"Undermined" Undermined
(a reply to Belot)

Carl Hoefer
ICREA Professor at University of Barcelona

v. 4, July 2018

Abstract:
In a recent article (Belot, 2016), Gordon Belot uses the so-called *undermining* phenomenon to try to raise a new difficulty for reductive accounts of objective probability, such as Humean Best System accounts. In this paper I will give a critical discussion of Belot's paper and argue that, in fact, there is no new difficulty here for chance reductionists to address.

Keywords: Humeanism, chance, objective probability, subjective probability, credence, Best Systems theory of laws and chance, actual frequentism.

1. Introduction

In a recent article (Belot, 2016), Gordon Belot attempts to raise a new difficulty for reductive accounts of objective probability, such as the numerous Humean Best System accounts that have flourished since 1994. On the basis of the well-known phenomenon of *undermining*, Belot argues for heretofore unrecognized unpleasant features of reductive accounts: they are forced to revise the probabilities assigned to certain events and histories, compared to what they would naturally be on a non-reductive theory of chance. They are also forced, counter-intuitively, to treat certain events and sequences of events as being probabilistically non-indepenent. Because of the unpleasantness of these alleged consequences, Belot suggests that non-reductive alternatives are preferable and deserve greater attention.

In this brief note I will show that Belot's argument is incorrect, and that there is no *new* issue for reductionists about chance to handle. This may seem like too minor a point to be the contribution of a whole paper. But because Belot's claims have the potential to mislead and distract many philosophers who are interested in reductionist accounts of objective chance, possibly even to spawn a pointless new sub-area in the chance literature, I feel it is important to have these criticisms available to philosophers of probability.

2. Belot's claims

The phenomenon of *undermining* is a well-known potential consequence of accounts of chance that reduce the facts about chance to facts about what events actually occur. It is easiest to see for the simplest form of chance reductionism: *actual frequentism*, the view that the objective probability of *A*-type events in a certain setup or reference class *S* is equal to the actual relative frequency with which *A*-type events occur in the relevant circumstances (if well defined). Consider Belot's toy-world example: a world which endures only 10 days and contains nothing but 10 chancy events, one each day, at each of which either the outcome *H* comes about, or the outcome *T*. In the full history of one such world, let's say that there are 6 *H* and 4 *T* events. Then, according to actual frequentism, the objective probability of *H* is 0.6, and the probability of *T* is 0.4. Now

imagine being in the toy world of this example, knowing that, according to actual frequentism, *Pr(H)* = 0.6, knowing that 9 of the 10 chancy events have already occurred, and knowing that there have already been 6 *H* outcomes. Knowing all this, one *knows* that the next chancy event will have outcome *T*. The outcome *H*, if it occurred, would *undermine* the very (frequentist) probability that we stipulated you know; it is incompatible with the stipulated facts, so it "can't happen", intuitively speaking.

More sophisticated, "Humean" reductionist accounts of chance tend to use the Mill-Ramsey-Lewis Best Systems approach (see (Lewis, 1994), (Hoefer, 2007)). Objective probability facts supervene on the non-modal, non-dispositional "occurrent" or "categorical" facts about events in the whole history of the world, but not in the simple way found in finite frequentism. On this approach, the objective probabilities for a given world are those found in an ideal chance theory, or set of chance laws, where the ideal theory has the best possible combination of simplicity, strength, and *fit* of the chances with actual event patterns and frequencies in that world. For certain forms of Best Systems accounts of chance, it is controversial whether undermining is a genuine possibility.[1] For pragmatic Best Systems accounts such as (Hoefer, 2007) and (Cohen & Callender, 2009), however, it is definitely a possibility that needs to be addressed. The correct chance rules in our world for a given type of event (such as spins of a roulette wheel, or radium decays) assign some non-zero probability to a future course of events (e.g. series of radium decays or roulette results) *F* which is such that no world could have the actual past history, *F* in the future, and also have the actual chance rules we posited to begin with. Given the Best System analysis of what objective chance *is*, this combination is not consistent.

What are the upshots of the undermining phenomenon, and what – if anything – should the chance reductionist do in response? The usual reaction (see (Lewis, 1994)) is to note that undermining is, to be sure, a peculiar, even unsettling phenomenon, but not a deal-breaker on its own; however, it creates outright contradictions for the chance reductionist in combination with the Principal Principle: contradictions in our rationally required *credences*, or subjective probabilities. Since one cannot tolerate one's analysis leading to contradictions, something has to give way. We will come back to this standard dialectic in the next section.

