = 4$, $p_h = 1$. Thus, all expansion paths must intersect there, which in turns implies that, within the unit simplex, all expansion paths are positively sloped (and do not intersect within the simplex).

Consider moving from some point (p_l, p_h) to (p_l', p_h) in the probability simplex, with $p_l < p_l'$. Denote the expansion path (p_l, p_h) is on as E_1 and the expansion path (p_l', p_h) is on as E_2. Then we can find points $(\hat{p}_l, 0)$ and $(\hat{p}_l', 0)$ such that the former is on expansion path E_1 and the latter is on expansion path E_2. Since the expansion paths cannot cross anywhere other than $(4, 1)$, $\hat{p}_l < \hat{p}_l'$. But we know from our previous reasoning that, regardless of the initial value of p_l, when increasing p_l and moving along the line $p_h = 0$, the slopes of the indifference curves decrease. So the slope of the indifference curve is lower at $(\hat{p}_l', 0)$ than $(\hat{p}_l, 0)$, meaning that the slope of the indifference curve must be lower at (p_l', p_h) than (p_l, p_h). Therefore, we get fanning out as p_l increases, regardless of p_h, so long as we are inside the probability simplex. □

Proof of Example 2. We consider the functional over (p_l, p_h) given by

$$V = -6p_l + p_l^2 + 7.82p_h - 3.2p_l p_h + 2.56p_h^2$$

Since $3.2^2 - 4 \times 2.56 = 10.24 - 10.24 = 0$, the indifference curves of V take the shape of *parabolas*, which have the same axis of symmetry. Thus all indifference curves either have lower contour sets that are (strictly) convex or upper contour sets that are (strictly) convex. In our case,

[34] See Lemmas A2.2-5 in their paper.

[35] For details, see Chew et al. (1991). Intuitively, the expansion paths all must intersect at center of the hyperbolas, or, in other words, at the point of intersection of the asymptotes.

[36] This follows from the fact that the asymptotes of the hyperbola must be on the same level set.

because the axis of symmetry has a positive slope and lies below the unit simplex, preferences have convex lower contour sets and hence satisfy quasi-convexity.

Moreover, $\frac{\partial V}{\partial p_l} = -6 + 2p_l - 3.2p_h$ and $\frac{\partial V}{\partial p_h} = 7.82 - 3.2p_l + 5.12p_h$. Thus, the slope of the indifference curves is $\mu(p_l, p_h) = -\frac{-6+2p_l-3.2p_h}{7.82-3.2p_l+5.12p_h}$.

Along the set of lotteries where $p_h = 0$, $\mu(p_l, p_h)$ reduces to $-\frac{-6.+2p_l}{7.82-3.2p_l}$. Taking the derivative of this with respect to p_l gives $\frac{0.347656}{(2.44375-p_l)^2} > 0$, so indifference curves are fanning in. This proves fanning in along the line $p_h = 0$.

In order to extend fanning in throughout the probability simplex, we use expansion paths in a similar way to Example 1. Since the indifference curves are parabolas, it is the case that the expansion paths are parallel.[37] Moreover, because the axis of symmetry of the indifference curves is an expansion path, the expansion paths have positive slopes.

Consider moving from some point (p_l, p_h) to (p_l', p_h), where $p_l < p_l'$. Denote the expansion path (p_l, p_h) is on as E_1 and the expansion path (p_l', p_h) is on as E_2. Then we can find points $(\hat{p}_l, 0)$ and $(\hat{p}_l', 0)$ such that the former is on expansion path E_1 and the latter is on expansion path E_2. Since the expansion paths cannot cross $\hat{p}_l < \hat{p}_l'$. But we know from our previous reasoning that, regardless of the starting value of p_l, when increasing p_l and moving along the line $p_h = 0$ the slope of the indifference curves increase. So the slope of the indifference curves is higher at $(\hat{p}_l', 0)$ than at $(\hat{p}_l, 0)$, which, in turns, implies that the slope of the indifference curves must be higher at (p_l', p_h) than (p_l, p_h). So we get fanning in as p_l increases, regardless of p_h. □

