

Brute Ignorance

0 Abstract

We know a lot about what the world is like. We know less, it seems, about what we know about what the world is like. According to a common thought, it is easier for us to come to know about the state of the world than to come to know about the state of our own knowledge. What explains this gap? An attractively simple hypothesis is that our ignorance about what we know is explained by our ignorance about the world. There are things we fail to know about what we know about the world because there are things we fail to know about the world. This hypothesis is often motivated by the idea that knowledge requires a margin-for-error (Williamson (1992, 1994, 2000, 2011)). In this paper, I'll argue that this simple hypothesis is inadequate. Not all our ignorance of our knowledge can be explained by our ignorance about the world. In this sense, at least some of our ignorance about what we know is brute.

1 Crowds & Classes

Consider two cases.

Crowd

Aron is at a theatre. Based on a quick glance at the crowd, Aron estimates that it contains 100 people. However, Aron's eyesight isn't perfect—he may have missed some people. And his attention is limited—he may have double-counted others. Given these facts, it seems Aron's knowledge will be inexact: there is no n such that he can know that the crowd contains exactly n people. However, it also seems Aron is not wholly ignorant about the crowd's size. For some k and k' , he knows that it contains at least k people and, similarly, that it contains at most k' people. (cf. Williamson (1994, §8.2))

Class

Nora is taking a class. The class meets 100 times and its location varies on an irregular basis between two rooms (Room A and Room B). Nora has a syllabus which says, for each day of class, which room it will meet in on that day. The day before the class first meets, Nora receives an email saying that the room in which it will meet has had to be switched on exactly one of the 100 days—however, the email does not specify which day. Given that the syllabus contains an inaccuracy, it seems Nora's knowledge will be inexact: she can't know where the class will meet on each of the 100 days. However, it also seems Nora is not wholly ignorant about where the class will meet when. For at least some days, Nora knows where the class will meet that

day.

Both cases illustrate a way we can gain incomplete knowledge from fallible sources. In each case, an individual has inexact knowledge of the state of the world. And yet there are many things they can, nevertheless, know. An imperfect view of a crowd is still good enough for knowledge that it is not very big or very small. Equally, a syllabus with one error is still accurate enough for knowledge of where the class will meet on at least some of the days about which it is correct.¹

Both cases also, arguably, illustrate a failure of **Introspection**.

Introspection If you know p , then you know that you know p .

In **Crowd**, there is some greatest k and least k' such that Aron knows there are between k and k' people in the crowd. However, whatever that range is, Aron does not seem capable of knowing that he knows the crowd's size is within that range. For any j and j' such that Aron knows that he knows there are between j and j' people in the crowd, either $j < k$ or $k' < j'$ (or both). In **Class**, there are some days for which Nora knows where the class will meet that day. However, Nora does not seem capable of knowing which days these are. In fact, arguably, there are no days for which Nora can know that she knows where the class will meet that day.

This is not news. **Introspection** is widely taken to admit counter-examples. Nevertheless, even if we are convinced—whether via argument, via appeal to intuition, or both—that knowledge fails to iterate freely, we would still like to have an explanation of why it fails. §2 turns to this question.

2 Margins

There is a canonical explanation of why Aron fails to know what he knows in **Crowd** (see Williamson (1992, 1994, 2000, 2011, 2013, 2014); Weatherson (2013); Srinivasan (2013); for critical discussion, see Weatherson (2004); Dutant (2007); Mahtani (2008); Goodman (2013); Cohen & Comesaña (2013); Goldstein (2022), among others). According to this explanation, his knowledge of the crowd's size is subject to a margin-for-error. Whatever the size of the crowd, Aron can't know it isn't slightly smaller and he can't know it isn't slightly larger. Moreover, this seems like something that Aron could know. So, in particular, (the explanation goes) Aron knows that, for any n : if there are n people in the crowd, then for all he knows, there are $n-1$ people in the crowd and for all he knows, there are $n + 1$.

To generate failures of **Introspection**, we'll assume that Aron's knowledge satisfies **Closure**.

¹Those who have hesitation about this judgment should feel free to increase the number of times the class is scheduled to meet in **Class** as much as they want.

Closure If you know each of a set of propositions, then you know any consequences of those propositions.

Closure involves an idealization away from our deductive limitations. But this idealization appears harmless. If **Introspection** fails in the absence of those limitations, we should not expect it to hold in their presence.

Now, consider the greatest known lower bound on the crowd's size. This is the greatest n such that Aron knows that there are at least n people in the crowd. Since it is the greatest, for all Aron knows there are exactly n people in the crowd. But he also knows that, if there are exactly n people in the crowd, he doesn't know that there aren't $n - 1$ people in the crowd. So, it follows that for all he knows he doesn't know that there aren't $n - 1$ people in the crowd. But, by assumption, Aron does know there are at least n people in the crowd. So there is something which Aron knows that he doesn't know he knows.

Crucially, on this picture, Aron's ignorance about what he knows about the size of the crowd is parasitic on his ignorance about the size of the crowd. There is something he knows about the size of the crowd which he fails to know that he knows only because he fails to know what size the crowd is. To see this, observe that, by reflecting on his reliability, Aron can come to know something about the size of the margin. Let m be the smallest number such that he knows that the margin is no greater than m . That is, m is the least number such that for all k , Aron knows that: if the crowd contains exactly k people, then he knows it contains at least $k - m$ people. Suppose that there are in fact n people in the crowd. By factivity, Aron knows that the crowd contains at least $n - m$ people. Of course, Aron doesn't know that he knows this. For all he knows, the crowd contains $n - 1$ people (in which case he can't rule out that, for all he knows, it contains $(n - 1) - m$ people). However, he is only unable to know that he knows it contains $n - m$ people because he is unable to know the exact size of the crowd. It is in this sense that the margin-for-error argument explains higher-order ignorance in terms of first-order ignorance.

Margins-for-error are often taken to provide a broad explanation of **Introspection** failure (Williamson (1994, §8.3)). In the next section, I'll argue that margin-for-error reasoning fails to generalize in the appropriate way. More broadly, I'll argue that there are examples of higher-order ignorance which is not parasitic on first-order ignorance at all. In cases which exhibit certain epistemic symmetries (such as **Class**), our ignorance about what we know cannot be explained by our ignorance about the world.

