Forthcoming in *Mind*

Daniel Immerman

How common knowledge is possible

Abstract. *The two of us commonly know a proposition just in case (i) we both know it, (ii) we both know that we both know it, (iii) we both know that we both know that we both know it, and so on. In a recent paper titled "Uncommon Knowledge", Harvey Lederman argues against the possibility of common knowledge. His argument rests on the empirical claim that there are minor individual variations in how we perceive things. This motivates a principle about perception: when two people are perceiving something and it perceptually appears a certain way to one of them, then, for all that person knows, it perceptually appears a slightly different way to the other. In this paper, I challenge Lederman's perceptual principle and thereby his argument. In particular, I argue that there are some circumstances in which things perceptually appear a certain way to me and nonetheless I know that they don't perceptually appear in a slightly different way to you. Indeed, I argue that not only are there exceptions to the perceptual principle, but they are widespread.*

You and I are walking along a busy street when you spot a llama on stilts. You point it out to me and we stop and stare.[1] We both know that there's a llama on stilts before us. In addition, we both know that we both know it. Furthermore, assuming we have reflected on the subject, we both know that we both know that we both know it.

A natural question arises: can we keep iterating forever? To make this question more precise, let us define some notions. First: common knowledge. The two of us *commonly know* a proposition just in case (i) we both know

[1]This case is, to some extent, modeled on [Lederman 2017, 1069].

it, (ii) we both know that we both know it, (iii) we both know that we both know that we both know it, and so on.

Second: being in a position to know. As I shall use the term in this paper, one is in a *position to know* something just in case one has all the empirical information one needs to know it. In other words, when one is in a position to know something, the only limits, if any, that keep one from knowing it are limits on one's reflective capacities (and knowledge thereof). Likewise, when you and I are in a position to commonly know something, the only limits, if any, that keep us from having this common knowledge are limits on our reflective capacities (and knowledge thereof).

Now the question becomes: are you and I in a position to commonly know that there's a llama on stilts before us?

There is a tradition in computer science, linguistics, and philosophy, of answering "yes" to this sort of question [Lederman 2017, 1069]. But in a recent paper titled "Uncommon Knowledge", Harvey Lederman argues that the answer is no; he thinks that in such circumstances we are not in a position to have common knowledge [Lederman 2017].

Lederman's argument rests on the empirical claim that there are minor individual variations in how we perceive things. This motivates a principle about perception: when two people are perceiving something and it perceptually appears a certain way to one of them, then, for all that person knows, it perceptually appears a slightly different way to the other. Lederman uses this principle to argue that these little differences add up and thereby serve

to rob us of common knowledge.

In this paper, I will be challenging this argument. The way I will challenge it is by arguing that the perceptual principle underlying it is false. In particular, there are some circumstances in which things perceptually appear a certain way to me and nonetheless I know that they don't perceptually appear in a slightly different way to you.

Even if a principle admits of exceptions, it might still hold in general. Thus perhaps Lederman could contend that most of the time, we are not in a position to have common knowledge. I explore this question as well, arguing that not only are there exceptions to his perceptual principle, but they are widespread.

My first section summarizes Lederman's argument. My second develops a counterexample. My third and fourth argue that counterexamples are widespread.

1 Lederman's argument

Let me start by laying out Lederman's argument. It works as follows: first, he constructs a scenario, second, he argues that in that scenario, the characters are not in a position to commonly know a certain proposition, and third, he generalizes this argument so that it applies to many other scenarios.

Here is Lederman's initial scenario:

SAILBOAT: Roman and Columba are ideal reasoners playing in

> a game show. Each contestant has a single button on a console in front of him or her. They have an unobstructed view of each other's faces, and of an area in the middle of the stage, where the hosts will place a sailboat. First, the hosts will bring out a toy sailboat (the 'test') with a 100 cm mast. They will then replace it with a sailboat chosen randomly from an array of sailboats of various sizes. If the mast of the new sailboat is taller than the test and both players press their respective buttons, they receive $1,000 each. If the mast is not taller than the test and both press, or if only one person presses their button, the person or people who pressed must pay the show $100. Today, the mast of the chosen boat is 300 cm tall. [Lederman 2017, 1075].