Proof of Proposition 6. For any distribution F over types, consider the strategies as specified in the Proposition. Type I voters are indifferent between all possible outcomes and hence will be indifferent between any randomization over p and q. Because we focus on the equilibria where no individuals play a weakly dominated strategy, it must be the case that types P vote for p and Q for q in equilibria. □

Before proceeding to the rest of the proofs, we denote the induced lotteries faced by individual i of type Γ given voting pattern m and distribution F by $p_i^*((\mathbb{V}^m), \Gamma, F)$ and $q_i^*((\mathbb{V}^m), \Gamma, F)$. We sometime refer to non-monotone types, that is, types $P2$, $Q2$, or I, by NM.

Proof of Proposition 7. First, by same arguments as in the proof of Proposition 6, it is clear that in any equilibrium, $P1$ and $Q1$ types will behave like expected utility maximizers, which implies points 2 and 3.

To show the existence of an anonymous equilibrium, notice that actions can't depend on an individual's identity, just their type. Thus $\alpha_i(\mathbb{V}^m, F) = \alpha(\mathbb{V}^m, F)$ and so $\beta_i^*(\mathbb{V}^m, \Gamma, F) = \beta^*(\mathbb{V}^m, \Gamma, F)$ for all i. We prove existence by contradiction, that is we will suppose no such equilibrium exists and show a contradiction occurs. We do this in several steps.

- Initially we suppose all NM types vote for p. Call this voting pattern $(1:1)$.[38] We will order the three NM types by increasing order of the threshold required to vote for q (given this voting pattern): $\mathbb{I}$, $\mathbb{II}$ and $\mathbb{III}$. Thus, if type $\mathbb{III}$ wants to switch their vote to q then all other NM types would as well. Since, by assumption, we are supposing this is not an equilibrium,

37 Again, see Chew et al. (1991).

38 In the proof we induct on the number of types (the number on the left), and within each type, on the number of individuals within it (the number on the right).

then at least one of the three NM types wants to deviate to voting for q. Clearly individuals of type $\mathbb{I}$ must want to switch (because of our ordering assumption).

We now order all possible individuals $1, 2, ..N$. We will consider each individual's strategy, conditional on him being of type $\mathbb{I}$ and induct on the order of the individuals. Begin with individual 1. By construction, in the proposed voting pattern, $\beta^*(\mathbb{V}^{(1:1)}, \mathbb{I}, F) > \beta(\mathbb{V}^{(1:1)}, \mathbb{I}, F)$ or, equivalently, $q^* \succ p^*$. So individual 1 in type $\mathbb{I}$ would prefer to switch to voting for q. Denote this voting pattern $(1:2)$.

Observe that under voting pattern $(1:2)$, we have that for all other individuals both $p^*(\mathbb{V}^{(1:2)}, \Gamma, F)$ and $q^*(\mathbb{V}^{(1:2)}, \Gamma, F)$ are closer to q than $p^*(\mathbb{V}^{(1:1)}, \Gamma, F)$ and $p^*(\mathbb{V}^{(1:1)}, \Gamma, F)$, respectively. Therefore, because all individuals in type $\mathbb{I}$ preferred to deviate from voting for p to voting for q under voting pattern $(1:1)$, it is now the case that $q^*(\mathbb{V}^{(1:2)}, \mathbb{I}, F)$ is strictly preferred to $p^*(\mathbb{V}^{(1:2)}, \mathbb{I}, F)$. Thus individual 2, if realized as type $\mathbb{I}$, will also have a strict incentive to switch his vote from p to q.

We continue by simply inducting on the number of individuals. After all individuals with index smaller than k have switched, we have voting pattern $(1:k)$. It is clear using the reasoning described above that all individuals in type $\mathbb{I}$ with index greater than k strictly prefer $q^*(\mathbb{V}^{(1:k)}, \mathbb{I}, F)$ to $p^*(\mathbb{V}^{(1:k)}, \mathbb{I}, F)$ and the same for those with index less than k, which guarantees that they will not switch back to vote for p. Thus, we conclude this step by having a potential anonymous equilibrium where of the NM types, types $\mathbb{I}$ vote for q and the other NM types vote for p.

