Search by Committee*

James Albrecht[†]
Georgetown University and IZA
albrecht@georgetown.edu

Axel Anderson
Georgetown University
aza@georgetown.edu

Susan Vroman
Georgetown University and IZA
vromans@georgetown.edu

JEL Codes: D72 and D83
Key Words: Search, Voting, Committees

Abstract

We consider the problem of sequential search when the decision to stop is made by a committee and show that a unique symmetric stationary equilibrium exists given a log concave distribution of rewards. We compare search by committee to the corresponding single-agent problem and show that committee members are less picky and more conservative than the single agent. We show how (i) increasing committee size holding the plurality fraction constant and (ii) increasing the plurality rule affect the equilibrium acceptance threshold and expected search duration. Finally, we show that unanimity is optimal if committee members are sufficiently patient.

*We thank Lones Smith for many helpful comments, Rob Shimer for spotting an error in an earlier version of the paper, Frits Beuker for providing a key step in one of our proofs, as well as two anonymous referees. We also received helpful comments at the UPenn Search and Matching Workshop, at the IZA Search conference, at the Tinbergen Institute, and at Georgetown University.

[†]Manuscript Correspondence: James Albrecht; albrecht@georgetown.edu; Telephone: (202)687-6105; Fax: (202)687-6102.

"A committee is a group that keeps minutes and loses hours." - Milton Berle

1 Introduction

In the classic sequential search problem, an individual makes one draw per period from an exogenous and known distribution. These draws are independently and identically distributed through time. After every draw, there is a decision to be made – to stop and accept the payoff given by the realization of the most recent draw or to continue searching. The benefit of further search is the expectation that a higher payoff will eventually be realized; the cost is that the searcher's enjoyment of this payoff will be delayed. Elaborations and applications of this optimal stopping framework abound in economics. See [13], [16], and [17] for a set of excellent surveys.

Almost all of the search literature to date has a common feature, namely, that the stopping decision is made by a single agent.[1] For many applications, this assumption is a good one, but often the decision to stop or to continue searching is made by a group of agents. Consider, for example, an academic department that brings a sequence of candidates to campus to interview for an open position, and suppose that after each visit, a decision is made whether to make an offer to the latest candidate or to continue the search. This decision is typically taken by a group of faculty. Or, consider a couple that is shown a sequence of rental properties. After each property is observed, the couple decides whether to accept the latest apartment they have seen or to continue their search. It is, of course, easy to think of other situations in which a group makes a stopping decision.

We model this group decision as a problem of *search by committee.* As in the single-agent problem, in each period, a group of agents (the committee) is presented with an option. The values that the committee members place on this option are draws from an exogenous and known distribution, and these draws are *iid* across committee members and over time. In each period, the committee votes whether to stop or to continue.

[1]The single exception that we know of is a 2008 unpublished paper by O. Compte and P. Jehiel, "Bargaining and majority rules: A collective search perspective," which we discuss below. At first glance, the problem of household search considered in the 2009 unpublished paper, "Joint-search theory: New opportunities and new frictions" by B. Guler, F. Guvenen and G. Violante might also seem to be an exception. However, they use a unitary model; that is, their household behaves as if it were a single agent.

Specifically, we consider a committee with N members and we suppose that at least M votes are required to stop. The voting game played by the committee aggregates its members' preferences. If at least M members find the current option acceptable, the search stops; otherwise, the search continues. Our approach thus combines two literatures, sequential search and private-values voting.

What do we learn from this combination? *First*, we show that the problem of search by committee is well posed. A symmetric stationary equilibrium exists and is unique given a log concavity assumption on the distribution of payoffs.

Second, we compare the outcomes of committee and single-agent search. Suppose the committee and the single agent face the same environment, i.e., the same distribution of payoffs and the same cost of delay. We show that committee members are always *less picky* than a single agent would be in the same environment in the sense that the acceptance threshold set by the committee members is always less than the threshold the single agent would set. In equilibrium, the acceptance threshold equals the discounted value of continuing to search; thus, the expected discounted payoff to a committee member is less than the corresponding value achieved by the single searcher. The fact that committee members set a lower acceptance threshold than a single agent need not imply, however, that expected search duration is shorter for a committee than it is for the single agent. In fact, the comparison between the expected duration of search for a committee and for a single searcher depends on the cost of delay in an interestingly non-monotonic way. Specifically, so long as unanimity is not required, the committee can expect to end its search faster than a single agent would if the cost of delay is either sufficiently low or sufficiently high. We also show that a standard result from the single-agent search literature, namely, that a single agent raises his or her acceptance threshold in response to a mean preserving spread in the distribution of rewards, can be reversed in the search-by-committee problem. In this sense, committee search is *more conservative* than is single-agent search. The two results, *less picky* and *more conservative*, follow from two fundamental elements of the search-by-committee problem. First, committee members impose externalities on each other that are by definition absent in the single-agent problem. (See [5] for a related point in a static voting game.) Second, the voting game played by its members defines a value function (the expected discounted payoff) for the committee. The value function in the single-agent problem is necessarily convex, but the committee value function is necessarily not convex.

Third, we examine the effects of changing the size of the committee on the acceptance threshold and on expected search duration. We do this holding the fraction of votes required to stop fixed; that is, we increase both M and N while holding M/N fixed. For example, we examine the effects of moving from a situation in which at least 2 out of 3 votes are required to stop to one in which at least 4 out of 6 votes are needed. We show that increasing committee size in this way leads to a decrease in the acceptance threshold. Equivalently, committee members are worse off as the size of the committee increases. We also show that, when unanimity is not required, increasing the size of the committee while holding the plurality fraction fixed decreases expected search duration so long as committee members are sufficiently impatient; on the other hand, increasing the size of the committee increases expected search duration when unanimity is required to stop.

Finally, we consider the effect of varying the plurality rule, M, holding committee size fixed. We show that expected search duration is always increasing in M. Starting with low values of M, the acceptance threshold increases as the required number of votes increases. However, if at some point an increase in M leads to a decrease in the acceptance threshold, then further increases in M also cause the acceptance threshold to fall. We also show that the welfare-maximizing choice of M increases as committee members become more patient and that unanimity is optimal for high enough (but bounded) rates of patience. The idea that unanimity can be optimal is in contrast to a standard result (e.g., [10]) from the common-values voting literature, albeit in a different context.

To make progress on a new problem, we have made simplifying assumptions. On the search side, we restrict our attention to the stationary sequential problem, and we assume that once an option is discarded it is lost forever to the committee (no recall). These assumptions are close to those of [14]. That is, we use a simple, one-sided search framework and do not embed the search-by-committee problem in a market environment in which the distribution of payoffs is endogenously generated by the actions of agents on the other side of the market. On the voting side, we restrict our attention to the private-values case, in which the values that committee members place on the option at hand are *iid* draws.[2] Thus, we do not allow for the possibility that voting can convey information. More fundamentally, we restrict the strategies available to committee members in the voting game. We consider Markovian strategies in which each committee member's vote (stop or continue) depends only on the option at hand. Thus our voting model harks back to the pioneering work of Hotelling and Black ([12] and [4]).[3]

As mentioned in footnote 1, Compte and Jehiel also consider a collective search problem. In both our paper and theirs, a committee considers a sequence of *iid* draws, and the problem is to find an equilibrium in the committee members' stopping rules. In our model, each committee member's reaction to an option is the realization of an idiosyncratic random variable. Compte and Jehiel consider our case, but only for a uniform distribution of rewards. They also allow the possibility that the committee members may have intrinsic differences in their tastes. However, this increased generality comes at a cost – almost all of their interesting results require the discount rate to go to one, i.e., the situation in which there is no time cost to search. In contrast, our results are not restricted to the limiting case. Indeed, some of our more surprising results depend critically on allowing for low discount rates (Propositions 3 and 4). Relative to their paper, our contribution is to give a much more general and complete analysis of the "symmetric" case. As this is the natural generalization of the canonical McCall sequential search problem, our primary contribution is to the search literature. In contrast, their contribution is primarily to the bargaining literature. In this sense, the two papers are natural complements.

In the next section, we describe our model and prove the existence of a unique symmetric stationary equilibrium. In Section 3, we compare search by committee to single-agent search and show that committee members are *less picky* and *more*

[2]Committee members who share a sense of purpose could be modeled as making positively correlated draws. A general approach to this problem would therefore posit affiliated values, as is often done in auction theory (see [15]). Two polar cases would then be perfectly positively correlated draws (common values) and our case of *iid* draws. The former coincides with the single-agent search problem. The case of *iid* draws affords two advantages. First, it is far more tractable than affiliated values. Second, it is the unexplored polar case, and the results in an affiliated values model are surely a blend of our new results and those of the well understood single-agent problem.

[3]We follow the terminology of Black, who also calls a collection of voters a *committee*. In his model, the committee decides between a proposal and the status quo. In our model, the proposal is the current option and the status quo is continuing to search. In [4], the status quo is exogenous, while the value of the "status quo" is endogenous in our model.

conservative. We elaborate on the two crucial aspects of the search-by-committee problem that drive these results, *externalities* and *non-convexities*, in the context of the simplest possible committee, namely, the case of $N = 2$. Section 4 explores how the acceptance threshold and expected search duration vary in committee size (holding the plurality fraction fixed) and in the number of votes required to stop. Finally, in Section 5, we conclude.

