

Independent Opinions?

Franz Dietrich & Kai Spiekermann¹

London School of Economics
October 2010

Abstract

Democratic decision-making is often defended on grounds of the ‘wisdom of crowds’: decisions are more likely to be correct if they are based on many independent opinions, so a typical argument in social epistemology. But what does it mean to have independent opinions? Opinions can be probabilistically dependent (threatening the ‘wisdom of crowds’) even if individuals form their opinion in causal isolation from each other. We distinguish four probabilistic notions of opinion independence. Which of them holds depends on how individuals are causally affected by environmental factors such as commonly perceived evidence. In a general theorem, we identify causal conditions guaranteeing each kind of opinion independence. These results have implications for whether and how ‘wisdom of crowds’ arguments are possible, and how truth-conducive institutions can be designed.

1 Introduction

When an individual forms opinions on factual questions and aims for correctness, it is often beneficial to adopt the opinion of an ‘expert’, a person with a high probability of being right. No doubt, such dependence on someone ‘who knows’ is a smart epistemic strategy from an individual perspective. At the same time, the dependence created by deference to experts will often have bad epistemic implications once we care about the ability of the *group* to form correct opinions. If many individuals follow experts, a set of *dependent* opinions is the result.

According to the popular idea of the ‘wisdom of crowds’, a decision based on many opinions is more likely to be correct than one based on just a few. Much of the trust in the judgment of large electorates, for instance, comes from the idea that a judgment is likely to be correct if it is approved by many voters. A single witness may well be mistaken, but twenty witnesses who all say the same may not, so the intuition. But these arguments assume that the group aggregates (sufficiently) *independent* individual opinions. Consider, for contrast, an extreme case of dependence: if all individuals were to adopt the opinion of the same expert, their opinions would be fully dependent, and the group would be no more competent than the expert. But if the individuals form independent opinions and are all at least somewhat better

¹The formal results presented in this paper are the work of Franz Dietrich.

than random at identifying the correct answer, then, as the famous Condorcet Jury Theorem tells us, the probability of the majority opinion being correct will increase with group size and converges to 1. Thus, the ‘wisdom of crowds’ hinges crucially on opinion independence.

Thinking about opinion independence reveals a systematic difference between individual and social epistemology. Generally speaking, individual epistemology recommends dependent opinions (in the form of positive correlation with experts), social epistemology recommends independent opinions. Pathologies of social opinion formation, such as informational cascades, biases and the influence of opinion leaders, are all instances of harmful opinion dependence. Therefore, if we are interested in collective rather than individual knowledge, we need institutions to ensure that individuals form independent opinions, possibly at a loss of individual epistemic success. (We also need a suitable aggregation procedure to use the independent opinions effectively, but aggregation is largely beyond the scope of this paper.²)

But what exactly does it mean to have independent opinions? First of all, one should not confuse causal and probabilistic independence. The former states that the opinions do not causally affect each other, the latter that they display no probabilistic dependencies (correlations). Ultimately, it is probabilistic—not causal—independence which matters for the ‘wisdom of crowds’, i.e., for the probability of correct aggregate opinions. Yet, as we argue, whether and in what way opinions are probabilistically (in)dependent is very much determined by causal interconnections between individuals and their environment. This paper distinguishes between different notions of probabilistic independence, and assesses which notion applies, depending on the underlying causal relations between individuals and other factors.

Assessing the epistemic quality of social practices and institutions is at the heart of social epistemology. Perhaps the simplest practice for forming collective opinions is voting. While we address voting, our analysis is much more general. The independence of opinions matters whenever collective opinion formation, of whatever sort, takes place. Group deliberation is a core example.

Some of the notions of independence developed in this paper lend themselves to *jury theorems*. Such theorems make independence and competence assumptions and conclude that ‘crowds are wise’, in the sense that the probability of a correct majority view increases as the group size increases (e.g., Grofman et al. 1983, Nitzan and Paroush 1984, List and Goodin 2001, Bovens and Rabinowicz 2006, Atkinson and Romeijn forthcoming, Hawthorne unpublished). The problem of opinion independence is tackled by several scholars. Among the many technical contributions are Boland (1989), Boland et al. (1989), Ladha (1992, 1993, 1995), Berg (1993) and Kaniovski (2010). But independence has also been discussed less formally in political philosophy, especially with reference to epistemic and deliberative democracy. Contributions include Grofman and Feld (1988), Estlund et al. (1989), Estlund (1994, 2008), Anderson (2006), Vermeule (2009, ch. 1) and Spiekermann and Goodin (2010). We approach the independence problem from a causal angle, which is crucial for gaining a deeper understanding of independence. To do so, we draw on causal networks.

²The Condorcet Jury Theorem, for instance, considers majority voting, a particular aggregation procedure for the case of dichotomous decisions.

Causal network reasoning has been employed before in the context of jury theorems (see Dietrich and List 2004, Dietrich 2008, Dietrich and Spiekermann 2010), but a general account of probabilistic independence in terms of causal interactions is still missing. In assuming an external standard for an opinion’s correctness, we are in line with the correspondence theory of truth and Alvin Goldman’s influential “veritism” approach in social epistemology (Goldman 1999, 59, 79ff.; 2004).

2 Four different independence conditions and their causal motivations

We assume that some individuals, labelled $i = 1, 2, 3, \dots$, must form opinions on a given issue.³ This might arise in the context of deciding between two alternatives, such as: convict or acquit the defendant in a court trial, predict that global warming will continue or that it will not, and so on. The opinions may, for instance, serve as votes in a formal voting procedure, or as inputs into group deliberation. It is easiest to think of an opinion as taking one of just two possible forms, such as a ‘yes’ or ‘no’ answer to a factual question, or a ‘guilty’ or ‘innocent’ vote in a jury. All our illustrations will follow this binary base-line case (as does the literature on jury theorems). But our model is much more general: opinions need not be binary and might for instance be sets of believed propositions (belief sets or judgment sets), numerical estimates (say, of the height of a mountain), or perhaps even degrees of belief. Exactly one opinion is ‘correct’, ‘right’ or ‘better’. Which opinion is correct is determined by an external fact, which is called the *state (of the world)* and is denoted x .⁴ For instance, the opinion ‘the defendant is innocent’ is correct if and only if the defendant has not committed the crime in question. Again, we shall mainly think of the state as taking one of only two possible forms, such as ‘yes’ or ‘no’ as the factually correct answer to a given question, following common practice in the literature (although the model places no restrictions on the possible forms of the state). What matters is that the state determines which opinion is correct, so that correctness supervenes on the state. Each person holds exactly one opinion; it might be correct or incorrect.