Lederman argues that in this scenario, the main characters lack common knowledge that the mast is taller than 100 cm. His argument rests on the idea that there is individual variation in the way things looks to different people, and thus for any way it looks to one of the contestants, for all that contestant knows, it may look very slightly differently to the other. Lederman thinks this motivates the following principle:

> Interpersonal Ignorance For all r, if it [the mast] looks to be r cm tall to one of the agents, then for all that agent knows, it looks to be $0.97r$ cm tall to the other. [Lederman 2017, 1081].

In order for his argument to work, Lederman needs this principle to be com-

mon knowledge. To ensure this, he adds to the scenario. In particular, he says that prior to playing the game, the players learn about perceptual variation among individuals. This yields him:

> CK Interpersonal Ignorance: For all r, Roman and Columba commonly know that if the mast looks to be r cm tall to one of them, then for all that player knows, it looks to be 0.97r cm tall to the other. [Lederman, 2017, 1081].

Lederman relies on this principle to argue that there is no common knowledge in Sailboat. Suppose the mast looks a certain height (say 300 cm) to Roman. Then, for all Roman knows, it looks slightly less tall (say 299 cm) to Columba. Roman knows that Columba is in a similar situation – that is – for all she knows, things look ever so slightly smaller to Roman than they do to her. Thus, for all Roman knows, for all Columba knows, the mast looks 298 cm tall. We can zig-zag back and forth to eventually get to the conclusion that for all Roman knows, for all Columba knows, for all Roman knows ... the mast is 100 cm tall. And thus, they do not commonly know that it is above 100 cm. [Lederman 2017, 1075].

Lederman next generalizes his argument. To do so, he argues that his point about perceptual variation carries beyond just variation about heights, but also to other continuous magnitudes, such as width, curvature, and color, and to other modalities, such as auditory modalities like pitch and modulation [Lederman 2017, 1086].

He uses this to motivate a principle similar to INTERPERSONAL IGNORANCE:

> WORLDLY IGNORANCE: For some constant parameter describing variations in Roman and Columba's visual perception, if Roman looks to be a certain way to Columba, then for all Columba knows, Roman's true hue, size and shape differ by the relevant constant parameter from how they appear to her.[2]

Lederman then uses this principle to argue that there are in principle barriers to common knowledge. Here, instead of moving from the mast height being 300 cm to being 100 cm, we instead move from Roman looking like a human to Roman looking like a rock. We thus get the conclusion that for all Roman knows, for all Columba knows, for all Roman knows ... for all Columba knows, Roman looks like a rock. If for all Columba knows, Roman looks like a rock, then for all she knows, he is a rock. Because rocks don't have knowledge, it follows that for all Roman knows, for all Columba knows, for all Roman knows ... for all Columba knows, Roman doesn't know anything. [Lederman 2017, 1086-7].

[2]I should note that in his paper, Lederman actually relies on an even stronger principle, which he calls CK WORLDLY IGNORANCE. ("CK" stands for common knowledge). This principle requires not merely that WORLDLY IGNORANCE holds, but also that Roman and Columba have common knowledge that it does. [Lederman 2017, 1086-7].

2 Challenging Lederman's arguments

Why does Lederman discuss SAILBOAT before going on to give his generalized argument? One reason is that he thinks SAILBOAT presents opportune circumstances for common knowledge. The lighting is good, the characters can clearly see each other, they can clearly see that the mast is greater than 100 cm, and so on. Thus, if he can show that the characters are not in a position to commonly know in this case, then he has provided some reason to think there we are never in a position to commonly know [Lederman 2017, 1075].

Is Lederman right? That is, does SAILBOAT provide ideal circumstances for common knowledge? I myself am somewhat skeptical. In this paper, I will be looking at some variants of this case, and arguing that, in these variants, Lederman's perceptual principles fail to apply. I will eventually work my way back to SAILBOAT, which I think presents a more subtle case.

2.1 Missing Boats

Recall that in SAILBOAT, Lederman makes use of a principle he calls INTERPERSONAL IGNORANCE, which runs as follows.