2 The Model

2.1 Assumptions

A *committee* is a pair (N, M), where N is the number of members and M is the number of votes required to end the search. Time is discrete, and all committee members discount the future at common rate $\delta \in (0, 1)$. In each period, the committee is presented with an option. Each committee member then draws a value for the option from a continuous cdf $F : [0, 1] \mapsto [0, 1]$ with positive density f. These values can represent von Neumann-Morgenstern utility payoffs or monetary payoffs for risk neutral agents. We assume value draws are *iid* both across time and across committee members. We rule out side payments, i.e., utility is non-transferable.

We assume that both $\int_0^z F(s)ds$ and $\int_0^z (1 - F(s))ds$ are log concave in z — for which it suffices that f be log concave, as is well-known. To see how we use these assumptions, define the truncated means:

$$\mu_h(z) \equiv E[X|X \geq z] \quad \text{and} \quad \mu_\ell(z) \equiv E[X|X < z].$$

The above log-concavity assumptions imply that $\mu_h'(z) \leq 1$ and $\mu_\ell'(z) \leq 1$ as is shown in [6].[4] We use these two upper bounds to establish uniqueness of equilibrium and also to sign some of our comparative statics results. Log concavity assumptions are common in many economic applications (search, signaling, mechanism design, etc.). See [3] for a survey of log concavity results and applications.

Each period is divided into two stages. In the first stage, the option arrives and each committee member's value is realized. In the second stage, the committee decides

[4]Similar results can be found in [11] and in [19].

whether to stop searching and accept the most recently observed option or to continue to the next period. By restricting voting to stopping with the most recent option or continuing to search we are ruling out *recall*; i.e., once an option is discarded it is lost forever.[5] We model this choice using a simple voting mechanism: the committee members simultaneously vote either to stop and accept the current option or to continue to search. The search ends if and only if at least M committee members vote to stop. Each committee member seeks to maximize his or her own discounted payoff.

A strategy for committee member i is a sequence of functions $\sigma_i = \{\sigma_i(t)\}_t$, such that $\sigma_i(t)$ maps from possible histories through time t to the set {continue, stop}. Player i employs a Markov strategy if $\sigma_i(t)$ is only a function of the most recently evaluated option. We restrict attention to symmetric stationary equilibria in which all players employ the same Markov strategy.

We assume that the above description of the model is common knowledge among the committee members. Given our Markovian restriction, whether individual draws are private or public information is immaterial. Further, whether an agent knows in advance that he or she is pivotal is also irrelevant.

2.2 Equilibrium

Once we assume no recall and stationary Markov strategies, cutoff strategies are optimal for the same reason they are optimal in single agent search problems: the continuation value is a constant with respect to the current draw. Suppose all other committee members set acceptance threshold z'. We define $W(z, z', N, M, \delta)$ to be the expected continuation value starting just before draws are made for a committee member who sets acceptance threshold z for all future time periods while all other committee members use threshold z'. Then we have

$$\begin{aligned} W(z, z', N, M, \delta) &= P(z', N-1, M-1)\delta W(z, z', N, M, \delta) + (1 - P(z', N-1, M))E[X] \\ &\quad + p(z', N-1, M-1)\int \max\{\delta W(z, z', N, M, \delta), x\} dF(x), \end{aligned}$$

[5] In single-agent search, the no-recall assumption is without loss of generality. However, search by committee is a game, and thus non stationary strategies can potentially be supported in equilibrium by conditioning on past history. With recall, the state variable is the entire past history of draws, so the no-recall assumption is important in the committee search problem. The assumption of Markov strategies, however, takes the bite out of the no-recall assumption in the committee search problem.

where $p(z', N-1, i) = \binom{N-1}{i}(1-F(z'))^i F(z')^{N-1-i}$ is the (binomial) probability that exactly i of the other $N-1$ members draw a value greater than or equal to z' and $P(z', N-1, i) = \sum_{j=0}^{i-1} p(z', N-1, j)$ is the probability that $i-1$ or fewer of the other $N-1$ members vote to stop. Notice that if $p(z', N-1, M-1) = 0$, i.e., if the probability that this committee member is pivotal equals zero, his or her payoff is independent of z. In this case, we impose the refinement that the committee member chooses z to solve $\max\{\delta W(z, z', N, M, \delta), x\}$ when the reward drawn is x.

We seek a symmetric stationary equilibrium, so we define $V(z, N, M, \delta) = W(z, z, N, M, \delta)$. Substituting and simplifying yields

$$V(z, N, M, \delta) = P(z, N, M)\delta V(z, N, M, \delta) + (1 - P(z, N, M))\Omega(z, N, M),$$

where $\Omega(z, N, M)$ is the expected payoff conditional on stopping. Rearranging, we have

$$V(z, N, M, \delta) = S(z, N, M, \delta)\Omega(z, N, M), \quad \text{where} \quad S(z, N, M, \delta) = \frac{1 - P(z, N, M)}{1 - \delta P(z, N, M)}.$$

An agent's payoff is only affected by the cutoff chosen when he or she is pivotal, i.e., when exactly $M-1$ of the remaining $N-1$ agents vote to stop. Thus, best responses are determined by considering a pivotal voter, who can choose either to stop and accept the most recently drawn reward, x, or instead continue. The pivotal agent solves $\max\{x, \delta V(z, \cdot)\}$. Since the continuation value, $V(z, \cdot)$, is a constant with respect to the current period's decision, the optimal strategy is an acceptance threshold, i.e., vote for an option *iff* its value exceeds some threshold. Thus, if all members set threshold z in the future, the acceptance threshold for the pivotal voter equates the values of stopping and continuing, given that he or she is pivotal now, i.e., $x(z) \equiv \delta V(z, \cdot)$.

For an equilibrium to be stationary and symmetric the following *equilibrium condition* must be satisfied:[6]

$$x^* = \delta V(x^*, \cdot).$$

In equilibrium, $V(x^*, \cdot)$ is the expected payoff for each committee member. Equivalently, since $V(x^*, \cdot)$ is proportional to x^*, the equilibrium acceptance threshold can be

[6] The equilibrium refinement discussed above excludes the trivial equilibria in weakly dominated strategies in which all agents set a threshold of zero or a threshold of one.

used as a measure of each committee member's welfare.

To establish existence and uniqueness and to characterize equilibrium, our first step is to show that $\Omega(x, N, M)$ is a weighted average of $\mu_h(x)$ and $\mu_\ell(x)$ and to bound its derivative with respect to x. These results are given in Lemma 1. The proof is given in the Appendix.

Lemma 1 *The stopping value $\Omega(x, N, M)$ is a weighted average of the truncated means μ_h and μ_ℓ, and $0 \leq \Omega_x(x, N, M) \leq 1$.*

Having characterized $\Omega(x, N, M)$, we now state the main result of this section.

Proposition 1 *A symmetric stationary Markov equilibrium exists and is unique.*

The details of the proof are in the Appendix. However, the basic idea is straightforward, as shown in Figure 1. Existence follows from $V(0, \cdot) = E[X] > 0$, $V(1, \cdot) = 0$, and $V(x, \cdot)$ continuous in x. To establish uniqueness, we need to show that $\delta V(x, \cdot)$ crosses the 45-degree line only once, i.e., that $\delta V_x(x, \cdot) < 1$. This last condition follows from the inequality $\Omega_x(x, N, M) \leq 1$ established in Lemma 1 and $\delta \in (0, 1)$.

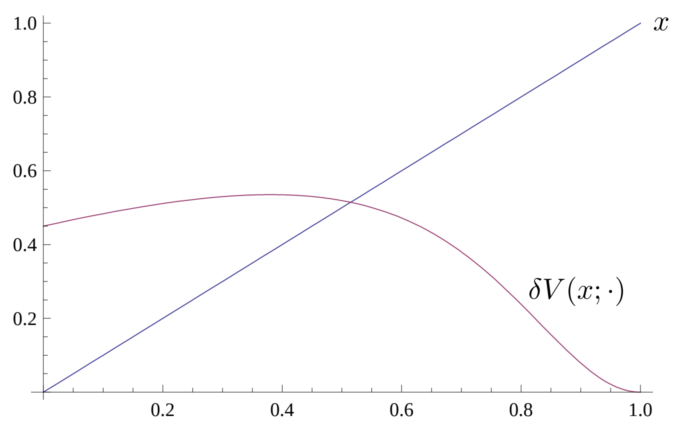

Figure 1: We have graphed $0.9V(x, 2, 2, 0.9)$ and X uniform on $[0, 1]$ to illustrate our equilibrium: $x^* \approx 0.515$.

3 Comparison to Single-Agent Search

In this section, we compare the symmetric stationary committee search equilibrium outcome with the outcome in the case of a single searcher. We first show that committee members are less picky than a single searcher would be in the same situation; i.e., the committee members set a lower acceptance threshold in equilibrium. We then show that committees conclude search more quickly for extreme (high or low) rates of patience but that the comparison can be reversed for intermediate rates of patience. Finally, we establish that committees are more conservative than single searchers in the sense that a mean preserving spread in F can lower the committee members' equilibrium acceptance threshold.

