Online Supplementary Appendix

Squid Voting Game

Rational Indecisiveness in Sequential Voting

Yuichiro Kamada[†] Yosuke Yasuda[‡]

December 30, 2022

B Inconsistent Strategies under General Asymmetric Voting

Consider the general voting model with $(p, K, L, \varepsilon_a, \varepsilon_b)$. Let

$$\bar{R} := 1 - \min\{p(1-\varepsilon_a)^N, (1-p)(1-\varepsilon_b)^N\}.$$

Note that in the symmetric case (i.e., $p = \frac{1}{2}$, $K = L = n+1$, and $\varepsilon_a = \varepsilon_b = \varepsilon$), $\bar{R}$ reduces to $\bar{R}_n$ that we defined in Section 4.1.

Consider a social choice function $f : \{a, b, \emptyset\}^{2n+1} \to \Delta(\{A, B\})$ that returns a probability distribution over outcomes for each realization of signals of all the N voters. Note that for any strategy profile σ in our original game, there is a social choice function f that achieves the same distribution over outcomes conditional on any realization of the state and the signals. Let F be the space of all social choice functions. Notice that a social choice function $f \in F$ determines the ex ante expected payoff to each voter. Let this payoff be $R(f)$.

Theorem 9. *Consider the general voting model with $(p, K, L, \varepsilon_a, \varepsilon_b)$.*

1. *For any $(p, K, L, \varepsilon_a, \varepsilon_b)$, there is an equilibrium σ that achieves the payoff $\bar{R}$.*
2. *For all $f \in F, R(f) \leq \bar{R}$.*
3. *(a) If $K \geq 3$, then there is $\bar{\varepsilon} > 0$ such that for all $\varepsilon_a \in (0, \bar{\varepsilon})$, there is an equilibrium that achieves the payoff of $p\varepsilon_a + 1 - p$.*

[†]Haas School of Business, University of California Berkeley, 2220 Piedmont Avenue, Berkeley, CA 94720-1900, USA, and University of Tokyo, Faculty of Economics, 7-3-1 Hongo, Bunkyoku, Tokyo, 113-0033, Japan. E-mail: y.cam.24@gmail.com

[‡]Osaka University, Department of Economics, 1-7 Machikaneyama, Toyonaka, Osaka 560-0043 Japan. E-mail: yosuke.yasuda@gmail.com

(b) If $L \geq 3$, then there is $\bar{\varepsilon} > 0$ such that for all $\varepsilon_b \in (0, \bar{\varepsilon})$, there is an equilibrium that achieves the payoff of $p + (1-p)\varepsilon_b$.

4. *The expected payoff from any pure strategy equilibrium is at least* $\min\{p+(1-p)\varepsilon_b, p\varepsilon_a + 1 - p\}$, *that is, it is*

$$\begin{cases} p + (1-p)\varepsilon_b & \text{if } p(1-\varepsilon_a) < (1-p)(1-\varepsilon_b) \\ p\varepsilon_a + 1 - p & \text{if } p(1-\varepsilon_a) > (1-p)(1-\varepsilon_b) \end{cases}.$$

Recall that the welfare for the consistent strategy equilibria depend on K and L (Corollary 2). Parts 2 and 4 of the above theorem show that the welfare bounds for inconsistent strategy equilibria do not depend on K or L.

Proof. **<u>Part 1:</u>** Fix $(p, \varepsilon_a, \varepsilon_b)$. Without loss of generality, assume $p(1-\varepsilon_a)^N \leq (1-p)(1-\varepsilon_b)^N$. For any K and L, the following strategy profile $\bar{\sigma}$ is an equilibrium and achieves the ex ante payoff $\bar{R}$.

- For any $k = 1, \ldots, K-1$, voter k chooses B if she receives a; otherwise she plays A.
- For any $k = K, \ldots, N$, voter k uses a consistent strategy. If he does not receive a signal, then (i) k plays B if he observes that voters $1, \ldots, K-1$ play A and voters $K, \ldots, k-1$ play B, and (ii) k plays A otherwise.