Belot makes no use of the Principal Principle, deliberately sets aside credences, and instead stays firmly in the realm of chance. His claim is that the chance reductionist is forced by undermining to *modify the objective probabilities* in odd ways that are unpleasantly counterintuitive.[2] We can see what Belot means using his toy world

[1] See (Loewer, 2004) and (Ismael, 2008) for arguments that undermining can not actually arise without presupposing illicit information from the future, which allows the issue to be set aside.

[2] Here it is appropriate to mention a complication of Belot's paper that I will otherwise ignore. Belot presents his new undermining challenge as one that affects *physical* chance theories of the kind which assign probabilities to entire world-histories. In his discussion of the 10-event toy worlds, in the first instance, the chance theories he has in view ascribe chances to each of the possible H/T sequences of length 10 – that is, to the possible complete world-histories. But Belot immediately moves to discussing the simplest such theory, one which assigns probability 1/1024 to each such world; such a theory is equivalent to one that simply says, of each chancy event, that it is independent of all others and has *Pr(H) = Pr(T)* = 0.5. From this point on in his discussion, the restriction to a special class of reductionist theories of chance falls out of sight and plays no role, as far as I can see. Nor does Belot say at any point that his undermining problem does *not* affect reductionist theories which do *not* proceed by assigning probabilities to entire world histories. So I will assume that Belot's alleged undermining problem is meant to

example, and presupposing actual frequentism again. Suppose one is considering the possible full world histories, of which there are 1024 (the set of all possible sequences of *H*'s and *T*'s of length 10). Suppose that one is told that *Pr(H) = Pr(T)* = 0.5. Initially that would seem to give each of the 1024 possible histories exactly the same probability, 1/1024. But on reflection, very many of the 1024 histories are, we might say, *kinematical* possibilities but not *dynamical* possibilities: under frequentism, they are incompatible with the stipulated probabilities. So in fact, Belot claims, those histories must be given probability zero, and this means that the probabilities of the non-ruled-out histories, as a collection, must be boosted to compensate (to jointly add up to 1).

Furthermore, for the same sort of reason, Belot claims that undermining induces failures of independence between distinct chancy events. In a *Pr(H) = Pr(T)* = 0.5 world one might initially suppose that each chancy event is independent of the others, as we commonly take coin flips to be. But this can't be so under frequentism: *Pr(H)* = 0.5, but *Pr(H in next event | first 5 events came out H)* = 0. And presumably, because (again) of the need to redistribute probability over the genuine possibilities after ruling out the undermining ones, even *Pr(H in next event | first 4 events came out H)* will not be equal to 0.5, but rather something considerably lower.

All these points are easy to grasp for the case of actual frequentism in finite toy worlds; but Belot spends several pages covering all the bases, arguing that this issue of modified objective chances arises in infinite worlds as well as finite worlds, and for Best System reductionism as well as actual frequentism. We need not look at how the arguments work for the other three cases, because the objection I will raise has to do with the very first move Belot makes: setting aside credence and the Principal Principle, and claiming that undermining forces reductionist accounts to amend the *chances* of events themselves. This is not the case. Undermining only forces a modification of our *credences* (and only in very bizarre thought examples), as Lewis (1994) maintained.

3. Not chance, but credence

In "Humean Supervenience Debugged" (1994), Lewis gave a very clear explanation of the problem that undermining presents to chance reductionists. Let's look at the way Lewis saw the issue. In the passage below, his discussion is framed from the perspective of an agent inside the world, part-way through world history; given history so far (*H*), what should we say about future courses of events that undermine the actual chances?

> Let *F* be some particular one of these alternative futures: one that determines different present chances than the actual future does. *F* will not come about, since it differs from the actual future. But there is some present chance of *F*. That is, there is some present chance that events would go in such a way as to complete a chancemaking pattern that would make the present chances different from what they actually are. The present chances *undermine* themselves. ...
>
> Could [a future like *F*] come to pass, given the present chances? Well, yes and no. It could, in the sense that there's non-zero present chance of it. It couldn't, in the sense that its coming to pass contradicts the truth about present chances. If it

apply to reductionist theories quite generally.