3.1 Committees are Less Picky

For fixed δ and F, we say that committee members are *less picky* than a single agent if the equilibrium acceptance threshold of the committee members is lower than the acceptance threshold that a single searcher would set. We establish the following in the Appendix.

Proposition 2 (Less Picky) *Committee members are less picky than a single searcher.*

The idea of the proof is straightforward. The single searcher can achieve at least as high an expected payoff as a committee member by mimicking the committee behavior but can then improve his or her payoff by optimizing. Since in equilibrium expected welfare is proportional to the acceptance threshold, committee members are less picky than a single searcher.

The intuition for Proposition 2 is as follows. A single searcher maximizes his or her continuation value, but a pivotal committee member cannot count on the committee doing so. More specifically, there are two negative externalities that committee members can impose on one another that do not arise in the single-searcher case. They can vote to stop when a committee member has drawn a low value or they can vote to continue when the committee member has drawn a high value. These externalities lead committee members to set a lower acceptance threshold than would a single searcher.[7]

[7]In the Börgers model of costly voting and voluntary participation ([5]), there is a single negative externality that arises because an individual's decision to participate makes it less likely that others

Note that this argument does not require independence of draws. Committees are less picky than single searchers even with correlated values, as long as the correlation is not perfect, at which point the single-agent and committee problems are identical.

3.2 Patient and Impatient Committees Conclude Search Faster

We know that committee members set an acceptance threshold that is below the corresponding threshold for a single searcher (holding F and δ constant), but this does not necessarily imply that the committee has a shorter expected search duration. In fact, the committee may take longer to search, as the following example illustrates.

Example 1 *Let X be uniform on $[0,1]$ and $\delta = 0.8$. Then a single searcher sets threshold 0.5, i.e., stops with probability 0.5. With $N = 5$ and $M = 4$, the equilibrium committee threshold is approximately 0.37, which yields a probability of stopping of approximately 0.39; i.e., the committee searches longer on average.*

The point illustrated by the above example is simple. We can, however, say something considerably less obvious about expected search duration. The comparison of expected search duration between a committee and a single searcher has an interesting non-monotonicity in δ. While expected search duration rises in δ for committees and for single searchers, the rate of change differs between committees and single searchers, so that the sign of the difference in expected search duration is not monotonic.

Our comparison of expected search duration between a committee and a single searcher relies on the following property of the binomial distribution:

Lemma 2 *If $N - n > M - m > 0$, $\exists\ \bar{x} \in (0,1)$ such that $P(x, N, M) - P(x, n, m)$ satisfies single crossing, negative for $x < \bar{x}$ and positive for $x > \bar{x}$.*

Note that the continuation probability for the single searcher is $F(x) \equiv P(x, 1, 1)$. The lemma is stated in a more general form because we use it again when we consider the implication for expected search duration of increasing committee size holding the plurality fraction constant.[8]

are pivotal and thus imposes a cost on them.

[8]Lemma 2 is related to a classic problem in probability theory. In 1693, the essayist Samuel Pepys asked Isaac Newton which of the following was most likely: (i) at least one “6” in 6 rolls of a fair die,

Proposition 3 *If $M < N$, then committees conclude search more quickly than individuals for sufficiently low* and *sufficiently high rates of patience.*

Formally, we show in the Appendix that there exists $0 < \delta_L \leq \delta_H < 1$ such that expected search duration is lower for the committee whenever $\delta \notin (\delta_L, \delta_H)$. When $\delta_L = \delta_H$, expected search duration is always lower for a committee.

To understand why committees conclude search faster for extreme rates of patience (and also why the comparison may be reversed for medium rates of patience), note that expected search duration depends on both the acceptance threshold (the threshold effect) and the probability of stopping given any acceptance threshold (the vote aggregation effect). Since committee members are less picky, the threshold effect always pushes committees toward concluding search faster. Thus, committees can only search longer if the vote aggregation effect is dominant and has the opposite sign of the threshold effect.

For low x, when $M < N$, Lemma 2 implies $P(x, N, M) < F(x)$. This means that whenever the committee and single-searcher acceptance thresholds are low enough, the vote aggregation effect reinforces the threshold effect, and the committee expects to conclude search faster. Low δ implies low acceptance thresholds, which implies that committees conclude search faster on average for low enough δ. This is illustrated in Figure 2. Above, we interpret low δ as impatience. Alternatively, since ours is a discrete-time model that does not specify the length of a period, we could interpret low δ as a situation in which there is a long interval between the arrival of options.

If δ is high, we cannot make the same argument. In fact, for high enough acceptance thresholds, the vote aggregation effect must have the opposite sign from the threshold effect (as long as $M > 1$). So the question becomes: which effect dominates? As $\delta \to 1$, the single-agent threshold goes to 1; thus, the probability of continuing goes to 1. As long as $M < N$, the committee threshold is bounded away from 1 as $\delta \to 1$ because as a weighted average of the truncated means, $\Omega < 1$. This in turn implies that the committee's probability of continuing is bounded away from 1. Thus, for sufficiently

(ii) at least two "6's" in 12 rolls of a fair die, or (iii) at least three "6's" in 18 rolls of a fair die. In our notation, if X is uniform on $[0, 1]$, the question is how $P(x, N, \alpha N)$ is ordered for $x = \alpha = 1/6$ and $N = 6, 12, 18$. As discussed in [18], the answer that Newton gave was correct, but his explanation only applies when the die is fair. More generally, it is interesting to know how $P(x, N, \alpha N)$ is ordered for all values of x, N and α. Lemma 2 addresses this question.

high δ, the threshold effect must dominate the vote aggregation effect. High δ can be interpreted either as patience or a situation in which options arrive quickly. For intermediate values of δ, the vote aggregation effect can be of the opposite sign and dominant as in Example 1. However, this intermediate range need not exist. For example, if $N = 2$ and $M = 1$ and values are distributed uniformly on $[0, 1]$, then the committee has a lower expected search duration than a single agent for *all* discount rates.

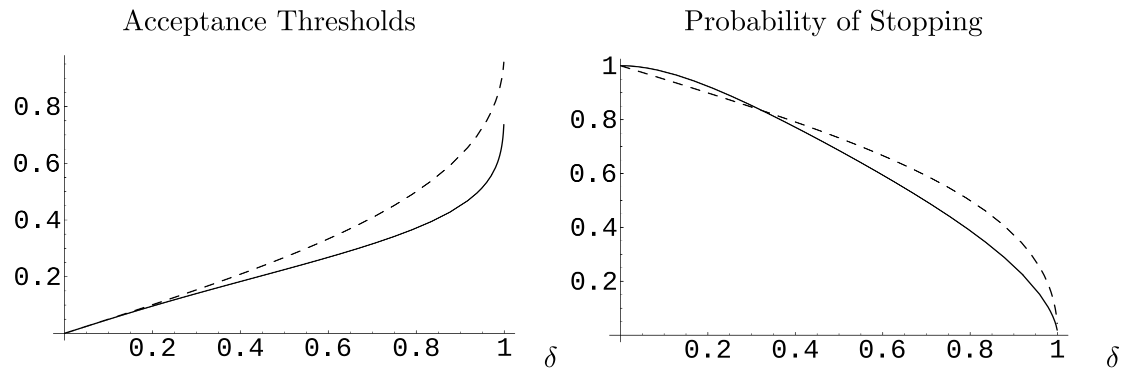

Figure 2: We compare the optimal single-agent solution to the committee equilibrium with $(N, M) = (5, 4)$ and X uniform on $[0, 1]$. In the graph on the left, we compare the single-agent acceptance threshold, $\tilde{x}$ (dashed line) to the committee threshold x^* (solid line). On the right, we compare the probability of stopping for a single agent, $1 - F(\tilde{x})$ (dashed line) to the probability of stopping for the committee, $1 - P(x^*, 5, 4)$ (solid line).

3.3 Committees are More Conservative

In this subsection, we show that committee members are *more conservative* than they would be were they single searchers. By more conservative we mean that a committee member may reduce his or her acceptance threshold in response to a mean preserving spread in the distribution of rewards, F. Equivalently, since the acceptance threshold gives each committee member's expected welfare, an increase in risk may make committee members worse off. In the single-agent search problem, however, mean preserving spreads in F are always good news, increasing continuation values and raising acceptance thresholds. To see that a mean preserving spread can reduce the acceptance threshold for a committee member, consider the case in which X is uniform on $[\underline{x}, \overline{x}]$. Let $\delta = 0.65$ and $M = N = 2$. If $[\underline{x}, \overline{x}] = [1, 3]$ then $x^* \approx 1.25$, while if $[\underline{x}, \overline{x}] = [0.5, 3.5]$ then $x^* \approx 1.22$.

The increase in the single searcher's acceptance threshold follows from considering the individual's value function ν which solves the recursion $\nu(x) = \max\{x, \delta E[\nu(X)]\}$. Since ν is the max of a constant and a linear function, ν is convex. As is well known, mean preserving spreads increase the expectation of a convex function, so that $E[\nu(X)]$ increases in mean preserving spreads. In turn, the acceptance threshold for single searchers (where $\widetilde{x} = \delta E[\nu(X)]$) must rise in mean preserving spreads.