3 Symmetry

In both cases above, we are interested in an agent's knowledge about a specific subject matter: what size the crowd is (in **Crowd**) and where the class will meet when (in **Class**). Corresponding to these subject matters, we can distinguish various scenarios. Each scenario can be thought of as a different way the world

could be with respect to that matter. In **Crowd**, the relevant scenarios correspond to different sizes of the crowd; in **Class**, to different pairings of rooms and days. Two worlds belong to the same scenario in a subject matter iff they agree on how the world is with respect to that subject matter. A proposition is about a subject matter iff each scenario in that subject matter entails either the proposition or its negation.

In **Crowd**, Aron's evidence favors some sizes the crowd could be over others. As a result, scenarios in the subject matter of the size of the crowd can be ordered into a series of epistemically better and worse cases. How much Aron can know about the size of the crowd decreases as its size departs further from what his evidence favors. If Aron knows he is subject to a margin-for-error, then he cannot rule out being in a scenario which is a slightly worse epistemic case than the scenario which in fact obtains. So his inability to know what size the crowd is can explain his inability to know what he knows about what size the crowd is.

Crucially, however, **Class** does not have this structure. Consider Nora's evidence about which room the class will meet in on each day. This evidence comprises what her syllabus says about where it will meet and that it contains exactly one error. Accordingly, any scenario in the relevant subject matter compatible with Nora's evidence will differ from what the syllabus says about which room the class will meet in when on exactly one day. Nora's evidence doesn't favor any one of these scenarios over any other. Given this, it seems hard to resist accepting certain symmetries in what she can know.

In particular, it seems that differences regarding the day on which the room has been switched shouldn't make a difference to Nora's ability to know a proposition about where the class will meet when, unless they make a difference to the truth of that proposition. If a proposition about where the class will meet when is entailed by each of a pair of scenarios, then knowing that proposition should be compatible with being in the first scenario iff it is compatible with being in the second.

For example, consider two scenarios each of which differs from Nora's syllabus regarding where the class will meet when on exactly one day. In the first, the room has been switched on 1st January. In the second, it has been switched on 2nd February. The two scenarios might differ with respect to what Nora can know (given her evidence). For instance, perhaps it is compatible with being in the first that she knows where the class will meet on 2nd February, even though knowing this isn't compatible with being in the second. However, they shouldn't differ with respect to whether Nora can know where the class will meet on 3rd March. If knowing this is compatible with being in the first, then it is compatible with being in the second.

Importantly, these symmetries can be expected to ramify into Nora's higher-order knowledge. Just as there are symmetries in what truths Nora can know across scenarios, there will likewise be symmetries in what truths she can know

that she knows across scenarios.

For any pair of scenarios compatible with what Nora knows, let p be a proposition about where the class will meet when which is entailed by both scenarios. Since p is a proposition about where the class will meet when, knowing p amounts to ruling out certain scenarios. Given that her evidence equally favors each scenario which differs from the syllabus on exactly one day, we should accept the following biconditionals:

- (+) Nora knows that [if she is in the first scenario, she knows p] iff she knows that [if she is in the second, she knows p].
- (−) Nora knows that [if she is in the first scenario, she doesn't know p] iff she knows that [if she is in the second, she doesn't know p].

To see why these symmetries are plausible, consider what would be required for (+) to fail. Specifically, there would need to exist three days (say: 1st January, 2nd February, and 3rd March) meeting the following conditions:

- i. For all Nora knows, the switch is on 1st January and for all she knows, the switch is on 2nd February.
- ii. Nora knows that if the switch is on 1st January, then she knows where the class will meet on 3rd March.
- iii. Nora doesn't know whether, if the switch is on 2nd February, then she knows where the class will meet on 3rd March.

In effect, this would amount to Nora knowing of certain scenarios and not of others that being in those scenarios is conducive to knowledge of some propositions which are true across both. By hypothesis, the only evidence Nora has about where the class will meet when is the syllabus and the email. Given that, it is unclear what could put Nora in a position to know this about some scenarios but not others.² The same considerations, *mutatis mutandis*, support (−).³

²An obvious asymmetry between the scenarios compatible with what Nora knows is that exactly one of them is actual (cf. Bacon (2014)). Since Nora doesn't know which scenario that is, though, it is unclear how it could make a difference to her ability to know which truths she can know in each scenario (even if it did make a difference to which truths she can know in each scenario).

³The point is not that there is no way of filling in details of the case which would motivate failure of these symmetries (given a particular background theory). Suppose the switch is necessary on day n because of planned construction work on that day in one of the rooms. Suppose Nora knows that if the work is planned for day n , then if it had been planned for a day other than n , it would have been planned for day k rather than day k' . Then safety-based theories could predict Nora knows that it is compatible with the switch being on day n that Nora knows where the class will meet on day k' but not that she knows where it will meet on day k . Or suppose Nora knows that the stakes of going to the right room are higher on day k than day k' (there is an exam on day k). Then pragmatic-encroachment theories could make

It is important to distinguish the observations above from the claim that there is no pair of worlds which differ with respect to the day of the switch (but are otherwise alike) such that, in exactly one, Nora is able to rule out being in the other. Below, I'll take seriously a view which has this consequence. In particular, on the view I'll defend, at each world, being in some scenarios would necessarily be more conducive to knowledge about the subject matter than being in others.⁴ However, even if this is the case, it seems clear that Nora shouldn't be in a position to know which those scenarios are.

3.1 Independence

To theorize about symmetries of this kind, we need to restrict our attention to what an agent knows about a particular subject matter. A subject matter, S , can be understood as a partition on over worlds: a set of a non-empty, disjoint subsets of the worlds in some set (Lewis (1988b,a); Yablo (2014); Plebani & Spolaore (2021)). Scenarios correspond to elements in this partition. Each $s \in S$ is a fully specific way things could be regarding S . For any world w , let s_w be the scenario in S to which w belongs. A proposition, p , is about a subject matter, S , iff p is the union of one or more scenarios in S . This provides us with a way of characterizing what an agent knows about what the world is like regarding some specific subject matter.

Given this setup, we can restate the higher-order symmetry principles about knowledge of a specific subject matter in more general terms. Consider positive symmetry, first:

Symmetry⁺ For any $s, s' \in S$ compatible with what you know at w and any p about S which is true throughout s and s' : You know in w that [either you aren't in s or you know p] iff you know in w that [either you aren't in s' or you know p].

