

Christian List and Adrian Vermeule
**Independence and interdependence:
lessons from the hive**

**Article (Accepted version)
(Refereed)**

Original citation:

List, Christian and Vermeule, Adrian (2014) *Independence and interdependence: lessons from the hive*. [Rationality and Society](#), 26 (2). pp. 170-207. ISSN 1043-4631
DOI: [10.1177/1043463114523713](https://doi.org/10.1177/1043463114523713)

© 2014 by [SAGE Publications](#)

This version available at: <http://eprints.lse.ac.uk/56678/>
Available in LSE Research Online: November 2014

LSE has developed LSE Research Online so that users may access research output of the School. Copyright © and Moral Rights for the papers on this site are retained by the individual authors and/or other copyright owners. Users may download and/or print one copy of any article(s) in LSE Research Online to facilitate their private study or for non-commercial research. You may not engage in further distribution of the material or use it for any profit-making activities or any commercial gain. You may freely distribute the URL (<http://eprints.lse.ac.uk>) of the LSE Research Online website.

This document is the author's final accepted version of the journal article. There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

Independence and Interdependence: Lessons from the Hive

Christian List and Adrian Vermeule*

Abstract: There is a substantial class of collective decision problems whose successful solution requires *interdependence* among decision makers at the agenda-setting stage and *independence* at the stage of choice. We define this class of problems and describe and apply a search-and-decision mechanism theoretically modeled in the context of honeybees and identified in earlier empirical work in biology. The honeybees' mechanism has useful implications for mechanism design in human institutions, including courts, legislatures, executive appointments, research and development in firms, and basic research in the sciences. Our paper offers a fresh perspective on the idea of "biomimicry" in institutional design and raises the possibility of comparative institutional analysis across species.

Keywords: Biomimicry, collective decision making, comparative institutional analysis, independence, information pooling, interdependence

For centuries, *homo sapiens* has learned tricks of design from other species, including both non-human animals and plants. In applied sciences such as engineering and aerodynamics, "biomimicry" exploits designs that arise from natural selection. The inventor of Velcro hook-and-loop fasteners, used in everyday clothes as well as high-tech products, copied the hooks by which cockleburrs snag the fur of passing animals; the shape of the Mercedes Benz bionic car mimics the boxfish to maximize aerodynamic efficiency; and Speedo's Fastskin body-hugging swimsuit, recently adopted by most Olympic swimmers, mimics the micro-features of sharkskin to minimize drag (Bushan 2009).

In the study of social organization, likewise, there is a long tradition of comparisons between humans and other animals. Of these the most famous may be Mandeville's *Fable of the Bees* (1714), but comparison and contrast between humans and social insects is much older.¹ Yet such analogies and disanalogies are typically either a literary conceit, as in Mandeville's case, or when meant seriously have

been pseudo-scientific. Indeed, there is an equally long history of politically motivated abuse of biological analogies, so much so that in some quarters the very notion is taboo (Rodgers 2009).

In recent years, however, scholars working at the intersection of biology and the social sciences have produced a growing body of research on collective decision making in the world of non-human animals. Herds of red deer appear to use a qualified majority rule for group decisions whether to move on or stay put, while among African elephants a majority of adult females decides (Conradt and Roper 2005). Some of the most striking findings concern social insects. Recent work has begun to put micro-foundations under the notion that insect colonies in some sense make collective decisions (e.g., List, Elsholtz and Seeley 2009, Franks et al. 2009, Seeley 2010); and it turns out that they make collective decisions extremely well, and through striking procedures. We suggest that *homo sapiens* can learn from these procedures, both in the theoretical sense and at the level of institutional design. In this paper, we develop some of these lessons from the hive for collective decision making.²

Broadly speaking, there are two ideal types of collective decision problems: first, those in which individuals have the same fundamental preferences, but different information or beliefs and (hence) different derived preferences; and second, those in which there is a bedrock conflict of fundamental preferences.³ The former are epistemic problems, the latter distributive ones. Of course, there are many mixed cases, but for clarity it is useful to focus on the extremes.

Both types of problems arise frequently in human groups. In what follows, we elicit some lessons from the hive for epistemic problems. Because worker bees do not reproduce themselves and their shared genetic interest is to help the mother queen survive and reproduce (e.g., Seeley 2005, 2010), conflict of fundamental preferences is minimal within the hive. As we will see, however, differences of information or belief are very much present, and social insects use intriguing mechanisms to sort out those differences and to settle upon a joint course of action.

The main lesson from the hive for epistemic decision making is the value of balancing *independence and interdependence*. In our central example, based on recent work in biology, honeybees can be

understood as using a special decision procedure for choosing nest sites: the bees prefer, all else equal, to assess options advertised by other bees, and in that sense decide interdependently, yet they assess those options in an independent manner, as formally defined below. The bees behave *interdependently* in setting the epistemic agenda, and *independently* in deciding whether to support – to “vote” for – any given option.⁴

By balancing independence and interdependence in this way, the honeybees have hit on an insight often overlooked by decision theorists. Teasing out the implications of Condorcet’s jury theorem, the literature on epistemic collective decision problems emphasizes the centrality and value of *independence* among the members of decision-making groups (for introductions, see Grofman, Owen and Feld 1983, List and Goodin 2001; in constitutional contexts, see Sunstein 2009, Vermeule 2009). Although it recognizes that complete independence is not always attainable, independence remains a key ideal and any kind of interdependence between decision makers is seen primarily as a risk.⁵ Moreover, the literature typically takes the agenda for epistemic decision making as exogenous, leaving it unspecified how groups with common preferences but dispersed information do or should decide what options they will decide among.

By contrast, our claim is that there is a substantial class of collective decision problems in which successful collective decision making requires interdependence at the stage of epistemic agenda-setting, as well as independence at the stage of choice. This generalizes and applies a mechanism theoretically modeled in the context of honeybees (List, Elsholtz and Seeley 2009) and identified in earlier empirical work in biology (e.g., Seeley, Visscher and Passino 2006, Lindauer 1955; for a recent overview, see Seeley 2010). We suggest that in a broad class of decision problems involving collective search, decision makers do best by striking a balance between independence and interdependence.

People are not bees. Likewise, automobiles are not boxfish; yet automotive engineers can design better cars by studying fish. Perhaps humans can design better decision procedures by studying insects. Our suggestion is emphatically *not* that humans should mindlessly copy the hive, or that any decision

procedure in use among social insects or other non-human animals can be directly transposed to the human world. Rather, certain *structural* features of decision-making environments, such as tradeoffs between speed and accuracy (to take an example discussed below), are common to decision making by humans and non-humans, in some settings. If natural selection has led to highly successful decision procedures for non-human animal groups in such environments, there is no reason not to examine those procedures to deepen our theoretical understanding and broaden our institutional repertoire.

Section I introduces Condorcet's jury theorem, explains the state of the literature on epistemic collective decision making, and shows that the literature both makes independence a central epistemic ideal, and also leaves agenda-setting exogenous – a mysterious black box. Section II introduces the problem of nest-site choice facing honeybee swarms, reviews a decision-making mechanism that combines interdependent agenda-setting with independent voting, and identifies general conditions under which such a mechanism will perform well. Section III applies the honeybees' mechanism to illuminate, and critique, a range of decision-making procedures in human institutions, including the certiorari process on the United States Supreme Court, agenda-setting by legislative committees, appointments in firms and universities, the choice of candidates by political parties, and the choice of research projects in commercial firms and in basic science. After detailing several applications, we consider in general terms how certain strategies of mechanism design – particularly “veiling mechanisms” that separate interdependent agenda-setting from independent evaluation of agenda items – can be helpful or even necessary to replicate some of the strengths of the bees' decision-making mechanism in human institutions. In the Conclusion, we return to the question whether the hive offers useful lessons.

I. Independence and the Jury Theorem

In an epistemic collective decision problem, a group of two or more individuals faces a choice among a number of options, where the individuals have common fundamental preferences but possibly different beliefs. That is, the individuals may disagree about the preferability of the options, but these disagreements are only informational. An omniscient observer would be able to rank the options in an

objective order that reflects the individuals' common preferences, and the individuals themselves would agree with this ranking, were they fully informed. We are looking for a mechanism by which the group can maximize its chance of choosing the best option or options.

Condorcet's jury theorem says that, under two conditions, majority rule provides such a mechanism, at least when there are only two options (e.g., Grofman, Owen and Feld 1983) (there are various ways of extending the theorem to more than two options, e.g., List and Goodin 2001). In the binary case, the conditions are the following. First, each individual has a better-than-random chance of identifying the best option (the "competence condition"); specifically, each individual has a probability greater than one-half of judging the first option to be best if this is the case, and of judging the second option to be best if that is the case. Secondly, the judgments of different individuals are mutually independent (the "independence condition"); that is, any individual's judgment about which option is best does not depend on any other individual's judgment on this question. Under these conditions, the probability that the majority supports the best option exceeds each individual's probability of doing so and approaches one – certainty – as the number of individuals increases.

Although simple to state in theory, the theorem's two conditions are often hard to meet in practice. The competence condition requires individual judgments to be positively correlated with the truth. Despite the initial plausibility of this condition, we can think of a number of cases in which, due to a lack of information, systematic bias, or the inherent difficulty of a judgmental task, individual judgments lack the required correlation with the truth; and when there is no such correlation, majority decisions are no better than random at picking the best option. Moreover, when there is a negative correlation, the reverse of Condorcet's effect kicks in: the probability that the majority supports the best option will then be smaller than each individual's probability and will approach zero with increasing group size.

Even more challenging than the competence condition, in many contexts, is the independence condition.⁶ If different individuals base their judgments on the same source of information or a limited number of sources, these judgments may become mutually dependent, and pooling them, as majority

voting does, cannot provide us with any new information beyond what was contained in the shared information they were based on. In the limiting case of perfect correlation between different individuals' judgments, the majority decision is no more reliable than the decision of any individual.