> INTERPERSONAL IGNORANCE For all r, if [the mast] looks to be r cm tall to one of the agents, then for all that agent knows, it looks to be 0.97r cm tall to the other. [Lederman 2017, 1081].

And recall that, in order for his argument to work, it not only has to hold,

it has to be common knowledge.[3] But as I shall now argue, this principle runs into problems in a number of variants of the original case. Here is an example:

> MISSING BOATS: Everything is the same as in SAILBOAT, except that the hosts, slightly before the game show begins, realize that a prankster has stolen all the sailboats with masts between 100 and 300 cm. The hosts announce this to the contestants. The contestants discuss amongst themselves and insist on playing the game despite the missing sailboats. They are now staring at a sailboat with a mast above 300 cm.

Even in advance of looking at the sailboat, the contestants know that there are certain ways it will not look. For example, they know that the mast will not look to be 200 cm. The reason is that there are no boats with masts even close to 200 cm. Similar things hold for 201 cm, 202 cm, and a bunch of other values.

Because of this there will be some range of ways the boat will not look to either participant. Ever so slightly above the top of this range will be a height such that a mas may well look that height to one of the contestants, but it will not look to be 1 cm less to the other. Call such a height "barely lookable". Thus, Columba knows the following: if the mast's height is barely lookable, then it won't look to be 1 cm lower to Roman.

[3]Note that Lederman himself doesn't need to endorse this; rather, he assumes it as part of an ad hominem argument. Thanks for an anonymous referee for pushing me to clarify this.

Interpersonal Ignorance says there is no such height. It implies that for every height, if the mast looks to be that height to one of the contestants, then for all that contestant knows, it will look 1 cm less to the other. And thus, in this case, Interpersonal Ignorance runs into problems.[4]

Let me pause to discuss how this criticism might be spelled out more fully. To do so, it will be helpful to focus on the following principle, which concerns the players' knowledge before they look at the mast:

> Intertemporal Interpersonal Ignorance: Prior to looking at the sailboat, the agents know the following: [for all r, if the mast will look to be r cm tall to one of the agents, then for all that agent will know, it will look to be 0.97r cm tall to the other].

It seems plausible that Lederman's claim that Interpersonal Ignorance holds is motivated by the claim that Intertemporal Interpersonal Ignorance holds, together with the assumption that when the agents look at the sailboat they do not lose knowledge in any interesting way.

[4]An anonymous referee suggests that perhaps Lederman's argument could be run with a principle called Mutual^n Interpersonal Ignorance:

> Mutual^n Interpersonal Ignorance For some r $<$ 1, for all x, if the mast looks to be xcm tall to one of the people, then for all the people mutually know it looks to be x$*$r cm tall to the other.

The referee adds that while perhaps a fan of common knowledge might say that using this principle in Lederman's argument would be question-begging, I haven't said anything to explicitly challenge it. I agree with the referee on both points. I don't currently see how to motivate it in a way that avoids a challenge like the one I give in this paper, but am open to the possibility that there is such a way of doing so. Thanks to the referee for this suggestion.

But in the case of, Missing Boats, Intertemporal Interpersonal Ignorance fails to apply. Prior to looking at the sailboat, for all the characters know, it will be barely lookable. And if it is indeed barely lookable, they will know it won't look to be 1 cm shorter to the other participant. Because Intertemporal Interpersonal Ignorance fails to apply in the case of Missing Boats, we lack motivation for the claim that Interpersonal Ignorance is known.[5]

Before I go into more details about this case, let me pause to be clear about the dialectic. While Lederman is committed to Interpersonal Ignorance holding in the case of Sailboat, he is not committed to it holding in other cases. Thus, it is consistent with what he says in his paper that Interpersonal Ignorance fails to hold in the case of Missing Boats. That said, as we shall see, I will be building on what I say regarding Missing Boats to argue (1) that his generalized argument against common knowledge fails and (2) that counterexamples to his perceptual principles are widespread.