Note that this is a pure strategy profile and is not consistent. Its default sequence is $(A, \ldots, A, B, \ldots, B)$, where A continues $K-1$ times and then B continues L times. To see that each voter chooses a best response, observe first that it is immediate that any voter who receives a signal is taking a best response. Also, observe that once there is a deviation from the default sequence, then voters take a best response given a belief that the first deviator from the default sequence has received a signal and followed the equilibrium strategy (such a belief can easily be shown to be consistent in the sense of Kreps and Wilson (1982)). So suppose that the voters 1 through $k-1$ have followed the default sequence, and suppose that voter k receives no signal. The posterior on α is $P(\alpha) = \frac{p(1-\varepsilon_a)^k}{p(1-\varepsilon_a)^k + (1-p)(1-\varepsilon_b)}$. If k follows the specified strategy, then her payoff is

$$P(\alpha)(1 - (1-\varepsilon_a)^{N-k}) + (1 - P(\alpha)) \cdot 1. \tag{10}$$

If k deviates, her payoff is $P(\alpha)$ if $k \geq K$. In this case, k is taking a best response if and only if

$$\begin{aligned} &P(\alpha)(1-(1-\varepsilon_a)^{N-k}) + (1-P(\alpha)) \geq P(\alpha) \\ \iff& 1-P(\alpha) \geq P(\alpha)(1-\varepsilon_a)^{N-k} \\ \iff& (1-p)(1-\varepsilon_b) \geq p(1-\varepsilon_a)^k(1-\varepsilon_a)^{N-k} \\ \iff& (1-p)(1-\varepsilon_b) \geq p(1-\varepsilon_a)^N. \end{aligned} \tag{11}$$

Since we assumed $p(1-\varepsilon_a)^N \leq (1-p)(1-\varepsilon_b)^N$, (11) holds if $(1-p)(1-\varepsilon_b) \geq (1-p)(1-\varepsilon_b)^N$. But this holds because $N \geq 1$. Hence, k is taking a best response.

If $k \leq K-1$, then her payoff from the deviation is

$$P(\alpha) + (1-P(\alpha))\tilde{Y},$$

where $\tilde{Y} := \text{Prob}(L-1 \text{ or more signals out of } N-1-k)$.

Thus, the payoff from playing A is no less than the payoff from playing B if and only if

$$P(\alpha)(1-(1-\varepsilon_a)^{N-k}) + (1-P(\alpha)) \cdot 1 \geq P(\alpha) + (1-P(\alpha))\tilde{Y},$$

or

$$P(\alpha) \leq \frac{1-\tilde{Y}}{1-\tilde{Y}+(1-\varepsilon_a)^{N-k}}.$$

Since $P(\alpha) = \frac{p(1-\varepsilon_a)^k}{p(1-\varepsilon_a)^k+(1-p)(1-\varepsilon_b)}$, this is equivalent to:

$$\frac{p(1-\varepsilon_a)^k}{(1-p)(1-\varepsilon_b)} \leq \frac{1-\tilde{Y}}{(1-\varepsilon_a)^{N-k}},$$

or

$$\frac{p}{1-p}\frac{(1-\varepsilon_a)^{N-1}}{(1-\varepsilon_b)} \leq 1-\tilde{Y}. \tag{12}$$

Now, note that $\tilde{Y}$ is equal to Y defined in the proof of Proposition 3 where we set $m=1$ and $i=k+1$, and the roles of A and B (and thus the roles of K and L, and of ε_a and ε_b) are reversed. Since that proof shows $(1-\varepsilon_a)^{L+m-i} \leq 1-Y$, we have

$$(1-\varepsilon_b)^{K+1-(k+1)} \leq 1-\tilde{Y},$$

or

$$(1-\varepsilon_b)^{K-k} \leq 1-\tilde{Y}.$$

Since $k \geq 1$ and $K \leq N$, we have $(1-\varepsilon_b)^{N-1} \leq (1-\varepsilon_b)^{K-k} \leq 1-\tilde{Y}$, which implies

$\frac{p}{1-p}\frac{(1-\varepsilon_a)^N}{1-\varepsilon_b} \leq 1-\tilde{Y}$ because we assumed $p(a-\varepsilon_a)^N \leq (1-p)(1-\varepsilon_b)^N$. This shows that k is taking a best response when $k \leq n$.

Finally, since the payoff is 0 under $\bar{\sigma}$ if and only if the state is α and no one receives a signal, the expected payoff is $1-p(1-\varepsilon_a)^N$.