In general, how mutual interdependence between different individuals' judgments affects the reliability of the resulting majority decision depends on the structure of the interdependence. To explain this, it helps to subdivide the conclusion of the jury theorem into two parts. The so-called "non-asymptotic" part states that the probability of a correct majority decision exceeds each individual's corresponding probability, while the so-called "asymptotic" part states that this probability approaches one with increasing group size. It can be shown that some forms of dependence between different individuals' judgments preserve both parts of the theorem's conclusion and only reduce the speed with which the probability of a correct majority decision converges to one. Such forms of dependence are in effect equivalent to a reduced group size and can be offset by increasing that size again. Dependence of this benign sort occurs, for example, when individuals derive their judgments from a mix of private information and signals received from others, whether from opinion leaders or from their peers (as discussed, e.g., in Ladha 1992, Estlund 1994).

By contrast, other forms of dependence between individuals have more dramatic consequences for the jury theorem, not merely reducing the speed with which the majority reliability converges to one with increasing group size, but undermining the second, asymptotic part of the jury theorem altogether. Suppose, for example, the members of a jury have epistemic access to the truth about a particular crime only via the shared body of evidence presented in the court room; none of the jurors has any private information that bypasses this single evidential route to the truth.⁷ Familiar rules of evidence impose precisely this constraint. The jury's reliability – its probability of convicting the defendant if and only if the defendant is guilty – may then still exceed the reliability of each individual juror, and so the non-asymptotic part of the jury theorem may continue to hold; together, the jurors may arrive at a more consistent interpretation of the evidence, for example. But the jury's reliability will never overcome the

epistemic bottleneck created by the jurors' mutual dependence on a single evidentiary route to the truth. If that shared evidence is limited or misleading, for example, they will never be able to transcend that limitation, regardless of how many jurors there are. In consequence, the reliability of the majority decision is subject to an upper bound at some threshold strictly below one – a maximal feasible level of reliability – which depends on the nature and quality of the evidence (for a revised jury theorem in this context, see Dietrich and List 2004; for recent, more general results, see Dietrich and Spiekermann 2013).

How pervasive is the kind of interdependence between different jurors' judgments that threatens the applicability of the jury theorem? This question is an empirical one, but even a cursory reflection on how opinion leaders and other epistemic bottlenecks can affect the formation of individual opinions in real-world settings suggests that violations of Condorcet's independence condition are frequent.

Furthermore, it has been argued that, although it is *possible* for Condorcet's two conditions to be simultaneously true, we can never obtain any evidence to *corroborate* their *joint* truth (Dietrich 2008). The reason, in very rough terms, is that to corroborate the competence condition, we must not focus on each individual's judgment in a single isolated decision problem, where we have no way of quantifying the individual's reliability, but we must average over a larger reference class of "similar" decision problems; some problems in that class will be easier, others harder, yet on average each individual may be shown to display the required competence. But once we look at such a larger reference class of decision problems, independence can no longer be corroborated; the reference class will inevitably exhibit some internal heterogeneity – as noted, some problems in it will be harder, others easier – and therefore judgmental performance is bound to be correlated across different individuals.

Some of the most pernicious violations of independence are those to which decision makers are most oblivious, when they misinterpret a situation in which individual judgments are not independent as one in which independence is satisfied. Individuals are then liable to draw false confidence from what they take to be the confluence of several independent sources of evidence, which are in fact highly correlated. This phenomenon underlies the "informational cascades" responsible for market bubbles, some instances of

mass hysteria, or the seemingly irrational spread of false beliefs in society (Bikhchandani, Hirshleifer and Welch 1992, Sunstein 2006). In an informational cascade, an accidental spell of support for some proposition or option is misinterpreted by other decision makers as evidence for the truth of the proposition or the quality of the option, thereby leading them to join the chorus of support. This, in turn, may be taken by others as even further evidence in support of the proposition or option and may thus trigger a snowball effect in which a small number of random signals can be amplified into a spurious consensus (see also List and Pettit 2004).

These considerations illustrate the risks associated with violations of independence and reinforce the centrality of independence as an epistemic ideal in collective decision making. Independent and competent assessments of the options seem to be the key conditions for efficient collective decisions. However, the work reviewed so far has taken the set of options as exogenously given, focusing only on the process by which the pattern of individual support for them is aggregated into a collective decision. Once agenda setting is taken into account as well, it turns out that there is an important class of decision problems in which interdependence between individuals can be put to good use, even from an epistemic perspective, provided interdependence is confined to the agenda-setting stage and carefully balanced with independence at the voting stage. To show this, we now look at the way honeybees choose nest sites.

II. The Mechanism and its Conditions

We begin with a brief empirical description of the mechanism by which honeybees choose their nest sites, as studied by Seeley et al. (e.g., 2004, 2006, 2010). We then review a simple theoretical model of the bees' collective decision process, drawing on recent collaborative work between social scientists and biologists (List, Elsholtz and Seeley 2009), which allows us to see the key determinants of the bees' decision-making performance. On this basis, we suggest which aspects of the bees' decision process may carry over to human collective decisions and under what conditions.

A. Background

At the end of spring or the beginning of summer, a honeybee colony that has grown too large tends to split. The queen bee leaves with approximately two-thirds of the worker bees, while a daughter queen stays in the maternal nest with the others. To survive, the bees that have left must quickly find a new nest. Empirical research has shown that they do so by means of a striking decision process (e.g., Seeley, Visscher and Passino 2006, Lindauer 1955, Seeley 2010). This involves a “search committee” of several hundred “scout bees” who roam the surrounding area in search for potential nest sites and then return to the swarm to draw the others’ attention to any good sites they have discovered.

In particular, after discovering a potential nest site, each scout bee performs a waggle dance whose orientation encodes the site’s location and whose duration encodes her assessment of the site’s quality. The better she perceives the site to be, the longer she dances. At first, the scout bees rely on discovering potential nest sites by chance, but once they observe other scouts dancing, they are more likely to investigate the sites advertised by those others. If they agree with the positive assessment of a site, they join the dance for it. In this way, sites supported by dancing bees are visited and inspected more often than other sites and, if not supported in error, tend to receive even more support. The process leads to a “consensus” relatively quickly – in one or two days – when the support rallies around one popular site. Once a critical threshold is reached, the swarm moves there. Crucially, when there are quality differences between different potential nest sites, the bees usually find one of the best ones (Seeley and Buhrman 2001).

There are three constraints that make the speed and accuracy of this decision process all the more surprising. First, the agenda of options is not straightforwardly given, unlike in the decision problems to which Condorcet’s jury theorem is usually applied. An indefinite number of places in the bees’ environment could in principle become candidates for nest sites, and a suitable method of agenda setting is needed to sort out the serious options from the non-starters. Second, although individual bees have some remarkable capacities, they are still fairly simple organisms, and a simultaneous and comparative

assessment of all potential nest sites is beyond any bee's capacities. For this reason, the bees' collective decision process must not place high cognitive demands on any individual bee. Third and relatedly, the bees' assessment of potential nest sites and their communication are subject to a significant amount of noise, and the decision process must therefore be error-tolerant. So which features of the bees' decision process explain its remarkable speed and accuracy, in light of these constraints? A stylized model of the process helps to reveal what drives its success. (An explanation of the underlying formal model is provided in the Appendix; here we merely indicate the main features.)

B. A Simple Model

The bees, like humans in a Condorcetian jury decision, face an epistemic collective decision problem, but, unlike in an ordinary jury decision, the set of options is more open-ended. So there is a group of individuals (in the present case several hundred scout bees) that has to make a choice among a number of options (different possible nest sites) but that number can be very large. As before, the individuals have common fundamental preferences but possibly different information. We can represent these preferences by assuming that each possible nest site has an objective, though unknown, quality level. The decision process extends over multiple time periods (where time, for simplicity, is discrete), and we model each individual scout bee's behavior over time (employing the formal model in List, Elsholtz and Seeley 2009, described in the Appendix). Technically, the model is an "agent-based" model, which is defined by specifying, first, what state each individual scout bee can be in during each time period, and second, how each scout bee changes her state from one time period to the next.

Let us begin by considering each scout bee's possible states. In any given time period, a scout bee can be in one of two states:

The non-dancing state. The bee is not dancing in support of any potential nest site, which can mean that she has not yet flown out to search, has not yet found any promising site, has ended a previous dance, is observing other bees, or is resting.

The dancing state. The bee is dancing in support of a potential nest site; this state of the bee is further specified by one parameter: the remaining dance duration.

We next specify how a scout bee changes her state from one time period to the next. There are two cases to consider:

The first case: the bee is in the non-dancing state in the given time period. The bee has some probability of remaining in that state in the next period – that is, of continuing to search, observe other bees, or rest – and a complementary probability of finding a potential nest site that she supports, and thereby of switching into the dancing state. Two factors determine whether she does so: her probability of finding and visiting one of the possible sites, and her assessment of that site. Whether the bee finds and visits a site depends on how easy it is to find it – effectively, its salience on the agenda. How she assesses the site’s quality depends on her epistemic competence and independence. At this juncture, we can plug different assumptions into the model, so as to compare their implications, as discussed below.

The second case: the bee is dancing in support of a potential nest site in the given time period. If the remaining dance duration is not yet over, she will continue to dance for that site in the next time period, that is, she will stay in the dancing state, with the remaining dance duration reduced by one period. If the remaining dance duration is over, she will switch back into the non-dancing state, that is, she will fly out to search afresh, observe other bees, or rest.

Having specified the possible states in which each scout bee can be as well as the way each bee changes her state from one time period to the next, we can use computer simulations to see how the states of a collection of scout bees change over time and how long it takes for a “consensus” – defined as a sufficiently large plurality of support – to emerge for a particular nest site. The computer simulation can be started by assuming that in the first time period all bees are in the non-dancing state. For present purposes, it suffices to summarize the findings in qualitative terms (for detailed quantitative results, see List, Elsholtz and Seeley 2009 and the summary in the Appendix below).