- Suppose again, continuing our contradiction, that this voting pattern (where of the NM types, types $\mathbb{I}$ vote for q and the other NM types vote for p) isn't an equilibrium. Denote this voting pattern by $(2:1)$. Now, we re-order the two remaining NM types that are voting for p under voting pattern $(2:1)$, calling them types $\mathbb{II}$ and $\mathbb{III}$.[39] Under our assumption that voting pattern $(2:1)$ is not an equilibrium, it must be the case that $\mathbb{II}$ types want to switch from voting for p to q.

  We now repeat the inductive process from the previous step but for individuals in type $\mathbb{II}$; order all individuals, and conditional on them drawing that type, switching them one by one from voting for p to voting for q. Observe that after individual k in type $\mathbb{II}$ switches from voting for p to q, that for all other individuals both $p^*(\mathbb{V}^{(2:k+1)}, \Gamma, F)$ and $q^*(\mathbb{V}^{(2:k+1)}, \Gamma, F)$ are both closer to q than $p^*(\mathbb{V}^{(2:k)}, \Gamma, F)$ and $q^*(\mathbb{V}^{(2:k)}, \Gamma, F)$ respectively. This means that (i) conditional on drawing type $\mathbb{II}$ no individual has an incentive to switch their votes, and (ii) conditional on drawing type $\mathbb{I}$ no individual would want to switch their vote back to p after any step in the inductive process. We conclude this step by having a potential equilibrium where of the NM types, types $\mathbb{I}$ and $\mathbb{II}$ vote for q and the type $\mathbb{III}$ vote for p.
- Lastly, we repeat the same exercise above, applying to type $\mathbb{III}$ voters. We will then conclude that we have an equilibrium in which all NM types vote for q, and have a strict preference to do so. This equilibrium is obviously anonymous, contradicting the assumption that no such equilibrium exists.

We now turn to proving the properties of the equilibrium. We have already proved parts 2 and 3. Suppose that an equilibrium exists with voting pattern $\mathbb{V}^m$ which induces, for each

[39] Types $\mathbb{II}$ and $\mathbb{III}$ need not correspond to the same groups as under voting pattern $(1:1)$; the ranking of the threshold to switch from p to q may be lower in one group under $(1:1)$ but higher under $(2, 1)$.

individual i, a pivot probability $\alpha(\mathbb{V}^m, F)$ and a threshold $\beta^*(\mathbb{V}^m, \Gamma, F)$. To see that 1 is true, observe that in the space of distributions F (using the weak* topology), generically $\beta^*(\mathbb{V}^m, \Gamma, F) \neq \beta_{i,\mathbb{V}^m,F}$. If in fact $\beta_i^*(\mathbb{V}^m, \Gamma, F) = \beta_{i,\mathbb{V}^m,F}$ then because of quasi-convexity the decision-maker still prefers not to randomize between the two. □

Before proceeding, we prove another useful Lemma.

Lemma 3. *For all $\epsilon > 0$ there exists an N^*, such that $N \geq N^*$ implies $0 < \alpha = p^* - q^* \leq \epsilon$.*

Proof of Lemma 3. Because F has full support and the fact that types $P1$ and $Q1$ vote for p and q respectively, it is always the case that any individual has a non-zero chance of being pivotal. However, as N goes to infinity the probability of being pivotal goes to 0, and thus α approaches to 0. □

Proof of Proposition 8. We prove each part in turn.