2.2 On perceptual differences across other modalities

In his paper, Lederman considers something like the point I have been pressing. In particular, he acknowledges the following possibility: even if, for all things look to Columba, the mast might look a little shorter to Roman, Columba might know, via some other modality, such as testimony, that it

[5]Thanks to an anonymous referee for suggesting this way of spelling out the challenge.

does not look this way.

Lederman then responds:

> But even if we consider everything a subject knows via all modalities, there will still be what we might call a 'total perceptual JND' [just noticeable difference], a variation in the sum total of perceptual appearances consistent with everything learned in any way (by any modality) whatsoever. These total perceptual JNDs will no doubt be smaller and more oddly behaved than the ordinary JNDs measured by psychophysicists in familiar paradigms. But it is plausible that they exist: the information we possess on the basis of perception – even all perception taken together – isn't perfectly accurate at very precise levels of detail. ... we can exploit variations in these total JNDs to move from the agents' actual perceptual appearances to quite 'distant', strange appearances. [Lederman 2017, 1088].

If what I've been saying is correct, then we have good reason to think what Lederman says here is incorrect. Put abstractly, the point I have been making is the following: (i) in gathering information about the world, I will come to know that there are certain ways things could not appear to certain people, (ii) but there will be other ways of appearing such that, for all I know, things appear those ways to these people, and thus, (iii) there is a border between ways of appearing. On the one side of this border are ways such that – for

all I know – things could appear that way. On the other side of the border are ways such that I know things don't appear that way. And this is where we find counterexamples to Lederman's idea: here, at the very edge of the border.

To illustrate, take the case of MISSING BOATS. Because we have expanded sensory modalities and are not merely talking about heights, the border is not merely a border between the heights that masts can appear to have. That is, it is not merely a border between the mast looking above 272.5 cm and below. We also have to take into account the way the host sounded when making the announcement about the prankster. Perhaps the host could have sounded slightly more high-pitched and slightly quieter. Perhaps the host could have looked slightly more embarrassed. And so on. The border is now along multiple dimensions. But that doesn't cause it not to exist; rather, it is now a multi-dimensional border.

Here is another way of putting my point. Lederman's idea was that, if things appear a certain way to me, then for all I know they could appear a slightly different way to you. My point is that this will not hold in the following sort of case: a case in which they appear in the most extreme way they could possibly appear, given how I know the circumstances to be. If they appear in the most extreme way they could possibly appear, given how I know the circumstances to be, then I know then they don't appear in an even more extreme way to you.

To develop this point even more concretely, here's an example.

> ARCHERY JUDGE: We are at an archery tournament and there is a judge (who I know to be honest) who is responsible for determining if an arrow landed within the target. The arrow is shot, the judge inspects the arrow, and then raises a green flag, which means: "this arrow is lying within the target". Some details: if it hadn't been within the target, she would have raised a purple flag. The rules say she has to hold the flag up for at least 10 seconds. She always holds it up for well over that amount of time. After she holds up the flag, I go to inspect the arrow myself.

Suppose, before I look at the arrow myself, but after I see her raise the flag, I think about what the judge knows. For all I know, there are various ways in which the arrow could appear to her. For instance, it could appear to be a bullseye. It could also appear to be within the target, but near the border. That said, I know there are other ways the arrow could not appear to her. For instance, it could not appear to be 10 cm outside the target. After all, she is an honest judge and raised the green flag, which signals that it appears to her as if the arrow is within the target. Given that there are some ways the arrow could appear, and other ways it could not, there must be some border of ways the arrow could appear to her.

There are also other borders along other dimensions; borders between how slowly it felt like she raised the flag, how long it felt like she held it up, and so on.

Now, suppose I take a look at the arrow and it looks to me at the extreme

edge of the borders of how it could look. That is, suppose the arrow looks to me at the maximal distance it could look (to me or her) from the center of the target, given that she raised the green flag, as opposed to the purple. Suppose all my other appearances are border as well. For instance, suppose that she had held the flag for what seemed (to me) the minimum amount of time she could, given that she's a conscientious judge and always holds it up well over 10 seconds.