C. The Determinants of the Bees' Collective Performance

As noted, the bees' collective performance depends on each scout bee's probability of finding and visiting one of the different potential nest sites and her competence and independence in assessing any such site once it has come to her attention. Let us compare different assumptions about each of these determinants of the process:

A scout bee's probability of finding and visiting each potential nest site. One theoretical possibility is that the probability that a particular site comes to a bee's attention depends only on the site's location and other exogenous factors. This would imply no communication or interdependence between the bees. Formally, each bee's unconditional probability of giving attention to a particular site would then be the same as her conditional probability of giving attention to it, given that one or more other bees have done so as well. Another possibility – the one supported by the empirical findings of field observation (e.g., Seeley, Visscher and Passino 2006) – is that while in the beginning a bee's probability of finding and visiting each possible site depends only on exogenous factors, so that finding a site is initially a random event, the probability increases as other bees start dancing for it. Thus the probability that any given site comes to a bee's attention – the site's "salience" on the agenda – is a weighted combination of an *ex ante* probability of finding it and the number of other bees advertising it. The weight of the second factor relative to the first can be taken to represent the level of interdependence between the bees. Once there is some interdependence, each bee's conditional probability of giving attention to any particular site, given that other bees have done so too, is higher than her unconditional probability of giving attention to it.

A scout bee's competence and independence in assessing any site. One theoretical possibility is that once a particular site comes to a bee's attention, she mimics other bees advertising it, so that her subsequent dance duration for the site is not determined by an independent assessment of its quality, but given randomly or by copying another bee's dance. Formally, if a bee mimics the dances of others, her probability of performing a dance of a particular duration for a given site, conditional on the site's having come to her attention, is unrelated to the site's quality. The empirical findings, however, support the

alternative possibility that a scout bee independently assesses a site that has come to her attention and that her dance duration for it then correlates positively (though imperfectly) with the site's quality. The strength of the correlation represents the bee's competence. Whereas a more sophisticated agent might be tempted to take the observed dance activity for a given site as a proxy for its quality and not to assess it independently at all, a bee's limited cognitive capacities prevent her from engaging in any such sophisticated epistemic free-riding.⁸ The technical sense in which a bee acts independently in assessing a site is that her probability of performing a dance of a particular duration for it, conditional on the site's having come to her attention and holding its quality fixed, remains the same irrespective of whether or not we also conditionalize on other bees' dance activity for it. To study the role played by this kind of independence, we can introduce, as a further model parameter, the probability that the bee's dance duration for any site is determined by an independent assessment of its quality rather than by mimicking other bees. This probability ranges from zero in the counterfactual case of no independence to one in the case of full independence.

So how does the bees' predicted decision-making performance vary as we vary these central model parameters – the bees' levels of interdependence, independence and competence? Computer simulations show the following. Assuming quality differences between different potential nest sites, *both* a certain level of interdependence in drawing each other's attention to promising sites *and* a certain level of independence in assessing the quality of any site once it has come to a bee's attention are needed to ensure that a "consensus" for a high-quality site will rapidly emerge. Further, given enough interdependence and independence, a moderate correlation between each bee's dance duration for her favored site and the site's actual quality – that is, a moderate individual competence – is sufficient to secure this outcome, and thus the decision process is error-tolerant.

In the hypothetical cases in which the bees lack either interdependence or independence, the decision process loses *either* its speed and decisiveness *or* its accuracy. Without interdependence, the bees fail to communicate to each other which sites are worth inspecting, and even good sites will only receive

attention from those (few) bees who stumble upon them randomly. The emergence of a consensus for any site – let alone a site that may be difficult to find – is therefore unlikely, and at least extremely slow. It is worth noting, however, that the opposite limiting case in which the bees set their agenda *only* interdependently is also suboptimal: if the bees consider only those sites advertised by others and do not randomly roam the area at all, there is no chance for them to find any good sites not yet discovered by others. Still, this negative effect comes into play only at very high levels of interdependence.

While the bees' interdependence is crucial for the speed and decisiveness of the decision process, their independence in assessing any sites that have come to their attention is crucial for its accuracy. Without independence, the bees are vulnerable to informational cascades, whereby any random fluctuation in the dance activity for some site can be amplified into a consensus for it regardless of its quality. The crucial link between dance activity and actual nest-site quality will then be compromised.

The computational results, and the bees' observed collective performance in identifying the best nest sites, suggest that the bees avoid the dual dangers of not giving enough attention to good sites on the one hand, and informational cascades on the other, through a finely balanced interplay of interdependence and independence: interdependence in communicating to each other which sites are worth inspecting – and thereby in setting the agenda – and independence in assessing the quality of any site they inspect (List, Elsholtz and Seeley 2009).

D. Conditions Favoring the Use of the Bees' Decision-making Mechanism

Although this decision-making mechanism has evolved in honeybees choosing nest sites, the structural features that make it work are transferable to other multi-agent systems. Any collection of agents that has to make fast and accurate decisions without an exogenously defined agenda can in principle implement the bees' decision-making protocol, as formally captured by the model we have described. What is needed is the ability to roam the space of possible options, to identify and independently rate potential options, however fallibly, and to draw each other's attention to options that are worth investigating. We have seen that organisms as simple as individual bees have this threefold ability, but nothing in this

package of skills is tied to a particular species, a particular decision problem, or a particular biological realization. We can view the bees' decision-making mechanism through a purely functionalist lens, abstracting away from the case of the bees, and ask in which decision-making environments a functionally similar mechanism would be useful.

In epistemic decision problems, three conditions seem to favor the use of such a mechanism: (1) an open-ended agenda, (2) (relatively) high stakes, and (3) (relatively) high opportunity costs of indecision. In the bees' case, as we have seen, all three conditions are clearly present. First, the agenda of potential nest sites is not straightforwardly given. Second, whether they find a good nest site, a mediocre site, or only a bad one can affect their survival and reproductive success. And third, indecision is not an option, since failing to reach a timely decision can threaten the swarm's survival.

Generally, we suggest that the less exogenously well-demarcated the agenda, the higher the stakes, understood as the utility differences between the options, and the higher the opportunity costs of indecision, the more a group can benefit from applying the bees' mechanism. How successful a group will be if it does so then depends on its ability to balance interdependence in signaling to each other which options are worth considering with independence in the individual assessments of those options, over and above each individual's epistemic competence.

III. Applications

We now turn to applications, using the mechanism of interdependence plus independence to assess a range of epistemic decision procedures in human institutions. Where the decisional environment most closely corresponds to the bees' environment of nest-site choice – exhibiting an open-ended agenda, high stakes, and high opportunity costs of indecision, in the presence of common fundamental preferences – a mechanism akin to the one used by the honeybees works to best advantage.

Our analysis is prescriptive and instrumental, not explanatory. We make no assumption that human institutions have evolved to efficiency, and thus do not seek to explain existing institutions by reference to the decision-making environment. We assume, in other words, that human institutions might or might

not work well, in any given setting, and that under certain conditions humans can learn from honeybees. If there are disanalogies between the bees' decision-making mechanism and any observed human procedure, this may sometimes give us some leverage to improve upon the latter. In particular, if a decision-making environment for humans has the features we have described, there may be strategies of mechanism design by which we could optimize human decision making in that environment. We take up the theme of mechanism design more generally after surveying a range of applications.

A. Agenda-Setting on the Supreme Court

If there is any governmental institution that routinely faces epistemic decision problems, it is the judiciary. Apart from a relatively small number of political cases, in which fundamental preferences differ along conservative and liberal lines, most cases present issues in which the judges have similar preferences but, at most, differing beliefs. This is true even at the level of the Supreme Court. Despite the strong selection pressure for hard cases to appear at the higher levels of the judicial hierarchy – easy cases are more likely to be settled, or never to be appealed – 47% of the Court's decisions were unanimous in the 2009 Term,⁹ and in the cases with dissents, some large fraction involved disagreements about facts or diverging predictions about the consequences of a ruling one way or another. At the agenda-setting stage, a recent study finds that “legal considerations strongly influence justices' agenda-setting behavior”, although ideological differences play a role as well (Black and Owens 2009, p. 1063). We bracket the latter point and assume, without too much distortion of reality, that the justices share common fundamental preferences about what types of cases to hear.

The Supreme Court is, however, unique in the broad control it enjoys over its own agenda. Since the Judiciary Act of 1925, the Court takes almost all of its cases by granting a petition for certiorari, and the Court possesses extremely broad discretion in deciding which petitions to grant. As a matter of practice, the Court is most likely to grant cases in which the lower courts have disagreed, or in which a federal statute has been invalidated on constitutional grounds.¹⁰ Yet in the end these practices are just rules of thumb or guidelines, which the Court follows or ignores according to circumstances. The Court's

discretion is increased by the sheer number of certiorari petitions that flood it every year – over 8,000 on average in recent years,¹¹ of which the Court usually grants slightly over 80 on average.¹² The problem for the Court as a body is to sift through the enormous mass of petitions to find the 1% that should be given the Court’s full attention.

The Court’s decision-making environment, in other words, combines nearly unlimited agenda control and high search costs. A great deal of the Court’s business lies in deciding what to decide (Perry, Jr. 1991). In this environment, the Court has developed, over time, a set of elaborate procedures for deciding what cases to take. The basic norm is a Rule of Four – the votes of four of nine Justices suffice to grant a certiorari petition for a full hearing. The voting takes place on petitions that are placed on a “discuss list” for the Justices’ weekly conferences; any Justice may place petitions on the list.

Yet how can the Justices sort through, in a single term, more than 8,000 certiorari petitions to decide which ones even to place on the list at all? To do so individually would consume most of the Justices’ time. Accordingly, most Justices – at the time of this writing, all but Justice Alito – participate in an institution that seeks to generate economies of scale at the preliminary stage of identifying plausible candidates for the discuss list. This institution is called the “cert[iorari] pool.” The law clerks (the scout bees) employed by the participating Justices divide up the mass of petitions among themselves and then circulate a memorandum to all the Justices in the cert pool. In some chambers, a clerk for Justice X will prepare a second memorandum, but often this is done, or done well, only for petitions that the initial writer of the pool memo has recommended to be granted. The pool memo writer, in other words, has some *de facto* leeway to shape the Court’s docket, especially by recommending denial. The institutional pressure to “deny cert” is enormous, and a recommendation of denial is rarely contested or closely scrutinized by clerks in other chambers, unless a case has obvious political import.