In our probabilistic framework, phenomena—such as the state and the opinions—are outcomes of random variables (with an underlying probability function denoted Pr). We thus consider a random variable \mathbf{x} generating the state of the world, and random variables $\mathbf{o}_1, \mathbf{o}_2, \dots$ generating the opinions of individuals 1, 2, ... Note our convention of using bold face letters to denote random variables. So, the variable \mathbf{x} ranges over the set of possible states, and each \mathbf{o}_i ranges over the set of possible opinions. What these sets are taken to be depends on the application; for instance, opinions might take the value ‘guilty’ or ‘innocent’, and the state the value ‘crime

³The total number of individuals does not matter for us. Technically, it may even be taken to be countably infinite.

⁴Different notions of objectivity are compatible with that assumption, as long as the fact about the correct opinion is not determined by the actual opinions of people. This excludes procedural notions of correctness, where an opinion is correct *if and because* it matches the opinion that arose collectively by applying an appropriate procedure.

committed’ or ‘crime not committed’.⁵

The simplest independence assumption one might come up with refrains from conditionalising on any background information:

Unconditional Independence (UI). The opinions $\mathbf{o}_1, \mathbf{o}_2, \dots$ are unconditionally independent.

Counterexamples to UI are easily constructed. In a context of voting between two alternatives, a and b , the information that the first ten voters vote for a plausibly raises the probability that voter 11 votes for a —formally, $\Pr(\mathbf{o}_{11} = a | \mathbf{o}_1 = \mathbf{o}_2 = \dots = \mathbf{o}_{10} = a) > \Pr(\mathbf{o}_{11} = a)$ —because the ten first votes strongly indicate that the state makes a correct (assuming that voters are reasonably competent in voting for the correct alternative), which in turn suggests that voter 11 also votes for a . But this is a violation of UI, whereby opinions should not be of any informational relevance to each other.

Note that this argument implicitly assumes the state \mathbf{x} has a causal effect on each opinion, as indicated in Figure 1. In this (and all following) plots, only two opinions




Figure 1: The state is a direct cause of the opinions.

are shown for simplicity. The arrows represent causal relationships, pointing from the causing variable to the affected variable. In Figure 1 the opinions are probabilistically dependent ‘through’ their common cause \mathbf{x} .

That UI is easily violated should not surprise scholars familiar with the Condorcet Jury Theorem, given that this theorem does *not* assume that votes are unconditionally independent but that they are state-conditionally independent. What is more surprising is that UI *does* hold in some circumstances, but we postpone this issue for now and turn to the more classical state-conditional notion of independence:

State-Conditional Independence (SI). The opinions $\mathbf{o}_1, \mathbf{o}_2, \dots$ are independent conditional on the state \mathbf{x} .⁶

⁵The way in which the correct opinion supervenes on the state is expressible by a ‘correctness’ function c from the set X of possible states to the set O of possible opinions, mapping each state x in X to the corresponding correct opinion $c(x)$.

⁶Conditional independence is defined like independence but with probabilities replaced by conditional probabilities. More precisely, as long as \mathbf{x} is discrete (e.g., binary), independence conditional on \mathbf{x} by definition means that for every value x which \mathbf{x} may take (with positive probability) there

This conditional notion of independence is the basis of Condorcet’s classical jury theorem (e.g., Grofman et al., 1983), which can be summarized as follows. Suppose a group performs a majority vote between two alternatives of which exactly one is correct. The correct alternative corresponds to our state \mathbf{x} , and the votes to our opinions $\mathbf{o}_1, \mathbf{o}_2, \dots$. Condorcet’s jury theorem makes an independence and a competence assumption. The former is simply SI. The latter is that, for every state x , the state-conditional probability of correctly voting for x exceeds $\frac{1}{2}$ (and is the same for each voter). Under these premises, the theorem concludes that the probability of a correct majority outcome increases in (odd) group size and converges to one.

State-Conditional Independence says that once we know the state of the world the opinions do not anymore bear any information on each other. The earlier objection to UI—namely that the opinions of some of the people tell us something about what the state is likely to be, and hence about what other people are likely to believe—does not work against SI because we can not learn anything new about the state if we have already conditionalised on it. The state plays the role of a common cause of the opinions. If the state is indeed the only common cause, SI is in line with Reichenbach’s famous *Common Cause Principle*, which is often understood roughly as follows: if a common cause fully explains the correlations between certain events, then these events are independent conditional on the common cause (Reichenbach 1956, 159-60). In slightly more general terms:

Common Cause Principle (stated informally). Phenomena which do not causally affect each other are probabilistically independent conditional on their common causes.

While the Common Cause Principle at first sight supports SI, it can be turned against SI once we consider other causal networks in which \mathbf{x} is not the only common cause of the opinions. Consider for instance the network in Figure 2. Here the opinions have two common causes, the state \mathbf{x} and another cause \mathbf{c} , which could be a factor like weather or room temperature. SI can now fail in much the same way as UI. Suppose for instance that weather has an effect on each juror in a court trial: the sunnier weather is, the more the jurors see the good in the defendant, and hence the more they are inclined to form the opinion that the defendant is innocent. Now, even after having conditionalised on the state of the world that the defendant is innocent, the opinions of the jurors are informative on each another, this time through the common cause of weather; for instance, ‘innocent’ opinions by the first ten jurors increase the probability that weather is sunny and hence the probability that juror 11 has the opinion ‘innocent’ too. In other words, the opinions are not state-conditionally independent but state-conditionally positively correlated, namely

is independence under the conditional probability function $\Pr(\cdot|x)$. Without discreteness restriction, the opinions are independent conditional on \mathbf{x} if they are independent under the conditional probability measure $\Pr(\cdot|x)$ for all values x that the random variable \mathbf{x} may take, except possibly from a set of values of \mathbf{x} that occurs with zero probability. (The clause ‘except...’ appears for technical reasons related to the general mathematical definition of conditional probabilities, which takes care of the case in which \mathbf{x} takes some or even all of its values with zero probability. We spare the reader the technicalities.)




