Quaestio facti. Revista Internacional sobre Razonamiento Probatorio
Quaestio facti. International Journal on Evidential Legal Reasoning
Sección: Ensayos
2025 I 9 pp. 1-19
Madrid, 2025
DOI: 10.33115/udg_bib/qf.i9.23151
Marcial Pons Ediciones Jurídicas y Sociales

ISSN: 2604-6202
Recibido: 06/03/2025 | Aceptado: 12/06/2025 | Publicado *online*: 29/06/2025

THE CONTENT OF LEGAL EVIDENCE

Timothy Williamson
University of Oxford and Yale University
timothy.williamson@philosophy.ox.ac.uk
Orcid 0000-0002-4659-8672

ABSTRACT: The boundary between old evidence and new evidence depends on how the content of evidence should be individuated. The paper explores conflicting pressures on the standard of individuation. Computational considerations and Frege cases of unknown co-reference both favour fine-grained individuation. The mathematical structure of probability theory, intensional and direct reference semantics, differences in format between verbal and perceptual evidence, the need for evidence to be transmitted from one context to another in memory and testimony, and the publicity of legal evidence all favour coarse-grained individuation. The paper argues that coarse-grained individuation is theoretically better motivated, and that pressures towards fine-grained individuation can be understood as resulting from our reliance on efficient but fallible disquotational heuristics for the ascription of agents' relations to propositions on the basis of their interactions with sentences expressing those propositions. Coarse-grained models can still be adapted *ad hoc* to understand more fine-grained phenomena.

KEYWORDS: Evidence; probability; form; content; intensionalism.

1. INTRODUCTION

We expect legal verdicts, judgments, and rulings to be supported by appropriate *evidence*. This relation of support is usually understood to hold primarily between the *content* of the verdict, judgment, or ruling and the *content* of the evidence, not just between sentence tokens or types, irrespective of such content. Two sentences, perhaps in different languages, may express the *same* content. Thus, contents are at a more abstract level than the sentences that express them. If one asks what these contents are, one may be referred to the philosophy of language. However, developments in modern philosophy of language raise contents to a level of abstraction so high that they may no longer seem fit for legal purposes. More specifically, contemporary semantic theories sometimes take two sentences that differ in cognitive significance to express the same content. But, for legal purposes, the cognitive significance of evidence is exactly what we are interested in. Thus, by ascending to the level of content, we risk abstracting away from the very thing we need. This paper will explain and explore such problems, for which there is no easy resolution.

All sorts of item are assessed against bodies of evidence: beliefs, claims, conjectures, theories, and propositions, as well as verdicts, judgments, and rulings. Which does our evidence *prove* or *refute*? Even if it is less decisive, which does it *confirm* or *disconfirm*?

For brevity, I will use "proposition" as a blanket term for whatever content is to be assessed against a body of evidence. Thus, we can ask how *probable* a proposition is on a given body of evidence.

When we update on new evidence, we ask a more dynamic question: does the new evidence make the proposition *more* or *less* probable than it was before?—or does it make no difference? If the ostensibly new evidence turns out to be just old evidence dressed up to look new, the probability should stay the same; updating was not needed. Some legal rules govern when new evidence may be introduced in a court case. If mere rewordings of old evidence counted as new evidence, such rules would create unnecessary problems. This suggests that evidence is understood as content, what is represented rather than what represents it.

The distinction between old and new evidence requires a standard for individuating evidence as the same or different. If we treat evidence as *encoded* by *representations* of some sort, we can ask: when do two representations encode the *same* evidence? Often, the representations are declarative sentences in contexts: the contexts are needed to fix the reference of indexicals and demonstratives such as "I", "he", "she", "they", "now", "then", "here", "there", "this", "that", and "those". If the evidence a sentence encodes in a context is just the content or proposition the sentence expresses in the context, the question becomes: when does a declarative sentence S1 in a context C1 express the same proposition as a declarative sentence S2 in a context C2?

Declarative sentences are not the only representations capable of encoding evidence. We gain evidence from sense perception too, encoded in our own perceptual states. A jury or judge may be shown photographs of the crime scene in evidence, and they may also gain evidence of what the defendant is physically capable or incapable of, by seeing them in the dock. Perhaps such perceptual evidence could not be adequately encoded in a declarative sentence, even a very long one. Contrary to what is sometimes assumed, this does not mean that perceptual evidence is non-propositional. Perceptual appearances can be veridical or non-veridical, accurate or inaccurate, whether or not they can be adequately put into words. If someone looks to me to be a certain way, that visual appearance is veridical if and only they are that way; otherwise, it is non-veridical. Such a distinction between, in effect, truth-conditions and falsity-conditions is enough for the visual appearance to have a proposition as content, whether or not it can be adequately expressed in words.

Even when written documents are submitted in evidence, the words in them do not exhaust their evidential significance. Various propositions *about* the documents matter do. For example, was it sent by post or by email? If it is handwritten, whose handwriting does it resemble? And so on.

Not all propositions constitute evidence. It is standardly assumed that any proposition inconsistent with the evidence is false, which presumes that all evidence is true. Thus, a false declarative sentence does not encode evidence. Of course, a witness may give false testimony in court, for example by saying "Mr Brown spent the night with me" (thereby providing him with an alibi) when Mr Brown did *not* spend the night with her. But the evidence that her testimony provides is not the false proposition that Mr Brown spent the night with her; rather, the evidence it provides is the true proposition that *she testified that* Mr Brown spent the night with her, or something like that.