This institutionalized process of search-and-agenda-setting bears an imperfect, but illuminating, resemblance to the honeybees’ decision procedure. The cert pool is like the bees’ scout committee. Just as interdependence among the bees means that bees are more likely to inspect nest sites advertised by

others, so too the effect of the cert pool is that Justices and clerks are more likely to pay close attention to petitions advertised by pool clerks as “certworthy,” or good candidates for a grant. Yet the evaluation of plausibly certworthy candidates is largely independent, both among the bees and among the clerks and their Justices. Once the cert pool writer has identified a plausible candidate for the discuss list, Justices and clerks in other chambers independently evaluate the petition to decide whether to place it on the Court’s agenda. Once a petition has been granted and a case has been given a full hearing, moreover, each Justice independently evaluates the legal claims. The Court’s process combines interdependence at a crucial preliminary stage – identifying, from the mass of petitions, plausible candidates for the agenda – with independence at all later stages of the decision-making process.

How well does the Court’s process work? On one level it is impossible to know, because we have no independent benchmark assessment that would tell us which petitions in fact warrant a full hearing. Yet we can offer some conditional conclusions, and one implication. If and to the extent that the Court’s decision-making environment at the certiorari stage is understood as having the same features as that of the honeybees – an open-ended agenda, real stakes, and real opportunity costs of inaction – then the Court’s process is well-engineered for that environment, subject to some improvements we describe below. Whether those conditions are met is a matter for debate. Many critics of the Court believe that the Court should hear more cases and thus should grant more petitions; presumably these critics believe that it is better for the Court to make more decisions than fewer, perhaps because decisions by the Court clarify the law and promote legal, economic and political certainty. Views such as these implicitly suppose, in other words, that there are high opportunity costs, from the social point of view, if the Court too often fails to reach consensus on which cases to hear. In this light, interdependence at the early stages of the Court’s agenda-setting process is desirable; perhaps even more interdependence than currently exists would be desirable. Conversely, however, the marked independence of the later stages of the process is desirable to the extent one thinks that the marginal stakes in the Court’s decisions are high, so

that information cascades and other phenomena associated with the lack of independence are especially harmful when they cause the Justices to reach consensus on the wrong (or worse) answer.

Suppose one believes that the Court's decision-making environment does present the combination of factors that make the honeybees' interdependence-independence mechanism useful. An implication is that the cert pool should incorporate most but not all of the Justices, to optimize the balance of interdependence and independence. Recall that although the honeybees preferably assess nest sites advertised by others, they nonetheless each retain a nonzero probability of stumbling upon potential nest sites on their own. This residual independence, even at the agenda-setting stage, is crucial to the efficiency of the mechanism. As noted above, where interdependence reaches its limiting maximal value, "there is not enough noise in the system for bees to discover any new sites not advertised by others. Small noisy deviations from perfect [interdependence] are necessary to permit the discovery of new sites" (List, Elsholtz and Seeley 2009, p. 758). Likewise, a cert pool containing all Justices and their clerks would in effect place in the hands of a single twenty-something law clerk a real measure of *de facto* power to set the Court's agenda, perhaps by burying certworthy cases. The existence of Justices who review all petitions independently of the pool is beneficial for the group, as it provides an independent check on the work of the pool clerks and creates a small amount of beneficial noise in the system, as the nonpool Justices and their clerks search for certworthy cases in parallel to the official search committee.

It is hard to say, based on these general considerations, what the optimal level of participation in the pool might be. The number has varied over time; when the pool began to operate in 1972, it had only five members, but its size has grown steadily over time, and at present all Justices except one participate. Whatever the optimal membership, it seems likely that zero non-participants is too little, and the Court is now uncomfortably close to that extreme. The larger point is that understanding the mechanism of interdependent search plus independent evaluation at least identifies the variables that determine the optimal setup of the pool.

B. Legislative Committees

Legislatures are not often thought of as epistemic decision-making institutions. Rather, legislatures often act as a kind of political marketplace for bargaining between the major political parties, who have different fundamental preferences over major policies. Yet *within* parties, it is entirely plausible that a great deal of epistemic decision making takes place. Although parties are themselves coalitions, especially under first-past-the-post voting systems, legislative co-partisans are much more likely to share fundamental preferences with one another than with members on the other side of the aisle. The co-partisans have common aims – perhaps to promote the public good, perhaps to stick the other party in the eye – and their problem is to aggregate differing information and beliefs so as to achieve their common aims.

Enter legislative committees. There are many different theories of committees, all of which seem to capture some truth; however, we will focus on the implications of one such theory, the partisan control account of Cox and McCubbins (2007). On this account, political parties control committees. In contrast to interest-group accounts, which picture legislators self-selecting onto committees, the partisan control view holds that party leaders select committee members to promote partisan interests. In contrast to informational accounts, which see committees as serving the interests of the median member of the whole legislature (who has decisive power under simple majority rule), committees serve the interest of the median member of the majority party, who has decisive power in selecting the internal legislative leaders, who in turn control the composition of committees.

It is at least compatible with this view to suppose that the partisan majorities who select and direct committees face an epistemic problem, conditional on the common preferences of co-partisans. Committee members search for policies that will promote the majority party's preferences, perhaps because the partisan majority believes those policies best for the nation, perhaps because they are politically constrained by constituents to adopt certain policies or block others, or perhaps to embarrass the other party or split the other party's internal coalition. Whatever the motivation, the task is epistemic

in that a partisan majority has (much of the time, on many issues) common preferences but dispersed information and differing beliefs about how to satisfy those preferences.

How can the party leaders identify and agree upon policies that will attain their ends? We suggest that the partisans might do well to imitate the bees, and in some respects already do so. They might set up a subgroup of the party membership to serve as a search committee. Individual members of this search committee would, in effect, roam the policy space to find politically useful legislative proposals. The members would then advertise any proposals identified as potentially useful, and would attract support from other co-partisans to the extent that, after inspecting the advertised candidates, they believe a given proposal is indeed politically valuable. If support reaches some critical threshold in the search committee, the committee would report out a bill embodying the proposed policy and, if politically feasible, the majority party would enact the bill into law.

As described, this process displays the combination of interdependence and independence characteristic of the bees' decision making. The co-partisans act interdependently at the agenda-setting stage in which useful candidate proposals are identified by individual members of the search committee. Committee members are more likely to give serious consideration to candidate policies identified by other committee members, rather than searching the policy space in a strictly individual fashion, without regard to the recommendations of others. Conditional on investigating alternatives proposed by other members, however, the members exercise independent judgment about the quality of those alternatives. The ultimate selection among the candidate policies is determined by independent assessment of alternatives generated interdependently.

Real-world legislatures are not so different from this model process. We can understand legislative committees as searching the policy space for politically advantageous proposals, and then exercising conditionally independent judgment on a set of alternatives generated in an interdependent fashion. Prescriptively, from the standpoint of the majority party, the key question about this decision-making mechanism is whether the costs of interdependence at the agenda-setting stage exceed the benefits. Recall

that the main cost is the possibility of premature herding towards a bad or at least suboptimal alternative because there has been insufficient exploration of alternatives. The main benefit is that insufficient levels of interdependence at the agenda-setting stage tend to produce a failure of consensus. Where the opportunity costs of inaction are high, as in the bees' environment, avoiding this failure of consensus becomes a collective imperative.

For party leaders, an important implication is that interdependence at the committee agenda-setting stage becomes more valuable as the costs of inaction increase. Imagine a political environment in which the majority party will suffer, politically, if it is perceived as running a "do-nothing Congress." At the early stages of the n-year legislative cycle, party leaders and committee chairs will do well to afford individual committee members more freedom to search out proposals that will put the opposing party in an awkward position, or will promote the majority's platform. This freedom will slow down the process of consensus formation, but result in an increase in the expected quality of the eventual consensus, from the majority party's point of view. As the election cycle nears its close, the opportunity costs of inaction increase, because the costs of being charged with running a do-nothing Congress increase. Party leaders should tighten up the process of search and resolution by requiring greater interdependence among committee members. Committee consideration should focus on the alternatives already identified, although committee members should be allowed to exercise independent judgment among those alternatives.

The comparison should not be pressed too far. How much are rank-and-file legislators really like scout bees? And do Senators bear anything more than a superficial resemblance to drones? A key difference between the insect hive and Capitol Hill is that legislatures are more specialized internally. Rather than having a single search committee to make a highly consequential collective decision (such as the choice of a new nest site), legislatures have multiple search committees each assigned to a different area within the total policy space. Moreover, these committees have partially overlapping jurisdictions, and the boundaries between their jurisdictions may be fuzzy.

Yet this jurisdictional fuzziness may work well, as judged against the bees' mechanism. As in the certiorari pool, where the optimal level of participation is not 100%, so too it is not desirable, from the standpoint of party leaders, that there be *complete* interdependence at the committee agenda-setting stage. Complete interdependence would eliminate all noise from the system, and thus eliminate any prospect for members to stumble upon new and highly advantageous proposals not found by others. Jurisdictional fuzziness can introduce some desirable noise by making it possible, although unlikely, that a member from another committee will stray into the policy space and stumble upon a valuable proposal that members of the principal committee have overlooked. Of course, we do not suggest that the jurisdictional overlap and fuzziness of typical legislative committee-structures is best explained on these grounds; it arises for a number of political and historical reasons, rather than on the basis of any considerations of optimal collective decision making. Yet in light of the bees' arrangement, jurisdictional overlap and fuzziness may be associated with some epistemic benefits.

C. Searching for Leaders: Executive Appointments in Firms

The bee's search-and-decision mechanism is most advantageous to the group when the decision-making environment combines high stakes (making independence at the voting stage valuable) with high opportunity costs of indecision (making interdependence at the agenda-setting stage valuable), while the agenda is relatively open-ended. Plausibly, the search for leaders in for-profit firms and non-profit organizations presents just such a decision-making environment. First of all, the set of possible candidates is not so easy to identify. The stakes of the choice are high, because leadership and charisma are scarce resources whose presence or absence can make or break institutions, and because firms and organizations tend to search for new leadership in periods of crisis, in which routine decision making is not viable and executive decisions are particularly consequential. The opportunity costs of failing to reach consensus on the selection of a new leader are also high under such conditions, because passivity and inaction are often the worst possible strategies for institutions in crises; it is better to have a strong hand at the helm than to drift in treacherous waters, even if it is unclear which way it is best to go.