Figure 2: Multiple direct common causes of opinions.

through the other common cause (‘weather’) on which SI fails to conditionalise.

This discussion suggests replacing SI by a notion of independence which conditionalises on all common causes of the opinions. By doing so we ‘control’ for all factors that causally affect more than one opinion, eliminating the dependence induced by such common factors. To state such a condition formally, let us extend the formal framework, which so far consists just of the state \mathbf{x} and the opinions $\mathbf{o}_1, \mathbf{o}_2, \dots$. Now we consider these and any number of additional random variables (representing phenomena which are directly or indirectly causally related to the opinions), and we consider a causal network over the variables. Formally, a causal network over some variables is a so-called *directed acyclic graph* over these variables, that is, a set of directed arrows between pairs of variables (representing causal relevance) such that there is no directed cycle of arrows.⁷ Figures 1 or 2 are examples of how the network might look like. Figure 3 is yet another example. Here, the state causally affects a




Figure 3: The state is an indirect common cause of opinions.

variable \mathbf{c} , which is interpretable as evidence (e.g., fingerprints, witness reports, etc.) and which in turn influences each opinion. Individuals are thus affected by the state only indirectly, via the ‘trace’ which the state leaves in the form of \mathbf{c} . The additional variables (such as the variable \mathbf{c} in Figures 2 and 3) may be binary or multi-valued. For instance, the variable weather may take the values ‘sunny’, ‘cloudy’, ‘rainy’ and so on; and the variable ‘body of evidence’ may take several forms as well. Some variables (such as room temperature) might even range over a continuum of values.

⁷For thorough discussions of causal networks, see Pearl 2000, ch. 1.

In the causal network, a variable \mathbf{a} is said to be a *direct cause* of another \mathbf{b} (and \mathbf{b} a *direct effect* of \mathbf{a}) if there is an arrow pointing from \mathbf{a} towards \mathbf{b} ($\mathbf{a} \rightarrow \mathbf{b}$). Further, \mathbf{a} is a *cause* of \mathbf{b} (and \mathbf{b} an *effect* of \mathbf{a}) if there is a directed path from \mathbf{a} to \mathbf{b} , i.e., a sequence of two or more variables starting with \mathbf{a} and ending with \mathbf{b} such that each of these variables (except from the last one) directly causes the next one. For instance, in Figure 3 the state \mathbf{x} directly causes \mathbf{c} , and indirectly causes the opinions. (When we use the verb “cause” we refer only to causal contribution; no sufficiency or necessity is implied.) A variable is a *common cause* (*effect*) of some variables if it is a cause (effect) of each of them. By a ‘common cause’ simpliciter we mean a common cause of (two or more) opinions. In all figures, such common causes are shown in grey. While in Figures 1-3 all causes of opinions are common causes, Figure 4 contains four *private causes* of opinions; they causally affect just one opinion. Note also that in Figure 4 some of the causes of opinions (namely, \mathbf{c}_2 , \mathbf{c}_4




Figure 4: Private and common causes of opinions.

and \mathbf{c}_6) are non-evidential: they are not related to the state. Although this might be viewed as ‘irrational’, individuals are often influenced by non-evidential causes such as room temperature (a *common* non-evidential cause) or the quality of one’s sleep last night (a *private* non-evidential cause).

Let us write χ (Greek ‘chi’) for the family of all common causes. In Figure 1 χ consists just of the state \mathbf{x} ; in Figures 2 and 3 of \mathbf{x} and \mathbf{c} ; and in Figure 4 of \mathbf{x} , \mathbf{c}_3 and \mathbf{c}_4 . In general, χ is a compound random variable with as many parts as there are common causes of opinions.⁸ We are now ready to state a new independence assumption, which is a direct application of the Common Cause Principle:

Common-Cause-Conditional Independence (CI). The opinions \mathbf{o}_1 , \mathbf{o}_2 , ... are independent conditional on the common causes χ .

This independence assumption may seem the most appealing one. It is backed by the Common Cause Principle and more generally by probabilistic theories of causality.

⁸The range of χ is the Cartesian product of the ranges of the common causes of opinions.

With CI, the independence of opinions is guaranteed as long as the opinions do not causally affect each other. It has, however, a weakness when applied in social epistemology, where we would like to feed opinion independence into an argument for the ‘wisdom of crowds’. The problem with CI is not so much that it is not sufficiently justified—CI is perhaps the most justifiable independence assumption—but rather that CI (like UI) is a premise which does not always lend itself to the kind of conclusions one is interested in.

To understand why, it is important to first realize that what matters ultimately is not independence of the *opinions* but rather independence of the *events of correct opinions*. The typical argument for the ‘wisdom of crowds’ says that a group whose members are independently more likely to get it right will quite probably get it right in majority. This argument involves independence of the events of holding correct opinions, not independence of the opinions (or votes) *simpliciter*. The kind of conclusion that independence of opinions *simpliciter* lends itself to is different and less relevant. One might reason that a group of jurors who are independently more likely to express the opinion ‘innocent’ will quite probably hold the majority opinion ‘innocent’. But what matters is less how likely a majority holds this opinion, but how likely it holds the *correct* opinion.