Symmetry⁺ implies that if you can't rule out failing to know something about S , then you can't rule out failing to know it while being in the scenario you are actually in. Call this **Independence**.

Independence For any p about S which you know at w : If for all you know at w [you are in s and don't know p], then for all you know at w [you are in s_w and don't know p].

To see why **Independence** follows from **Symmetry**, consider an arbitrary world $v \in s$ which is compatible with what you know at w . Since $v \in s$ is compatible with what you know at w , any proposition about S which you know at w is true throughout s . Since that proposition is something you know at w , it

the same prediction.

The point is that our judgments that **Introspection** fails do not depend on filling in the details of the case in this way. So these theories cannot give us a general explanation of Nora's higher-order ignorance.

⁴At least on a natural story about how indeterminacy interacts with modality (cf. Bacon (2018)).

is also true throughout \mathbf{s}_w (by factivity). So suppose that you don't know it at v . Then for all you know w , you are in \mathbf{s} and don't know p . So, by **Symmetry**⁺, it follows that for all you know in w , you are in \mathbf{s}_w and don't know it.

Independence implies that your ignorance about your knowledge about a subject matter is not parasitic on your first-order ignorance about that subject matter. If **Introspection** fails for what you know about \mathbf{S} , then there must be a world, v , compatible with what you know such that there something you in fact know about \mathbf{S} which you fail to know at v . **Independence** implies that if there is a world of this kind, then there is a world of this kind which belongs to the actual scenario. So there is nothing you know about the subject matter which you fail to know that you know only because you fail to know exactly which scenario you are in.

In this sense, positive symmetries imply that Nora's second-order ignorance is brute. Her failure to know what she knows about where the class will meet when cannot be attributed to her failure to know where the class will meet when.

3.2 Marginal Tolerance

In fact, the considerations above suggest something stronger: in cases exhibiting symmetries, failures of **Introspection** cannot be explained by a margin-for-error principle at all. Consider negative symmetry:

Symmetry[−] For any $\mathbf{s}, \mathbf{s}' \in \mathbf{S}$ compatible with what you know at w and any p about \mathbf{S} which is true throughout \mathbf{s} and \mathbf{s}' : You know at w that [either you aren't in \mathbf{s} or you don't know p] iff you know at w that [either you aren't in \mathbf{s}' or you don't know p].

Let \sim be a reflexive, symmetric relation over scenarios. Intuitively, $\mathbf{s} \sim \mathbf{s}'$ iff \mathbf{s} and \mathbf{s}' differ marginally. According to the margin-for-error principle, if two scenarios differ only marginally, then you don't know in one that you aren't in the other.

Margins For any $\mathbf{s}, \mathbf{s}' \in \mathbf{S}$ and any $w \in \mathbf{s}$: if $\mathbf{s} \sim \mathbf{s}'$, then for all you know at w , you are in \mathbf{s}' .

Together, however, **Margins** and **Symmetry**[−] imply **Marginal Tolerance**.⁵

Marginal Tolerance If for all you know at w , you are in \mathbf{s} and $\mathbf{s} \sim \mathbf{s}'$, then for all you know at w you are in \mathbf{s}' .

Marginal Tolerance says that if you can't rule out being in a scenario, then you can't rule out being in any scenario marginally different from that scenario. To see why it follows from **Symmetry**[−] and **Margins**, suppose $\mathbf{s} \sim \mathbf{s}'$ and for all you know at w you are in \mathbf{s} . Let p be the proposition that \mathbf{s}' does not obtain (i.e., $p = \{u | u \notin \mathbf{s}'\}$). By **Margins**, for each $v \in \mathbf{s}$, for all you know at v , you are in \mathbf{s}' . So there is no $v \in \mathbf{s}$ at which you know p . So, you know at w that either you aren't in \mathbf{s} or you don't know p . But observe that p is about \mathbf{S}

⁵Assuming that **Margins** is known and that it is known which scenarios differ marginally.

(since it doesn't cross-cut any scenarios in \mathcal{S}). And it is compatible with what you know at w that you are in \mathbf{s}_w (by factivity). So, by **Symmetry**⁻, it follows that you know at w that you either you aren't in \mathbf{s}_w or you don't know p . So, (by factivity again) it follows that you don't know p at w . So for all you know at w , you are in \mathbf{s}' .

Marginal Tolerance has drastic implications. Suppose there is a sequence $\mathbf{s}_1, \dots, \mathbf{s}_n$ such that for all $i < n : \mathbf{s}_i \sim \mathbf{s}_{i+1}$. An immediate consequence of the principle is that if \mathbf{s}_1 is compatible with what you know at w , then \mathbf{s}_n is, too.⁶ That is, if you can't rule out being in a scenario, then you can't rule out being in any scenario connected to it by a sequence of marginally differing scenarios.

More generally, **Marginal Tolerance** implies that any failure of knowledge about the subject matter to iterate cannot be explained by **Margins**. If **Introspection** fails for knowledge about \mathcal{S} , there must be a pair of scenarios, $\mathbf{s}, \mathbf{s}' \in \mathcal{S}$, such that you can't rule out being in \mathbf{s} , you can rule out being in \mathbf{s}' , but you can't rule out being in \mathbf{s} and being unable to rule out being in \mathbf{s}' . **Marginal Tolerance** implies that any such counter-example must involve a pair of scenarios which are *not* linked by a chain of marginally different scenarios. As a result, the constraints on knowledge imposed **Margins** can't be what explains the failure of **Introspection** in cases in which **Symmetry**⁻ holds.

4 Indeterminacy

To hold that higher-order ignorance is brute is not to hold that higher-order ignorance lacks explanation. It is merely to hold that it is not explained by first-order ignorance. What, then, can explain it instead? In this section I'll argue that, in many cases, our ignorance about our own epistemic state can be explained by its indeterminacy.

As observed above, there are many things Nora knows. She is not entirely ignorant about where the class will meet when. More specifically, there are some days for which she knows which room the class will meet in on that day. She also knows, for each day, what the syllabus says about where the class will meet on that day and that, on some day, the class will not meet where the syllabus says it will.

However, she doesn't know everything. In particular, she doesn't know on which day the rooms were switched. Assuming that Nora's knowledge satisfies **Closure**, it follows that there is at least one day on which the rooms weren't switched such that she doesn't know where the class will meet on that day. That's because which day the rooms were switched on follows from the true propositions about where the class will meet on each other day together with her other background knowledge.⁷

⁶The base case (where $n = 2$) is an instance of **Marginal Tolerance**. The proof proceeds by induction on n .