> came to pass, the truth about present chances would be different. Although there is a certain chance that this future will come about, there is no chance that it will come about while still having the same present chance it actually has. It's not that if this future came about, the truth about the present would change retrospectively. Rather, it would never have been what it actually is, and would always have been something different.
>
> This undermining is certainly very peculiar. But I think that, so far, it is no worse than peculiar.
>
> ... [A]t first sight the Principal Principle says nothing against undermining. It concerns, rather, the connection between chance and credence.
>
> But look again, and it seems that the Principal Principle does rule out undermining. ... We recall that F had some present chance of coming about, so by the Principal Principle, $C(F|E) \neq 0$. But F is inconsistent with E, so $C(F|E) = 0$. Contradiction. I could tolerate undermining as merely peculiar. But not contradiction! [482–483]

The Principal Principle is a constraint on reasonable or rational credences, i.e., subjective or personal probabilities, and the basic idea is that in the absence of "better" information, if one knows the chance of some event A happening then one's credence in A should equal the chance. PP reads:

$$C(A|XE) = x,$$

where E is the background knowledge of a rational agent, and is presumed to contain no information about whether or not a future chancy event A will obtain, and X is the proposition stating that the objective chance of A is x. Our undermining event F can be substituted for A, and PP tells us: have credence in F coming about that is equal to the objective chance of F.[3] On the other hand, given a reductionist analysis of chance, X, E and F are not jointly consistent. Therefore, by the rules of the probability calculus (which rational credences should respect), $C(F|XE) = 0$. This is the real threat of undermining, according to Lewis and most commentators: reductionism about chance, plus the PP, lands us in contradictory prescriptions.

The key part of Lewis' exposition, however, for our purposes, is the first half, which ended with something "very peculiar". Can an undermining future event F come about? Yes and no; *"It could, in the sense that there's non-zero present chance of it.* It couldn't, in the sense that its coming to pass contradicts the truth about present chances.*"* The first claim, in italics, is crucial: *there is a non-zero present chance of* F *coming about*. This is a simple statement of fact about what the correct theory of chances for our world asserts. That theory, or the relevant part of it, asserts: F has objective chance greater than zero, and in that sense, F is indeed a physical possibility. Belot seems to think that because of undermining, the theory of chances for our world must change its mind, so to speak, or "take back" its claim that $Pr(F) > 0$. But on what grounds? Belot gives no argument that this must be the case, he simply assumes it after presenting the undermining issue for the 10-event toy world theories and actual frequentism, and continues to assume it when he turns to other forms of reductionism and to non-finite worlds.

[3] In his exposition, Lewis has built X into the agent's background knowledge E.

For example, discussing the possible frequentist chance theories for the toy 10-event worlds, Belot claims that, under frequentism, imposing a statistical postulate (such as the chance rule that $Pr(H) = Pr(T) = 0.5$)

> . . . has propositional content of the most straightforward sort (to assert the statistical postulate is to eliminate some worlds as candidates to be the actual world). But contracting the space of worlds in this way inevitably distorts relations of probability. Consider the case of [10-event theories]. Under a straight reading of [the 0.5 chance theory], there are one thousand and twenty-four equiprobable histories; and the chance of H on Day Six is .5, no matter what the states are on Days One through Five. Under the frequentist interpretation of this theory, there are only two hundred and fifty-two possible histories: more than three quarters of the histories in T10 are excluded as being inconsistent with the frequentist understanding of the claim that the chance of H is even each day. As a consequence of this contraction of the space of histories, events that are considered probabilistically independent on a straight reading of the theory come out as dependent on the frequentist reading—for example, according to a frequentist, the chance of the state being H on Day Six is zero, conditional on the state having been H on all [5] preceding days. (pp. 785 – 786).