1. Because F has full support and the fact that types $P1$ and $Q1$ vote for p and q respectively, it is always the case that any individual has a non-zero chance of being pivotal. Let $\bar{z}^*$ indicate the highest value of z_Γ^* across all NM types, which means $\bar{z}^* \in (0, 1)$.[40] For a large enough f_{P1}^* and large enough N, in any voting pattern it is very likely, for each individual i, that p is chosen whenever i is not pivotal. Thus, for all individuals $p^*(\mathbb{V}^m, \Gamma, F)$ and $q^*(\mathbb{V}^m, \Gamma, F)$ are both in $(\bar{z}^*, 1)$, meaning that all NM individuals will choose p^*. The same reasoning must also be true for $f_{P1} > f_{P1}^*$. We can conduct a similar exercise for f_{Q1}.
2. We prove the result in two steps. First, we show that it holds for all anonymous equilibria. Recall that in all anonymous equilibria, all individuals of the same type take the same action. For large enough N, the proportion of each type in the total number of voters is known with near certainty (equals f_Γ). Moreover, fixing an equilibrium it is known exactly what action each type takes. This means that with near certainty we know what proportion of the total number of voters choose p and what proportion choose q, and hence the outcome of the voting game is known with near certainty; in other words, for all individuals $\beta_{i,\mathbb{V}^m,F}$ is arbitrarily close to either 1 or to 0. Without loss of generality suppose $\beta_{i,\mathbb{V}^m,F}$ is arbitrarily close to 1. Then $p^*(\mathbb{V}^m, \Gamma, F)$ is arbitrarily close to p and since N is large, Lemma 3 implies that $q^*(\mathbb{V}^m, \Gamma, F)$ is also arbitrarily close to $p^*(\mathbb{V}^m, \Gamma, F)$ (and so to p). This immediately implies that for any Γ, $q^*(\mathbb{V}^m, \Gamma, F)$ and $p^*(\mathbb{V}^m, \Gamma, F)$ are both greater than z_Γ^*, and so all NM types prefer to choose p. This proves the second part in the case of anonymous equilibria.
   The next step is to prove that with large N, generically all equilibria are anonymous. Consider two different individuals, i and j, who are considering their strategies, conditional on drawing the same type Γ. For large enough N, even if they choose different strategies, $\alpha_i(\mathbb{V}^m, F)$ is arbitrarily close to $\alpha_j(\mathbb{V}^m, F)$ (and both are arbitrarily close to 0). Moreover, $\beta_i^*(\mathbb{V}^m, \Gamma, F)$ is arbitrarily close to $\beta_j^*(\mathbb{V}^m, \Gamma, F)$. Thus, for large enough N if $p_i^*(\mathbb{V}^m, \Gamma, F) \succ q_i^*(\mathbb{V}^m, \Gamma, F)$ then $p_j^*(\mathbb{V}^m, \Gamma, F) \succ q_j^*(\mathbb{V}^m, \Gamma, F)$. We can iterate this argument over all individuals of a given type, and we obtain an anonymous equilibrium.

[40] Recall that for any type Γ, z_Γ^* is such that V_Γ is strictly decreasing on $[0, z_\Gamma^*]$ and strictly increasing on $[z_\Gamma^*, 1]$.

Thus, the only situation where we may have non-anonymous equilibria is where we have an (infinite) sequence of N along which $p_i^*(\mathbb{V}^m, \Gamma, F) \sim q_i^*(\mathbb{V}^m, \Gamma, F)$ holds. But using similar arguments as given above, it can be shown that this generically will not happen.

3. Suppose $f_{P1} > \epsilon > 0$. Suppose we need for a proportion of at least T^* people to vote for q in order for it to be chosen. But even if all NM types vote for q, as T^* goes to 1 the probability that the proportion of votes for q is greater than T goes to 0. Thus $p^*(\mathbb{V}^m, \Gamma, F)$ and $q^*(\mathbb{V}^m, \Gamma, F)$ both go to p, so p^* is preferred over q^* by all NM types. Thus in equilibrium all NM must vote for p. Clearly, for $T > T^*$ the same reasoning holds. □

Proof of Proposition 9. Proof of the first part: Each voting pattern generates a $\beta_{i,\mathbb{V}^m,F}$. Observe that since types $P1$ and $Q1$ always vote for p and q, respectively, as the proportion of NM types goes to 0 it is the case that $\beta_{i,\mathbb{V}^m,F}$ approaches some $\hat{\beta}$ regardless of the voting pattern of the NM types. Similarly, $\beta^*(\mathbb{V}^m, \Gamma, F)$ approaches $\beta^*(\Gamma)$. But since generically it is not the case that $\beta^*(\Gamma) = \hat{\beta}$ (for example, any small change in the distribution F which shifts small weight from $P1$ types to $Q1$ types will alter it), each NM type will have a unique best response regardless of the strategy of any other NM type.