Now, independence of the opinions *implies* independence of the correct opinion events once we have conditionalised on the state \mathbf{x} . It is easy to see why: conditional on the state being x , if the opinions $\mathbf{o}_1, \mathbf{o}_2, \dots$ are independent then the events that \mathbf{o}_1 matches x , \mathbf{o}_2 matches x , ... are independent, too. In other words, the assumption of State-Conditional Independence implies what is needed for ‘wisdom of crowds’-type arguments, namely that events of correct opinions are independent across people. Hence, the independence assumption SI does not have the flaw of not pertaining to the events of correct opinions (which is why Condorcet’s classical jury theorem is based on SI rather than on UI). Similarly, if it so happens that the state \mathbf{x} features among the common causes χ —as it does indeed in all of the above Figures 1-4—then Common-Cause-Conditional Independence does not have the flaw either, as it then implies independence of the correct opinion events conditional on the common causes. But there are plausible situations in which the state \mathbf{x} is not a common cause. Figure 5 is one such case. Here the common cause \mathbf{c} affects both the state and the opinions.




Figure 5: The state is not a cause of opinions.

As a plausible example, imagine a murder case in which the jurors learn that the

defendant has bought cyanide (represented by \mathbf{c}). This fact is a common cause of the opinions of the jurors, who take murder (and guilt) to be more likely if the defendant has bought cyanide in advance. Since having bought cyanide facilitates poisoning, the purchase causally affects not just the opinions but also whether the murder takes place; hence the network of Figure 5. Note that state \mathbf{x} is not a cause of any opinions.

Whenever the state \mathbf{x} is not a common cause, CI does not conditionalise on it, and therefore does not lend itself to arguments about the probability of correct majority judgments, i.e., about the ‘wisdom of the crowds’. So we have to add the state into the conditionalisation, just as Condorcet’s jury theorem conditionalises on the state by using SI rather than UI. In total, we thus have to conditionalise on all common causes plus the state. But what does this mean? Following Dietrich (2008) and Dietrich and Spiekermann (2010), the *decision problem* faced by the group can be conceptualized as being a description of two things:

- the fact to find out about, conceptualized as the state of the world;
- the circumstances (environment) in which people search, conceptualized as the common causes influencing the opinions.

If we conditionalise on the decision problem, we include the state by default, thus making sure that not only the opinions but also the events of correct opinions are independent. We therefore have to conditionalise on the decision problem in the just-defined sense. Formally, let us write $\boldsymbol{\pi}$ for the decision problem defined as a family containing the state \mathbf{x} and all common causes. Clearly, the problem $\boldsymbol{\pi}$ reduces to the common causes $\boldsymbol{\chi}$ if the state is a common cause (as in Figures 1-4). In general $\boldsymbol{\pi}$ is isomorphic to the state-circumstances pair $(\mathbf{x}, \boldsymbol{\chi})$.⁹ We are now in a position to state our final independence condition:

Problem-Conditional Independence (PI). The opinions $\mathbf{o}_1, \mathbf{o}_2, \dots$ are independent conditional on the problem $\boldsymbol{\pi}$.

3 The causal foundation of each independence condition: a general theorem

While the last section has given informal causal motivations for the four independence conditions, this section turns to a formal result. The result gives us precise sufficient (and in fact essentially necessary) conditions on causal interconnections for each independence condition to hold. Given this result, once we know the individuals’ causal environment we can infer which kinds of opinion independence should (not) be assumed. And if a social planner can design the environment, he can do so to induce the kind of independence he aims for.

To be able to infer probabilistic features from causal interconnections, one must of course assume that probabilities are compatible with the causal network. What such compatibility amounts to has been settled precisely in the theory of causal (and Bayesian) networks (e.g., Pearl 2000). Formally, probabilities (more precisely: the

⁹This pair contains the state twice if the state is among the common causes, but such a redundancy poses no problem.

the prediction of a bank run might cause the bank run. Though interesting and sometimes very real, such cases violate one of the core assumptions of many theories of social epistemology (at least among those committed to a correspondence theory of truth): the assumption that an external fact determines correctness and not vice versa. ‘Self-fulfilling prophecies’ are ruled out.

In Figure 6b, by contrast, the state is a common effect of private causes of opinions. To show the relevance of such a setup, we need a more complex example. Suppose an intelligence agency observes two different subjects at different ends of the town. The agency knows from reliable sources that if and only if both subject 1 leaves the house at noon (\mathbf{c}_1) and subject 2 leaves the house at noon (\mathbf{c}_2), the two subjects will have a conspiratory meeting (\mathbf{x}). One agent observes subject 1, another subject 2, and for security reasons they cannot directly communicate with each other. Both agents form opinions on whether the meeting will take place. Each agent’s opinion is influenced only by his own observation (either \mathbf{c}_1 or \mathbf{c}_2), so that these two causes influence both the opinions and the state. This example shows that 6b is a plausible causal setup, but like for 6a the condition in clause (b) of Theorem 1 is violated and Problem-Conditional Independence should not be assumed. Indeed, PI is intuitively violated: conditional on the state, we can infer something about an agent’s opinion if we learn about other opinions. For instance, if we know that the conspiratory meeting does not take place (we condition on \mathbf{x} being ‘no meeting’) and we learn that agent 1 believes that the meeting will take place, we can infer that subject 1 has left the house. But since there is no meeting, we also infer that subject 2 stays at home and that agent 2 holds the corresponding opinion. We have learned something about 2’s opinion from 1’s opinion, a violation of PI. A plausible notion of opinion independence in cases like 6a and 6b must not conditionalise on the state. Therefore, only Common-Cause-Conditional Independence and hence—somewhat surprisingly, as there are no common causes—Unconditional Independence hold.

Although Theorem 1 seems to deal only with two of our independence conditions (CI and PI), an immediate corollary of part (a) gives us causal conditions for our other two independence conditions:

Corollary. Suppose probabilities are compatible with the causal network, and no opinion is a cause of any other opinion. Then:

- State-Conditional Independence holds if only the state is a common cause;
- Unconditional Independence holds if there are no common causes at all.

Notice how strong the causal conditions for SI and UI are. Among the above figures, only Figure 1 satisfies the condition for SI that the state is the only common cause, and only Figures 6a and 6b satisfy the ‘no common cause’ condition for UI. It might surprise that there exist plausible causal interconnections for which UI holds. Figure 6b is such a plausible network, as discussed above.