As Belot sees it, the mere fact that a chance theory is asserted "under frequentism" (or under some other reductive account of chance) means that the chance theory itself now *says* that certain *prima facie* possibilities are in fact impossibilities, and hence have objective probability zero. But this claim involves equivocating between what can be *inferred* from certain stipulated facts, including the chance facts as given in a theory, and what the chance theory literally *says*, about objective probabilities. The equivocation is beguiling and easy to fall into, but is a mistake nonetheless. The real undermining problem is all on the side of rational agents and their credences, and what those credences should be if the agents are apprised of both a reductive account of chance and what the chances for a given world are. Belot attempts to set aside credence and PP, and somehow make trouble for chance reductionists on the side of their objective chances alone; but there is simply no reason why a chance reductionist should agree with Belot's claims about what their chance theories *in themselves* say, nor follow him down the path of modifying the content of the chance rules in a quest to make undermining go away.[4] There never

[4] Although Belot does not explicitly supply an argument for his claim that reductionist chance theories literally "say" or "claim" that undermining events are impossible and therefore have chance zero, I can see how one might arrive at this view if one starts from the widely held, but incorrect, assumption that the domain of propositions over which a chance theory is defined and assigns probabilities should be wide enough to include *itself*; that is, that chance theories should or must assign objective probabilities to their own truth or falsehood. Let T be the proposition that asserts that in our world, the chances are given by theory $T@$. Let H be the world's history up to some point of time, and F be an undermining future event (given H and $T@$), as usual. Let $P_{T@}(\_)$ be the chance function corresponding to $T@$, and let T be in the domain of $P_{T@}(\_)$. If $P_{T@}(\_)$ obeys the axioms of probability, as it should, then $P_{T@}(F|HT)$ must be equal to zero, since F, H and T are jointly contradictory (assuming, as always, a reductive analysis of chance). Now, if we think of the conditioned 'T' as being redundant in $P_{T@}(F|HT) = 0$ – surely a sensible thing to think, after all, how can a theory of chance ($T@$) not be seen as asserting that the chances are as it, itself, says (T)? – then we might see it as trivial to re-write that equation as:

was any need to do this because, as Lewis said, on the side of *chance alone*, undermining is "merely peculiar."

To see this clearly, let's return to the "no" part of Lewis' "yes and no": "It couldn't, in the sense that its coming to pass contradicts the truth about present chances." What sort of claim is being made here? Is it a claim about physical impossibility *as delineated by our chance theory*, i.e., a claim that our chance theory says *F* has chance equal to zero? I think it is clear that it is not; it is a claim about what we can coherently believe. Given that the chances are as stated by *X*, and given *E* (or *H*), and understanding the reductive analysis of chance (whether actual frequentism, or Lewis' Best Systems account, or my pragmatic Best Systems account), we cannot coherently conceive of all three being true in actuality (*F, X, E*). If we take ourselves to know *X* and *E*, then we can infer *F won't* come about, and thus – in a sense, given *X* and *E* – "couldn't" come about. But this is simply a logical inference on our part, and it does not entail that *F* is physically impossible or that *F* has objective chance equal to zero.

Undermining does give rise to contradiction, but the problem is entirely on the side of credence, on what we propositions we can coherently believe to be true in actuality. Belot supposes that because of undermining, we must interpret the theory of chance for our world as *itself* "saying" that *F* cannot come about (or that, given that history up to now is as *E* asserts, *F* cannot come about). But the theory of chances need say no such thing; it is just a set of rules or laws that prescribe objective probabilities to certain sorts of events. The *starting point* of the undermining issue is noting that the theory ascribes an objective chance greater than zero to *F*, and that fact remains the case, even after we notice the peculiarity of undermining, and even after we notice that our *credences* concerning *F* are overdetermined in a bad way.

Before we wrap up, it may be interesting to think through what would follow, for reductive chance theories, if Belot's claim that they must skew the chances were correct. Chance reductionism involves giving some prescription for how the objective chance facts follow from (either full, or partial) facts about what events occur in actuality (future as well as past and present). Notice that if undermining forced chance theories to "change their minds" about the probabilities of certain potential events like *F*, that prescription would no longer be correct. So, what would the correct reductive prescription be? It would have to still be some prescription for how the objective chance facts follow from the occurrent facts (else, the reductionist would have given up their thesis); but such prescriptions, generically, allow undermining! In which case, perhaps the prescription will need to be modified again; the spectre of an endless regress looms.

$P_{T@}(F|H) = 0$. And this is what Belot, in effect, asserts: that undermining histories like *FH* are assigned chance equal to zero under reductionist accounts.