⁷There are revisionary accounts of **Class** on which **Closure** fails. Most notably, someone

Of which days does Nora know where the class will meet that day? Given the symmetries in her evidence, there is nothing to favor any particular answer. For each day on which the rooms weren't switched, her evidence regarding where the class will meet on that day is the same. It is hard, then, to see what could make it the case that she knows where the class will meet on some of those days but fails to know where it will meet on others, when the syllabus is equally accurate about both.

We face a puzzle. Any choice appears arbitrary. And yet, given what we have said, some choice is compulsory. This arbitrariness is typical of indeterminacy. In the absence of anything to settle which of an exhaustive set of alternatives obtain, it seems that all we can say is that while it is determinate that some alternative obtains, there is no alternative that determinately obtains.

The idea that indeterminacy accompanies arbitrariness is most familiar from the sorites. Similar arguments have been made for indeterminacy of counterfactuals (Stalnaker (1980)); of reference (Breckenridge & Magidor (2012); Barnes (2014)); of causation (Swanson (2016); Bernstein (2016)); and of moral obligation (Dougherty (2013); Barnes (2014)). In each case, indeterminacy is taken to be motivated by the absence of anything which could settle the relevant facts in a non-arbitrary fashion.

In the present case, for each non-switch day, there is no non-arbitrary way of settling whether it is one of the days of which Nora fails to know where the class will meet that day. Accordingly, there is no non-switch day of which it is determinate Nora knows where the class will meet that day and no non-switch day of which it is determinate she fails to know where the class will meet that day. The only day for which it is not indeterminate whether Nora knows where the class will meet that day is the day of the switch (since she determinately fails to know where it meets on that day).

Indeterminacy precludes knowledge: If it is indeterminate whether p , then you don't know p and you don't know $\neg p$ (Greenough (2003); Williams (2008); though see Dorr (2003) for dissent). Moreover, this ignorance is, in an important way, ineliminable. A proposition is about imprecise matters iff it remains indeterminate regardless of how things are with respect to precise matters. Acquiring knowledge about precise matters cannot help to resolve ignorance about imprecise matters.

Accordingly, the indeterminacy of what we know can account for our inability to know what we know. §5 develops a model of the interaction between ignorance and indeterminacy. This model draws on the approach to theorizing about precision and imprecision in Bacon (2018), supplementing it to theorize

could hold that, while Nora cannot know where the class will meet on any particular day, she can know any disjunction of two or more distinct true claims on the syllabus. This view gives up **Closure**, since taken together, these disjunctions imply where the class will meet on each day on which the switch didn't occur. I will not argue for **Closure** here; my restricted aim is to explore the options for those who want to retain **Closure** in cases like **Class**.

about what can be known in cases with different kinds of structure. Within the model, our ignorance about what we know remains brute—it is not explained by ignorance about the world. But it is not inexplicable. Instead, it is explained by the indeterminacy of what we know about the world.

5 A Model of Indeterminacy and Knowledge

A world is a way of fully settling how things are. To play its role, a world must settle both how things are with respect to precise matters and how things are with respect to imprecise matters.

Let W be a domain of worlds. We can represent the distinction between precise and imprecise matters with a function, $|\cdot|$, which maps each world to a partition over W (Bacon (2018)). $|w|$ is the set of maximally specific precise ways things could be (relative to what counts as precise at w). If $p \in |w|$ and $v, u \in p$, then v and u agree on all matters which are precise at w . A proposition $p \in \mathcal{P}(W)$ is precise at w iff it is the union of one or more elements of $|w|$. Precise propositions settle precise matters only. $|\cdot|$ is not assumed to be constant function, since whether a difference between two worlds is a precise or imprecise matter may itself be an imprecise matter.⁸

What is determinate at a world is settled by what is precise at that world. It is determinate that p at w iff there is some $q \in |w|$ such that $w \in q \subseteq p$. Put another way, the determinate truths at w are the truths which follow from how things are precisely at w (relative to w). It is indeterminate whether p at w iff neither p nor \bar{p} is determinate at w .

What is known at a world is represented by a relation, R , over worlds. It is known that p at w iff $R(w) \subseteq p$. We assume that what is known is constrained by what is precise.

Ignorance $\forall p \in |w| : \text{either } p \subseteq R(w) \text{ or } p \subseteq \overline{R(w)}$.

Ignorance says that worlds which differ only with respect to imprecise matters are epistemically indiscriminable. If v and u agree on all matters which are precise at w , then v is compatible with what is known at w iff u is compatible with what is known at w . As a result, at any world the strongest known proposition is precise at that world. An immediate consequence is that if it is indeterminate

⁸The framework is intended to be neutral on how imprecision is understood. There is an interpretation in terms of ontic vagueness on which each $w \in W$ corresponds to maximally specific metaphysically possible way for reality to be (though cf. Bacon (2018, §14)). However, this interpretation is not obligatory. On an epistemicist interpretation, each $w \in W$ will correspond to an epistemically possible way for reality to be (but there will be at most one world in each $p \in |w|$ which corresponds to a metaphysically possible way for things to be). On a semantic interpretation, each $w \in W$ will correspond to a metaphysically possible way for reality to be along with a specification of the semantic facts about some vague language. Readers who prefer to reserve the term ‘world’ for metaphysically possible total specifications of reality should feel free substitute its use here for a term they find more neutral (such as ‘index’ or ‘point’).

whether p at w (i.e. neither p nor \bar{p} is determinate at w), then it is unknown that p at w .

The forgoing provides a general framework for theorizing about the interaction of knowledge and indeterminacy. To represent specific cases we need to consider particular subject matters about the precise and imprecise.

5.1 A Model of Class

In **Class**, there is some set of precise ways the world could be such that: (i) for each of member of the set, Nora has strong evidence that things are not that way; but (ii) Nora has overwhelming evidence that the world is one of the ways in that set. Cases with this structure can be represented by imposing additional constraints on models.

To represent the relevant ways the world could be, we introduce a subject matter $\mathbf{S} = \{\mathbf{s}_1, \dots, \mathbf{s}_n\}$. \mathbf{S} is a partition over a subset of W comprising n cells. We assume that each $\mathbf{s}_i \in \mathbf{S}$ is precise at each world. That is, for all $w \in W$ and $\mathbf{s}_i \in \mathbf{S}$: \mathbf{s}_i is the union of one or more elements of $|w|$. In the case of **Class**, each $\mathbf{s}_i \in \mathbf{S}$ corresponds to the set of worlds at which the switch occurred on the i th day.