Proof of the second part: Observe that if $\beta_{i,\mathbb{V}^m,F}$ is arbitrarily close to 1 then all individuals will vote p. Similarly if it is arbitrarily close to 0, all individuals will vote q. If the proportion of NM types goes to 1 and all NM types vote for p, then $\beta_{i,\mathbb{V}^m,F}$ goes to 1 and so we have an equilibrium. Similar logic applies if all NM types vote for q. □

References

Abeler, Johannes, Falk, Armin, Goette, Lorenz, Huffman, David, 2011. Reference points and effort provision. Am. Econ. Rev. 101 (2), 470–492.

Agranov, Marina, Goeree, Jacob K., Romero, Julian, Yariv, Leeat, 2017. What makes voters turn out: the effects of polls and beliefs. J. Eur. Econ. Assoc. 16 (3), 825–856.

Austen-Smith, David, Banks, Jeffrey S., 1996. Information aggregation, rationality, and the Condorcet jury theorem. Am. Polit. Sci. Rev. 90 (01), 34–45.

Baisa, Brian, 2013. Essays in Economic Theory. Yale University.

Barseghyan, Levon, Molinari, Francesca, O'Donoghue, Ted, Teitelbaum, Joshua C., 2013. The nature of risk preferences: evidence from insurance choices. Am. Econ. Rev. 103 (6), 2499–2529.

Bartels, Larry M., 1988. Presidential Primaries and the Dynamics of Public Choice. Princeton University Press.

Brown, Robert, 1986. Social Psychology, second edition. Free Press.

Callander, Steven, 2007. Bandwagons and momentum in sequential voting. Rev. Econ. Stud. 74 (3), 653–684.

Callander, Steven, 2008. Majority rule when voters like to win. Games Econ. Behav. 64 (2), 393–420.

Camerer, Colin F., Ho, Teck-Hua, 1994. Violations of the betweenness axiom and nonlinearity in probability. J. Risk Uncertain. 8 (2), 167–196.

Carbajal, Juan Carlos, Ely, Jeffrey, 2012. Optimal Contracts for Loss Averse Consumers. Technical report. University of Queensland, School of Economics.

Carbajal, Juan Carlos, Ely, Jeffrey C., 2016. A model of price discrimination under loss aversion and state-contingent reference points. Theor. Econ. 11 (2), 455–485.

Chew, Soo Hong, Epstein, Larry G., Segal, Uzi, 1991. Mixture symmetry and quadratic utility. Econometrica, 139–163.

Chew, Soo Hong, Epstein, Larry G., Segal, Uzi, 1994. The projective independence axiom. Econ. Theory 4 (2), 189–215.

Chew, Soo Hong, Konrad, Kai, 1998. Bandwagon effects and two-party majority voting. J. Risk Uncertain. 16 (2), 165–172.

Coombs, Clyde H., Huang, Lily C., 1976. Tests of the betweenness property of expected utility. J. Math. Psychol. 13 (3), 323–337.

Crawford, Vincent P., 1990. Equilibrium without independence. J. Econ. Theory 50 (1), 127–154.

Dekel, Eddie, 1986. An axiomatic characterization of preferences under uncertainty: weakening the independence axiom. J. Econ. Theory 40 (2), 304–318.

Delquié, Philippe, Cillo, Alessandra, 2006. Expectations, disappointment, and rank-dependent probability weighting. Theory Decis. 60 (2–3), 193–206.

Eisenhuth, Roland, 2019. Reference-dependent mechanism design. Econ. Theory Bull. 7 (1), 77–103.