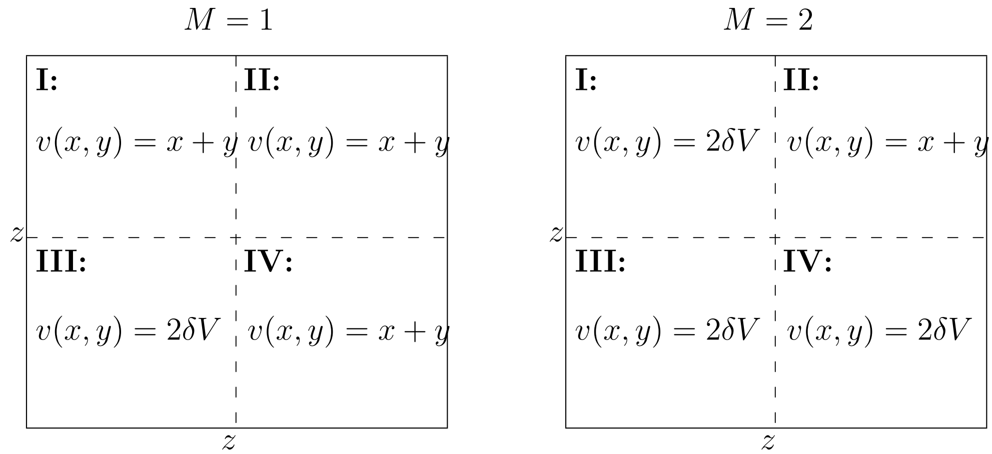

Figure 3: Joint value function. We show $v(x, y)$ in each region, given symmetric threshold z and continuation value V.

Why does this simple logic not carry over to the committee problem? To provide insight, let $N = 2$ and consider an analogous value function approach to the committee problem. Let $v(x, y)$ be the sum of the two committee members' payoffs given that they both set threshold z and draw values (x, y). Our continuation value function is related to $v(x, y)$ as follows:

$$V(z, \cdot) = \frac{1}{2} \int \int v(x, y) f(x) f(y) dx dy, \quad (1)$$

where $v(x, y) = x + y$ when the committee stops searching and $2\delta V(z, \cdot)$ otherwise. In Figure 3, we consider the function v. Any symmetric threshold z divides the unit square into four regions. If $M = 1$, then $(x, y) < (z, z)$ is the continuation region in which $v(x, y) = 2\delta V(z, \cdot)$, while $v(x, y) = x + y$ elsewhere. When $M = 2$, then $v(x, y) = x + y$ on the stopping region $(x, y) \geq (z, z)$, and $v(x, y) = 2\delta V(z, \cdot)$ elsewhere.

Negative externalities obtain in Regions I and IV. For example, when $M = 1$, in Region I, member 2 forces a conclusion to the search problem despite a relatively low draw by member 1. When $M = 2$, it is member 1 who imposes a negative externality on member 2 in Region I, forcing a continuation of search even though member 2 has drawn a high value.

Now consider how changes in F affect the continuation value and the threshold x^* by examining equation (1). As in the single-agent problem, continuation values and acceptance thresholds must be increasing in first order stochastic dominance changes in F since v is increasing in both x and y. What about mean preserving spreads in F? If v were convex in x and y, then we would get monotonicity in mean preserving spreads in F, as in the single-agent problem. However, it turns out that v is not convex in x and y.

To see the non-convexities, fix $M = 1$ and again consider Figure 3. Take any three pairs $(x, y), (x', y'), (x'', y'')$ and any $\lambda \in (0, 1)$ such that $(x, y) = \lambda(x', y') + (1 - \lambda)(x'', y'')$, $x' < x < z < x''$, and $y'' < y < z < y'$, i.e. (x, y) in Region III, (x', y') in Region I and (x'', y'') in Region IV. Then $v(x, y) = 2\delta V = 2z$ (in equilibrium) and

$$\lambda v(x', y') + (1 - \lambda)v(x'', y'') = \lambda(x', y') + (1 - \lambda)(x'', y'') = x + y.$$

Finally, $2z > x + y$ since z exceeds x and y individually. Thus, when $M = 1$, v is not everywhere convex in (x, y). To see that v is not convex in (x, y) when $M = 2$, choose pairs such that (x', y') is in Region I, (x'', y'') is in Region IV, and (x, y) is in Region II. Then $v(x', y') = v(x'', y'') = 2z < v(x, y) = x + y$.

While v is not convex in (x, y), given our independence assumption we could still prove monotonicity in mean preserving spreads in F if $v(x, y)$ were convex in one variable holding the other fixed (bi-convex). However, this weaker convexity condition is also not met. We show this in Figure 4, in which we graph $v(x, y)$ for $M = 1$ and $y < z$ (left) and $v(x, y)$ for $M = 2$ and $y > z$ (right). These non-convexities have an intuitive externality interpretation. For example, consider the $M = 1$, $y < z$ case. At

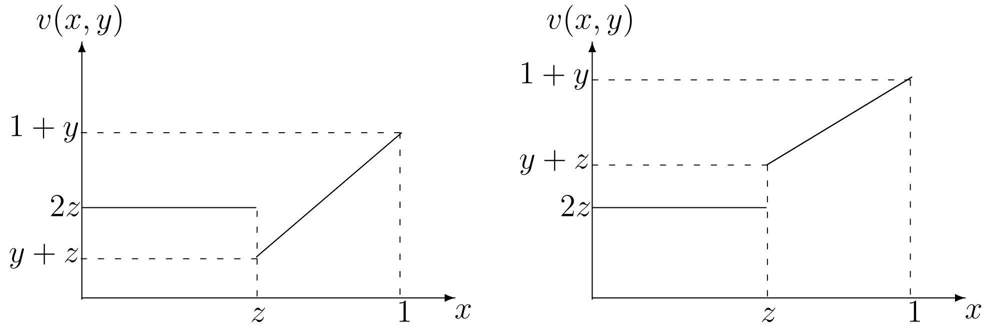

Figure 4: v is not bi-convex. We graph $v(x, y)$ as a function of x holding y fixed. In the graph on the left: $M = 1$ and $y < z$, while on the right: $M = 2$ and $y > z$.

the boundary $x = z$, agent 1 is just indifferent between continuing and stopping, while agent 2 strictly prefers continuing ($y < z$). Thus, agent 1 confers a negative externality on agent 2, and the joint payoff jumps down discontinuously. Similarly, when $M = 2$ and $y > z$, agent 2 strictly prefers to stop, so that agent 1 confers a negative externality on agent 2 by forcing continuation when $x < z$, and the joint payoff discontinuously jumps upward at $x = z$.

4 Effects of Committee Size and the Plurality Rule

In this section, we investigate the effects of committee size (N) and the plurality rule (M) on the equilibrium acceptance threshold and on expected search duration. Since, as we noted in Section 2, the equilibrium acceptance threshold is proportional to the expected payoff for each committee member, our results have both positive and normative implications.

4.1 Committee Size

We first examine the effects of changes in committee size. We do this holding the fraction of votes required to stop, $\alpha = M/N$, constant, recognizing that the committee problem is only defined when αN is an integer. This assumption means that as we increase committee size, we are proportionally increasing the number of votes required to stop. An increase in committee size, holding α constant, decreases the acceptance threshold; i.e., individual committee members have lower expected payoffs. When it comes to expected search duration, there is again both a threshold effect and a vote aggregation effect. The threshold effect unambiguously acts to decrease expected search duration as the size of the committee grows, while the vote aggregation effect reinforces the threshold effect at low values of x (equivalently, for low values of δ) but counteracts it at higher values of x. The pattern exhibited by the vote aggregation effect follows from Lemma 2. Our final result on the effects of a change in committee size is that, when unanimity is required to stop, increasing the size of the committee causes expected search duration to rise. We summarize the above discussion in the following Proposition.

Proposition 4 *If $M = \alpha N$, then x^* falls in N. If $\alpha < 1$, there exists a $\delta_L > 0$ such that expected search duration falls in N for $\delta < \delta_L$. If $\alpha = 1$, expected search duration rises in N.*

The result with respect to x^* presented in Propostion 4 is related to our result that committees are less picky. In a sense, the latter is more general than the result with respect to x^* in Proposition 4 in that it compares the single searcher (a committee with M and N equal to 1) to committees with any values for M and N, i.e., it does not require $\alpha = M/N$ constant.

The intuition for why x^* falls when N increases holding M/N fixed is the same as that underlying Proposition 2; namely, additional committee members imply additional externalities. The method of proof, however, is more involved. The idea behind our proof of Proposition 2 was simple – the single searcher can always mimic committee behavior. Since the single searcher can in fact do better than that, he or she must achieve a higher payoff; equivalently, $\widetilde{x} > x^*$. That logic cannot be carried over to our proof of Proposition 4. Suppose we compare a situation in which at least one vote out of two is required to stop to a situation in which at least two votes out of four are required. In order to mimic the behavior of the larger committee, the two members of the smaller committee would have to coordinate their behavior. Thus, we have to examine directly how $V(x^*, N, \alpha N, \delta) = S(x^*, N, \alpha N, \delta)\Omega(x^*, N, \alpha N)$ varies with N. The key step, which is nontrivial, is to show that $\Omega(x^*, N, \alpha N)$ falls in N. That is, as the size of the committee increases while holding the plurality fraction constant, the expected payoff per committee member conditional on stopping falls. The details are given in the appendix.