Although truth is necessary for being part of the evidence, it is of course not sufficient. A truth is part of *your* evidence only if you *have* it, in other words, you stand in an appropriate epistemic relation to it. Elsewhere, I have argued that the appropriate epistemic relation is *knowing* (Williamson, 2000, p. 184-208). Likewise, a truth is part of a *court's* evidence only if the court has it, in other words, the court stands in an appropriate epistemic relation to that truth, arguably, a knowledge-like relation.

This paper focuses not on whether being true or known is required for being evidence, but instead on the nature of the evidence itself, specifically, on how it is individuated, and, relatedly, whether the identity or distinctness of one's evidence must be *transparent* to one. Suppose the jury hears some evidence, but it is opaque to them whether it is genuinely new evidence, or just old evidence in a new guise. Then it may also be opaque to the jury how it should update on this evidence: whether it enhances the support for some proposition, so their confidence in it should increase, or whether it is just some old evidence supporting in a new form, so their confidence in the proposition should stay the same. Thus, individuating evidence in ways opaque to us may render us unable to work out how to comply with an epistemic

or legal norm, such as one of proportioning one's confidence in a proposition to the evidence. What our evidence *appears* to support may conflict with what it really does support, in ways we are in no position to adjudicate, since we do not know which is reality and which appearance.

Someone might respond: if so, at least for legal purposes, our evidence had better be individuated in a way that *is* transparent. But perfect transparency may be impossible, for general epistemological reasons (Williamson, 2000, p. 93-113). The most systematic framework for thinking rigorously about evidence may imply gross failures of transparency, not just slight failures. This paper works through these issues.

2. HOW IS EVIDENCE INDIVIDUATED?

The central choice-point for individuating evidence is how fine-grained or coarse-grained the individuation is to be: the finer the grain, the more it takes for two representations to encode the same evidence; the coarser the grain, the less it takes for them to encode the same evidence.

To get more precise, we need a more informative conception of evidence. We will assume that evidence is *propositional*, in the sense of being true or false in possible circumstances. As already noted, that it is propositional does not mean that it is linguistic or quasi-linguistic in form; a sensory perception may be veridical or non-veridical, and have a content true or false in various possible circumstances, without being linguistic or quasi-linguistic in form. Evidence must be propositional, since propositions are compatible or incompatible with it; compatibility and incompatibility are relations between propositional items. In what follows, each piece or body of evidence is treated as a proposition. I have argued for the propositionality of evidence in detail elsewhere (Williamson, 2000, p. 194-200). Thus, the individuation of evidence boils down to the individuation of propositions.

The most salient coarse-grained approach to individuating propositions is *intensionalism*, on which propositions are identical if and only if they are necessarily equivalent, in other words, they cannot differ in truth-value. In each possible world, they are both true or both false. For present purposes, we take this impossibility to be metaphysical, rather than logical or epistemic; not much here will turn on the exact understanding of "metaphysical".

That identical propositions are necessarily equivalent is trivial, for each proposition p is necessarily equivalent to itself, so by Leibniz's Law of identity, if p and q are identical, p is necessarily equivalent to q. What is controversial in intensionalism is the converse claim, that necessarily equivalent propositions are identical, which more fine-grained approaches deny. The latter are *hyperintensional*, holding that some pairs of propositions are distinct but necessarily equivalent.

Of course, different philosophers may use the theoretical term "proposition" with different senses, perhaps differing in the associated fineness of grain. For present pur-

poses, the question would then be: which of those senses is best-suited to understanding the function of evidence, at least in a legal context? A key issue for individuating evidence is whether there are distinct but necessarily equivalent pieces of evidence. The next section considers this issue in the setting of probability theory, since its rigorous mathematical framework is conducive to clarity. Although evidence in legal contexts is not usually discussed in such precise quantitative terms, it sometimes is. An expert witness may testify about the probability that a trace of DNA found on an item of clothing at the murder scene came from the defendant. Thus, a theory of legal evidence should at least be able to make sense of probabilistic relations of evidence to propositions.

3. PROBABILITY SPACES AND COARSE-GRAINED INDIVIDUATION

As just noted, one role for evidence is to provide a basis on which to assess the probability of propositions. We often need to ask how probable a proposition is *on our evidence*[1]. An account of evidence that cannot make sense of probabilities on our evidence is to that extent dysfunctional. That is not to say that *all* evidential relations are probabilistic, just that some of the most important and best-understood of them are. A theory of evidence that does not engage properly with probability theory is inadequate to the modern world.

The cognitive power of probabilistic thinking derives from its mathematical structure. That structure unfolds within a mathematical framework made formally precise by the mathematical definition of a *probability space*. That definition reveals how evidence is individuated in standard probability theory.

A probability space is based on an underlying set Ω, whose members are called "outcomes". They are sometimes informally conceived as the possible outcomes of an experiment, though that way of thinking is not very relevant in a legal context. What matters is that they are mutually exclusive and jointly exhaustive; each outcome is meant to be maximally specific in all relevant respects. Subsets of Ω are called "events." The probability space also includes a probability measure Pr, mapping events to real numbers in the closed interval [0, 1], conceived as their probabilities. The impossible empty event {} has probability 0; the necessary event Ω has probability 1. The probability of the union of any finite or countably infinite collection of mutually incompatible events is required to be the sum of the probabilities of those events. Consequently, the probabilities of any event and its complement event sum to 1. To handle combinatorial difficulties arising if Ω is infinite, the standard definition of a probability space avoids requiring probabilities to be assigned to *all* subsets of Ω (with the device of an "σ-field»)[2]. For simplicity, we can ignore that restriction;

[1] On such evidential probabilities see Williamson (2000, p. 209-237)

[2] See Grimmett and Stirzaker (2001, ch. 1).

it makes no difference to the issues discussed in this paper. In a standard legal setting, the relevant possibilities can presumably be treated as finite in number.