Eliaz, Kfir, Ray, Debraj, Razin, Ronny, 2006. Choice shifts in groups: a decision-theoretic basis. Am. Econ. Rev. 96 (4), 1321–1332.

Ellis, Andrew, 2016. Condorcet meets Ellsberg. Theor. Econ. 11 (3), 865–895.

Ericson, Keith M. Marzilli, Fuster, Andreas, 2011. Expectations as endowments: evidence on reference-dependent preferences from exchange and valuation experiments. Q. J. Econ. 126 (4), 1879–1907.

Feddersen, Timothy J., Pesendorfer, Wolfgang, 1996. The swing voter's curse. Am. Econ. Rev. 86 (3), 408–424.

Feddersen, Timothy, Pesendorfer, Wolfgang, 1997. Voting behavior and information aggregation in elections with private information. Econometrica, 1029–1058.

Feddersen, Timothy, Pesendorfer, Wolfgang, 1998. Convicting the innocent: the inferiority of unanimous jury verdicts under strategic voting. Am. Polit. Sci. Rev. 92 (01), 23–35.

Feddersen, Timothy J., Pesendorfer, Wolfgang, 1999. Abstention in elections with asymmetric information and diverse preferences. Am. Polit. Sci. Rev. 93 (02), 381–398.

Freeman, David J., 2019. Expectations-based reference-dependence and choice under risk. Econ. J. 01.

Ghosal, Sayantan, Lockwood, Ben, 2009. Costly voting when both information and preferences differ: is turnout too high or too low? Soc. Choice Welf. 33 (1), 25–50.

Gilboa, Itzhak, Schmeidler, David, 1989. Maxmin expected utility with non-unique prior. J. Math. Econ. 18 (2), 141–153.

Gill, David, Prowse, Victoria, 2012. A structural analysis of disappointment aversion in a real effort competition. Am. Econ. Rev. 102 (1), 469–503.

Goeree, Jacob K., Yariv, Leeat, 2015. Conformity in the lab. J. Econ. Sci. Assoc. 1 (1), 15–28.

Goidel, Robert K., Shields, Todd G., 1994. The vanishing marginals, the bandwagon, and the mass media. J. Polit. 56 (03), 802–810.

Granovetter, Mark, 1978. Threshold models of collective behavior. Am. J. Sociol., 1420–1443.

Heidhues, Paul, Kőszegi, Botond, 2008. Competition and price variation when consumers are loss averse. Am. Econ. Rev. 98 (4), 1245–1268.

Heidhues, Paul, Kőszegi, Botond, 2014. Regular prices and sales. Theor. Econ. 9 (1), 217–251.

Herweg, Fabian, Müller, Daniel, Weinschenk, Philipp, 2010. Binary payment schemes: moral hazard and loss aversion. Am. Econ. Rev. 100 (5), 2451–2477.

Hong, Chew Soo, 1983. A generalization of the quasilinear mean with applications to the measurement of income inequality and decision theory resolving the Allais paradox. Econometrica, 1065–1092.

Hung, Angela A., Plott, Charles R., 2001. Information cascades: replication and an extension to majority rule and conformity-rewarding institutions. Am. Econ. Rev. 91 (5), 1508–1520.

Kahneman, Daniel, Tversky, Amos, 1979. Prospect theory: an analysis of decision under risk. Econometrica, 263–291.

Karni, Edi, Safra, Zvi, 1989. Dynamic consistency, revelations in auctions and the structure of preferences. Rev. Econ. Stud. 56 (3), 421–433.

Kőszegi, Botond, Rabin, Matthew, 2007. Reference-dependent risk attitudes. Am. Econ. Rev., 1047–1073.

Krishna, Vijay, Morgan, John, 2011. Overcoming ideological bias in elections. J. Polit. Econ. 119 (2), 183–211.

Krishna, Vijay, Morgan, John, 2012. Voluntary voting: costs and benefits. J. Econ. Theory 147 (6), 2083–2123.