To represent the indeterminacy of what is known, we introduce a second subject matter, $\mathbf{I} = \{\mathbf{i}_{j_1, \dots, j_d} | 1 \leq j_1 < \dots < j_d \leq n\}$. Here, $d < n$ is a constant representing the maximum achievable level of exactness of what is known about \mathbf{S} . The idea is that d corresponds to the minimum number of elements of \mathbf{S} which cannot be ruled out regardless of how things are with respect to precise matters.⁹ We assume that, for each $\mathbf{i}_{j_1, \dots, j_d} \in \mathbf{I}$, there is no precise way the world could be which would settle whether $\mathbf{i}_{j_1, \dots, j_d}$. For all $w \in W$ and all $p \in |w|$: for all $\mathbf{i}_{j_1, \dots, j_d} \in \mathbf{I}$: $\mathbf{i}_{j_1, \dots, j_d} \cap p \neq \emptyset$.

Intuitively, the idea is that each element of \mathbf{I} corresponds to a way of settling which elements of \mathbf{S} cannot be ruled out at any world. We impose this interpretation via the following principle.

Plenitude For any $1 \leq j_1 < \dots < j_d \leq n$ and any $v \in \mathbf{i}_{j_1, \dots, j_d} \in \mathbf{I}$: if $x \leq d$: $R(v) \cap \mathbf{s}_{j_x} \neq \emptyset$.

Plenitude says that each $\mathbf{i}_{j_1, \dots, j_d} \in \mathbf{I}$ comprises worlds at which $\mathbf{s}_{j_1}, \dots, \mathbf{s}_{j_d}$ cannot be ruled out. Since, for each $\mathbf{i}_{j_1, \dots, j_d} \in \mathbf{I}$ and $w \in W$, it is indeterminate whether $\mathbf{i}_{j_1, \dots, j_d}$ at w , it follows that it is indeterminate which elements of \mathbf{S} cannot be ruled out. Moreover, since there is no $\mathbf{s}_i \in \mathbf{S}$ and $w \in W$ such that it is determinate at w that \mathbf{s}_i can be ruled out, there is no (non-trivial) proposition about \mathbf{S} which is known to be known at any world.

⁹Since it is plausible that it is indeterminate what level of exactness is achievable, this involves an idealization away from higher-order indeterminacy. A more realistic model would treat the value of d as a function of worlds which can differ across worlds which agree on all precise matters. This complication is suppressed for the purposes of exposition, since it makes no difference to the results below.




Figure 1

We establish the adequacy of the model via the following result.¹⁰

Fact 1. Let R be the weakest¹¹ relation satisfying **Ignorance** and **Plenitude**. Then it follows that:

- (i) **Symmetry**⁺ and **Symmetry**⁻ hold.¹²
- (ii) **Introspection** fails.¹³

For a concrete example, consider model in **Figure 1** satisfying **Ignorance** and **Plenitude**. In this model, $n = 3$ and $d = 2$. Each colored point corresponds to a world. Worlds which belong to the same element of S are located in the same dashed region. Worlds which belong to the same element of I are shaded the same color. $|\cdot|$ maps each world to S . R -accessibility is represented by the

¹⁰**Symmetry**⁺ holds iff for all $w \in W$ and $s, s' \subseteq X \subseteq S$: $\forall v \in R(w) \cap s : R(v) \subseteq \bigcup X$ iff $\forall u \in R(w) \cap s' : R(u) \subseteq \bigcup X$. **Symmetry**⁻ holds iff for all $w \in W$ and $s, s' \subseteq X \subseteq S$: $\forall v \in R(w) \cap s : R(v) \not\subseteq \bigcup X$ iff $\forall u \in R(w) \cap s' : R(u) \not\subseteq \bigcup X$. **Introspection** holds iff R transitive.

¹¹I.e., for all $w \in W$, $R(w)$ is the smallest set compatible with the constraints.

¹²*Proof:* consider any $s_k, s_{k'} \in S$ such that $R(w) \cap s_k \neq \emptyset$ and $R(w) \cap s_{k'} \neq \emptyset$. Consider any $X \subseteq S$ such that $s_k \cup s_{k'} \subseteq \bigcup X$. Next, take an arbitrary $v \in R(w) \cap s_k$. Consider the element $i_{j_1, \dots, j_d} \in I$ such that $v \in i_{j_1, \dots, j_d}$. i_{j_1, \dots, j_d} has a non-empty intersection with each $p \in |w|$. By **Ignorance**, it follows that $R(w) \cap s_{k'} \cap i_{j_1, \dots, j_d} \neq \emptyset$. By **Plenitude**, for any $u \in i_{j_1, \dots, j_d} \cap s_k$ and any $s_x \notin X$: $R(u) \cap s_x \neq \emptyset$ iff $R(u) \cap s_{k'} \neq \emptyset$. But v was arbitrary. So, for **Symmetry**⁺, it follows that for any $v \in R(w) \cap s_k$: if $R(v) \not\subseteq \bigcup X$, then there is some $u \in R(w) \cap s_{k'}$ such that $R(u) \not\subseteq \bigcup X$. Conversely, for **Symmetry**⁻, it follows that for any $v \in R(w) \cap s_k$: if $R(v) \subseteq \bigcup X$, then there is some $u \in R(w) \cap s_{k'}$ such that $R(u) \subseteq \bigcup X$.

¹³*Proof:* consider an arbitrary $s_k \in S$. Observe that since $d < |S|$, there is some $i_{j_1, \dots, j_d} \in I$ and some $w \in i_{j_1, \dots, j_d}$ such that $R(w) \cap s_k = \emptyset$. By **Ignorance**, for all $i_{j'_1, \dots, j'_d} \in I$: $R(w) \cap i_{j'_1, \dots, j'_d} \neq \emptyset$. But, by **Plenitude**, there is some $i_{j'_1, \dots, j'_d} \in I$ such that for all $v \in i_{j'_1, \dots, j'_d}$: $R(v) \cap s_k \neq \emptyset$. So there is some $v \in R(w)$ such that $R(v) \cap s_k \neq \emptyset$. Hence, R is non-transitive.

reflexive transitive closure of each color of arrow. So, each world R -accesses: (i) itself; (ii) any world reachable by one or more arrows of the same color; and (iii) no other worlds.