4 Conclusion

The widespread conceptual confusion about the notion of independence has hindered progress in assessing whether and when ‘crowds are wise’. It largely explains the disagreements in social epistemology about whether jury theorems have any relevance. Only by understanding causal interrelations is a proper analysis of opinion independence possible.

This paper distinguishes between four distinct notions of independence, as summarized in Table 1. This table highlights our two dimensions of categorization. To

Table 1: Different notions of independence.

		explicit state-conditionalization?	
		yes	no
explicit common-	yes	PI	CI
cause-conditionalization?	no	SI	UI

make a notion of independence realistic, the conditionalisation has to include the common causes; to make it suitable for jury theorems (or ‘wisdom of crowds’ arguments), the conditionalisation has to include the state of the world. State-Conditional Independence is the commonly used notion in orthodox statements of Condorcet’s jury theorem. Common-Cause-Conditional Independence is new and most generally defensible from the perspective of the theory of causal networks: it applies *always* as long as the opinions do not causally affect each other. Problem-Conditional Independence was recently introduced in Dietrich and Spiekermann (2010). If Unconditional Independence does not appear in the literature, it is probably because most scholars take it to be obviously false; but surprisingly we find plausible causal setups in which this simplest of all independence assumptions is justified. Theorem 1 and its corollary give formal causal-network-theoretic foundations for each of the four independence assumptions. These results suggest that the causal conditions for the classical State-Conditional Independence assumption are quite special and of limited real-world significance. One will usually have to go beyond classical independence to make sound arguments in support of the ‘wisdom of crowds’.

As already indicated, those independence assumptions which conditionalise on the state of the world \mathbf{x} (and perhaps on more)—State-Conditional and Problem-Conditional Independence—lend themselves to formal arguments for the ‘wisdom of crowds’, i.e., jury theorems. Such arguments require a competence assumption compatible with its independence assumption. State-Conditional Independence is the basis of Condorcet’s classical jury theorem, explained above. As for Problem-Conditional Independence, Dietrich and Spiekermann (2010) prove a corresponding jury theorem, with a problem-conditional rather than state-conditional competence assumption. Much further work is needed to develop the causal approach. We leave it as a future challenge to develop jury theorems for the aggregation of non-binary opinions, such as judgment sets or degrees of belief.

References

- Anderson, E. (2006), 'The Epistemology of Democracy', *Episteme* 3(1-2), 8-22.
- Atkinson, D. & Romeijn, J. (forthcoming), 'A Condorcet jury theorem for unknown juror competence', *Politics, Philosophy, and Economics*.
- Berg, S. (1993), 'Condorcet's jury theorem, dependency among voters:', *Social Choice and Welfare* 10, 87-95.
- Boland, P. J. (1989), 'Majority systems and the Condorcet jury theorem', *Journal of the Royal Statistical Society. Series D (The Statistician)* 38(3), 181-189.
- Boland, P. J.; Proschan, F. & Tong, Y. (1989), 'Modelling dependence in simple and indirect majority systems', *Journal of Applied Probability* 26(1), 81-88.
- Bovens, L. & Rabinowicz, W. (2006), 'Democratic Answers to Complex Questions - an Epistemic Perspective', *Synthese* 150, 131-153.
- Conradt, L. & Roper, T. J. (2005), 'Consensus decision making in animals', *Trends in Ecology & Evolution* 20(8), 449 - 456.
- Dietrich, F. (2008), 'The premises of Condorcet's jury theorem are not simultaneously justified', *Episteme* 58(1), 56-73.
- Dietrich, F. & List, C. (2004), 'A model of jury decisions where all jurors have the same evidence', *Synthese* 142, 175-202.
- Dietrich, F. and Spiekermann, K. (2010), 'Epistemic Democracy with Defensible Premises', Unpublished manuscript.
- Estlund, D. M. (1994), 'Opinion leaders, independence, and Condorcet's Jury Theorem', *Theory and Decision* 36(2), 131-162.
- Estlund, D. M. (2008), *Democratic Authority: A Philosophical Framework*, Princeton University Press, Princeton.
- Estlund, D.; Jeremy, W.; Grofman, B. & Feld, S. L. (1989), 'Democratic Theory and the Public Interest: Condorcet and Rousseau Revisited', *American Political Science Review* 83(4), 1317-1340.
- Goldman, A. I. (2004), 'Group Knowledge versus Group Rationality: Two Approaches to Social Epistemology', *Episteme* 1, 11-22.
- Goldman, A. I. (1999), *Knowledge in a Social World*, Clarendon, Oxford.
- Grofman, B. & Feld, S. L. (1988), 'Rousseau's General Will: A Condorcetian Perspective', *American Political Science Review* 82(2), 567-576.
- Grofman, B.; Owen, G. & Feld, S. L. (1983), 'Thirteen Theorems in Search of the Truth', *Theory and Decision* 15, 261-278.
- Hawthorne, J. (unpublished) Voting in Search of the Public Good: the Probabilistic Logic of Majority Judgments, University of Oklahoma.
- Kaniovski, S. (2010), 'Aggregation of correlated votes and Condorcet's Jury Theorem', *Theory and Decision* 69(3), 453-468.
- Ladha, K. K. (1995), 'Information pooling through majority-rule voting: Condorcet's jury theorem with correlated votes', *Journal of Economic Behavior & Organization* 26(3), 353-372.
- Ladha, K. K. (1993), 'Condorcet's jury theorem in light of de Finetti's theorem', *Social Choice and Welfare* 10(1), 69-85.
- Ladha, K. K. (1992), 'The Condorcet Jury Theorem, Free Speech, and Correlated

- Votes', *American Journal of Political Science* 36(3), 617–634.
- List, C. & Goodin, R. E. (2001), 'Epistemic Democracy: Generalizing the Condorcet Jury Theorem', *Journal of Political Philosophy* 9(3), 277–306.
- Nitzan, S. & Paroush, J. (1984), 'The significance of independent decisions in uncertain dichotomous choice situations', *Theory and Decision* 17(1), 47–60.
- Pearl, J. (2000), *Causality: models, reasoning and inference*, Cambridge University Press, Cambridge.
- Reichenbach, H. (1956), *The direction of time*, University of California Press, Berkeley.
- Spiekermann, K. & Goodin, R. E. (2010), 'Courts of Many Minds', manuscript under review.
- Vermeule, A. (2009), *Law and the Limits of Reason*, Oxford University Press, Oxford.