The implication is that leadership searches in firms and organizations, especially in crisis conditions, should attempt to combine interdependence and independence in roughly the ways we have outlined. We will focus on the structure and procedures of executive search committees in universities. When universities select new leadership – say, the university’s President – the process typically involves a search committee. At Harvard University, for example, the search committee that selected Drew Gilpin Faust as President consisted of the six members of the University Corporation, plus three members of the University’s Board of Overseers. University search committees often operate in a secretive fashion, so it is difficult to know how they make decisions. But the pool of suitable candidates is usually not so easily defined, and in the environment in which Faust was selected – in early 2007, as the financial crisis became ever more severe – the stakes were high and the opportunity costs of deadlock serious.

In such an environment, leadership search committees should engage in interdependent agenda-setting and independent assessment. Interdependent agenda-setting will mean that individual members of search committees tend to focus their attention on candidates previously proposed by other members of the search committee, rather than roaming the space of candidates purely on their own. The committee members who propose candidates early in the process may have predominant influence in setting the agenda. Yet this does not entail interdependence at the stage of evaluation. To the contrary, committee members should decide with strict independence whether the candidates on the agenda meet the threshold set by the group’s established criteria. Procedural mechanisms such as the use of secret ballots within the committee can maximize independence under certain conditions.

In this picture, there are twin evils to be avoided. On the one hand, insufficient interdependence would result in an excessively protracted search process, as committee members would spend too much time searching for new candidates to put on the agenda, and devote too little attention to evaluating the candidates others have put forward. On the other hand, insufficient independence would produce informational cascades that might settle on a bad candidate. The optimum is a process that allows agenda-

setters to structure the pool of candidates, yet subjects those candidates to fully independent evaluation – maximizing the chances of settling on a good candidate, with reasonable expedition.

Similar lessons apply to committee decisions on the award of symbolically important prizes, such as Nobel Prizes and other national or international recognitions of merit. In such decisions, the pool of candidates tends to be open-ended; the stakes are high due to the cultural, intellectual, political or sometimes commercial repercussions of an award, and not meeting the deadline for an award is normally not an option. Here, too, committees are often secretive about their procedures, but our discussion suggests that – at least in cases of common fundamental preferences – they would do well by balancing interdependence in arriving at a list of nominated candidates with an independent assessment of these candidates' merits.

D. Research and Development in Firms

Just as the honeybees' mechanism can be applied to identify candidates for certain positions or awards, so it can also be applied to identify projects worth developing. Imagine a firm with a large staff of research experts – scientists, engineers, or other knowledge workers – who engage in two sorts of tasks, dividing time between them in some proportion. One task is to independently search the space of technically feasible innovations for potentially profitable innovations, and then to promote them to other researchers. Another task is to assess potentially profitable innovations promoted by colleagues. In the second task, assessment is independent, but the choice of innovations to be assessed is interdependent. By promoting an innovation, researchers set the epistemic agenda for colleagues, who suspend their independent search and decide whether the proposal meets some threshold of plausibility. Proposals that attain sufficient support in the research group are kicked upstairs for further assessment by higher management.

How should this process be structured to maximize the firm's expected utility? From the standpoint of the firm, the optimal allocation of time by each individual researcher will not be either of the corner solutions – either the one in which each researcher spends all her time independently searching for profitable innovations, or the one in which each researcher spends all her time assessing innovations

proposed by others. In the latter case, there will be no innovations to assess – not everyone can be reactive, or there will be nothing to react to – while in the former case, no consensus will form and the firm’s collective resources will be scattered too widely across different projects. The optimal time allocation balances independence and interdependence.

There are several examples of highly creative firms that use an optimizing mechanism of this sort. In the 1950s, 3M allowed its research staff to devote 15% of their time to independent projects, whose results would belong to the firm. Famous innovations resulted, including Post-It Notes and masking tape. More recently, Google has a similar policy at the 20% level, which has been credited with producing Gmail and Google News.¹³

This picture is a heroic simplification, because for-profit firms are often cited as examples of groups of actors with conflicting fundamental preferences. Standard principal-agent models of such firms begin with the premise that the lower-level agents have preferences that diverge from the preferences of the principal – the firm’s leadership, somehow defined – so that researchers may want to slack off, or to research questions that are of maximal interest to them rather than of maximal expected utility to the firm, and so on. But the picture we advance has some utility to the extent that compensation mechanisms, incentive schemes, informal norms, or selection and screening at the hiring stage align the interests of researchers with the interests of the firm and its principal(s). Where that is so, it is not impossible to understand the optimal research and development process in for-profit firms along the lines we suggest, and to reconfigure the research and development process within actual firms accordingly.

E. Basic Research

To illustrate the limits of the analysis, we suggest that there is a large domain in which the social utility of the bees’ decision-making mechanism is more limited: basic research, especially in the natural sciences. By basic research we mean research that has no currently foreseeable applications or direct payoff for applied sciences, such as engineering. Although society does well to fund a portfolio of basic research,

some small fraction of which will pay off handsomely in the long run, there is no expectation that the payoff will materialize in the short run. Theoretical research in physics is the standard example.

Importantly, the process has both an epistemic dimension and a dimension of collective choice. The epistemic dimension is that we envisage basic research as a process of searching for theories about (some aspect of) the world, and we assume that some theories are objectively correct while others are objectively incorrect. The collective choice component can be understood either from the standpoint of funding institutions, or from the standpoint of scientists themselves. For funders, typically panels or other groups, the problem is to fund a portfolio of basic research that maximizes net present value to the funding institution or to society generally. For scientists, the problem is to generate a set of possible theories and then to reach consensus on the ones that are true, while collectively rejecting those that are false. As we will see, these twin collective aims – theory generation and theory sifting – trade off against one another.

In basic research, the bees' mechanism would entail that researchers focus their attention primarily on questions identified by other researchers, and then proceed to address them independently. This has indeed been the tendency in basic research, which is increasingly conducted in teams across the natural sciences. Such teams in effect focus on questions identified by a leader or head, or on questions that have become fashionable, having received attention from other researchers; team members do not ask whether the question is the right one to pursue, only how, if at all, it can be answered. Although there are powerful institutional and individual incentives to form such teams, their social utility is an open question.

Our analysis suggests that an excessively high degree of interdependence at the agenda-setting stage of basic research is undesirable. The reason is that the opportunity cost of failing to reach consensus on basic theories is relatively low, where one assumes as we do that consensus on those theories will have no immediate payoff. Basic research is a long-run enterprise, in which it is better that things be settled right, eventually, than that things be settled today. The cost of failing to reach consensus – the main cost of low interdependence – is typically a lesser concern; the greater concern is, or should be, that excessive

interdependence may leave some very promising theories sitting about undiscovered, because no one has been searching for them.

Our point is not that basic research teams are inaccurate, or unreliable. With the bees' combination of high interdependence and high independence, basic research teams will make accurate assessments of theories put on their scientific agenda by team leaders, yet some excellent theories may go unconsidered by anyone. With lower interdependence, more scientists will roam the theoretical space alone or in smaller teams; more theories will be explored, but by fewer people in each case. That retards the generation of consensus on the theories considered, yet reduces the number of true theories overlooked altogether.

By the nature of the case, it is hard to know whether there are many excellent theories waiting to be found, and if so where they might be (if one knew those things, one would already be in position to find the theories). But we can motivate our view by pointing to the domination of string theory within theoretical physics, as a plausible example of the costs of interdependence. In the current generation, "it is virtually impossible (in the U.S.) for someone not working within [the string theory] paradigm to be hired as an assistant professor [of physics] at a major research university" (Elster 2009, pp. 19-20). According to critics, theoretical physics is focused to the point of obsession on evaluating and expanding one particular theory or family of theories while other approaches go unexplored. Meanwhile, a growing number of people have begun to question whether string theory is even a scientific enterprise at all, given the difficulty of using the theory to generate implications that are both testable and unique to the theory (Elster 2009, p. 20; Woit 2006). If the criticism is right, theoretical physics suffers from excessive interdependence of the research agenda. Similar social dynamics among researchers may be responsible for the excessive disciplinary rigidity of which mainstream economics is sometimes accused.¹⁴

F. Statistical Groups: Ratings and Individual Choices

We may also consider an extension, from the judgments made by actual groups to the virtual judgments made by statistical or notional groups, and to individual choices that rely upon the judgments of those

statistical groups. An example of what we have in mind is the website “Rate My Professors,”¹⁵ where students can see, for any given professor, an “overall quality” score that averages all ratings, as well as composite scores for “easiness” and “hot(ness).” Students who use websites of this sort to decide where to allocate their course time are implicitly relying upon the collective judgment of a statistical group. Importantly, not all professors have ratings. Which professors will be rated depends on the decentralized choices of other students, who in effect advertise the (high or low) quality of a professor by choosing to rate them. That choice in turn influences the choices of later students to take or not to take a given course; the students who do so may then record an independent assessment of the professor’s quality, which will in turn influence the choices of yet later students, and so on. There is interdependence among students at the stage of deciding which course to take, yet independence (ideally, at least) at the stage of rating.

More generally, a similar combination of interdependence and independence can appear in a broad range of individual choices influenced by the decisions of earlier participants to advertise (including in negative terms) the quality of the choices. The relevant category here is the *rating system*, which can be distinguished from a *ranking system*; the former is decentralized while the latter is inherently centralized. Under a ranking system, some individual or group attempts to reach a synoptic overview of the relevant choices – an exogenously defined set – and compares them all with one another in order to arrive at an overall ranking. Under a rating system, by contrast, collective judgments emerge from the decentralized action of participants who search the space of possible choices, advertise for or against candidates they like or dislike, and thereby influence the choices of other participants whether or not to consider or sample the relevant goods. Although individuals decide which options to consider or pursue by considering the aggregate judgment of a virtual group, their assessment of those options is (ideally) independent. In this respect, the growing prevalence of rating systems for consumer goods – particularly experience goods like films,¹⁶ hotels, vacations,¹⁷ and gourmet restaurants¹⁸ – shows that *homo sapiens* has already begun to do as the hive does, although many millennia after the hive perfected its mechanism of decentralized search.