There are two mistakes in this train of thought. First, the assumption that *T* should or must be in the domain of $P_{T@}(\_)$ can be rejected (see (Hoefer, 2007) and (Ismael, 2008) for reasons to reject this assumption). Second, it is wrong to think that a chance theory like *T@* asserts the truth of *T*. In general, for reductionists, *T@* neither asserts *T* nor assigns it probability conditional on H equal to one, except perhaps at or near the end of time. In general, if it does exist, $P_{T@}(T|H) \neq 1$, and therefore $P_{T@}(F|HT) \neq P_{T@}(F|H)$. The former, if defined, must equal zero, but the latter need not.

Belot may be thinking that after self-revision to ensure that the original underminers get chance zero, no further undermining can then arise. But it is not clear what would justify such an assumption. After all, if our revisions zero out some possible events that (initially) had non-zero probability, we must redistribute that probability elsewhere. Qualitatively, what this means is that some or all of the nice, normal-looking events or event sequences get a bit of a "boost" to their chances.[5] But now, if future history can be indefinitely big and long, we see scope for a new kind of undermining to arise. Let *A* be one of those events (e.g. some long sequence of individual chancy events coming out a certain way) that gets a boost, from having chance x to having chance $x + \varepsilon$. And suppose that in that big, long future, a huge number of chances for *A* to either occur or not occur arise, and *A* in fact occurs with frequency equal to x (as is surely reasonable, given the original reductive recipe that led to the initial chance theory). This is a deviation of the frequency of *A* type events from its (new) objective chance – and hence, potentially an underminer. So there would seem to be at least a potential problem here to be addressed, if one were to follow Belot and assume that reductive chance theories must modify themselves so as to avoid undermining.[6]

Fortunately, there is no need to accept Belot's prescription that reductive chance theories must be self-effacing and correct themselves.

4. Conclusion

Belot's "Undermined" tries to show that chance reductionism gets into trouble just on the side of what such theories say about chances, before we even consider what happens when we consider the relationship between objective chance and credence captured by the PP. But his claims can be rejected. Reductivist chance theories do not, and need not, assign chance zero to undermining histories; so they do not need to correct themselves, retracting the chances they (so to speak) initially seemed to be asserting.

That does not mean that chance reductionists have no undermining problem; we certainly do! The problem continues to be the threat of contradictions on the credence side of things. See Hoefer (2018, chapter 5) for a full discussion of how the undermining problem should be resolved.

References

Belot, G. (2016). Undermined. *Australasian Journal of Philosophy*, 1–11. https://doi.org/10.1080/00048402.2016.1139604

Cohen, J., & Callender, C. (2009). A better best system account of lawhood.

[5] Structurally, this is just like the gambler's fallacy: being sure that too-unusual streaks (like seeing 10 coin flips in a row land Heads) are bound to be corrected sooner or later (and thus ascribing probability greater than ½ to the next flip landing Tails after seeing several Heads in a row).

[6] How plausible it is that such neo-undermining arises for a once-modified chance theory depends on the type of reductive account. For a pragmatic account such as mine (2007/2018), it is not likely that such 2nd level undermining could arise, except in toy worlds, because a reasonable amount of deviation of actual frequencies from chances is built in from the start. The correction amount ε, for reasonable (non-toy) worlds, is always going to be too small to make potential undermining trouble.

Philosophical Studies, *145*(1).
Hall, N. (1994). Correcting The Guide to Objective Chance. *Mind*, *103*(412), 505–517. https://doi.org/10.1093/mind/103.412.505
Hoefer, C. (1997). On Lewis's objective chance: "humean supervenience debugged." *Mind*, *106*(422).
Hoefer, C. (2007). The Third Way on Objective Probability: A Sceptic's Guide to Objective Chance. *Mind*, *116*(463). Retrieved from http://mind.oxfordjournals.org/content/116/463/549.short
Hoefer, C. (2018). *Chance in the World: a Humean Guide to Objective Chance*. Forthcoming, Oxford University Press.
Ismael, J. (2008). Raid! Dissolving the Big, Bad Bug. *Nous*, *42*(2), 292–307. https://doi.org/10.1111/j.1468-0068.2008.00681.x
Lewis, D. (1994). Humean Supervenience Debugged. *Mind*, *103*(412).
Loewer, B. (2004). David Lewiss Humean Theory of Objective Chance. *Philosophy of Science*, *71*(5), 1115–1125. https://doi.org/10.1086/428015