By Lemma 2, the probability of continuing exhibits a single crossing as N increases holding M/N constant. Thus, as in Proposition 3, the result for expected duration involves a comparison of two opposing effects – the vote threshold effect and the aggregation effect. For low δ, these reinforce each other so that expected duration falls as N increases. For high δ, however, the two effects are opposed and the resulting effect on expected duration is unclear.

4.2 Plurality Rule

In this subsection, we consider how the acceptance threshold, x^*, and expected search duration vary with the number of votes required to stop, M, holding committee size, N, constant. We first show that expected search duration is increasing in M. Second, we show that starting with low values of M, the acceptance threshold increases as M increases. However, if at some point an increase in M leads to a decrease in the acceptance threshold, then further increases in M also cause the acceptance threshold to fall. In short, the acceptance threshold is either everywhere increasing in M or is hump-shaped in M. Since, as we showed above, the acceptance threshold is proportional to the expected payoff for committee members, our results for x^* have implications for the optimal plurality rule. We explore these in Proposition 6 below. First, however, we summarize the effects of an increase in M on the acceptance threshold and expected search duration, holding N fixed.

Proposition 5 *Expected search duration is increasing in M. There exists a $k \in (0, N]$ such that x^* is increasing in M when $M < k$, while x^* decreases in M when $M > k$.*

The fact that increasing M raises expected search duration is not as obvious as it might appear at first glance. Of course, increasing M would necessarily increase expected search duration were x^* held fixed. However, particularly for high values of M, the acceptance threshold can fall as the plurality requirement increases. That is, we need to show that the vote aggregation effect, which is positive for all M, necessarily overwhelms the threshold effect when the two are opposed.

The intuition for the potentially non-monotonic effect of M on the acceptance threshold has to do with the two externalities in the search-by-committee problem that were discussed in Section 3. Changing a parameter of the problem – in this case the plurality rule – can alleviate one of these externalities at the cost of exacerbating the other. When M is low, the important externality is that relatively many committee members (as many as $N-M$) may be forced to stop when they would prefer to continue. The other externality – that relatively few committee members (fewer than M) may be forced to continue when they would prefer to stop – is less important. When M is high, the relative importance of these two effects is reversed. Increasing M reduces the effect of the first externality (the one that is especially costly when M is low) and increases the effect of the second externality (the one that is relatively unimportant when M is low). Thus, when M is low, an increase in the plurality rule makes committee members better off; equivalently, increases x^*. When M is high, the effect of an increase in the plurality rule on the second externality may become more important, i.e., x^* can fall.

The Optimal Plurality Rule

Since the acceptance threshold varies in a systematic way with the plurality rule, and since the expected payoff for committee members is proportional to x^*, we can set M to maximize the committee members' expected payoffs. We show that the welfare-maximizing choice of M increases with δ and that unanimity is optimal for high enough (but bounded) δ.[9]

Proposition 6 *The welfare-maximizing plurality rule, M, is weakly rising in the discount rate, δ. Given sufficient patience, unanimity is welfare maximizing.*

As above, we can gain some intuition by considering an externality interpretation. The relative costs of the two external effects (forcing your fellow committee members to stop when they would prefer to continue; forcing them to continue when they would prefer to stop) vary with the discount rate. When δ is low, the cost of being forced to stop on a low draw is low relative to the cost of being forced to continue on a high draw, so the optimal M is low. When committee members are patient, i.e., when δ is high, these relative costs are reversed. The optimal M is thus increasing in δ. If δ is high enough, the cost of being forced to continue on a high draw is low enough relative to the cost of being forced to stop on a low draw that the optimal plurality rule is unanimity. We have illustrated Propositions 5 and 6 in Figure 5.

[9] A similar result for the special case in which X is uniform on $[-1, 1]$ is proven in the Compte and Jehiel paper.

M	$\delta = 0.5$	$\delta = 0.7$	$\delta = 0.9$	$\delta = 0.99$
1	0.250004	0.350079	0.450767	0.496838
2	**0.250038**	0.350695	0.455635	0.507872
3	0.249935	**0.351811**	0.467617	0.539919
4	0.247398	0.347614	**0.474651**	0.591816
5	0.234549	0.325010	0.450359	**0.607018**
6	0.201928	0.276230	0.388368	0.552019
7	0.144056	0.198564	0.290138	0.445081

Figure 5: This table lists x^* for X uniform on $[0, 1]$ and $N = 7$ for different values of M and δ. In every column we highlight the highest value of x^* for the given δ.

We can compare our results on the optimal plurality rule to the extensive literature on voting with common values and private information, which began with Condorcet ([8]). To fix ideas, assume there are two choices, $\{c_1, c_2\}$, and two states of the world, $\{\omega_1, \omega_2\}$. Everyone agrees that choice c_i is optimal in state ω_i, but everyone has private information about the state of the world (think of a jury that would like to convict *iff* the defendant is guilty). In this context, Condorcet argued that for a "large population" (formally as $N \to \infty$), majority rule yields the correct choice with probability 1 (the Condorcet Jury Theorem), although he assumed sincere voting. Austen-Smith and Banks ([2]) proved the Condorcet Jury Theorem for strategic voting, i.e., assuming that pivotal voters correctly process the information contained in being pivotal. Feddersen and Pesendorfer ([10]) showed that unanimity is the *uniquely* suboptimal rule. That is, requiring any plurality fraction $\alpha < 1$ results in the correct action with probability arbitrarily close to 1 for high enough N, while if unanimity is required, the probability of a correct choice does not converge to 1.

Clearly, we have a very different model, one with private values and search externalities, while in the information aggregation literature, values are common and suboptimal decisions result from information externalities. Thus, we do not want to push the comparison too far. To the extent that the information aggregation literature's main message has been that requiring unanimity is (uniquely) suboptimal, we offer a contrasting message; namely, with search externalities, unanimity can be optimal.

5 Conclusion

In this paper, we analyze a new type of search problem, search by committee, in which the decision to stop or to continue searching is made by a group of agents. First, we show that the problem is well posed in that a symmetric stationary equilibrium exists and is unique given a log concavity assumption on the distribution of payoffs. We then show that agents in a committee are less picky than they would be were they searching on their own; that is, they set a lower acceptance threshold. This does not necessarily imply a lower expected search duration. We find that the expected search duration of a committee versus that of a single agent varies with the discount rate in a non-monotonic way. We also show that committee members are more conservative than a single-agent searcher in the sense that a mean preserving spread in the distribution of returns may make them worse off.

We then examine the effects of varying committee size and the plurality rule. We show that increasing the size of the committee holding the plurality fraction M/N constant decreases the acceptance threshold; equivalently the expected payoff per committee member falls. Assuming $M < N$, expected search duration is shorter for larger committees if the committee members are sufficiently impatient, but when unanimity is required, expected search duration increases with N. We also show that, holding N constant, expected search duration is increasing in M and that the acceptance threshold is either increasing in M or "hump shaped" in M. Finally, we find that the welfare-maximizing plurality rule M increases with the discount rate and that unanimity is optimal for sufficiently high (but bounded) discount rates.

The single-agent sequential search problem has been extended in many directions. One could do the same in the search-by-committee problem. We leave these extensions for later research since our aim here is to introduce and analyze the search-by-committee model in its most basic form.

Appendix: Proofs

Lemma 1 *The stopping value $\Omega(x, N, M)$ is a weighted average of the truncated means μ_h and μ_ℓ and $0 \leq \Omega_x(x, N, M) \leq 1$*

Step 1: Ω as a weighted average of (μ_h, μ_ℓ).

$$\Omega(x,N,M) \equiv \sum_{i=M}^{N} \frac{p(x,N,i)}{1-P(x,N,M)} \left[\frac{i}{N}\mu_h(x) + \frac{N-i}{N}\mu_\ell(x) \right]$$
$$\equiv w(x)\mu_h(x) + (1-w(x))\mu_\ell(x),$$

where

$$w(x) = \sum_{i=M}^{N} \frac{i\ p(x,N,i)}{N(1-P(x,N,M))} \equiv \sum_{i=M}^{N} \left(\frac{i}{N}\right) w_i(x)$$

STEP 2: $w'(x) < 0$.