Part 2: Fix any $f \in F$. Let $f(\emptyset,\ldots,\emptyset)(A)=q$. Then, conditional on the signal realization $(\emptyset,\ldots,\emptyset)$, the voters' payoff is q under state α and $1-q$ under state β. The probability that the signal profile $(\emptyset,\ldots,\emptyset)$ realizes is $(1-\varepsilon_a)^N$ under state α and $(1-\varepsilon_b)^N$ under state β. Hence, the ex ante expected payoff is at most

$$
\begin{aligned}
&p(1-(1-\varepsilon_a)^N)\cdot 1+(1-p)(1-(1-\varepsilon_b)^N)\cdot 1+\max_{q\in[0,1]}\left(pq(1-\varepsilon_a)^N+(1-p)(1-q)(1-\varepsilon_b)^N\right)\\
&=(1-p(1-\varepsilon_a)^N-(1-p)(1-\varepsilon_b)^N)+\max\{p(1-\varepsilon_a)^N,(1-p)(1-\varepsilon_b)^N\}\\
&=1+\\
&\quad\max\{p(1-\varepsilon_a)^N-p(1-\varepsilon_a)^N-(1-p)(1-\varepsilon_b)^N,(1-p)(1-\varepsilon_b)^N-p(1-\varepsilon_a)^N-(1-p)(1-\varepsilon_b)^N\}\\
&=1+\max\{-(1-p)(1-\varepsilon_b)^N),-p(1-\varepsilon_a)^N\}\\
&=1-\min\{(1-p)(1-\varepsilon_b)^N),p(1-\varepsilon_a)^N\}\\
&=\bar{R}.
\end{aligned}
$$

Part 3: We only provide the proof for part (3b). The proof for part (3a) is symmetric.

Consider the following strategy profile, which we denote by $\underline{\sigma}$.

- Voter $k=1,\ldots,K-1$: If all the actions observed so far are A, then play A irrespective of the signal. Otherwise, play a consistent strategy in which A is played if no signal is received.

- Voter $k=K,\ldots,N-1$: If all the actions observed so far are n times of A followed by $k-1-n$ times of B, then play B irrespective of the signal. Otherwise, play a consistent strategy in which A is played if no signal is received.

- Voters $k=N$: Play a consistent strategy in which A is played if no signal is received.

Note that this strategy profile has the ex ante payoff of $p+(1-p)\varepsilon_b$. We show that, if $n\geq 2$, $\underline{\sigma}$ constitutes an equilibrium with a belief that, after any deviation, any voter who has not received a signal assigns probability 1 to α (such a belief can easily be shown to be consistent in the sense of Kreps and Wilson (1982)).

Consider voter k. Fix any history of actions by the previous voters. First, suppose that k receives signal a. Then, since the outcome will be A if k follows $\underline{\sigma}_k$, playing $\underline{\sigma}_k$ induces the expected payoff of 1, which is the highest possible payoff in this game. Hence, $\underline{\sigma}_k$ is a best response.

Second, suppose that k receives signal b.

- Suppose that the voters so far have followed $\underline{\sigma}$.
  - If $1 \leq k \leq K-1$, then if she follows $\hat{\sigma}_k$ then her payoff is ε_b. If instead she plays B then at least $L-1$ voters from the set of subsequent voters have to receive signal b in order for k to expect the payoff of 1, and otherwise she receives the payoff of 0. Hence, her payoff is $O(\varepsilon_b^{L-1})$. Since $L-1 \geq 2$, this implies that there exists $\bar{\varepsilon} > 0$ such that for all $\varepsilon_b < \bar{\varepsilon}$, following $\underline{\sigma}_k$ is a best response.
  - If $K \leq k \leq N-1$, then playing A ends the game with outcome A, so it induces the expected payoff of 0. Hence, playing B is a best response.
  - If $k = N$, then playing B induces the payoff of 1 with probability 1, so it is a best response to play B.
- Suppose that there is a voter who has deviated from $\underline{\sigma}$. Then, the play by the subsequent voters will not be affected by the action taken by voter k. Hence, it is a best response for voter k to provide additional vote for B. Therefore, playing B is a best response.

Third, suppose that voter k did not receive a signal.