G. Interdependence, Independence and Institutional Design: Veiling Mechanisms

In any decisional environment that satisfies the three conditions we have described – an open-ended agenda, high stakes, and high opportunity costs of indecision – the challenge for institutional designers is to balance interdependence and independence. The former is beneficial at the stage of agenda-setting, in order to coordinate on suitable agenda items, while the latter is beneficial at the stage of evaluation of those items, in order to maximize the accuracy of the resulting decision.

A difficulty, however, is that interdependent agenda-setting might spill over into the stage of evaluation. Suppose, for example, that in an executive search committee one committee member is known to have placed a candidate on the agenda because he or she considers the candidate to be of very high quality. That knowledge might then affect the others' substantive evaluations of the candidate, reducing their independence. Indeed, taking into account the information embodied in the agenda proposals of others may be individually rational for any given participant, so long as the benefits of incorporating the information are greater than the additional cognitive costs.¹⁹

There is thus a potential disanalogy between bees and humans. In the bees' mechanism, individuals do not directly observe others' full quality assessment, which is encoded in the duration of their dances. Dance duration indirectly affects the probability that other bees will investigate a given site, but there is no direct spillover of information from the agenda-setting stage to the evaluation stage. Observing duration may exceed the cognitive capacities of any individual bee. Thus the bees are indirectly advantaged by their low cognitive capacities, combined with the structure of their decision-making mechanism. Paradoxically, humans' superior cognitive capacity enables them to infer information from agenda proposals, compromising independence; humans thus face the spillover problem.

The proper response to the problem, however, is not to declare the bees' mechanism irrelevant to human collective choice in epistemic contexts. Rather, the challenge is to replicate the conditions for the bees' success by employing strategies of mechanism design. The same human ingenuity that creates the spillover problem can also ameliorate it, by enabling human institutional designers to develop

mechanisms that shield the stage of independent evaluation – in whole or at least in part – from the information required for interdependent agenda-setting.

The precise mechanisms that can enforce such shielding are highly context-specific. But we will describe one general class of *veiling mechanisms*²⁰ that in one way or another limit the information held by participants when evaluating proposals others have placed on the agenda. Such mechanisms will create a barrier between interdependence and independence to prevent spillovers that might compromise the latter. By depriving human decision makers of information they might use – quite rationally – in ways that compromise independence, veiling mechanisms indirectly replicate the bees’ lower cognitive capacities, with the paradoxical result of improving group performance.

Here are some examples. In executive search committees, one might establish an impartial officer to serve as a depository for proposals to place a candidate on the agenda. The officer will inform members of the bare fact that some other member has proposed consideration of the candidate, but nothing more. Whereas under ordinary procedures, each member of the committee will know who proposed the candidate and will hear a formal or informal presentation of reasons in the candidate’s favor – information that might compromise independence at the stage of evaluation – the laundering of proposals through an intermediary acts as a partial veil that reduces the flow of compromising information.

Likewise, the Supreme Court might experiment with veiling procedures. Under current practice, the pool memorandum that recommends a grant of certiorari often contains a detailed statement of the case and of the parties’ arguments, and lists the name of the authoring law clerk. Although we doubt that this information has a large compromising effect on the independence of evaluation by Justices and other clerks, especially since cases are ultimately decided only after full briefing and argument, one might eliminate the law clerks’ names and at least some of the other information to minimize the risk of spillover. In these and other settings, spillover is not only a problem, but also an opportunity to design mechanisms that replicate and exploit the advantages of the bees’ procedure.

Conclusion

We conclude by underscoring the central programmatic implications of our analysis for the study of institutional design. Whatever the merits of the honeybees' collective-choice mechanism, it reveals two major gaps in the literature on epistemic collective choice: the role of agenda-setting and the importance of time. More broadly, the honeybees' mechanism illustrates a mode of arbitrage – from the evolved decision-making strategies of non-human animals to the design of human institutions – that amounts to a form of comparative institutional analysis across species. We offer a few remarks on each point in turn.

Epistemic agenda-setting. Machiavelli ([ca. 1513] 1996) observed that “a multitude without a head is useless.” In his motivating example, a group of plebians who threatened to secede from Rome proved entirely incapable of negotiating with the patricians, because the plebians had no leader to make proposals to them and speak for them. Although Machiavelli did not clearly distinguish between the aggregation of judgments and preferences, his observation holds in either setting, insofar as he is pointing out that it is often costly for decision-making groups to structure their own agendas. Even when a decision-making group shares all fundamental preferences in common, it must make choices between alternatives, and the alternatives must come from somewhere. Given realistic constraints on the time and cognitive capacities of the group, it is not feasible for all members of the group to put as many proposals as they see fit on the agenda, and then for the group to vote on all proposals. Instead, some member or members must act as epistemic agenda-setters who narrow the range of options.

From another standpoint, however, Machiavelli's claim is misleading. The epistemic agenda-setter on a given issue need not be a “head” or leader, in the sense of an individual member, specified *ex ante*, who possesses agenda-setting authority across the board. Rather the agenda-setter can be a rank-and-file member, or the member of a search committee, who advertises the quality of a given option to other members, who in turn coalesce around a particular option and influence the choice of the group as a whole. In the honeybees' procedure, there are epistemic agenda-setters, but they are otherwise

unremarkable rank-and-file bees. It is pragmatically necessary that there be epistemic agenda-setters, but that is a different topic than leadership.

Many issues lurk here, and there is much to be explored. The honeybees' procedure merely illustrates that work on collective epistemic decision making, in economics and rational choice theory more generally, has neglected the issue of epistemic agenda-setting, and has generally rested content with models that treat the options for decision as exogenous. The next generation of epistemic models should relax this assumption.

Time, truth-tracking, and collective search. A second and related issue is that Condorcetian models of collective epistemic decision making are excessively static. Those models illuminate the idea that collective decisions can produce correct or incorrect judgments, relative to the common aims of the group; in that modest sense, the models attempt to show conditions under which group decisions can "track the truth" (Grofman, Owen and Feld 1983, Estlund 1993, List and Goodin 2001). Yet the Condorcetian models show little appreciation of the brute fact that searching for the truth takes time. Sensible groups will trade off the benefits of obtaining the very best answer against the opportunity costs of information exchange, deliberation, and possible failure to reach consensus, resulting in no group decision at all (in effect a decision for the status quo, which may be untenable or the worst option of all). By contrast, models of search processes in economics, behavioral economics and biology are acutely sensitive to the opportunity costs of search and to tradeoffs between speed and accuracy. Yet those models frequently involve a single decision maker, perhaps a consumer, and thus abstract away from the crucial collective epistemic problem: multiple decision makers have different beliefs and information, which must somehow be aggregated through an optimizing procedure.

The honeybees' decision procedure lies at the intersection of these problems. The honeybees face a problem of collective search in an epistemic context. Some options are much better than others, given the common fundamental preferences of the group, yet different members of the group have different information than others, and the group's problem is to pick an optimal strategy for collective search. Of

course, the honeybees do not pick a strategy, either individually or collectively; yet natural selection has produced individual-level behaviors that cause the bees to behave in a collectively optimal fashion. For human purposes, the bees' as-if collective search strategy is worth considering in any decision-making environment that resembles the bees' environment. As we have tried to show, humans face similar environments in a range of institutional settings.

Comparative institutional analysis across species. The last point explains the sense in which, and the conditions under which, the hive provides useful lessons for humans. Precisely because individual bees have such low cognitive capacities, the apparent efficiency of their collective search strategy, in their environment, is all the more striking. To the extent that their aggregation mechanism exhibits efficiency, the human problem is to describe the features of their environment in suitably abstract terms – the terms of economic theory and decision theory – and then to ask whether there are similar human environments or decision problems to which the bees' mechanism might be carried over. Nothing in this process of analysis, abstraction and transposition requires drawing dubious analogies between humanity and the hive. Drawing lessons from the hive is merely a form of comparative institutional analysis *across* rather than *within* species. Just as institutional designers may observe other human organizations in the social world, to find institutional forms that would otherwise never have occurred to them, so too designers may observe the products of natural selection to broaden their repertoires.

References

- Akçay, E., J. Roughgarden, J. Fearon, J. Ferejohn and B. Weingast. 2010. Biological institutions: The political science of animal cooperation. Working paper, Stanford University.
- Austen-Smith, D. and T. J. Feddersen. 2006. Deliberation, preference uncertainty, and voting rules. *American Political Science Review* 100: 209–217.
- Berend, D. and L. Sapir. 2007. Monotonicity in Condorcet's Jury Theorem with dependent voters. *Social Choice and Welfare* 28(3): 507-528.

- Bhushan, B. 2009. Biomimetics: lessons from nature – an overview. *Philosophical Transactions of the Royal Society A* 367(1893): 1445-1486.
- Bikhchandani, S., D. Hirshleifer and I. Welch. 1992. A Theory of Fads, Fashions, Custom, and Cultural Change as Informational Cascades. *Journal of Political Economy* 100: 992-1026.
- Black, Ryan C. and Ryan J. Owens. 2009. Agenda setting in the Supreme Court: The collision of policy and jurisprudence. *Journal of Politics* 71: 1062-1075.
- Boland, P. J. 1989. Majority Systems and the Condorcet Jury Theorem. *The Statistician* 38: 181-189.
- Conradt, L. and C. List, eds. 2009. *Group decision making in humans and animals*. Theme issue of *Philosophical Transactions of Royal Society B* 374: 717-852.
- Conradt, L. and T. J. Roper. 2005. Consensus decision making in animals. *Trends in Ecology and Evolution* 20: 449-456.
- Cox, G. W. and M. D. McCubbins. 2007. *Legislative Leviathan: Party Government in the House*, 2nd edition. Cambridge: Cambridge University Press.
- Depew, David J. 1995. Humans and other political animals in Aristotle's *History of Animals*. *Phronesis* 40: 156-181.
- Dietrich, F. 2008. The premises of Condorcet's jury theorem are not simultaneously justified. *Episteme* 5(1): 56-73.
- Dietrich, F. and C. List. 2004. A Model of Jury Decisions Where All Jurors Have the Same Evidence. *Synthese* 142: 175-202.
- Dietrich, F. and K. Spiekermann. 2013. Epistemic Democracy with Defensible Premises. *Economics and Philosophy* 29(1): 87-120.
- Dyer, J., A. Johansson, D. Helbing, I. D. Couzin and J. Krause. 2009. Leadership, consensus decision making and collective behaviour in humans. *Phil. Transactions of the Royal Society B* 364: 781-789.
- Elster, Jon. 2009. Excessive Ambitions. *Capitalism and Society* 4(2): Article 1.