Note that $w(x)$ is the $w_i(x)$-weighted average of the function i/N, which is increasing in i. Thus, if we can show that increasing x causes a first-order stochastic decrease in the weights $w_i(x)$, we are done. As is well known, the monotone likelihood ratio property implies first-order stochastic dominance. That is, if $w_i(x)/w_j(x)$ increases in x for all $i < j$, then $w'(x) < 0$.

$$\frac{w_i(x)}{w_j(x)} = \frac{\binom{N}{i}}{\binom{N}{j}} F(x)^{j-i}(1-F(x))^{i-j} \Rightarrow$$
$$\frac{\partial(w_i(x)/w_j(x))}{\partial x} = (j-i)f(x)\frac{\binom{N}{i}}{\binom{N}{j}} F(x)^{j-i-1}(1-F(x))^{i-j-1} > 0 \quad \forall\ i < j$$

STEP 3: LOG-CONCAVITY COMPLETES THE PROOF THAT $\Omega_x(x,N,M) \leq 1$.

$$\Omega_x(x,N,M) = w(x)\mu_h'(x) + (1-w(x))\mu_\ell'(x) + w'(x)(\mu_h(x) - \mu_\ell(x))$$
$$\leq w(x)\mu_h'(x) + (1-w(x))\mu_\ell'(x) \quad \text{(by } w'(x) \leq 0 \text{ and } \mu_h > \mu_\ell\text{)}$$
$$\leq 1 \quad \text{(by Log Concavity).}$$

Proposition 1 *A symmetric stationary Markov equilibrium exists and is unique.*

STEP 1: EXISTENCE.

We have $V(x,\cdot)$ continuous in x, $\delta V(0,\cdot) = \delta \int x f(x)dx > 0$, and $\delta V(1,\cdot) = 0$.

Step 2: Single Crossing $\Rightarrow$ Uniqueness.

Given $\delta V(0,\cdot) > 0$ and continuity of $V(x,\cdot)$ in x, if there are multiple equilibria, $\delta V_x(x^*,\cdot) \geq 1$ for at least one of them. Suppose that $\delta V_x(x,N,M,\delta) \geq 1$ and note that

$$V_\delta(x,\cdot) = \frac{P(x,\cdot)V(x,\cdot)}{1-\delta P(x,\cdot)}.$$

Then

$$V_{\delta x}(x,\cdot) = \frac{P(x,\cdot)V_x(x,\cdot)}{1-\delta P(x,\cdot)} + \frac{V(x,\cdot)P_x(x,\cdot)}{(1-\delta P(x,\cdot))^2} \geq 0,$$

where the inequality follows from the assumption that $V_x \geq 1$ and the fact that the continuation probability is increasing in x (i.e. $P_x \geq 0$). Thus, V_x is maximized at $\delta = 1$, but since $\delta V_x(x,N,M,1) = \delta\Omega_x(x,N,M) < 1$ (by Lemma 1) and $\delta \in (0,1)$, we have a contradiction. Thus, $\delta V_x(x^*,\cdot) < 1$ and there cannot be multiple equilibria. □

Proposition 2 (Less Picky) *Committee members are less picky than a single searcher.*

Proof: The single searcher can achieve at least as high an expected payoff as a committee member by mimicking the committee behavior. That is, assume a symmetric committee threshold x^* and let the single searcher commit to the following strategy:

1. Generate $N-1$ draws from a standard uniform distribution.
2. Stop whenever M or more of these draws exceed $1-F(x^*)$.
3. Continue if fewer than $M-1$ of these draws exceed $1-F(x^*)$.
4. Employ threshold x^* otherwise.

The single searcher, however, can do better. For example, suppose the single searcher draws $x < x^*$ and that M of the $N-1$ draws exceed $1-F(x^*)$. The single searcher can increase his or her expected payoff by continuing. The expected payoff for a single searcher thus exceeds the expected payoff for a committee member: $\widetilde{x} > x^*$. □

Lemma 2 *If* $N-n > M-m > 0$, $\exists\ \bar{x} \in (0,1)$ *such that* $P(x,N,M) - P(x,n,m)$ *satisfies single crossing, negative for* $x < \bar{x}$ *and positive for* $x > \bar{x}$.

Step 1: Likelihood Ratio Simplification.

We use the relationship between the binomial and beta cdfs (see, e.g., [7, p. 82]), namely,

$$P(x, N, M) = \frac{N!}{(N-M+1)!M!}\int_0^{F(x)} t^{N-M-1}(1-t)^M dt$$

$$\Rightarrow \quad P_x(x, N, M) = f(x)\frac{N!}{(N-M+1)!M!}F(x)^{N-M-1}(1-F(x))^M.$$

Thus, the ratio $P_x(x, N, M)/P_x(x, n, m)$ is proportional to

$$h(F) \equiv F(x)^{(N-n)-(M-m)}(1-F(x))^{M-m}.$$

Note that $h(0) = h(1) = 0$ and that $h(F)$ has a unique maximum in $(0, 1)$ at $F = ((N-n)-(M-m))/(N-n)$.

Step 2: The Single Crossing Result.

$$\lim_{F\to 0} h(F) = 0 \text{ and } P(0, \cdot) = 0.$$

By the continuity of $h(F)$, this implies that for x close to 0 (F close to 0), $P(x, N, M) < P(x, n, m)$. Similarly,

$$\lim_{F\to 1} h(F) = 0 \text{ and } P(1, \cdot) = 1 \Rightarrow P(x, N, M) > P(x, n, m) \text{ for } x \text{ near } 1.$$

Thus, by continuity, $P(x, N, M) = P(x, n, m)$ for at least one $x \in (0, 1)$, and if $P(x, N, M)$ and $P(x, n, m)$ cross more than once in $(0, 1)$, they must cross at least 3 times. Call these three crossing points $0 < F_1 < F_2 < F_3 < 1$. By the endpoint conditions, we must have $h(F_1) > 1$, $h(F_2) < 1$, and $h(F_3) > 1$. But h is continuous and unimodal when $N - n > M - m > 0$, a contradiction. □

Proposition 3 *If $M < N$, then committees conclude search more quickly than individuals for sufficiently low* and *sufficiently high rates of patience.*

Step 1: If $M = 1$, committees conclude search more quickly $\forall \delta \in (0, 1)$.

Trivially, $P(x, N, 1) < F(x) = P(x, 1, 1)$ for each $x \in (0, 1)$. Denote the single agent acceptance threshold by $\widetilde{x}(\delta)$. By Proposition 2, $x^*(N, 1, \delta) < \widetilde{x}(\delta)$, so $P(x^*(N, 1, \delta), N, 1) < F(\widetilde{x}(\delta))$.

Step 2: When $N > M \geq 2$, Committees search more quickly for low δ.

Given $N > M \geq 2$, $\exists\, \bar{x} \in (0, 1)$ s.t. $P(x, N, M) < F(x)$ $\forall x < \bar{x}$ by Lemma 2. Since $\widetilde{x}(\delta)$ is strictly increasing in δ, $\widetilde{x}(0) = 0$, and $\lim_{\delta \to 1} \widetilde{x}(\delta) = 1$, we can define δ_L by $\widetilde{x}(\delta_L) = \bar{x}$. Thus, for all $\delta < \delta_L$

$$\begin{aligned} \widetilde{x}(\delta) < \bar{x} \;&\Rightarrow\; P(\widetilde{x}(\delta), N, M) < F(\widetilde{x}(\delta)) \quad \text{(by Lemma 1)} \\ &\Rightarrow\; P(x^*(N, M, \delta), N, M) < F(\widetilde{x}(\delta)) \quad \text{(by Proposition 2).} \end{aligned}$$

Step 3: Committees conclude search more quickly for high δ.

If $M < N$, then $\Omega(x, N, M) < 1$ for all x. Thus, $x^*(N, M, \delta) < 1$ for all δ. This implies

$$\lim_{\delta \to 1} P(x^*(N, M, \delta), N, M) < 1 = \lim_{\delta \to 1} F(\widetilde{x}(\delta)),$$

Since $P(x^*(N, M, \delta), N, M)$ and $F(\widetilde{x}(\delta))$ are continuous in δ, $\exists\, \delta_H < 1$ such that $P(x^*(N, M, \delta), N, M) < F(\widetilde{x}(\delta))$ $\forall\, \delta > \delta_H$. $\square$

Proposition 4 *If $M = \alpha N$, then x^* falls in N. If $\alpha < 1$, there exists a $\delta_L > 0$ such that expected search duration falls in N for $\delta < \delta_L$. If $\alpha = 1$, expected search duration rises in N.*

Proof Preliminaries: While our committee problem is only defined for non negative integer values of M and N such that $M \leq N$, we can use the equivalence between the binomial and beta cdfs to write all our functions as continuous functions of N and M. Note that since we have set $M = \alpha N$ in this proposition, the partial derivatives that we discuss, $V_N(x^*, N, \alpha N, \delta)$, $S_N(x^*, N, \alpha N, \delta)$, and $\Omega_N(x^*, N, \alpha N, \delta)$, are calculated taking into account the dependence of M on N, i.e., we take the partial derivative of the function with respect to N plus α times the partial derivative with respect to M.

Step 1: $\Omega_N(x, N, \alpha N) \leq 0$.