Structurally, a probability space is like a simple Kripke model for modal logic. Instead of "worlds", one speaks of "outcomes"; instead of "propositions", conceived as sets of worlds, one speaks of "events", conceived as sets of outcomes; instead of "W", the set of all worlds, one speaks of "Ω", the set of all outcomes. To complete the analogy, one can treat the universal relation, which any outcome has to any outcome, as an accessibility relation R for the probability space, so any world is possible from the perspective of any world.

We also want to discuss the probabilities of events *conditional* on an event E, where we in effect exclude all outcomes outside E and recalibrate: the conditional probability of an event H on E is the proportion (weighted by Pr) of the E-region of the probability space that overlaps the H-region. When Pr(E) is nonzero, we can therefore define the conditional probability Pr(H|E) of H on E as the ratio Pr(H∩E)/P(E). When Pr(E) = 0 but E is non-empty, we may still be able to make sense of the conditional probability Pr(H|E) by treating conditional probability as basic and undefined, though it must still satisfy the ratio equation when Pr(E) is nonzero. We can conceive E here as new information on which to update. Then Pr(H|E) is the probability of the proposition H *on* the evidence. For present purposes, we make the simplifying assumption that whenever E is evidence, Pr(E) is nonzero (it is not crucial to the arguments of this paper). The probability space gives us just what we need for evidential probabilities.

To reach that satisfactory conclusion, we had to treat both the proposition H and the evidence E as events, sets of outcomes, in order for their unconditional and conditional probabilities to be defined in the model. In effect, events in a probability space are individuated intensionally: E = F whenever, for each possible outcome o, E is true in o (o∈E) if and only if F is true in o (o∈F). "Possible" here corresponds to whatever modal standard is built into the probability space, which need not be metaphysical possibility; many structural effects are the same whatever flavour of modality is involved. In particular, for any such modality, every equation that follows from the axioms for Boolean algebra corresponds to an identity of events. For example, for any two events E and F, E = E(F∩E), even if what E∩(F∪E) is "about" (both E and F) seems to differ from what E is "about" (just E). Consequently, the conditional probabilities P(E|E) and P(E|(E∩F)∪E) are both 1, even though the entailment from the right to the left of the conditional bar "|" is easier to see in the first case than in the second. The same applies when the Boolean equation is vastly harder to verify: in general, computational differences between two such Boolean-equivalent expressions disappear at the level of events. Thus, the mathematical framework of probability theory automatically enforces much of the flattening effect of intensionalism (Builes (2020) discusses other apparently hyperintensional probabilistic phenomena).

Some philosophers will conclude that standard probability theory needs revising. Such methodological advice is potentially disastrous. No small tweak would

suffice to respect computational differences, nor is it clear what larger ones would be needed. A major revision of mathematical probability theory could easily make it computationally intractable, or pitifully weak; either way, it would undermine the scientific value of the theory. In a legal setting, computationally intractable probabilities are likely to be useless.

A subtler move is to keep the mathematics unchanged, but modify the interpretation of natural language over the probability space. A sentence "S" of natural language such as "The DNA came from the defendant" is mapped to an event ["S"], a subset of Ω. Normally, it is taken for granted that the linguistic operations of conjunction, disjunction, and negation are to be interpreted by the set-theoretic operations of intersection, union, and complementation in Ω, respectively. For instance, ["S and T"] = ["S"]∩["T"]. This is no mere convention, since those linguistic operations are used to define the corresponding set-theoretic operations. For instance, the set E∩F is defined as the set to which an outcome belongs if and only if it belongs to E *and* belongs to F. But other approaches are conceivable. In particular, we could define ["S"] as the set of outcomes *o* in which the sentence "S" is true *as interpreted in o itself* (this is like the use of diagonal propositions to interpret some beliefs compatibly with intensionalism in Stalnaker (1999)). Thus, to give a toy example, if "and" in *o* is interpreted as *or*, *o* belongs to ["S and T"] if and only if *o* belongs to ["S"] *or o* belongs to ["T"]. Thus, if *o* belongs to ["S"] but not to ["T"], then *o* belongs to ["S and T"] but not to ["S"]∩["T"], so ["S and T"] ≠ ["S"] ∩ ["T"]. Consequently, sentences of natural language equivalent to each other in propositional logic may be mapped to different events in the probability space. The flattening effect of intensionalism is thereby avoided.

But the reinterpretation strategy only shifts the wrinkle in the carpet. Although the pure mathematical theory of probability remains intact, all the complications return as soon as we try to apply it to a problem stated in natural language. We lose our understanding of the structural relations amongst the events under discussion, and so cannot make proper use of the mathematics. The reinterpretation was in any case ill-motivated, because it models a situation where the logical connectives of natural language are not clearly understood, but such semantic confusions are not the main source of our computational limitations[3]. Mathematical probability theory provides a forceful case for the coarse-grained intensional individuation of evidence.

4. FORMATS OF EVIDENCE

Intensionalism allows us to model the content of evidence as a set of possible worlds, but it does not imply that we must encode that content by a list of labels of

[3] See Williamson (2024, p. 244-248)

possible worlds or in any other specific way. It is largely neutral on the form in which the content is expressed.