Intuitively, the model in [Figure 1](#) represents a simplified, three-day version of **Class**. On this interpretation, each $s_i \in \mathbf{S}$ corresponds to the worlds at which the rooms were switched on day i . Each $i_{j,k} \in \mathbf{I}$ corresponds to the worlds at which s_j and s_k cannot be ruled out. For example, suppose $w \in s_i \cap i_{j,k}$. Since R is reflexive, the strongest proposition known at w is $s_i \cup s_j \cup s_k$. As a result, strictly more will be known where $i = k$ or $i = j$ than where $i \neq j$ and $i \neq k$. Among the worlds in $i_{j,k}$, more is known in s_j and s_k than in other scenarios. That's because worlds in $i_{j,k}$ are those in which s_j and s_k cannot be ruled out, whatever scenario obtains. As can be checked, the model satisfies **Symmetry**⁺ and **Symmetry**⁻ while allowing for **Introspection** failure.

6 Inexactness

In **Class**, Nora's evidence for each of a set of claims is on a par, yet she has overwhelming evidence that some claim in the set is false. Cases with this structure are abundant. Most prominently, this structure is present in many versions of the preface paradox ([Makinson \(1965\)](#)). However, cases with the same structure equally afflict our testimonial knowledge ([Hawthorne \(2003, 48-49\)](#); [Smith \(2016, 72-75\)](#)), inductive knowledge ([Goodman & Salow \(2021, 176\)](#)) and perceptual knowledge ([Littlejohn & Dutant \(2020, 1597\)](#), [Carter & Goldstein \(2021, 2523\)](#)).

How common these cases are depends partly on how similar our evidence for claims must be for them to be on a par. Take any set of claims for which our evidence is fallible and is on a par. As long as we can find a sufficiently large set, our evidence that some claim in the set is false will be overwhelming. By the reasoning above, our higher-order knowledge about which of these claims is true will exhibit positive and negative symmetries. Our ability to proliferate such cases will be limited only by our ability to find sufficiently many claims for which our evidence on a par (cf. [Littlejohn & Dutant \(forthcoming\)](#)).

Not all cases of **Introspection** failure have this structure, though. Claims about the size of the crowd are not all on a par for Aron, given his evidence. What implication then, if any, does the preceding discussion have for cases like **Crowd**?

It is possible that our higher-order ignorance has multiple sources. Perhaps some failures of introspection are attributable to indeterminacy and others to a margin-for-error. Still, it is worth considering whether a unified explanation is available.¹⁴

¹⁴One reason to consider alternative explanations arises from concern about why we should think that Aron's knowledge is subject to a margin-for-error. A common thought is that Aron could easily have formed a false belief in a relevantly similar proposition about the size of the

In **Crowd**, Aron’s knowledge is inexact. There is some k and $k' > k$ such that, for any $k \leq i \leq k'$, Aron can’t know the crowd doesn’t contain i people. Moreover, his knowledge of how inexact his knowledge is also appears inexact. Intuitively, he can’t know what the greatest such k and least such k' are (cf. Dutant (2007); Rosenkranz & Dutant (2020); Williamson (2011, 2014)).

Why not? A natural explanation is that it is indeterminate how inexact Aron’s knowledge is. There appears to be nothing to settle the least k and greatest k' such that Aron knows the size of the crowd is between k and k' . Any choice is arbitrary. As we saw, this arbitrariness is symptomatic of indeterminacy. It suggests that there is no k for which it is determinate that k is the greatest lower bound (or the least upper bound) on the crowd’s size known by Aron.

If it is indeterminate how inexact Aron’s knowledge is, then it is indeterminate what Aron knows. This indeterminacy can act as a barrier to Aron’s ability to know what he knows about the crowd’s size. And, arguably, it gives us a fully adequate characterization of Aron’s higher-order ignorance without the need for a margin-for-error. I’ll demonstrate this point by considering a variant of the model in §5.

6.1 A model of Crowd

In **Crowd**, there is some set of precise ways the world could be which can be ordered with respect to how accurate Aron’s evidence would be if the world were that way. As with **Class**, cases with this structure can be represented by imposing additional constraints on models.

To represent the relevant ways the world could be, we introduce a subject matter $\mathbf{S}^* = \{\mathbf{s}_1, \dots\}$. We assume that each $\mathbf{s}_i \in \mathbf{S}^*$ is precise at each world. In the case of **Crowd**, each $\mathbf{s}_i \in \mathbf{S}^*$ corresponds to the set of worlds at which there are exactly i people in the crowd. Let e be an integer value representing the agent’s best estimate (given how things appear) about which element of \mathbf{S}^* obtains.

To represent the indeterminacy of what is known, we introduce a second subject matter $\mathbf{I} = \{\mathbf{i}_1, \dots\}$. The idea is that each \mathbf{i}_k corresponds to the set of worlds at which $[x - k, x + k]$ is the maximum achievable level of exactness of what is known about \mathbf{S}^* . At any world in \mathbf{i}_k , there are at least $2k$ elements of \mathbf{S}^* which cannot be ruled out (and at some worlds in \mathbf{i}_k , there may be more)

Unlike in the model of **Class**, we will not assume that for each $\mathbf{i}_k \in \mathbf{I}$, whether \mathbf{i}_k is not settled by how things are precisely. At any world there will presumably be some values of k such that it is determinate that the maximum level of exactness of what is known about the size of the crowd is less than k or greater than k .

crowd using the same method(s) (Williamson (1994, 2009); Sainsbury (1997); Manley (2007); Dutant (2016), cf. Mahtani (2008)). But the question of what counts as the same method and what propositions are relevantly similar shouldn’t be expected to be resolved in a non-circular fashion, independent of the question of what Aron knows (Williamson (2000)). As a result, it is far from clear whether this can give us independent grounds for thinking that Aron’s knowledge is subject to a margin-for-error.

In this case, how things are precisely will imply \mathbf{i}_k does not hold.