Recall from Lemma 1, that $\Omega(x, N, \alpha N) = w(x)\mu_h(x) + (1 - w(x))\mu_\ell(x)$. Thus, we wish to show that $w(x)$ is decreasing in N when $M = \alpha N$. Let $q \equiv 1 - F(x)$.

$$
\begin{aligned}
w(q) &= \left[\sum_{i=M}^{N} \frac{i}{N}\binom{N}{i} q^i (1-q)^{N-i}\right] / \left[\sum_{i=M}^{N} \binom{N}{i} q^i (1-q)^{N-i}\right] \\
&= \left[\sum_{i=M}^{N} \binom{N-1}{i-1}\left(\frac{q}{1-q}\right)^i\right] / \left[\sum_{i=M}^{N} \binom{N}{i}\left(\frac{q}{1-q}\right)^i\right] \\
&= \frac{q}{1-q}\left[\sum_{i=M-1}^{N-1} \binom{N-1}{i}\left(\frac{q}{1-q}\right)^i\right] / \left[\sum_{i=M}^{N} \binom{N}{i}\left(\frac{q}{1-q}\right)^i\right] \\
&= \frac{M}{N}\frac{G(1, M-N, M, -q/(1-q))}{G(1, M-N, M+1, -q/(1-q))}.
\end{aligned}
$$

The final equality follows from the identity

$$\sum_{i=M}^{N} \binom{N}{i}(-z)^i = (-z)^M \frac{N!}{(N-M)!M!} G(1, M-N, M+1, z),$$

where G is Gauss' hypergeometric function, satisfying[10]

$$G(1, b, c, z) = \sum_{k=0}^{\infty} \frac{(b)_k}{(c)_k} z^k \quad \text{with} \quad (b)_k = \frac{(b+k-1)!}{(b-1)!} \quad \text{and} \quad (c)_k = \frac{(c+k-1)!}{(c-1)!}.$$

Let $z = -q/(1-q) < 0$ and recall that $M = \alpha N$. Then we complete the proof by showing: $G(1, (\alpha-1)N, \alpha N, z)/G(1, (\alpha-1)N, \alpha N + 1, z)$ falls in N.[11]

The hypergeometric series G satisfies the identity:

$$(z-1)G(1, b, c, z) + \left(\frac{b-c}{c}\right) zG(1, b, c+1, z) + 1 = 0,$$

or in our notation

$$(z-1)G(1, (\alpha-1)N, \alpha N, z) - (\frac{1}{\alpha})zG(1, (\alpha-1)N, \alpha N + 1, z) + 1 = 0$$

[10] There are many references for hypergeometric functions, e.g., [9] and [1].

[11] The proof that this ratio is falling in N was provided by Frits Beukers, Mathematics faculty, Utrecht University.

$$\Rightarrow \quad \alpha(1-z)\frac{G(1,(\alpha-1)N,\alpha N,z)}{G(1,(\alpha-1)N,\alpha N+1,z)} = -z + \alpha G(1,(\alpha-1)N,\alpha N+1,z)^{-1}.$$

Note that the LHS of the previous equality equals $(1-z)w(q)$. We can show that the RHS of the equality is decreasing in N by showing that $G(1,(\alpha-1)N,\alpha N+1,z)$ is increasing in N. We do this by using the Euler integral:

$$G(1,(\alpha-1)N,\alpha N+1,z) = \alpha N \int_0^1 (1-s)^{\alpha N-1}(1-zs)^{(1-\alpha)N} ds = \int_0^1 (1-z(1-y^{1/\alpha N}))^{(1-\alpha)N} dy$$

where the second equality follows from change of variable $y=(1-s)^{\alpha N}$. Simplification yields

$$w(q) = q + \alpha(1-q)\left[\int_0^1 (1+\frac{q}{1-q}(1-y^{1/\alpha N}))^{(1-\alpha)N} dy\right]^{-1}.$$

Differentiating the integrand in $w(q)$ gives

$$\frac{\partial[(1+(\frac{q}{1-q})(1-y^{1/\alpha N}))^{(1-\alpha)N}]}{\partial N}$$
$$= \frac{(1-\alpha)h(N,y)^{(1-\alpha)N-1}}{\alpha N}\left(\alpha N\ h(N,y)\log(h(N,y)) + (\frac{q}{1-q})y^{1/\alpha N}\log[y]\right) > 0,$$

where $h(N,y) = 1+(\frac{q}{1-q})(1-y^{1/\alpha N}) \geq 1$ given $\alpha N > 0$, $q \geq 0$, and $y \in [0,1]$. Thus, $w(q)$ decreases in N and so $\Omega_N(x,N,\alpha N) < 0$.

Step 2: x* falls in N.

Using the equilibrium condition, $x^* = \delta V(x^*,N,\alpha N,\delta) = \delta S(x^*,N,\alpha N,\delta)\Omega(x^*,N,\alpha N)$, we get the following comparative static:

$$\frac{\partial x^*}{\partial N} = \frac{\delta V_N(x^*,\cdot)}{1-\delta S_x(x^*,\cdot)\Omega(x^*,\cdot)-\delta S(x^*,\cdot)\Omega_x(x^*,\cdot)} \tag{2}$$

The denominator is positive as $S_x < 0$ and $1-S\Omega_x > 0$ (by Lemma 1), and so x^* falls in N if $V_N(x^*,\cdot) < 0$.

We now show that $V_N(x^*,\cdot) < 0$. Consider first

$$V_\delta(x,\cdot) = \frac{(1-P(x,\cdot))P(x,\cdot)}{(1-\delta P(x,\cdot))^2}\Omega(x,\cdot) = \frac{P(x,\cdot)}{(1-\delta P(x,\cdot))}V(x,\cdot).$$

This implies

$$V_{N\delta}(x,\cdot) = V_{\delta N}(x,\cdot) = \frac{P(x,\cdot)}{(1-\delta P(x,\cdot))} V_N(x,\cdot) + \left[\frac{P(x,\cdot)}{(1-\delta P(x,\cdot))}\right]^2 V(x,\cdot).$$

Thus, $V_N(x,\cdot) \geq 0$ implies $V_{N\delta}(x,\cdot) > 0$, and we have

$$V_N(x, N, \alpha N, \delta) \geq 0 \Rightarrow V_N(x, N, \alpha N, 1) > 0 \quad \forall \delta < 1.$$

This is equivalent to

$$V_N(x, N, \alpha N, 1) \leq 0 \Rightarrow V_N(x, N, \alpha N, \delta) < 0 \quad \forall \delta < 1.$$

Finally, since $S(x, N, \alpha N, 1) = 1$ and $S_N(x, N, \alpha N, 1) = 0$, we have

$$V_N(x, N, \alpha N, 1) = \Omega_N(x, N, \alpha N) \leq 0 \quad \text{(by step 1)}.$$

Since $\delta \in (0,1)$, x^* must fall with N.

Step 3: $\exists \delta_L > 0$ such that Search duration falls in N for $\delta < \delta_L$.

By Lemma 2, P satisfies single crossing in x. Let $\bar{x}$ be the threshold defined in Lemma 2. That is, for all $x < \bar{x}$, we have $P(x, N, \alpha N) < P(x, n, \alpha n)$ when $N > n$. Since $x^*(n, \alpha n, \delta)$ is continuous in δ and $\lim_{\delta \to 0} x^*(n, \alpha n, \delta) = 0$, there exists a $\delta_L > 0$ such that $x^*(n, \alpha n, \delta) < \bar{x}$ for all $\delta < \delta_L$. From Step 2, we have that $x^*(N, \alpha N, \delta) < x^*(n, \alpha n, \delta)$. Therefore, $P(x^*(N, \alpha N, \delta), N, \alpha N) < P(x^*(n, \alpha n, \delta), n, \alpha n)$ for all $\delta < \delta_L$.

Step 4: Search duration rises in N when $\alpha = 1$.

When $\alpha = 1$, $\Omega(x, N, N) = \mu_h(x)$ is constant in N. Further our equilibrium condition is $x^* = \delta S(x^*, \cdot)\Omega(x^*, \cdot)$, and so when $\alpha = 1$, x^* and $S(x^*, \cdot)$ must move in the same direction in N. By Step 2, x^* falls in N, and so $S(x^*, \cdot)$ must fall as well. Since S is inversely related to P, and search duration is given by $(1 - P(x^*, N, \alpha N))^{-1}$, search duration must rise in N. $\square$

Proposition 5 *Expected search duration is increasing in M. There exists a $k \in (0, N]$ such that x^* is increasing in M when $M < k$, while x^* decreases in M when $M > k$.*

Step 1: Search Duration Rises in M.

Assume $M < M'$ and $x^*(N, M, \delta) < x^*(N, M', \delta)$. Since the probability of continuing, P, increases in x and M: $P(x^*(N, M, \delta), N, M) < P(x^*(N, M', \delta), N, M')$.