Since evidence is propositional, it is easily taken to be expressed in the form of declarative sentences. But that form is inappropriate for much evidence from present or past perception. Apart from the direct role of non-verbal evidence in court, the court may also have to weigh non-verbal evidence from outside the court. For example, the police report that a witness picked out the defendant from an identity parade as the man, previously unknown to her, whom she saw standing outside the victim's house on the night of the murder. The court must determine how reliable her identification was; to do so, it must assess the evidential value of her memory several weeks later of the person she saw in poor light from the other side of the road. How good a look did she get at him, and how clear was her subsequent memory? The memory on which her identification was based was primarily non-verbal. She may remember how he looked; she may even articulate the content of her memory to herself by thinking "He looked like *that"*, where "*that"* is her memory demonstrative, pointing to how he looked (as far as she remembers). The trouble is that most of the cognitive work is done by the perceptual memory itself, at which the verbal memory demonstrative is gesturing. Without the connection to her memory, the bare sentence is of little use. Thus, the court is forced to engage with questions about pieces of non-verbal evidence.

Some legal evidence is verbalized; some is not. Evidence is not restricted in form to any one natural language, or to any one sense modality. To think in probabilistic terms about the underlying evidential relations among all the propositions and pieces of evidence in play, we must notionally map all these representations in diverse formats onto propositions (events, subsets of Ω) in a single probability space, to render them mutually commensurable. The same applies when we have to assess the logical or epistemically modal relations among those propositions and pieces of evidence, though the set of which the propositions are subsets may be called "W" rather than "Ω". Thus, propositions are a "common currency" for cognitive interaction, or, more accurately, abstract exchange values. Propositions are not themselves representations. They are nothing like sentences or sensory images. For propositions conceived as sets of worlds, worlds too are non-representational; they are better understood as something like ways for things to be. But this coarse-grained approach to evidence and probability faces some serious challenges, which the next section starts to explore.

5. FREGE PUZZLES OF LEGAL EVIDENCE

A predictable challenge to the coarse-grained framework for evidence and probability comes from Frege puzzles, like that of Hesperus and Phosphorus.

For the sake of argument, let us pretend that Dr Jekyll, the pillar of the establishment, and Mr Hyde, the brutal killer, are real people, indeed, the *same* real person. Mr Hyde is on trial for murder, under that name. The lawyers in the case and the

members of the jury are rational people, but have no idea that the names "Dr Jekyll" and "Mr Hyde" co-refer. They have encountered both names, and associate them with very different descriptions. Clearly, witness testimony given using the name "Mr Hyde" will be treated, and in some sense *should* be treated, very differently from exactly similar testimony but with the name "Mr Hyde" replaced throughout by the name "Dr Jekyll" (the latter testimony would be dismissed as irrelevant to the case). In a subjective Bayesian framework, one would expect this difference to be reflected in the jury's credences (subjective probabilities or degrees of belief) for the event expressed by the sentence "The defendant is guilty" after updating on the new evidence. Yet, on the predominant semantics for proper names in contemporary philosophy of language, a proper name is a *directly referential* singular term: it contributes only its bearer to determining the proposition expressed by a sentence in which the name occurs, used in the normal way (not mentioned within quotations marks or the like), according to the standard compositional semantics for the language. Thus, for example, the sentences "Mr Hyde was present" and "Dr Jekyll was present" express the same proposition, the same content, in a given context; likewise for the sentences "Mr Hyde is guilty" and 'Dr Jekyll is guilty. Consequently, if evidential and probabilistic relations depend only on content, substituting one sentence in a pair for the other makes no difference to those evidential and probabilistic relations. But that conclusion seems absurd.

Of course, the work of the court can include identifying potential Frege cases and resolving them. But that is no guarantee of success. Even if the court is better at it than any single individual present, a Frege case may still go unidentified. If things go badly, the legal proceedings may come to a conclusion in ignorance of the co-reference between the two names.

Frege puzzles are no mere artefacts of natural language semantics, but genuine cognitive phenomena obviously relevant to our uses of evidence, including legal evidence. Theories of evidence cannot in good conscience just brush Frege puzzles aside. How should they be understood?

An alternative semantics for proper names would not meet the generality of the problem, since Frege puzzles can arise for any kind of singular term, for example, perceptual demonstratives: "she" (with the reference fixed by visual memory from one occasion) and "she" (with the reference fixed by visual memory from another occasion) may non-obviously refer to the same woman. All kinds of directly referential singular term raise the same problem.

At this point, some philosophers will naturally suggest a Fregean solution to Frege puzzles, appealing to a distinction between sense and reference. Thus, the names "Dr Jekyll" and "Mr Hyde" have the same referent but different senses. Since Fregean senses are individuated cognitively, at first sight that approach looks better equipped to capture the evidential difference between co-referential singular terms.

The trouble is that Frege puzzles can arise from mere differences of word, even when the words are paradigms of synonymy, as Saul Kripke (1979) pointed out. I will use one of his examples. The words "furze" and "gorse" are natural kind terms for the very same kind of shrub. They are not associated in English with different descriptions, except trivially with metalinguistic descriptions, such as "shrub whose name begins with an 'f'" or "shrub whose name begins with a 'g'". Although an individual speaker may idiosyncratically associate them with different non-metalinguistic descriptions, that is not a feature of their linguistic meaning. But even a native speaker who understands both terms by normal standards and is reasonably confident that they co-refer may still be *slightly* less confident of "Furze is gorse" than of "Furze is furze", and so rationally assign a slightly lower probability to the former than to the latter. It is not irrational to be less than perfectly confident that one has not confused the meanings of the terms, even if one has *in fact* not confused them. For such a speaker, learning "Mary saw no furze on that hill" may be slightly better evidence for "There is no furze on that hill" than learning "Mary saw no gorse on that hill", with corresponding differences in their rational credences. These differences in cognitive significance arise simply from the difference of word.