However, we will assume that there is no \mathbf{i}_k and $v \in \mathbf{i}_k$ such that it is determinate that \mathbf{i}_k at v . This reflects the idea that it is indeterminate how inexact knowledge about \mathbf{S}^* is. More specifically, let $d \geq 1$ be some integer constant representing the level of determinacy regarding the maximum achievable exactness of knowledge about \mathbf{S}^* .¹⁵ We assume that at $w \in i_k$, it is indeterminate whether i_x at w , for any $x \geq 1$ such that $k - d \leq x \leq k + d$. Accordingly, there is no i_k and w such that it is determinate that i_k at w .

To model cases like **Crowd**, we impose two conditions on R , in addition to **Ignorance**.

Inexactness If $w \in \mathbf{i}_k$ and $|j - e| \leq k$, then $R(w) \cap \mathbf{s}_j \neq \emptyset$.

Inexactness says that at any \mathbf{i}_k -world, there is no size of the crowd which is within k of e which is known not to obtain. This simply implements the idea that elements of \mathbf{I} correspond to ways of settling the maximum achievable level of exactness of knowledge about \mathbf{S}^* .

Convexity If $k \leq x \leq k'$, $R(w) \cap s_k \neq \emptyset$ and $R(w) \cap s_{k'} \neq \emptyset$, then $R(w) \cap s_x \neq \emptyset$.

Convexity says that if it is not ruled out that the crowd contains k people and it is not ruled out that the crowd contains k' people, then it is not ruled out that the crowd contains x people, for any $k \leq x \leq k'$.

In order to evaluate **Margins**, we simply assume that there is some relation \sim over \mathbf{S}^* which relates each $\mathbf{s}_i \in \mathbf{S}^*$ to \mathbf{s}_{i-1} and \mathbf{s}_{i+1} (and perhaps some other elements of \mathbf{S}^* , too). **Margins** says that if $w \in \mathbf{s}_i$ and $\mathbf{s}_i \sim \mathbf{s}_j$, then $R(w) \cap \mathbf{s}_j \neq \emptyset$.

We can establish the adequacy of the model via the following observation:

Fact 2. Let R be the weakest relations satisfying **Ignorance**, **Inexactness**, and **Convexity**.

Then it follows that both **Introspection** and **Margins** fail.¹⁶

Let R be the weakest relation satisfying **Ignorance**, **Inexactness**, and **Convexity**. Not only does **Introspection** fail in the model, but it fails at every

¹⁵Since it is plausible that it is indeterminate what levels of exactness of knowledge are determinately unachievable, this is an idealization away from higher-order indeterminacy. A more plausible model would treat the value of d as a function of worlds which can differ across worlds which agree on all precise matters. This complication is suppressed for the purposes of exposition, since the primary result, below, is independent of its introduction.

¹⁶*Proof:* for **Introspection**, consider $w \in \mathbf{s}_n \cap \mathbf{i}_k$, where $n \leq e$. By **Ignorance**, $R(w) \cap \mathbf{i}_{k+1} \neq \emptyset$. But by **Inexactness**, for all $v \in \mathbf{i}_{k+1}$, $R(v) \cap \mathbf{s}_{e+(k+1)} \neq \emptyset$. But, since $n \leq e$, $R(w) \cap \mathbf{s}_{e+(k+1)} = \emptyset$. So R is non-transitive.

For **Margins**, consider $w \in \mathbf{s}_n \cap \mathbf{i}_k$, where $n \leq (e - k)$. Since R is the weakest such relation: if $R(w) \cap \mathbf{s}_i \neq \emptyset$, then $i \geq n$. But $\mathbf{s}_n \sim \mathbf{s}_{n-1}$. So **Margins** fails.

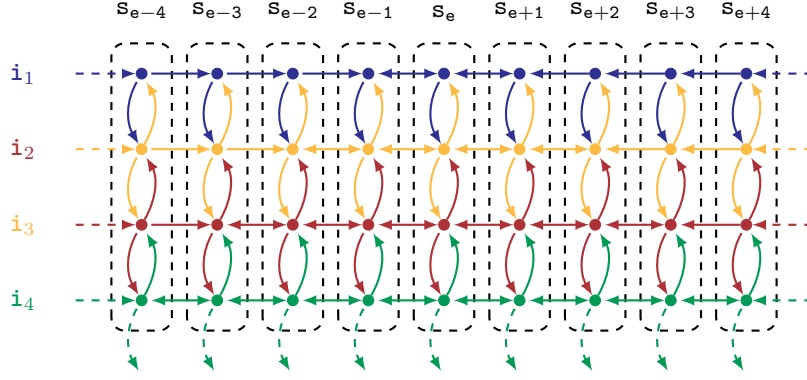


Figure 2

world. For each $w \in W$ there will be some $v \in R(w)$ such that $R(v) \not\subseteq R(w)$. This is because, for any world, whatever the strongest proposition known at that world, it is not determinate at that world that that proposition is known. Accordingly, by **Ignorance**, at every world there will be some proposition which is known but which is not known to be known.

As a concrete example, consider the model in **Figure 2**, which satisfies all three constraints (where $d = 1$). As above, accessibility corresponds to the transitive reflexive closure of each color of arrow. As can be easily checked, the model permits ‘cliff-edge’ knowledge: for some s_n and $w \in s_n$. Suppose $w \in s_n \cap i_k$, where $n \leq e$ and $k < n$. Then for any $j < n$, then $R(w) \cap s_j = \emptyset$ (Hellie (2005); Cohen & Comesaña (2013); Bonnay & Égré (2008, 2009); Goldstein (2022)). So, **Margins** is violated. Despite this, **Introspection** fails at every world. Moreover, second-order ignorance about what is known about the size of the crowd is not parasitic on first-order ignorance about the size of the crowd: for any $w \in s_n$, there is some $v \in R(w) \cap s_n$ such that $R(v) \not\subseteq R(w)$.

That indeterminacy can explain our second-order ignorance does not imply that it does. Perhaps, in cases like **Crowd**, our knowledge is constrained both by indeterminacy and by a margin for error. Still, we should be cautious about entertaining unnecessary explanations. Indeterminacy in what we know is independently necessary to explain our ignorance about how inexact our knowledge is. What I have aimed to show is that, once we posit indeterminacy about the inexactness of what we know, this suffices to explain our ignorance of what we know.