A Appendix: Proof of Theorem 1

Assume that probabilities are compatible with causal interconnections (in the sense of the Parental Markov Condition) and no opinion is a cause of another opinion. We first prove part (a) and then part (b). The informal idea in both proofs is that dependence between opinions can only arise if information can travel along a path in the network without the path being 'blocked' by the variables on which one conditionalises. The formal definition of 'blocking' (or 'd-separating') and the theorem whereby such blocking implies conditional independence are borrowed from the theory of causal networks, where they play a central role. Throughout we write \mathcal{C} for the set of common causes.

Proof of part (a). We have to show that the opinions are independent conditional on \mathcal{C} . By the 'blocking theorem' in the theory of causal networks (e.g., Pearl, 2000, Theorem 1.2.4) it suffices to show that \mathcal{C} blocks every path from an opinion to another opinion, in the usual technical sense that for any such path

- (i) *either* the path contains a chain ' $\mathbf{a} \rightarrow \mathbf{b} \rightarrow \mathbf{c}$ ' or fork ' $\mathbf{a} \leftarrow \mathbf{b} \rightarrow \mathbf{c}$ ' such that \mathbf{b} is in \mathcal{C}
- (ii) *or* the path contains a collider ' $\mathbf{a} \rightarrow \mathbf{b} \leftarrow \mathbf{c}$ ' such that \mathbf{b} is not in \mathcal{C} and does not cause any variable in \mathcal{C} .

To show this, consider any path from an opinion \mathbf{o}_i to another opinion \mathbf{o}_j . Call the path $(\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_m)$, where m (≥ 2) is the number of variables in the path and $\mathbf{a}_1 = \mathbf{o}_i$ and $\mathbf{a}_m = \mathbf{o}_j$. By definition of a path, any two neighbours $\mathbf{a}_t, \mathbf{a}_{t+1}$ are connected by an arrow, of the form ' $\mathbf{a}_t \rightarrow \mathbf{a}_{t+1}$ ' or ' $\mathbf{a}_t \leftarrow \mathbf{a}_{t+1}$ '.

Case 1: the arrow between \mathbf{a}_1 and \mathbf{a}_2 points towards \mathbf{a}_2 (' $\mathbf{a}_1 \rightarrow \mathbf{a}_2$ '). It is impossible that between all neighbouring variables $\mathbf{a}_t, \mathbf{a}_{t+1}$ the arrow points towards \mathbf{a}_{t+1} , since otherwise \mathbf{o}_i would be a cause of \mathbf{o}_j . Let \mathbf{a}_t be the earliest variable in the path such that an arrow points from \mathbf{a}_{t+1} to \mathbf{a}_t . Notice the collider ' $\mathbf{a}_{t-1} \rightarrow \mathbf{a}_t \leftarrow \mathbf{a}_{t+1}$ '. It is impossible that \mathbf{a}_t is in \mathcal{C} or causes a variable in \mathcal{C} , since otherwise \mathbf{o}_i (which causes \mathbf{a}_t) would cause other opinions. Therefore \mathcal{C} blocks the path via clause (ii).

Case 2: the arrow between \mathbf{a}_1 and \mathbf{a}_2 points towards \mathbf{a}_1 ($\mathbf{a}_1 \leftarrow \mathbf{a}_2$). It is impossible that between all neighbouring variables $\mathbf{a}_t, \mathbf{a}_{t+1}$ the arrow points towards \mathbf{a}_t , since otherwise \mathbf{o}_j would be a cause of \mathbf{o}_i . Let \mathbf{a}_t be the earliest variable in the path such that an arrow points from \mathbf{a}_t to \mathbf{a}_{t+1} . Notice the fork $\mathbf{a}_{t-1} \leftarrow \mathbf{a}_t \rightarrow \mathbf{a}_{t+1}$.

Subcase 2.1: $\mathbf{a}_t \in \mathcal{C}$. Then \mathcal{C} blocks the path via clause (i).

Subcase 2.2: $\mathbf{a}_t \notin \mathcal{C}$. Then \mathbf{a}_t (which already causes \mathbf{o}_i) cannot also cause \mathbf{o}_j . So we do not have a chain $\mathbf{a}_t \rightarrow \mathbf{a}_{t+1} \rightarrow \dots \rightarrow \mathbf{a}_m$. Choose \mathbf{a}_s as the earliest variable among $\mathbf{a}_t, \mathbf{a}_{t+1}, \dots, \mathbf{a}_{m-1}$ such that the arrow between \mathbf{a}_s and \mathbf{a}_{s+1} points towards \mathbf{a}_s . Note the collider $\mathbf{a}_{s-1} \rightarrow \mathbf{a}_s \leftarrow \mathbf{a}_{s+1}$. The variable \mathbf{a}_s neither belongs to \mathcal{C} nor causes a member of \mathcal{C} , since otherwise the variable \mathbf{a}_t (which causes \mathbf{a}_s) would belong to \mathcal{C} . Therefore \mathcal{C} blocks the path via clause (ii).

Proof of part (b). Now suppose that \mathbf{x} is not a common effect of any opinions or private causes thereof. We have to show that the opinions are independent conditional on $\{\mathbf{x}\} \cup \mathcal{C}$. Again by the ‘blocking theorem’ (e.g., Pearl, 2000, Theorem 1.2.4), it suffices to show that $\{\mathbf{x}\} \cup \mathcal{C}$ blocks every path from an opinion to another opinion, i.e., that for every such path

- (i*) either the path contains a chain $\mathbf{a} \rightarrow \mathbf{b} \rightarrow \mathbf{c}$ or fork $\mathbf{a} \leftarrow \mathbf{b} \rightarrow \mathbf{c}$ such that \mathbf{b} is in $\{\mathbf{x}\} \cup \mathcal{C}$
- (ii*) or the path contains a collider $\mathbf{a} \rightarrow \mathbf{b} \leftarrow \mathbf{c}$ such that \mathbf{b} is not in $\{\mathbf{x}\} \cup \mathcal{C}$ and does not cause any variable in $\{\mathbf{x}\} \cup \mathcal{C}$.