Here, in this case, there is no gap. There is no extra little bit of distance away from the target the arrow could look to her. There is no shorter amount of time it could have felt to her. That is, there are cut-offs along every dimension: the farthest from the center of the target, the shortest the amount of time the flag can be held, and so on. If we're at the extremes along every dimension then I know that things cannot appear to her to be even farther along the extremes.

2.3 On Lederman's generalized argument

I have just been arguing that the principle underlying Lederman's SAILBOAT argument – INTERPERSONAL IGNORANCE – fails to be known in some cases, such as MISSING BOATS. But, as I have noted, this is consistent with Lederman's argumentative strategy. He didn't say this principle holds in every case; he just said that it holds in SAILBOAT. What he did say is that the principle that he uses for his generalized argument – WORLDLY IGNORANCE – is known in every case. Recall that this general principle runs as follows.

> WORLDLY IGNORANCE: For some constant parameter describing variations in Roman and Columba's visual perception, if Roman looks to be a certain way to Columba, then for all Columba knows, Roman's true hue, size and shape differ by the relevant constant parameter from how they appear to her.

Unfortunately for Lederman, this principle, too, admits of exceptions. Indeed, it fails to hold in SAILBOAT, Lederman's flagship case.

Here, I rely on the following feature of SAILBOAT, which Lederman himself included: the contestants have an unobstructed view of each other's faces. Per the rules of the game, the contestants know that they will be playing the game with another human (as opposed to a rock). Thus, there are certain ways that the other human could not look. For instance, the human could not look exactly like a rock. The human could not look like a rock, but with a human head. And so on. As before, choose a way things could look every-so-slightly outside of this range and call it "barely lookable".

This serves to undermine the motivation for knowledge of WORLDLY IGNORANCE. If Roman looks barely lookable to Columba, then it's not the case that, for all she knows, his true hue, size, and shape differ ever so slightly, so that he looks every-so-slightly more rock-like. After all, Columba knows she is playing with a human and not a rock. Thus, knowledge of WORLDLY IGNORANCE fails to be motivated in the case of SAILBOAT. (Similar things apply to the case of MISSING BOATS.)

3 How common is knowledge of gaps

If what I have said is right, then INTERPERSONAL IGNORANCE faces problems in MISSING BOATS and WORLDLY IGNORANCE faces problems in SAILBOAT and MISSING BOATS. The reason is that the characters in these cases have what I call a "knowledge of gaps" – they know that they won't find themselves in certain sorts of circumstances. And this allows them to prevent perceptual differences from robbing them of common knowledge.

In this section, I will argue that knowledge of gaps is fairly common. To do so, I distinguish three varieties of knowledge of gaps and note some everyday examples.

3.1 Knowledge that certain possibilities won't obtain

The first sort of knowledge of gaps is one we have already seen. It involves knowledge that certain possibilities won't obtain. For example, in the case of MISSING BOATS, our characters had knowledge that the mast wouldn't be a certain height. Sometimes, in everyday life, we have this sort of knowledge. Here is an example:

> DINNER PARTY I am throwing a dinner party for a small group of close friends. Everyone has arrived except for two: Big John and Little John. As their names suggest, they differ in height; Big John is taller than average (six feet) and Little John is shorter than average (five feet). The doorbell rings, so the two people

nearest the door, Alice and Bob, open it and greet the new arrival.[6]

Here, then, is an everyday case in which characters have knowledge that certain possibilities won't obtain. In particular, given that they know who is coming to the party, Alice and Bob know that there are certain heights that the guest will not have. In particular, they know that the guest will not be five feet six inches tall.

3.2 Knowledge that there will be additional perceptual clues

A different sort of knowledge of gaps is the following: knowledge that if a certain possibility does occur, things will appear differently to the participants. So, for example, consider the following way variant of SAILBOAT:

> DISTINCTIVE FACIAL EXPRESSIONS: Everything is consistent with the way SAILBOAT is described. Notably, as with SAILBOAT, the contestants have an unobstructed view of each other's faces. As it turns out, Roman has distinctive facial expressions. If he can easily tell that the mast is above 100 cm, he grins broadly. If he is unsure, he frowns and furrows his brow. He never does anything in between.