Some differences in cognitive significance cut finer than the words themselves, for example when the words are context-sensitive. The word "she", used as a perceptual demonstrative, differs in cognitive significance depending on the speaker's state of perceptual attention. Thus, differences in cognitive significance cut at least as finely as ordered pairs of a linguistic expression and a context.

The Fregean explanation of a Frege puzzle depends on a difference in sense without a difference in reference; that is what motivates the distinction between sense and reference. Thus, since a mere difference of word gives rise to a Frege puzzle, the Fregean explanation implies that the mere difference of word gives rise to a difference in sense. But, if the senses of words are individuated as finely as the words themselves (paired with contexts), the motivation for postulating a separate category of senses is undermined: one might as well make do with the words themselves (paired with contexts), and cut out senses as a redundant layer of middle management between language and the world. On this view, a representation does not *encode* evidence; it *is* the evidence.

Such an approach can be implemented for testimonial evidence. As already noted, when the witness testifies in court "Dr Jekyll was present", what becomes part of the evidence is not the proposition that Dr Jekyll was present, since the witness may be lying or mistaken, rather than the proposition that the witness testified that Dr Jekyll was present. One could cut the grain more finely by ruling that what primarily becomes part of the evidence is not such an indirect speech report but a direct speech report, the proposition that the witness testified "Dr Jekyll was present". If the witness had instead testified "Mr Hyde was present", the proposition that the witness testified "Mr Hyde was present" would instead have become part of the evidence. Those two direct speech reports are by no means mutually truth-conditionally equi-

valent, since the former entails that the witness used the name "Dr Jekyll" and does not entail that she used the name "Mr Hyde", whereas the latter entails that she used the name "Mr Hyde" and does not entail that she used the name "Dr Jekyll". By contrast, there is no such barrier to the mutual truth-conditional equivalence of the two indirect speech reports, since their truth does not depend on providing the very words used by the witness. Does this metalinguistic approach fine-grain the evidence in the required way?

The trouble is that, if we treat testimony as the mere uttering of words, we lose its evidential relevance to the verdict. An account that ignores the proposition the witness expressed by uttering the words "Dr Jekyll was present' is ill-equipped to make sense of the evidential relations at issue. But once we reintroduce propositions to understand the testimony"s evidential bearing on the case, we are back with the problem that "Dr Jekyll was present" and "Mr Hyde was present" express the *same* proposition.

Testimony in court is a special case of transmitting evidence across linguistic contexts. An obvious example is that the word "I" refers to the witness when she uses it but not when anyone else does. Often, when evidence is transmitted across contexts, what matters is preserving *reference*, not preserving anything like Fregean sense: the Fregean approach is much less suited to the legal setting than one might initially expect. The next section will reinforce that point.

6. TRACKING REFERENCE

In moving across contexts, we update to keep track of things, otherwise we may lose evidence. This process can be used to illustrate how identity of reference, not identity of sense, let alone identity of representation, is often what matters in our handling of evidence.

Here is an example. A kidnapped hostage is trying to keep track of what has been happening to him. He remembers: "Today I heard raised voices in the corridor". To preserve what matters in that piece of evidence the next day, he needs a different sentence: he must remember "Yesterday I heard raised voices in the corridor". The word "yesterday" as the hostage used it one day has nothing like the same sense as the word "today" as he used it the previous day; it simply has the same reference.

That example involved the same agent at different times. Here is a similar example that involves different agents at the same time. A kidnapper tells the hostage: "You'll be released soon". To capture what matters in that piece of evidence, the hostage needs a different sentence: he must recognize "I'll be released soon", or at least "The kidnapper said that I'll be released soon". The word "I" as used by the hostage has nothing like the same sense as the word "you" as used by the kidnapper; it simply has the same reference.

In both examples, when things go well, the intended recipient comes to *know* the proposition expressed by the new sentence, otherwise it could hardly form part of

the recipient's evidence, on the speaker's knowledgeable testimony[4]. To acquire that knowledge, the hostage needs something like implicit knowledge of the coreference between "today" uttered one day and of the coreference between "you" uttered by the kidnapper and "I" uttered by him. But all of that is quite consistent with the coarse-grained individuation of the evidence itself, as it is transmitted from one context to another.

Similar phenomena occur for perceptual and memory demonstratives. Here is an example. A robber buries her loot in a forest, intending to return and dig it up again once the coast is clear. She knows "I buried it by this tree", using "this tree" as a perceptual demonstrative. With a memory image of the tree, she may preserve the knowledge "I buried it by that tree", using "that tree" as a memory demonstrative, once the tree is no longer in sight. But her vivid memory image of the tree will be of little use if she cannot remember where in the forest the tree is, or how to reach it. Of course, preserving the right reference is not enough. The user must in effect *know* that reference has been preserved. Moreover, the reference must be preserved in some relevantly usable form. But that does not mean that the usable form must be built into the content.

The transmission of evidence across epistemic contexts or perspectives is crucial to many of its uses: from one time to another in memory, from one agent to another in testimony. Evidence is not typically private; it does not consist in the mental states of a particular individual at a particular time. Obvious examples come from evidence in scientific, political, social, and of course legal contexts. The evidence on which the court bases its verdict should be shared by all the relevant parties in court, otherwise justice is not seen to be done. Thus, individuating evidence in terms of a single epistemic context or perspective is inappropriate for central functions of evidence.