- Estlund, D. 1993. Making truth safe for democracy. Pp. 71-100 in D. Copp, J. Hampton and D. Copp, eds. *The Idea of Democracy*. New York: Cambridge University Press.
- Estlund, D. 1994. Opinion Leaders, Independence, and Condorcet's Jury Theorem. *Theory and Decision* 36: 131-162.
- Feddersen, T. and W. Pesendorfer. 1996. The Swing Voter's Curse. *American Economic Review* 86(3): 408-424.
- Franks, N. R. et al. 2009. Speed versus accuracy in decision-making ants: expediting politics and policy implementation. *Philosophical Transactions of the Royal Society B* 364: 845-852.
- Gressman, Eugene et al. 2007. *Supreme Court Practice*. 9th edition.
- Grofman, B., G. Owen and S. L. Feld. 1983. Thirteen theorems in search of the truth. *Theory and Decision* 15: 261-278.
- Hobbes, Thomas. 1651. *Leviathan*.
- Kaniovski, S. 2008. Aggregation of correlated votes and Condorcet's Jury Theorem. *Theory and Decision* 69(3): 453-468.
- Kitson, M., ed. 2005. *Economics for the Future*. Special issue, *Cambridge Journal of Economics* 29(6).
- Ladha, K. 1992. The Condorcet Jury Theorem, Free Speech and Correlated Votes. *American Journal of Political Science* 36: 617-634.
- Lindauer, M. 1955. Schwarmbienen auf Wohnungssuche. *Zeitschrift für vergleichende Physiologie* 37: 263-324.
- List, C., C. Elsholtz, C. and T. D. Seeley. 2009. Independence and interdependence in collective decision making: an agent-based model of nest-site choice by honeybee swarms. *Philosophical Transactions of the Royal Society B* 364: 755-762.
- List, C. and R. E. Goodin. 2001. Epistemic Democracy: Generalizing the Condorcet Jury Theorem. *Journal of Political Philosophy* 9: 277-306.

- List, C. and P. Pettit. 2004. An Epistemic Free Riding Problem? In P. Catton and G. Macdonald, eds. *Karl Popper: Critical Appraisals*. London: Routledge.
- List, C. and P. Pettit. 2011. *Group Agency: The Possibility, Design and Status of Corporate Agents*. Oxford: Oxford University Press.
- Lorenz, J., H. Rauhut, F. Schweitzer and D. Helbing. 2011. How social influence can undermine the wisdom of crowd effect. *Proceedings of the National Academy of Sciences* 108 (28): 9020-9025.
- Machiavelli, N. [ca. 1513] 1996. *Discourses on Livy*, translated by H. C. Mansfield and N. Tarcov. Chicago: University of Chicago Press.
- Mandeville, Bernard. 1714. *The Fable of the Bees; or, Private Vices, Publick Benefits*.
- Perry, Jr., H.W. 1991. *Deciding to Decide: Agenda Setting in the United States Supreme Court*. Cambridge, Massachusetts: Harvard University Press.
- Roberts, Jr., John G., Chief Justice. 2006-2009. *2006 Year-End Report on the Federal Judiciary* 9; *2007 Year-End Report on the Federal Judiciary* 9; *2008 Year-End Report on the Federal Judiciary* 8; *2009 Year-End Report on the Federal Judiciary* 2. All available at: <http://www.supremecourt.gov/publicinfo/year-end/year-endreports.aspx>
- Rodgers, Diane M. 2009. *Debugging the Link Between Social Theory and Social Insects*. Baton Rouge, LA: Louisiana State University Press.
- Seeley, T. D. 2005. *The Wisdom of the Hive: The Social Physiology of Honey Bee Colonies*. Cambridge, MA: Harvard University Press.
- Seeley, T. D. 2010. *Honeybee Democracy*. Princeton, NJ: Princeton University Press.
- Seeley, T. D. and Buhrman, S. C. 2001. Nest-site selection in honey bees: how well do swarms implement the ‘best-of-N’ decision rule? *Behavioral Ecology and Sociology* 49: 416-427.
- Seeley, T. D. and P. K. Visscher. 2006. Group decision making in nest-site selection by honey bees. *Apidologie* 35: 101–116.

Seeley, T. D., P. K. Visscher and K. M. Passino. 2006. Group decision making in honey bee swarms. *American Scientist* 94: 220-229.

Sunstein, C. R. 2006. *Infotopia: how many minds produce knowledge*. New York: Oxford University Press.

Sunstein, Cass R. 2009. *A Constitution of Many Minds*. Princeton: Princeton University Press.

Vermeule, Adrian. 2007. *Mechanisms of Democracy: Institutional Design Writ Small*. New York: Oxford University Press.

Vermeule, Adrian. 2009. *Law and the Limits of Reason*. New York: Oxford University Press.

Woit, Peter. 2006. *Not Even Wrong: The Failure of String Theory and the Search for Unity in Physical Law*. New York: Basic Books.

Appendix: A Formal Model

In what follows, we give a brief summary of the formal model and simulation results from List, Elsholtz and Seeley (2009), on which we draw in the main text.

Model

There are

- n scout bees, labeled $1, 2, \dots, n$;
- k potential nest sites, labeled $1, 2, \dots, k$, where $q_j \geq 0$ is the objective quality of site j ;
- discrete time periods, labeled $1, 2, 3, \dots$

The state of each bee i at time t is represented by a pair $x_{i,t} = (s_{i,t}, d_{i,t})$, where

- $s_{i,t} \in \{0, 1, 2, \dots, k\}$ is the nest site for which the bee dances at the given time ($s_{i,t} = 0$ means “no dance”), and
- $d_{i,t} \geq 0$ is the remaining duration of the dance (measured in number of time periods).

The initial state of the model is $x_{i,1} = (0,0)$ for all i ("no dancing at the beginning"). The state of each bee in each time period depends on her own state and that of the other bees in the previous period. We consider two cases.

The first case: bee i is in the non-dancing state at time t . Formally, $s_{i,t} = 0$. Bee i then commences a dance for one of the sites $1, 2, \dots, k$ or remains in the non-dancing state at time $t+1$ with probabilities $p_{1,t+1}, p_{2,t+1}, \dots, p_{k,t+1}$ and $p_{0,t+1}$, respectively. This determines $s_{i,t+1}$. The probabilities (for $j > 0$) are given by the formula

$$p_{j,t+1} = (1 - \lambda)\pi_j + \lambda f_{j,t},$$

where π_j is the *ex ante* probability of finding site j , $f_{j,t}$ is the proportion of bees dancing for site j at time t , and λ is the *interdependence* parameter, capturing how much the bees influence each other through signaling. The remaining dance duration at time $t+1$ is given by the formula

$$d_{i,t+1} = \begin{cases} q_j \exp(T_\sigma) & \text{with probability } 1 - \mu \text{ ("independent assessment")} \\ K \exp(T_\sigma) & \text{with probability } \mu \text{ ("non-independent mimicking")}, \end{cases}$$

where T_σ is a normally distributed random variable with mean 0 and standard deviation $\sigma \geq 0$ (the *reliability* of bee i – the lower, the more reliable), and K is some strictly positive constant, not related to the quality of nest site j . The parameter μ is a proxy for the independence between bees in their assessment of any nest site, once they have identified it. This completes the specification of $x_{i,t+1} = (s_{i,t+1}, d_{i,t+1})$.

The second case: bee i is dancing for one of the sites at time t . Formally, $s_{i,t} > 0$. Here

$$x_{i,t+1} = \begin{cases} (s_{i,t}, d_{i,t} - 1) & \text{if } d_{i,t} > 1 \\ (0,0) & \text{otherwise.} \end{cases}$$

Now, a *consensus* between the bees is reached when the total number of bees dancing for one of the sites at some time t ,

$$n_{j,t} = |\{i : s_{i,t} = j\}|,$$

meets an appropriate “quorum” criterion, such as the following illustrative ones:

- site j receives more dancing support than any other site at time t , i.e., $n_{j,t} > n_{h,t}$ for any $h \neq j$ with $n_{h,t} \neq 0$ (a “weak consensus” criterion);
- site j receives more than twice as much support as the second most supported site at time t , i.e., $n_{j,t} > 2n_{h,t}$ for any $h \neq j$ with $n_{h,t} \neq 0$, and more than 20% of the scout bees are engaged in dancing (a “strong consensus” criterion).

Of course, other criteria could be used as well, but the computational results were relatively robust.

Results

Using this model, List, Elsholtz and Seeley (2009) provide computational support for the following hypotheses:

- For non-extremal values of σ and λ (and $\mu = 0$), the bees choose the best site (*Hypothesis 1*).
- The bees’ independence in assessing the sites’ quality and their interdependence through signaling are necessary and sufficient for the reliability of the decision process (*Hypothesis 2*).

In the simulations reported in the cited paper, the number of scout bees was fixed at $n = 200$ and the number of nest sites at $k = 5$. For illustrative purpose, the sites’ quality levels q_1, \dots, q_5 were fixed at 3, 5, 7, 9, 10, capturing differences across sites as well as closeness between the top sites. These parameter values were motivated by Seeley’s empirical studies. It was also assumed that the *ex ante* probabilities of finding one of the sites was just 25%, evenly distributed across the five sites, so that $\pi_1 = \dots = \pi_5 = 5\%$ and $\pi_0 = 75\%$. The bees’ behavior was simulated over 300 time periods.