Consider any path $(\mathbf{a}_1, \dots, \mathbf{a}_m)$ from an opinion \mathbf{o}_i ($= \mathbf{a}_1$) to another opinion \mathbf{o}_j ($= \mathbf{a}_m$).

Case 1: the arrow between \mathbf{a}_1 and \mathbf{a}_2 points towards \mathbf{a}_2 ($\mathbf{a}_1 \rightarrow \mathbf{a}_2$). Construct a collider $\mathbf{a}_{t-1} \rightarrow \mathbf{a}_t \leftarrow \mathbf{a}_{t+1}$ as in Case 1 of part (a). Again, \mathbf{a}_t neither is in \mathcal{C} nor causes a variable in \mathcal{C} .

Subcase 1.1: \mathbf{a}_t neither is \mathbf{x} nor causes \mathbf{x} . Then, in summary, \mathbf{a}_t neither belongs to $\{\mathbf{x}\} \cup \mathcal{C}$ nor causes a variable in $\{\mathbf{x}\} \cup \mathcal{C}$. So, $\{\mathbf{x}\} \cup \mathcal{C}$ blocks the path via clause (ii*).

Subcase 1.2: \mathbf{a}_t is \mathbf{x} or causes \mathbf{x} . We cannot have arrows $\mathbf{a}_t \leftarrow \mathbf{a}_{t+1} \leftarrow \dots \leftarrow \mathbf{a}_m$, since otherwise \mathbf{a}_t , and hence \mathbf{x} , would be a common effect of the opinions \mathbf{a}_1 ($= \mathbf{o}_i$) and \mathbf{a}_m ($= \mathbf{o}_j$). Let \mathbf{a}_s be the earliest variable among $\mathbf{a}_t, \mathbf{a}_{t+1}, \dots, \mathbf{a}_{m-1}$ with an arrow $\mathbf{a}_s \rightarrow \mathbf{a}_{s+1}$. Note the fork $\mathbf{a}_{s-1} \leftarrow \mathbf{a}_s \rightarrow \mathbf{a}_{s+1}$ (see Figure A.1).

Subsubcase 1.2.1: $\mathbf{a}_s \in \mathcal{C}$. Then we are done as $\{\mathbf{x}\} \cup \mathcal{C}$ blocks the path via clause (i*).

Subsubcase 1.2.2: $\mathbf{a}_s \notin \mathcal{C}$. Then we do not have arrows $\mathbf{a}_s \rightarrow \mathbf{a}_{s+1} \rightarrow \dots \rightarrow \mathbf{a}_m$, since firstly \mathbf{a}_s does not commonly cause the opinion \mathbf{a}_m ($= \mathbf{o}_j$) and another opinion because $\mathbf{a}_s \notin \mathcal{C}$, and secondly \mathbf{a}_s does not *privately* cause the opinion \mathbf{a}_m because otherwise \mathbf{a}_t , and hence \mathbf{x} , would be a common effect of two variables (namely \mathbf{o}_i and \mathbf{a}_s) which are or privately cause distinct opinions. Given that we do not have arrows $\mathbf{a}_s \rightarrow \mathbf{a}_{s+1} \rightarrow \dots \rightarrow \mathbf{a}_m$, there must be a variable \mathbf{a}_r among $\mathbf{a}_{s+1}, \dots, \mathbf{a}_{m-1}$ with a




Figure A.1: The path in Subcase 1.2.

collider ' $\mathbf{a}_{r-1} \rightarrow \mathbf{a}_r \leftarrow \mathbf{a}_{r+1}$ '. Let \mathbf{a}_r be the *last* such variable among $\mathbf{a}_{s+1}, \dots, \mathbf{a}_{m-1}$. Note that we either have arrows ' $\mathbf{a}_r \leftarrow \dots \leftarrow \mathbf{a}_m$ ' or arrows ' $\mathbf{a}_r \leftarrow \dots \leftarrow \mathbf{a}_p \rightarrow \dots \rightarrow \mathbf{a}_m$ ' for some $r < p < m$ (see Figure A.2). We may assume without loss of generality that in the second case \mathbf{a}_p does not belong to \mathcal{C} , since otherwise $\{\mathbf{x}\} \cup \mathcal{C}$ would block the path via clause (i*). Now,

(*) \mathbf{a}_r does not belong to \mathcal{C} (hence, does not cause a member of \mathcal{C}),

In the case of ' $\mathbf{a}_r \leftarrow \dots \leftarrow \mathbf{a}_m$ ' this is because otherwise the opinion $\mathbf{a}_m (= \mathbf{o}_j)$ would cause another opinion, and in the case of ' $\mathbf{a}_r \leftarrow \dots \leftarrow \mathbf{a}_p \rightarrow \dots \rightarrow \mathbf{a}_m$ ' it is because otherwise \mathbf{a}_p would belong to \mathcal{C} . Notice further that \mathbf{a}_r is an effect either of the opinion $\mathbf{a}_m = \mathbf{o}_j$ (in the case of ' $\mathbf{a}_r \leftarrow \dots \leftarrow \mathbf{a}_m$ ') or of a private cause of this opinion (in case of ' $\mathbf{a}_r \leftarrow \dots \leftarrow \mathbf{a}_p \rightarrow \dots \rightarrow \mathbf{a}_m$ '). From this it follows that

(**) \mathbf{a}_r is not \mathbf{x} and does not cause \mathbf{x} ,

as otherwise \mathbf{x} would be a common effect of on the one hand (via \mathbf{a}_r) the opinion \mathbf{o}_j or a private cause thereof, and on the other hand (via \mathbf{a}_t) the opinion $\mathbf{o}_i (= \mathbf{a}_1)$. By (*) and (**), \mathbf{a}_r neither belongs to $\{\mathbf{x}\} \cup \mathcal{C}$ nor causes a member of $\{\mathbf{x}\} \cup \mathcal{C}$. So, the path is blocked via clause (ii*).