Ledyard, John O., 1984. The pure theory of large two-candidate elections. Public Choice 44 (1), 7–41.

Ledyard, John O., et al., 1981. The Paradox of Voting and Candidate Competition: a General Equilibrium Analysis. Technical report.

Machina, Mark J., 1982. "Expected utility" analysis without the independence axiom. Econometrica, 277–323.

Machina, Mark J., 1984. Temporal risk and the nature of induced preferences. J. Econ. Theory 33 (2), 199–231.

Masatlioglu, Yusufcan, Raymond, Collin, 2016. A behavioral analysis of stochastic reference dependence. Am. Econ. Rev. 106 (9), 2760–2782.

Moreno, Bernardo, Ramos-Sosa, María del Pino, 2017. Conformity in voting. Soc. Choice Welf. 48 (3), 519–543.

Nakajima, Daisuke, 2011. First-price auctions, Dutch auctions, and buy-it-now prices with Allais paradox bidders. Theor. Econ. 6 (3), 473–498.

Neilson, William S., 1994. Second price auctions without expected utility. J. Econ. Theory 62 (1), 136–151.

Niemi, Richard G., Bartels, Larry M., 1984. The responsiveness of approval voting to political circumstances. PS Polit. Sci. Polit. 17 (03), 571–577.

Nordhöy, Frode, 1962. Group Interaction in Decision Making Under Risk. PhD thesis. Massachusetts Institute of Technology.

Palfrey, Thomas R., Rosenthal, Howard, 1983. A strategic calculus of voting. Public Choice 41 (1), 7–53.

Palfrey, Thomas R., Rosenthal, Howard, 1985. Voter participation and strategic uncertainty. Am. Polit. Sci. Rev. 79 (01), 62–78.

Pruitt, Dean G., 1971. Choice shifts in group discussion: an introductory review. J. Pers. Soc. Psychol. 20 (3), 339.

Quiggin, John, 1982. A theory of anticipated utility. J. Econ. Behav. Organ. 3 (4), 323–343.

Rosato, Antonio, 2016. Selling substitute goods to loss-averse consumers: limited availability, bargains, and rip-offs. Rand J. Econ. 47 (3), 709–733.

Roux, Nicolas, Sobel, Joel, 2015. Group polarization in a model of information aggregation. Am. Econ. J. Microecon. 7 (4), 202–232.

Sarver, Todd, 2018. Dynamic mixture-averse preferences. Econometrica 86 (4), 1347–1382.

Schkade, David, Sunstein, Cass R., Kahneman, Daniel, 2000. Deliberating about dollars: the severity shift. Columbia Law Rev., 1139–1175.

Segal, Uzi, 1987. Some remarks on Quiggin's anticipated utility. J. Econ. Behav. Organ. 8 (1), 145–154.

Sobel, Joel, 2014. On the relationship between individual and group decisions. Theor. Econ. 9 (1), 163–185.

Stoner, James A.F., 1968. Risky and cautious shifts in group decisions: the influence of widely held values. J. Exp. Soc. Psychol. 4 (4), 442–459.

Stoner, James Arthur Finch, 1961. A Comparison of Individual and Group Decisions Involving Risk. PhD thesis. Massachusetts Institute of Technology.

Sunstein, Cass R., 2005. Why Societies Need Dissent. Harvard University Press.

Sunstein, Cass R., Hastie, Reid, Payne, John W., Schkade, David A., Viscusi, W., 2002. Punitive Damages: How Juries Decide. University of Chicago Press, Chicago.

Sydnor, Justin, 2010. (Over)insuring modest risks. Am. Econ. J. Appl. Econ. 2 (4), 177–199.

Taylor, Curtis R., Yildirim, Huseyin, 2010. A unified analysis of rational voting with private values and group-specific costs. Games Econ. Behav. 70 (2), 457–471.

Wakker, Peter, 1994. Separating marginal utility and probabilistic risk aversion. Theory Decis. 36 (1), 1–44.

Wakker, Peter P., 2010. Prospect Theory: For Risk and Ambiguity. Cambridge University Press.