Assume instead $x^*(N, M', \delta) < x^*(N, M, \delta)$. Expected search duration is inversely related to S, so we must show that $S(x^*(N, M, \delta), N, M, \delta)$ falls in M. Define $x' = x^*(N, M', \delta)$ and $x = x^*(N, M, \delta)$, and abuse notation by suppressing arguments N and δ, then

$$
\begin{aligned}
& S(x, M)\Omega(x, M) - S(x', M')\Omega(x', M') = x - x' \\
\Leftrightarrow\ & S(x, M)\Omega(x, M) - S(x', M')\Omega(x', M') + S(x', M')\Omega(x, M) - S(x', M')\Omega(x, M) = x - x' \\
\Leftrightarrow\ & [S(x, M) - S(x', M')]\,\Omega(x, M) + S(x', M')\,[\Omega(x, M) - \Omega(x', M')] = x - x' \\
\Leftrightarrow\ & [S(x, M) - S(x', M')]\,\Omega(x, M) = x - x' + S(x', M')\,[\Omega(x', M') - \Omega(x, M)]\,.
\end{aligned}
$$

We wish to show that $S(x, M) > S(x', M')$, i.e. that the left hand side is positive. Since the expected payoff conditional on stopping, Ω, is increasing in M (see the proof of Lemma 1, Step 1), the right hand side weakly exceeds:

$$
\begin{aligned}
x - x' + S(x', M')\,[\Omega(x', M) - \Omega(x, M)] &\geq x - x' + [\Omega(x', M) - \Omega(x, M)] \\
&= x - x' - \int_{x'}^{x} \Omega_x(y, M)dy > 0,
\end{aligned}
$$

which gives the desired result. □

Step 2: $\exists k \in (0, N]$ such that for all $M < k$, $x^*(M + 1) > x^*(M)$, while for all $M > k$, $x^*(M + 1) < x^*(M)$.

Henceforth we abuse notation and let $x^*(M) = x^*(N, M, \delta)$.

Step 2-A: A Useful Single Crossing Property.

Define $\phi(x, M) \equiv P(x, \cdot)x + (1 - P(x, \cdot))\Omega(x, \cdot)$. Since, $x^* = \delta V(x^*, \cdot)$, we have $V(x^*(M), N, M, \delta) = \phi(x^*(M), M)$. As with $\delta V(x, \cdot) - x$, $\Omega_x \leq 1$ implies that $\delta\phi(x, M) - x$ satisfies single crossing, positive for $x < x^*(M)$ and negative otherwise. In turn this implies:

$$x^*(M + 1) < x^*(M) \Leftrightarrow \phi(x^*(M), M + 1) < \phi(x^*(M), M).$$

Routine algebra establishes that:

$$\phi(x, M+1, \cdot) - \phi(x, M, \cdot) = p(x, N, M)\left[x - \psi(x, M)\right],$$

where $\psi(x, M) \equiv \frac{M}{N}\mu_h(x) + (1 - \frac{M}{N})\mu_\ell(x)$, i.e. the expected payoff when exactly M members vote to stop with threshold x.[12] Altogether we may conclude:

$$\begin{array}{ccc} x^*(M+1) < x^*(M) & \Leftrightarrow & x^*(M) < \psi(x^*(M), M) \\ x^*(M+1) > x^*(M) & \Leftrightarrow & x^*(M) > \psi(x^*(M), M) \end{array} \tag{3}$$

Step 2-B: $x^*(1) > x^*(0)$.

When $M = 0$: $P = 0$ and $\Omega = E[X]$, and we have:

$$0 < \delta E[X] = x^*(0) > \mu_\ell(x^*(0)) = \psi(x^*(0), 0).$$

Thus, $x^*(1) > x^*(0)$ by condition (3).

Step 2-C: $x^*(M+1) \leq x^*(M) \Rightarrow x^*(M+2) < x^*(M+1)$.

$$\begin{aligned} x^*(M+1) \leq x^*(M) &\Rightarrow x^*(M) \leq \psi(x^*(M), M) \quad \text{(By condition (3))} \\ &\Rightarrow x^*(M+1) \leq \psi(x^*(M+1), M) \quad \text{(By } \psi_x(x, M) \leq 1\text{)} \\ &\Rightarrow x^*(M+1) < \psi(x^*(M+1), M+1) \quad \text{(By } \psi \uparrow M\text{)} \\ &\Rightarrow x^*(M+2) < x^*(M+1) \quad \text{(By condition (3))}. \end{aligned}$$

Step 2-B asserts that x^* is initially increasing in M, and Step 2-C asserts that once x^* stops increasing in M, it strictly decreases for higher values of M. □

Proposition 6 *The welfare-maximizing plurality rule, M, is weakly rising in the discount rate, δ. Given sufficient patience, unanimity is welfare maximizing.*

[12]For an intuition: We are comparing the value when requiring $M+1$ versus M votes, fixing the threshold and continuation value at x. The difference in values is then the probability that exactly M values exceed the threshold, $p(x, N, M)$, times the difference in value from continuing x and stopping with exactly M values above the threshold $\psi(x, M)$.

Step 1: The welfare maximizing M weakly rises in δ.

Suppose not. Let $M_H < M_L$ be welfare maximizing for $\delta_H > \delta_L$, and assume M_L does not maximize welfare at δ_H, i.e. $x^*(N, M_H, \delta_H) > x^*(N, M_L, \delta_H)$. Since M_L is welfare maximizing for δ_L, $x^*(N, M_L, \delta_L) \geq x^*(N, M_L - 1, \delta_L)$, which implies

$$\begin{aligned} & x^*(N, M_L - 1, \delta_L) \geq \psi(x^*(N, M_L - 1, \delta_L), M_L - 1) \quad \text{(by condition (3))} \\ \Rightarrow \quad & x^*(N, M_L - 1, \delta_H) \geq \psi(x^*(N, M_L - 1, \delta_H), M_L - 1) \quad \text{(by } \psi_x(x, M) \leq 1 \text{ and } x^* \uparrow \delta\text{)} \\ \Rightarrow \quad & x^*(N, M_L, \delta_H) \geq x^*(N, M_L - 1, \delta_H) \quad \text{(by condition (3))} \\ \Rightarrow \quad & x^*(N, M_L, \delta_H) \geq x^*(N, M_H, \delta_H) \quad \text{(by Prop. 5 and } M_L > M_H\text{)}, \end{aligned}$$

which contradicts M_L not being welfare maximizing for δ_H.

Step 2: For sufficiently high δ, $x^*(N, N, \delta) > x^*(N, M, \delta)$ for all $M < N$.

Assume $M < N$. Then $\Omega(x, N, M) > \psi(x, M)$ and $\lim_{\delta \to 1} x^* = \Omega(x^*, \cdot)$

$$\begin{aligned} \Rightarrow \quad & \lim_{\delta \to 1} [x^*(N, M, \delta) - \psi(x^*(N, M, \delta), M)] > 0 \\ \Rightarrow \quad & \lim_{\delta \to 1} x^*(N, N, \delta) - x^*(N, M, \delta) > 0 \quad \text{(by condition (3))} \\ \Rightarrow \quad & \exists \delta^* < 1 \ \text{ s.t. } \ x^*(N, N, \delta) > x^*(N, M, \delta) \ \forall \ \delta > \delta^*, \end{aligned}$$

where the last implication follows from the continuity of x^* in δ and ψ in x. □

References

[1] M. Abramowitz, I. Stegun, Handbook of Mathematical Functions, Dover, New York, 1965.

[2] D. Austen-Smith, J. Banks, Information aggregation, rationality, and the Condorcet Jury Theorem, The American Political Science Review, 90 (1996), 34–45.

[3] M. Bagnoli, T. Bergstrom, Log-concave probability and its applications, Economic Theory, 26 (2005), 445–469.

[4] D. Black, The decisions of a committee using a special majority, Econometrica, 16 (1948), 245–261.

[5] T. Börgers, Costly voting, American Economic Review, 94 (2004) 57–66.

[6] K. Burdett, Truncated means and variances, Economics Letters, 52 (1996), 263–267.

[7] G. Casella, R. Berger, Statistical Inference, Second Edition, Pacific Grove: Duxbury,2002.

[8] M. D. Condorcet, Essay on the Application of Analysis to the Probability of Majority Decisions, Paris,1785.

[9] A. Erdélyi, Higher Transcendental Functions, McGraw-Hill, New York,1953.

[10] T. Feddersen, W. Pesendorfer, Convicting the innocent: The inferiority of unanimous jury verdicts under strategic voting, American Political Science Review, 92 (1998), 23–35.

[11] A. Goldberger, Abnormal selection bias, in: S. Karlin, T. Amemiya, L. Goodman (Eds.), Studies in Econometrics, Time Series, and Multivariate Statistics, Academic Press, New York, 1983, pp. 67–84.

[12] H. Hotelling, Stability in competition, Economic Journal, 39 (1929), 41–57.

[13] S. Lippman, J. McCall, The economics of job search: A survey, Economic Inquiry, 14 (1976), 155–189, 347–68.

[14] J. McCall, Economics of information and job search, Quarterly Journal of Economics, 84 (1970), 113–126.

[15] P. Milgrom, R. Weber, A theory of auctions and competitive bidding, Econometrica, 50 (1982), 1089–1122.

[16] D. Mortensen, Job search and labor market analysis, in: O. Ashenfelter, R. Layard (Eds.), Handbook of Labor Economics, North-Holland, Amsterdam, 1986, pp. 849–920.

[17] R. Rogerson, R. Shimer, R. Wright, Search-theoretic models of the labor market: A survey, Journal of Economic Literature, XLIII (2005), 959–988.

[18] S. Stigler, Isaac Newton as probabilist, Statistical Science, 21(2006), 400–403.

[19] S.B. Vroman, No-help-wanted signs and the duration of job search, Economic Journal, 95 (1985), 767–773.