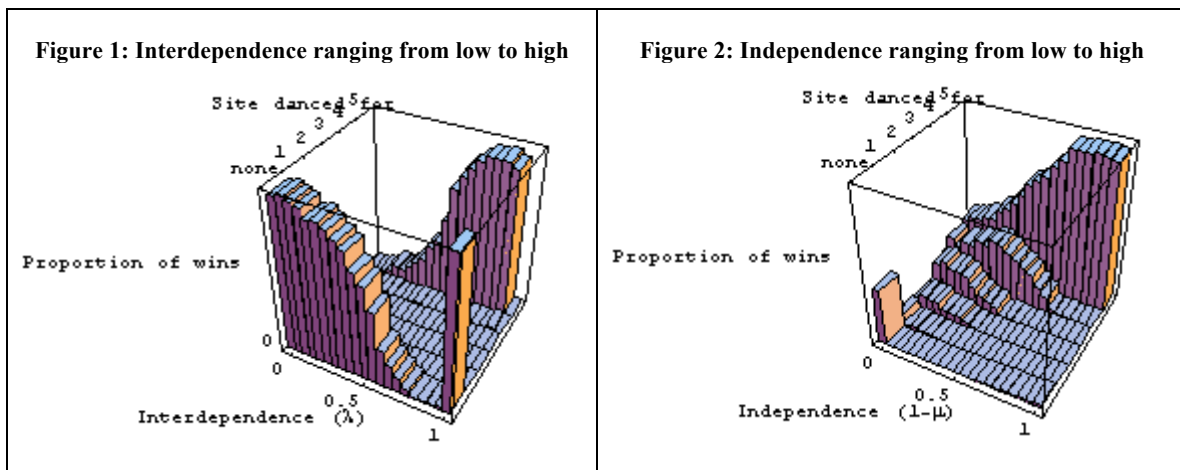
Table 1, from List, Elsholtz and Seeley (2009), shows how often each nest site was chosen in 250 runs of the simulation, for different levels of individual reliability and interdependence, using the two consensus criteria defined above. The results are broadly consistent with Hypothesis 1.

		High individual reliability ($\sigma = 0.2$)		Low individual reliability ($\sigma = 1$)	
		Strong consensus criterion	Weak consensus criterion	Strong consensus criterion	Weak consensus criterion
High interdependence ($\lambda = 0.8$)	1 st best site	246 (98.4%)	250 (100%)	199 (79.6%)	237 (94.8%)
	2 nd best site	0 (0%)	0 (0%)	5 (2%)	12 (4.8%)
	3 rd best site	0 (0%)	0 (0%)	0 (0%)	0 (0%)
	None	4 (1.6%)	0 (0%)	46 (18.4%)	1 (0.4%)
Medium interdependence ($\lambda = 0.5$)	1 st best site	104 (41.6%)	226 (90.4%)	94 (37.6%)	220 (88%)
	2 nd best site	0 (0%)	22 (8.8%)	1 (0.4%)	28 (11.2%)
	3 rd best site	0 (0%)	0 (0%)	0 (0%)	0 (0%)
	None	146 (58.4%)	2 (0.8%)	155 (62%)	2 (0.8%)
Low interdependence ($\lambda = 0.2$)	1 st best site	11 (4.4%)	176 (70.4%)	7 (2.8%)	190 (76%)
	2 nd best site	1 (0.4%)	63 (25.2%)	0 (0%)	58 (23.2%)
	3 rd best site	0 (0%)	1 (0.4%)	0 (0%)	0 (0%)
	None	238 (95.2%)	10 (4%)	243 (97.2%)	2 (0.8%)

Table 1: Simulation Results for Hypothesis 1 (from List, Elsholtz and Seeley 2009)

To separate out the effects of independence and interdependence, List, Elsholtz and Seeley (2009) further ran a number of simulations with the level of interdependence ranging from low ($\lambda = 0$) to high ($\lambda = 1$), while assuming high individual reliability and independence ($\sigma = 0.2$ and $\mu = 0$), and a number of simulations with the level of independence ranging from low ($\mu = 1$) to high ($\mu = 0$), while assuming high interdependence and individual reliability in case a bee checks a site's quality ($\lambda = 0.8$ and $\sigma = 0.2$).

Figures 1 and 2 display the proportion of wins for each of the five sites, ordered by quality from worst (site 1) to best (site 5) (in 250 simulations for each set of parameter values), for different levels of interdependence and independence, where the stronger consensus criterion is used for determining the chosen site. The results are broadly consistent with Hypothesis 2.



Figures 1 and 2: Simulation Results for Hypothesis 2 (from List, Elsholtz and Seeley 2009)

For further details, readers are directed to List, Elsholtz and Seeley (2009).

Notes

* C. List, London School of Economics, Departments of Government and Philosophy; A. Vermeule, Harvard Law School. We are grateful to Jacob Gersen, Aziz Huq, Daryl Levinson, Martha Minow, Eric Posner, participants at a conference at the University of Chicago Law School in September 2010, and the anonymous reviewers for comments, and, in addition, to Larissa Conrads, Franz Dietrich, Christian Elsholtz and Thomas Seeley for helpful conversations in related collaborations. We also thank Janet Kim for helpful research assistance.

¹ Famous examples include Hobbes's *Leviathan* (1651, ch. 17), and Aristotle's *History of Animals*, cited and quoted in Depew (1995, p. 156).

² The idea of a fruitful dialogue between research on human and non-human collective decisions has been suggested in a recent symposium, edited with an introductory survey (pp. 719-742) by Conrads and List (2009). A noteworthy interdisciplinary study included in this symposium is the one by Dyer et al. (2009). In a subsequent working paper, Akçay et al. (2010) suggest that biologists can learn from institutional analysis in political science. We reverse their emphasis by exploring what human institutional designers can learn from non-human animals.

³ We here employ the notions of “preferences” and “beliefs” as they are used in decision theory, micro-economic theory, and philosophy more generally. Preferences and beliefs are different kinds of intentional attitudes of an agent. Beliefs are representational attitudes encoding what the agent takes the world to be like. Preferences are conative attitudes encoding what the agent wants the world to be like. A rational agent, very roughly speaking, pursues his or her preferences in accordance with his or her beliefs. For further details on this notion of agency, see, e.g., List and Pettit (2011, ch. 1).

⁴ The terminology of “independence” and “interdependence” was introduced in relation to honeybee decisions in List, Elsholtz and Seeley (2009), building on earlier empirical work (e.g., Seeley, Visscher and Passino 2006, Lindauer 1955).

⁵ On the independence condition and various relaxations of the condition, see, e.g., Boland (1989), Ladha (1992), Estlund (1994), Dietrich and List (2004), Berend and Sapir (2007), Kaniowski (2008), Dietrich (2008), and Dietrich and Spiekermann (2013). Outside the specific literature on Condorcet’s jury theorem, there is, of course, a sizeable body of work on how rational agents update their beliefs in response to signals received from others, but, unlike the present paper, that work does not focus on epistemic agenda setting, and so the kind of interdependence between agents analyzed in that literature is different from the one discussed here. See, among many others, Austen-Smith and Feddersen (2006). There is also social-scientific experimental evidence suggesting that the wisdom-of-crowds effect can be undermined by social influences. See, e.g., Lorenz, Rauhut, Schweitzer and Helbing (2011).

⁶ The following discussion draws on List and Pettit (2011, ch. 4). See also the earlier references on various relaxations of independence.

⁷ Formally, jurors are no longer independent conditional on the original truth about the crime here; they are at most independent conditional on the shared body of evidence (Dietrich and List 2004).

⁸ On the notion of epistemic free-riding, see List and Pettit (2004).

⁹ To be sure, this point is only consistent with epistemic voting; it does not necessarily demonstrate its existence. Unanimity does not logically entail that there is no conflict of fundamental preferences in the

case at hand. Under unusual circumstances, it is possible that there are two blocs of Justices who have opposed fundamental preferences but also have opposing beliefs, and who thus share identical derived preferences, although for completely different reasons. But unanimity is at least compatible with, and provides some evidence of, lack of deep preference conflict. The most casual glance at the Court's decisions in any Term, especially the ones issued between November and (say) April, will show a large number of unanimous or near-unanimous decisions in humdrum technical cases, on matters of regulation, taxation and court procedure, in which it is implausible that there are fundamental conflicts.

¹⁰ Gressman et al. (2007), pp. 242-50 (“The Supreme Court often, but not always, will grant certiorari where the decision of a federal courts of appeals, as to which review is sought, is in direct conflict with a decision of another court of appeals on the same matter of federal law or on the same matter of general law as to which federal courts can exercise independent judgments. One of the primary purposes of the certiorari jurisdiction is to bring about uniformity of decisions on these matters among the federal courts of appeals.”) (emphases in original deleted); *id.*, at 264-67 (“Where the decision below holds a federal statute unconstitutional or where a federal statute is given an unwarranted construction in order to save its constitutionality, certiorari is usually granted because of the obvious importance of the case.”)

¹¹ From October Term 2004 to October Term 2008, the average number of petitions was 8170.6 (7496 cases were filed in the 2004 term, 8521 in 2005, 8857 in 2006, 8241 in 2007, and 7738 in 2008). See Roberts, Jr. (2006-2009).

¹² In the years 2004 to 2008, an average of 82.8 cases were argued before the Court. (87 in 2004, 87 in 2005, 78 in 2006, 75 in 2007, and 87 in 2008). See Roberts, Jr. (2006-2009).

¹³ See, e.g., <http://www.scottberkun.com/blog/2008/thoughts-on-googles-20-time/>, accessed 29/8/2010.

¹⁴ See, e.g., a special journal issue on “Economics for the Future”, edited with introduction (pp. 827-835) by Kitson (2005).

¹⁵ At <http://www.ratemyprofessors.com/>.

¹⁶ E.g., <http://www.rottentomatoes.com/>.

¹⁷ E.g., <http://www.tripadvisor.com/>.

¹⁸ E.g., the *Guide Michelin* and <http://www.yelp.com/>.

¹⁹ We might describe this as the “evaluator’s curse.” Cf. Feddersen and Pesendorfer (1996).

²⁰ For an analysis of veiling mechanisms in constitutional design, see Vermeule (2007), pp. 27-71.