Of course, we can treat the court itself at a given time as constituting a context in its own right and as having its own collective epistemic perspective. But that perspective must not be confused with the epistemic perspective of any one individual present. For example, a witness may know whether she is lying while the court does not. More generally, evidence available to the court will do little good unless it is also available to the relevant individuals in court, which already raises the issue of cross-perspectival epistemic availability.

A related example from the social sciences comes in Anna Mahtani's work (2017; 2021; 2023) on the interpretation of probability in welfare economics. She shows convincingly that although the probabilities used to define an agent's expected welfare under a policy are officially interpreted as subjective credences, the definition has trouble with Frege cases, where the same agent is presented under several non-equivalent guises, so the agent's expected welfare should be guise-dependent, but

[4] For the equation of one's total evidence with one's total knowledge see Williamson (2000, p. 200-208)

the theory treats it as guise-independent. Mahtani makes an impressively careful attempt to adapt theories of welfare economics to a Fregean interpretation of the probabilities, but the result is a heavy increase in theoretical complexity and, as she admits, no fully satisfactory solution. A better approach may be to understand the probabilities more objectively, as guise-independent, making Frege cases irrelevant, and thereby retain the (slightly reinterpreted) original definitions and theories (Williamson, 2024, p. 212-221)[5].

This perspective-independent approach to evidence is consonant with now-standard frameworks for the semantics of natural languages, as already described, on which singular terms are directly referential. Contrary to what one might first have expected, semantic theories with nothing like Fregean senses individuated by cognitive significance may be better adapted than those with Fregean senses to cognitive purposes, once account is taken of the diachronic and social aspects of cognition.

7. IS EVIDENCE TRANSPARENT?

The direct reference semantics tends to make agents' credences (degrees of belief) opaque to the agents themselves. Is that a problem?

Imagine quizzing members of the jury on their credences, and in particular on their conditional credences. For the sake of argument, we may pretend that they understand the mathematics of probability perfectly, and are very careful not to confuse using a word with mentioning it. You make it quite clear to them that you are interested in their conditional probabilities about non-metalinguistic matters. First, you ask them "What is your credence that Mr Hyde did the murder, given that Mr Hyde did the murder?" They give the trivially mathematically correct answer "1". Then you ask them "What is your credence that Mr Hyde did the murder, given that Dr Jekyll did the murder?" Their first instinct may be to offer 0 or some number close to it, since they assume that the two names refer to different men. On reflection, however, it may occur to them that "Dr Jekyll" and "Mr Hyde" could be names of the same man. They may even accept a direct reference semantic theory, and appreciate its potential relevance to the case at hand. But they are still in no position to answer "1" to the second question, since, for all they know, the two names do *not* co-refer. If they do accept the direct reference theory, they may have to answer "I don't know": from their perspective, their conditional probability may be 1 (if the two names co-refer) and it may be close to 0 (if the two names do not co-refer). On the other hand, if they accept something more like a Fregean semantic theory, they may give an answer close to 0, though that answer will be incorrect if the direct reference theory is in fact right. If one tries to measure their credences by their betting behaviour instead, the results will be similar, because the bets are still offered to them in linguistic form.

[5] For more on Frege cases for credences see Braun (2016).

In short, given the direct reference theory, in Frege cases, agents are at best ignorant and at worst mistaken about their own conditional credences on non-metalinguistic matters. Such credences are not in general transparent to the agent.

Switching to credences on metalinguistic matters does not automatically restore transparency. Just as names of people can non-obviously co-refer, so can names of sentences. To use another of Kripke's examples, someone may not realize that "Paderewski was Polish" uttered in a conversation about pianists is the same sentence as "Paderewski was Polish" uttered in a conversation about politicians (Kripke, 1979). Moving the objects of credence into the meta-language only postpones the problem.

Opponents of direct reference semantics may regard all this as constituting a strong argument for a more Fregean approach. We have already seen some of the grave difficulties for Fregean semantics. In particular, one should not imagine that Fregeanism can restore the transparency of evidential relations to the agent. The attempt to make them transparent was what collapsed difference of sense into difference of word (which is itself not fully transparent, as the Paderewski case). Despite its initial promise, the use of differences in cognitive significance as a guide to differences in meaning has turned out badly for philosophers of language. Epistemologists and philosophers of law should not repeat their mistake.

Coarse-grained semantics is far from the only obstacle to the wishfully imagined transparency of evidence. I have argued elsewhere on general epistemological grounds that *no* non-trivial condition is such that, whenever it obtains, one is in a position to know that it obtains[6]. The argument applies in particular to the condition that a given proposition is part of one's evidence, and to the condition that the proposition is *not* part of one's evidence. Coarse-grained semantics adds a further class of (especially vivid) examples of non-transparency.

Should we "re-engineer" the "concept" of evidence, that is, change the meaning of the word "evidence" and of words with similar meanings in English and other languages, in order to make its application more transparent? That would involve making it stand for something more fine-grained. Given the arguments of previous sections, that would in turn require making it more sensitive to the form of representations, and so metalinguistic. We could indeed gerrymander and stipulate such a new metalinguistic meaning for the word "evidence", at least for theoretical contexts over which we had some control, but the benefits would be unlikely to outweigh the costs. By the general argument mentioned above, it would still not achieve full transparency. Indeed, it promises to be of much less theoretical and practical use than the word "evidence" as currently used. More specifically, the arguments about the capacity to interface with the mathematical structure of probability theory, the diversity of formats in which evidence is presented, and the diachronic and social availability of evidence, do not just reflect the way the word "evidence" happens to be currently used: they concern our more general epistemic and cognitive needs, to

[6] For this anti-luminosity argument see Williamson (2000, p. 93-113)

which the coarse-grained engineering is already well-adapted. It would make the term more cumbersome and difficult to use, the preserve of a philosophical elite with time on their hands. It would be as misguided and ineffective as most philosophers' proposals to "re-engineer concepts."