[6] Anonymous editors suggested that I clarify whether Alice and Bob have been told the heights of Big John and Little John or whether their knowledge comes merely from observation. I personally am hoping that nothing I say in the paper hinges on understanding the case in one way, rather than the other. But for the sake of precision: please understand the case in the second way.

This sort of case provides a different way to avoid Lederman's arguments. Here, the characters don't know that there are certain heights that the mast will not be. Rather, the characters know that they will have evidence that sharply distinguishes the case where the mast is clearly above 100 cm and the case in which it is not.

This sort of knowledge of gaps is fairly common. It is often present, even when we lack the first type of knowledge of gaps. So, for example, let us return to the dinner party example. Let us alter it slightly, so that guests are allowed to bring plus ones. At this point, Alice and Bob no longer know that there won't be someone who shows up looking halfway in between Big John and Little John. But they know that if this were to happen there would be additional perceptual clues that would allow them to distinguish this guest from Big John and Little John. For instance, the person would introduce themselves as a guest, they would arrive with a known attendee, and so on.

3.3 Knowledge that people will probe in borderline cases

Another final sort of knowledge of gaps derives from the fact that people will often probe in borderline cases by seeking new evidence to help distinguish between various cases. So for example, consider the following variant of Sailboat:

> Talk It Out: Everything is consistent with the way Sailboat is described. Unfortunately, Roman lacks distinctive facial expressions, so it is often hard to tell what he's thinking. Columba

> doesn't want to have to pay $100. So if she hits a borderline case, she will ask Roman whether he agrees that the mast is above 100 cm. That way, she'll avoid the case in which she presses the button and he doesn't.

This sort of case provides yet another way to avoid Lederman's argument. Here, the characters lack the sort of evidence that they had in the previous case. But they can choose to communicate so as to distinguish cases that are along the border.

Again, it is worth emphasizing that this sort of knowledge is fairly common. Indeed, this is often how we would expect characters to behave if the other sorts of knowledge of gaps is absent. So, for example, suppose we alter the dinner party example so that in addition to Big John being invited, his twin, whom none of the other guests have ever met, is expected to make an appearance and is arriving solo. In this case Alice and Bob lack evidence that allows them to distinguish the way Big John looks from the way his twin looks. So, if someone shows up and it's not clear visually whether it's Big John or his twin, someone – either the guest or Alice and Bob themselves will presumably do something to clarify the situation. And at this point, they will acquire evidence that allows them to clearly distinguish Big John from his twin.

4 How robust must the knowledge of gaps be to avoid Lederman's argument?

In my last section, I made the case that knowledge of gaps is fairly common. But I haven't said much about how robust this knowledge of gaps is. And it might be thought that, unless it is fairly robust, Lederman will be able to repair his argument so as to evade my challenge.

4.1 An example

To develop this idea, let us return to the example of dinner party. As a reminder, here is the case:

> Dinner Party: I am throwing a dinner party for a small group of close friends. Everyone has arrived except for two: Big John and Little John. As their names suggest, they differ in height; Big John is taller than average (six feet) and Little John is shorter than average (five feet). The doorbell rings, so the two people nearest the door, Alice and Bob, open it and greet the new arrival.

Before Alice and Bob open the door, it seems they have knowledge of gaps. They know that the person ringing the bell will either be Big John or Little John, and thus that there are certain ways that he will not look. For instance, he will not look to be five feet six inches.

Let us focus in on their knowledge regarding Big John. Alice and Bob know various things about the way the doorbell-ringer will appear, if it is

indeed Big John. So, for instance, they know that there is a range of heights he might look. In addition, they know there is a range of heights he will not look. As per usual, right above the border is a height that I will call "barely lookable". This is a height such that, for all they know, he will look to be that height, but they know he will not look a little shorter.

Now consider the following possibility: Big John is the guest, but he is barely lookable to Alice. It might be thought that if he appears and looks this way, she will lose some of her knowledge of gaps. Perhaps, for instance, she will lose her knowledge that the doorbell-ringer is one of the Johns. After all, sometimes people show up at your house unexpectedly. And if the doorbell-ringer appears barely lookable to Alice, Alice might start to wonder if it is not Big John, but instead someone else. For instance, she might start to wonder whether the doorbell-ringer is a UPS delivery person. This motivates an idea like the following:

> Lost Knowledge At The Border: If the guest barely looks like Big John to Alice, then she doesn't know it's Big John at the door.