- Suppose that the voters so far have followed $\underline{\sigma}$. Then, the posterior belief on α is $\frac{1}{2}$.
  - Suppose $k \leq K-1$. If k follows $\underline{\sigma}_k$, then her payoff is $\frac{1}{2} + \frac{1}{2}\varepsilon_b$. If she does not follow $\underline{\sigma}_k$, then her payoff is $\frac{1}{2} + \frac{1}{2}O(\varepsilon_b^{L-1})$. Since $L-1 \geq 2$, following $\underline{\sigma}_k$ is a best response.
  - Suppose $K \leq k \leq N-1$. If k follows $\underline{\sigma}_k$, then her payoff is $\frac{1}{2} + \frac{1}{2}\varepsilon_b$. If she does not follow $\underline{\sigma}_k$, then her payoff is $\frac{1}{2}$. Hence, following $\underline{\sigma}_k$ is a best response.
  - Suppose $k = N$. Then, k is indifferent between the two actions. So following $\underline{\sigma}_k$ is a best response.
- Suppose that there is a voter who has deviated from $\underline{\sigma}$ before k's move. Then, the play by the subsequent voters will not be affected by the action taken by voter k. Hence it is a best response for k to provide additional vote for A because k's belie assigns probability 1 to α.

This completes the proof of part 3.

Part 4: The proof we present below is analogous to the one in the symmetric case (the proof of Proposition 2).

Fix a pure strategy equilibrium σ. Take the default sequence of σ. Suppose first that the last voter on this default sequence, whom we denote voter k^*, plays A.

First, we show the following lemma.

Lemma 7. *Suppose that the votes by voters* $1, \ldots, k^* - 1$ *have followed the default sequence of* σ*. Then, voter* k^* *votes for* B *if and only if she receives signal* b*.*

Proof of Lemma 7. Consider the history in which voters $1, \ldots, k^* - 1$ have followed the default sequence of σ. Given the assumption that k^* plays A given no signal, it suffices to check the cases when she receives a signal.

1. First, we show that k^* plays B if she receives signal b. We show this using induction on k. To do this, fix $k' \in \{k^*, \ldots, N\}$. Suppose, as an induction hypothesis, that for every $k > k'$, voter k plays B if he receives signal b.[1] Suppose that voter k' receives signal b. Given this signal, her posterior probability on β is 1. If k' votes for A, then her expected payoff is 0. If k' votes for B, then there is a strictly positive probability that all the subsequent voters receive signal b and, in that case, the game ends with outcome B by the induction hypothesis. Hence, her expected payoff is strictly positive. Thus, playing B is a unique best response for k', and hence, voter k' plays B if she receives signal b. This shows that voter k^* plays B if she receives signal b.

2. Next, we show that k^* plays A if she receives signal a. To see this, suppose to the contrary that k^* plays B if she receives signal a. In this case, the outcome must be A with probability 1 after k^* plays B because otherwise, playing B would give k^* a payoff strictly less than 1, while she would get the payoff of 1 if she played A, making her choice B suboptimal. This implies that the outcome will be A if k^* chooses B when the state is α, no matter what her signal is. Now, suppose that k^* did not receive a signal. Let $P(\alpha)$ be the posterior of k^* at such an information set. Then, the expected payoff of voter k^* is $P(\alpha)$ if she plays A. If she instead plays B, then the outcome will be A with probability 1 if the state is α as we have concluded. The argument in item 1 above implies that if the state is β and all the subsequent voters receive signal b, which happens with probability $\varepsilon^{2n+1-k^*} > 0$, then the outcome is B. Hence, the expected payoff of voter k^* is at least $P(\alpha) + (1 - P(\alpha))\varepsilon^{2n+1-k'}$ if she plays B, and this is strictly greater than $P(\alpha)$ because $1 - P(\alpha) > 0$ and $\varepsilon > 0$. Hence, playing A is

[1] Note that this hypothesis is vacuously true when $k' = N$.

suboptimal for k^* when she does not receive a signal, which contradicts the assumption that she votes for A given no signal. Thus, k^* plays A if she receives signal a.

□

Given Lemma 7, once k^* plays B, the subsequent voters assign posterior probability 1 to state β, so the outcome will be B with probability 1.