8. DISQUOTATIONAL HEURISTICS

If the semantics of our language is really coarse-grained, why does it *seem* so fine-grained to us? Why is it so tempting to judge that, for example, one could have good evidence that Mr Hyde did the murder without having good evidence that Dr Jekyll did the murder? The thought process is not hard to reconstruct. Faced with a question of the form "How good is S's evidence that P?", by default we try to put ourselves in S's epistemic position, asking ourselves "How good is my evidence that P?", answering as best we can from that perspective, and then treating that as an answer to the original question. From a perspective without effective access to the identity "Dr Jekyll is Mr Hyde" (in those very words), a rational person might well answer "I have good evidence that Mr Hyde did the murder" and "I do not have good evidence that Dr Jekyll did the murder" in response to the corresponding questions. We then treat those answers as *true* of the epistemic perspective at issue. This presumes that the first-person present-tense perspective is decisive for assessing evidential relations. Pre-theoretically, such judgments of evidential relations are not *obviously* wrong even in Frege cases; we have no obvious incentive to do better, and in any case no idea *how* to do better. The problems with a Fregean approach to semantics have only gradually become clear through decades of reflection in the philosophy of language. Direct reference semantics highlights one way in which judgments about evidential relations in a given epistemic position can sometimes be corrected from outside that position, but there are many other ways too. A court of law does not give witnesses the final say on how probable a proposition is on their evidence.

When we are led into thinking of evidence as fine-grained in content, we are in effect guided by a *heuristic*, an efficient but fallible way of answering a question. It is not just a specific heuristic dedicated to the word "evidence". Rather, there is a far more general kind of *disquotational* heuristic, by which one attributes to an agent S a cognitive relation to the proposition that P on the basis of S's actual or potential interaction (or lack of it) with the *sentence* "P", where the relevant kind of interaction depends on the relation to be attributed (Williamson, 2024). Since two sentences can express the same proposition, and S may interact differently with those two sentences because S fails to treat them as expressing the same proposition, we may end up equivocating as to whether S has the given relation to the proposition they both express. Having good evidence for a proposition is one of many such relations.

Describing the operation of the heuristic as "disquotation" is strictly misleading, for in the simplest case the user only has to repeat the speaker's words; there are no quotation marks to strip off. Often, something more is required, such as converting

the speaker's first-person pronoun to the user's second-person or third-person pronoun or a proper name, and likewise for temporal indexicals such as "now" and "today". But that is usually done unreflectively, with minimal effort, and is in any case a departure from disquotation, imposed by the nature of indexicality. Nevertheless, I will continue to use the word "disquotational" for want of a better.

Arguably, such a disquotational heuristic is active in Kripke's notorious puzzle about belief, where Pierre interacts differently with the English sentence "London is pretty" and the French sentence "Londres est jolie", even though they express the same proposition (Kripke 1979). In that case, the propositional attitude is belief, and the relevant kind of interaction is assent; the specific mindreading heuristic there is for judging what people believe by what they say. Kripke himself does not present his disquotational principles for belief attribution as merely heuristic in status, but they are arguably better understood as restricted cases of such a heuristic (Williamson, 2021; 2024, p. 181-192).

Heuristics are studied extensively in contemporary psychology. They may be seen in a negative light, as forms of irrationality, and described as "cheap and dirty", in the tradition of Daniel Kahneman (Kahneman *et al.*, 1982), or in a more positive light, as forms of bounded rationality, and described as "fast and frugal", in the tradition of Gerd Gigerenzer (Gigerenzer *et al.*, 2011). But both traditions agree that fallible heuristics play a major role in human cognition. It is therefore not at all surprising that we rely on fallible heuristics in attributing cognitive relations. Indeed, the work of Kahneman and Tversky on heuristics and biases has been applied to understanding judicial decisions (Peer and Gamliel, 2013).

Some heuristics are built into our perceptual systems; for example, colour boundaries in the visual field are used as a guide to the edges of three-dimensional objects; that heuristic explains why camouflage is effective. Disquotational heuristics are presumably built into our mindreading system, which seems to be humanly universal, like perceptual heuristics. In both cases, we typically rely on them unreflectively, not consciously aware that we are relying on a heuristic, though we can become aware of our reliance on them through conscious reflection. Of course, our mindreading capacity uses far more than disquotational heuristics, which are specific to mindreading on the basis of *linguistic* behaviour. We also attribute propositional attitudes on the basis of non-linguistic behaviour to non-human animals and very young children, and indeed to human adults too, to explain that behaviour.

Arguably, just as visual illusions result from heuristics built into the human visual system, so philosophical paradoxes result from heuristics built into more general human cognitive systems (Williamson, 2024). By contrast, not all our heuristics are humanly universal; a bias may be culture-specific. Some heuristics may also be conscious, such as a doctor's idiosyncratic heuristic for interpreting X-rays.

Our reliance on fallible heuristics in perception does not justify scepticism about our knowledge of the external world; likewise, our reliance on fallible heuristics in mindreading does not justify scepticism about our knowledge of other minds. But

we do need to take a somewhat more critical attitude to our data. In particular, we should be wary of dismissing powerful explanatory frameworks just on the basis of confident judgments—for example, about Frege cases—that can be explained as products of efficient but fallible heuristics which are predicted to go wrong in just such cases.