In short, even if we grant that Alice and Bob have knowledge of gaps prior to getting a good look at the person at the door, perhaps this knowledge will go away if they were presented with certain evidence, viz. the doorbell-ringer being barely lookable. Perhaps this is enough to motivate an analog to Interpersonal Ignorance:

> Doorbell-ringer Ignorance: For all r, if the doorbell ringer looks to be r cm tall to one of the agents, then for all that agent knows, the doorbell ringer looks to be 0.97r cm tall to the other.

After all, so the thought goes, suppose we have a case in which r is a height that is barely lookable. Suppose, e.g. that it is Big John at the door, but he looks much shorter than Alice is used to. In such a case, maybe Alice will start to wonder whether it is indeed Big John at the door. This will cause her to lose her knowledge of gaps – she will be forced to allow for the possibility that the doorbell-ringer is not one of the Johns, but instead someone else, such as a UPS delivery person. And then, for all she knows, such a person will appear even shorter – 0.97r cm tall – to Bob.[7]

4.2 How common is this sort of case

How common is this sort of case? That is, how common are cases in which one's knowledge of gaps will not survive evidence that one is at the border?

There are some reasons for thinking that this sort of case will be relatively infrequent. As I noted in the previous section, there are several different ways of acquiring knowledge of gaps. And there are reasons to think that in a typical case, multiple of these sources will be available, thus allowing one's knowledge to survive evidence that one is at the border.

[7]Note that this is consistent with the relevant intertemporal principle holding (and being known) providing the defeat happens (and is known in advance to happen) in the right way at the right time. Thanks to an anonymous referee for raising this worry. Thanks also to the same referee, in later comments, for encouraging me to add this footnote, which I have quoted nearly verbatim from their comments.

As I noted, one sort of knowledge of gaps involves perceptual clues that allow us to distinguish between borderline cases. It is natural to expect this sort of evidence to be present in a case like the one I just discussed. That is, if the person at the door were a UPS deliver person, there would presumably be perceptual clues that would allow Alice and Bob to distinguish him from Big John. Presumably the person would be carrying a package as opposed to a bottle of wine, wearing a uniform as opposed to a dinner jacket, refuse the hug that Alice offered, and so on.

Likewise, suppose that Big John is barely lookable (say he looks 5 ft 9 in) and Alice starts to wonder whether he looks even shorter to Bob (say 5 ft 6 in). Again, in a typical case, Alice would have evidence to rule out this possibility. For instance, as I noted, people will often probe if they find themselves in confusing or unusual circumstances. So, for instance, if Big John had looked way shorter than normal to Bob, Alice would presumably expect Bob to have already remarked on it.

In short, there are reasons to think that in many, and perhaps nearly all cases, there will be knowledge of gaps – of one form or another – that is robust enough to handle borderline cases.

5 Conclusion

If what I have argued in this paper is correct, then Lederman's argument challenging common knowledge fails. Not only do his key principles face counterexamples, they do so in a broad range of circumstances. And thus,

for all Lederman has shown, common knowledge is frequently available.

It is worth noting that my counterexamples do not focus on Lederman's original case, viz. SAILBOAT. Thus a question remains: what should be said about this case? Lederman had initially presented it as a paradigm case of public information. If what I have said in this paper is right, then several options present themselves. One is that my arguments can be extended to show that INTERPERSONAL IGNORANCE should be rejected even in this case. A second is that, if we manage to sharpen the case in a way that it avoids my arguments, it will no longer appear to be a paradigm case of public information. A final option is that it will manage to evade my arguments while maintaining its status as a paradigm case of public information. I myself don't have immediate views about which of these three possibilities is correct, but think it is well worth investigating further.[8]

References

[Lederman 2017] Lederman, Harvey. 2017. 'Uncommon knowledge.' *Mind*, 127(508):1069–1105.

[8]Thanks to an anonymous referee for pushing me to add this paragraph.