Case 2: the arrow between \mathbf{a}_1 and \mathbf{a}_2 points towards \mathbf{a}_1 (' $\mathbf{a}_1 \leftarrow \mathbf{a}_2$ '). Construct the fork ' $\mathbf{a}_{t-1} \leftarrow \mathbf{a}_t \rightarrow \mathbf{a}_{t+1}$ ' as in Case 2 of part (a).

Subcase 2.1: $\mathbf{a}_t \in \{\mathbf{x}\} \cup \mathcal{C}$. Then \mathcal{C} blocks the path via clause (i*).

Subcase 2.2: $\mathbf{a}_t \notin \{\mathbf{x}\} \cup \mathcal{C}$. Since in particular $\mathbf{a}_t \notin \mathcal{C}$, we can construct a collider ' $\mathbf{a}_{s-1} \rightarrow \mathbf{a}_s \leftarrow \mathbf{a}_{s+1}$ ' as in Subcase 2.2 of part (a) (see Figure A.3), and again \mathbf{a}_s neither belongs to \mathcal{C} nor causes a member of \mathcal{C} .

Subsubcase 2.2.1: \mathbf{a}_s neither is nor causes \mathbf{x} . Then, in summary, \mathbf{a}_s neither is in nor causes a variable of $\{\mathbf{x}\} \cup \mathcal{C}$, and hence $\{\mathbf{x}\} \cup \mathcal{C}$ blocks the path via clause (ii*).




Figure A.2: The path in Subsubcase 1.2.2.

Subsubcase 2.2.2: \mathbf{a}_s is or causes \mathbf{x} . We cannot have arrows ' $\mathbf{a}_s \leftarrow \mathbf{a}_{s+1} \leftarrow \dots \leftarrow \mathbf{a}_m$ ', since otherwise \mathbf{x} would be a common effect (via \mathbf{a}_s) of the opinion $\mathbf{o}_j (= \mathbf{a}_m)$ and the private cause \mathbf{a}_t of the opinion $\mathbf{o}_1 (= \mathbf{a}_1)$. If we have arrows ' $\mathbf{a}_s \leftarrow \dots \leftarrow \mathbf{a}_q \rightarrow \dots \rightarrow \mathbf{a}_m$ ' for some $s < q < m$, then \mathbf{a}_q must be in \mathcal{C} since otherwise \mathbf{a}_q would be a *private* cause of the opinion $\mathbf{a}_m (= \mathbf{o}_j)$ so that \mathbf{x} would be a common effect of the private causes \mathbf{a}_q and \mathbf{a}_t ; and since ' $\mathbf{a}_{q-1} \leftarrow \mathbf{a}_q \rightarrow \mathbf{a}_{q+1}$ ' is a fork with $\mathbf{a}_q \in \mathcal{C}$ we are done by clause (i*). Now assume the remaining case that we neither have arrows ' $\mathbf{a}_s \leftarrow \dots \leftarrow \mathbf{a}_m$ ' nor arrows ' $\mathbf{a}_s \leftarrow \dots \leftarrow \mathbf{a}_q \rightarrow \dots \rightarrow \mathbf{a}_m$ '. There must be a variable \mathbf{a}_r among $\mathbf{a}_{s+1}, \dots, \mathbf{a}_{m-1}$ such that we have a collider ' $\mathbf{a}_{r-1} \rightarrow \mathbf{a}_r \leftarrow \mathbf{a}_{r+1}$ '. Choose \mathbf{a}_r to be the *latest* variable among $\mathbf{a}_{s+1}, \dots, \mathbf{a}_{m-1}$ with the collider property. Note that we either have arrows ' $\mathbf{a}_r \leftarrow \dots \leftarrow \mathbf{a}_m$ ' or arrows ' $\mathbf{a}_r \leftarrow \dots \leftarrow \mathbf{a}_p \rightarrow \dots \rightarrow \mathbf{a}_m$ ' for some $r < p < m$ (see Figure A.4). We may assume without loss of generality that in the second case $\mathbf{a}_p \notin \mathcal{C}$, as otherwise we would be done by clause (i*). Then

(***) \mathbf{a}_r does not belong to \mathcal{C} (so, does not cause a member of \mathcal{C}).

In the case of arrows ' $\mathbf{a}_r \leftarrow \dots \leftarrow \mathbf{a}_m$ ' the reason for (***) is that otherwise the opinion $\mathbf{a}_m (= \mathbf{o}_j)$ would cause another opinion; and in the case of arrows ' $\mathbf{a}_r \leftarrow \dots \leftarrow \mathbf{a}_p \rightarrow \dots \rightarrow \mathbf{a}_m$ ' the reason is that otherwise \mathbf{a}_p would belong to \mathcal{C} . Note further that \mathbf{a}_r




Figure A.3: The path in Subcase 2.2.

is caused either by the opinion $\mathbf{a}_m = \mathbf{o}_j$ (in the case of arrows ' $\mathbf{a}_r \leftarrow \dots \leftarrow \mathbf{a}_m$ ') or by a private cause of this opinion (in case of arrows ' $\mathbf{a}_r \leftarrow \dots \leftarrow \mathbf{a}_p \rightarrow \dots \rightarrow \mathbf{a}_m$ '). It follows that

(****) \mathbf{a}_r is not \mathbf{x} and does not cause \mathbf{x} ,

since otherwise \mathbf{x} would be a common effect firstly (via \mathbf{a}_r) of the opinion \mathbf{o}_j or a private cause thereof, and secondly (via \mathbf{a}_s) of the private cause \mathbf{a}_t of the opinion \mathbf{o}_i ($= \mathbf{a}_1$). By (***) and (****), \mathbf{a}_r does not belong to or cause a member of $\{\mathbf{x}\} \cup \mathcal{C}$. So, the path is blocked via clause (ii*). ■




Figure A.4: The path in Subsubcase 2.2.2.