9. COARSE-GRAINED CONTENT AND EPISTEMIC NORMATIVITY

Given the non-transparency of evidence to the agent, we can expect epistemic norms to be less than fully *operational*: one is not always in a position to know whether one is complying with a given epistemic norm. When one is not in a position to know what one's evidence includes, one is also not in a position to know which beliefs are proportionate to one's evidence. This is a general feature of the human cognitive predicament, and indeed of the cognitive predicament of just about any finite creature. It is something we have to learn to with. Attempts to get around it by fiddling with the definition of "evidence" or the wording of the norms are futile; they fail to restore transparency and they lead to bad epistemology, as I have argued elsewhere (Williamson, 2000). That should be no surprise in a legal setting, where the epistemic bearing of given evidence is often in dispute. Complying with evidential norms is often *difficult*, and rational agents are not always in a position to know whether they have succeeded.

In particular cases, we can often find more or less *ad hoc* ways to finesse the difficulty. For example, we can often model Frege cases within the standard framework of mathematical probability framework by tweaking the semantics of direct reference, since the overall mathematical framework does not by itself enforce that semantics. We can model the epistemic predicament of the jury by treating the names "Dr Jekyll" and "Mr Hyde" as non-rigid designators, varying in reference from one outcome to another independently of each other, so that the sentences "Dr Jekyll was present" and "Mr Hyde was present" express distinct events and differ in truth-value at some outcomes. Alternatively, one can achieve a similar effect by treating the names as directly referential and so rigid singular terms, each constant in reference across all outcomes, but referring to *distinct* objects. Although such models do not capture the genuine metaphysical possibilities, they may often simulate the epistemic predicaments well enough for the purpose at hand. They may clarify for us the thinking, and especially the decision-making, of an agent who is unclear about the relevant identities. By such *ad hoc* means, we can make sense of various ways of deviating from the real evidential relations.

There is a temptation to try to generalize and unify all such modelling devices into a "universal" framework for modelling evidential relations, or at least evidential probability, with the flexibility to include all such deviations. That is a methodologi-

cal error, and a dead end: the result will be too weak to be useful. It is better to work with a strong theory of genuine evidential relations, and to understand the hard cases as local deviations from the latter, perhaps the results of illusions of genuine evidential relations. Some modelling assumptions are theoretically fruitful without showing anything about the nature of the target phenomena. For instance, the highly successful Lotka-Volterra model of predator-prey population dynamics uses differential equations, and so treats the number of predators (say, foxes) and the number of prey (say, rabbits) as varying continuously, even though we know that the numbers are discrete (Weisberg, 2013).

In theorizing about evidential relations, we do better to retain the systematic explanatory power of standard probability theory and standard intensional semantics, while acknowledging that to understand the messy complications of real cognitive life we must often track fine-grained differences at the level of representational form rather than represented content, in ways adapted to the case at hand. This kind of *ad hoc* adaptation to the messy complications of real life makes theoretically minded philosophers uneasy, but it is surely no stranger to legal practice.

ACKNOWLEDGMENTS

I thank participants in the 2024 Girona workshop on "Knowledge and Legal Standards of Proof" for valuable discussion of a talk from which this paper derives, and two anonymous referees for this journal for helpful comments on a draft of this paper.

REFERENCES

Braun, D. (2016). The objects of belief and credence. *Mind*, *125*, p. 469-497. *https://doi.org/10.1093/mind/fzv137*

Builes, D. (2020). A paradox of evidential equivalence. *Mind*, *129*, p. 113-127. *https://doi.org/10.1093/mind/fzz046*

Gigerenzer, G., Hertwig, R., and Pachur, T. (eds.) (2011). *The Foundations of Adaptive Behaviour*. Oxford University Press.

Grimmett, R., and Stirzaker, D. (2001). *Probability and Random Processes* (3rd ed.). Oxford University Press.

Kahneman, D., Slovic, P., and Tversky., A. (eds.) (1982). *Judgment under Uncertainty: Heuristics and Biases*. Cambridge University Press.

Kripke, S. (1979). A puzzle about belief. In A. Margalit (ed.), *Meaning and Use* (p. 239-283). Reidel.

Mahtani, A. (2017). The *ex ante* Pareto Principle. *Journal of Philosophy*, *114*, p. 303-323. *https://doi.org/10.5840/jphil2017114622*

Mahtani, A. (2021). Frege's puzzle and the *ex ante* Pareto Principle. *Philosophical Studies*, *178*, p. 2077-2100. *https://doi.org/10.1007/s11098-020-01524-w*

Mahtani, A. (2023). *The Objects of Credence.* Oxford University Press.

Peer, E., and Gamliel, E. (2013). Heuristics and biases in judicial decisions. *Court Review: The Journal of the American Judges Association*, *49*, p. 114-118. *https://digitalcommons.unl.edu/cgi/viewcontent.cgi?article=1428&context=ajacourtreview*

Stalnaker, R. (1999). *Context and Content: Essays on Intentionality in Speech and Thought*. Oxford University Press.

Weisberg, M. (2013). *Simulation and Similarity: Using Models to Understand the World*. Oxford University Press.

Williamson, T. (2000). *Knowledge and its Limits*. Oxford University Press.

Williamson, T. (2021). Epistemological consequences of Frege puzzles. *Philosophical Topics*, *49*, p. 287-319. *https://doi.org/10.5840/philtopics202149226*

Williamson, T. (2024). *Overfitting and Heuristics in Philosophy*. Oxford University Press.