7 Conclusion

Indeterminacy comes in different kinds: semantic, metaphysical, and epistemic. What kind of indeterminacy afflicts our knowledge is an important question. Crucially, however, indeterminacy of any kind is a barrier to knowledge (Greenough (2003); Williams (2008)). Accordingly, the proposal above can remain neutral on whether indeterminacy is best understood as metaphysical or semantic. Indeed, it is even open that indeterminacy is to be understood epistemically, as ineliminable ignorance about what is known about a subject matter. What is crucial is only that such ineliminable ignorance is brute: it cannot be explained parasitically by first-order ignorance about that same subject matter.

References

- Bacon, Andrew. 2014. Giving Your Knowledge Half a Chance. *Philosophical Studies*, 1–25.
- Bacon, Andrew. 2018. *Vagueness and Thought*. Oxford University Press.
- Barnes, Elizabeth. 2014. Fundamental Indeterminacy. *Analytic Philosophy*, **55**(4), 339–362.
- Bernstein, Sara. 2016. Causal and Moral Indeterminacy. *Ratio*, **29**(4), 434–447.
- Bonnay, Denis, & Égré, Paul. 2008. Margins for Error in Context. *Pages 103–107 of: Carpintero, G., & Koelbel, M. (eds), Relative Truth*. Oxford University Press.
- Bonnay, Denis, & Égré, Paul. 2009. Inexact Knowledge with Introspection. *Journal of Philosophical Logic*, **38**(2), 179–227.
- Breckenridge, Wylie, & Magidor, Ofra. 2012. Arbitrary Reference. *Philosophical Studies*, **158**(3), 377–400.
- Carter, Sam, & Goldstein, Simon. 2021. The Normality of Error. *Philosophical Studies*, **178**(8), 2509–2533.
- Cohen, Stewart, & Comesaña, Juan. 2013. Williamson on Gettier Cases and Epistemic Logic. *Inquiry: An Interdisciplinary Journal of Philosophy*, **56**(1), 15–29.
- Dorr, Cian. 2003. Vagueness without Ignorance. *Philosophical Perspectives*, **17**(1), 83–113.
- Dougherty, Tom. 2013. Vague Value. *Philosophy and Phenomenological Research*, **89**(2), 352–372.
- Dutant, Julien. 2007. Inexact Knowledge, Margin for Error and Positive Introspection. *Pages 118–24 of: Samet, Dov (ed), Proceedings of Tark XI*.
- Dutant, Julien. 2016. How to Be an Infallibilist. *Philosophical Issues*, **26**(1), 148–171.
- Goldstein, Simon. 2022. Fragile Knowledge. *Mind*, **131**(522), 487–515.
- Goodman, Jeremy. 2013. Inexact Knowledge without Improbable Knowing. *Inquiry: An Interdisciplinary Journal of Philosophy*, **56**(1), 30–53.
- Goodman, Jeremy, & Salow, Bernhard. 2021. Knowledge from Probability. *Electronic Proceedings in Theoretical Computer Science*, **335**(June), 171–186.
- Greenough, Patrick. 2003. Vagueness: A Minimal Theory. *Mind*, **112**(446), 235–281.
- Hawthorne, John. 2003. *Knowledge and Lotteries*. Oxford University Press.
- Hellie, Benj. 2005. Noise and Perceptual Indiscriminability. *Mind*, **114**(455), 481–508.
- Lewis, David. 1988a. Relevant Implication. *Theoria*, **54**(3), 161–174.
- Lewis, David K. 1988b. Statements Partly About Observation. *Philosophical Papers*, **17**(1), 1–31.

- Littlejohn, Clayton, & Dutant, Julien. 2020. Justification, Knowledge, and Normality. *Philosophical Studies*, **177**(6), 1593–1609.
- Littlejohn, Clayton, & Dutant, Julien. forthcoming. What Is Rational Belief? *Noûs*.
- Mahtani, Anna. 2008. Williamson on Inexact Knowledge. *Philosophical Studies*, **139**(2), 171–180.
- Makinson, David C. 1965. The Paradox of the Preface. *Analysis*, **25**(6), 205.
- Manley, David. 2007. Safety, Content, Apriority, Self-Knowledge. *Journal of Philosophy*, **104**(8), 403–423.
- Plebani, Matteo, & Spolaore, Giuseppe. 2021. Subject Matter: A Modest Proposal. *Philosophical Quarterly*, **71**(3), 605–622.
- Rosenkranz, Sven, & Dutant, Julien. 2020. Inexact Knowledge 2.0. *Inquiry: An Interdisciplinary Journal of Philosophy*, 1–19.
- Sainsbury, R. M. 1997. Easy Possibilities. *Philosophy and Phenomenological Research*, **57**(4), 907–919.
- Smith, Martin. 2016. *Between Probability and Certainty: What Justifies Belief*. Oxford University Press UK.
- Srinivasan, Amia. 2013. Are We Luminous? *Philosophy and Phenomenological Research*, **90**(2), 294–319.
- Stalnaker, Robert C. 1980. A Defense of Conditional Excluded Middle. *Pages 87–104 of: Harper, William, Stalnaker, Robert C., & Pearce, Glenn (eds), Ifs*. Reidel.
- Swanson, Eric. 2016. Inderterminacy in Causation. *Philosophical Quarterly*, **67**(268), 606–624.
- Weatherson, Brian. 2004. Luminous Margins. *Australasian Journal of Philosophy*, **82**(3), 373–383.
- Weatherson, Brian. 2013. Margins and Errors. *Inquiry: An Interdisciplinary Journal of Philosophy*, **56**(1), 63–76.
- Williams, J. Robert G. 2008. Ontic Vagueness and Metaphysical Indeterminacy. *Philosophy Compass*, **3**(4), 763–788.
- Williamson, Timothy. 1992. Inexact Knowledge. *Mind*, **101**(402), 217–242.
- Williamson, Timothy. 1994. *Vagueness*. Routledge.
- Williamson, Timothy. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.
- Williamson, Timothy. 2009. Probability and Danger. *Pages 1–35 of: Amherst Lecture in Philosophy*, vol. 4.
- Williamson, Timothy. 2011. Improbable Knowing. *In: Dougherty, T. (ed), Evidentialism and Its Discontents*. Oxford University Press.
- Williamson, Timothy. 2013. Gettier Cases in Epistemic Logic. *Inquiry: An Interdisciplinary Journal of Philosophy*, **56**(1), 1–14.
- Williamson, Timothy. 2014. Very Improbable Knowing. *Erkenntnis*, **79**(5), 971–999.
- Yablo, Stephen. 2014. *Aboutness*. Princeton University Press.