Lemma 8. *Under σ, if voter 1 receives signal a, then she expects that the outcome will be A with probability 1.*

Proof of Lemma 8. We use induction. Fix $k \leq k^*$. Suppose as an induction hypothesis that for every $k' \in \{k+1, \ldots, k^*\}$, if the actions by voters $1, \ldots, k'-1$ have followed the default sequence and k' receives signal a, then the outcome will be A with probability 1. Then, suppose that the actions by voters $1, \ldots, k-1$ have followed the default sequence and k receives signal a. If k plays the action specified in the default sequence, then either (i) no subsequent voters receive a signal, or (ii) there is at least one subsequent voter who receives signal a. In case (i), by the definition of the default sequence, the outcome will be A with probability 1. In case (ii), the first voter who receives signal a expects that the outcome will be A, which follows from the induction hypothesis. Hence, the outcome will be A in either case. Therefore, k's expected payoff is 1 if k plays the action specified in the default sequence. Since k under σ must be doing at least as good as playing any action, this implies that, under σ, k expects the payoff of 1, and hence she expects that the outcome will be A with probability 1. This completes the induction argument. Therefore, we have shown that, if voter 1 receives signal a, then she expects that the outcome will be A with probability 1. □

Lemma 9. *Under σ, if any voter receives signal b when the votes so far followed the default sequence, then she expects that the outcome will be B with probability at least ε_b.*

Proof of Lemma 9. We use induction. Fix $k \leq k^* - 1$. Suppose as an induction hypothesis that for every $k' \in \{k+1, \ldots, k^*-1\}$, if the actions by voters $1, \ldots, k'-1$ have followed the default sequence and k' receives signal b, then the outcome will be B with probability at least ε_b. Then, suppose that the actions by voters $1, \ldots, k-1$ have followed the default sequence and k receives signal b. If k plays the action specified in the default sequence, then either (i) no subsequent voters in $\{k+1, \ldots, k^*-1\}$ receive a signal, or (ii) there is at least one subsequent voter in $\{k+1, \ldots, k^*-1\}$ who receives signal b. In case (i), by Lemma 7, voter k^* will play B if she receives signal b, and thus there is probability ε_b that the outcome will be B as we concluded after Lemma 7. In case (ii), the first voter who receives signal b expects that the outcome will be B with a probability of at least ε_b, which follows from the induction hypothesis. Hence, the outcome will be B with at least ε_b probability in either case. Therefore, k's expected payoff under σ must be at least ε_b if she plays the action specified in the default sequence. Since k under σ must be doing at least as good as playing any action, this implies that, under σ, k expects a payoff of at least ε_b, and hence she expects that the outcome will be B with probability at least ε_b. This completes the induction argument. Therefore, we have shown that, if any voter receives signal b, then she expects that the outcome will be B with probability at least ε_b. □

Finally, consider voter 1. If the state is α, then either (i) no voters in $\{1, \ldots, k^*\}$ receive a signal, or (ii) there is at least one voter in $\{1, \ldots, k^*\}$ who receives signal a. In case (i), the outcome will be A with probability 1 by the choice of the default sequence. In case (ii), we have shown that the outcome will be A with probability 1. Hence, voter 1's expected payoff is 1 conditional on this event.

If the state is β, then either (i) no voters in $\{1, \ldots, k^* - 1\}$ receive a signal, or (ii) there is at least one voter in $\{1, \ldots, k^* - 1\}$ who receives signal b. In case (i), voter k^* will have seen the history that has followed the default sequence. If she receives no signal, which happens with probability $1 - \varepsilon_b$, then she plays A by assumption, and thus the outcome will be A with probability 1. If she receives signal b, which happens with probability ε_b, then the outcome will be B with probability 1 as we concluded after Lemma 7. Overall, voter 1 expects the probability that the outcome will be B to be ε_b conditional on (i). In case (ii), take the first voter who has received signal b. We have shown that this voter assigns a probability of at least ε_b that the outcome will be B. Thus, conditional on the event that (ii) realizes, voter 1 expects that the probability that the outcome will be B is at least ε_b. Therefore, if the state is β, the outcome will be B with at least ε_b probability.

Overall, since the prior on the states α and β are p and $1 - p$, respectively, voter 1's expected payoff is at least

$$p \cdot 1 + (1 - p) \cdot \varepsilon_b = p + (1 - p)\varepsilon_b.$$

Now, if the last voter on the default sequence plays B, a symmetric proof shows that voter 1's expected payoff is at least $p\varepsilon_a + 1 - p$. Overall, the minimum equilibrium payoff is $\min\{p + (1 - p)\varepsilon_b, p\varepsilon_a + 1 - p\}$.

Since $p + (1 - p)\varepsilon_b < p\varepsilon_a + 1 - p$ if and only if $p(1 - \varepsilon_a) < (1 - p)(1 - \varepsilon_b)$, the proof is complete. □