The Chances of Choices

Reuben Stern*

Forthcoming in *The British Journal for the Philosophy of Science*

Abstract

It is sometimes thought that if we treat decision-theoretic options as interventions, then we can use evidential decision theory to vindicate causal dominance reasoning. This is supposed to be guaranteed by a causal modeling axiom that implies that interventions are probabilistically independent of their non-effects—namely, the Causal Markov Condition. But there are two concerns for this line of reasoning. First, the Causal Markov Condition doesn't imply that an agent should regard their intervention as probabilistically independent from its non-effects when the agent has "exotic evidence"—i.e., evidence about some variable that they regard as causally downstream from their intervention. Second, the Causal Markov Condition is not plausible when we interpret it as implying constraints on subjective probability distributions, because there are cases where it is rational for an agent to regard variables as causally independent but subjectively probabilistically dependent. In this paper, I argue that interventionists can answer these challenges by adopting a conception of choice according to which there are significant constraints on the objective probabilities for decision-theoretic options.

1 Introduction

When you make a real choice, the existing chances don't compel or incline you to choose in any particular way. For were the chances to nudge you toward some

*Duke University – reuben.stern@duke.edu – https://sites.google.com/view/reubenstern/home.

option(s), then the choice wouldn't be fully yours. This means that if you're currently confronted with a genuine choice, then the chance of your choosing any particular way must, immediately prior to settling for yourself what to do, equal the chance of choosing any other way.

To some, this line of reasoning may seem completely anodyne. After all, when you make a genuine choice, it's reasonable to think that your choice must be up to you in the sense that any inclination or disposition to choose in any particular way must arise from within you, and not from antecedently determined chances that incline you toward some subset of your option(s). But to the trained philosopher, this line of reasoning may sound alarm bells. Anyone who has taken a course on free will knows that we often come to decision-making contexts disposed to choose in particular ways—because of our upbringing, because of our genetic makeup, because of our personality, etc.—and it thus may seem overly strict to maintain that the chances of choices cannot compel us to choose in any particular way.

Nevertheless, I will argue here that roughly this commonsensical (if naïve) view pays serious decision-theoretic dividends. Specifically, I will argue that if the chances mustn't distinguish between an agent's options, then we can rescue Meek and Glymour's (1994) idea that we can use the machinery of evidential decision theory to attain causal-decision-theoretic recommendations by treating the decision-maker's options as *interventions*. This is a desirable result not only insofar as it identifies a path for those with causal-decision-theoretic sympathies to square the rationality of their favored choices with the evidentialist's insight that we should choose in a way that makes the desired outcomes most probable,[1] but also insofar as it promises to reshape the dispute between causal decision theorists and evidential decision theorists. For if we can rescue the Meek and Glymour (MG) idea, then we can fruitfully view the dispute between causal decision theorists and evidential decision theorists as turning on how we should represent genuine agency, rather than on which of two irreconcilable norms is more intuitive.[2]

[1] Papineau (2001, p. 244) maintains that the evidentialist's insight is self-evident: "Evidential theory simply recommends that agents perform those actions that make desired results most probable. This recommendation doesn't seem to need any further justification. Doesn't everybody want it to be probable that they will get what they want?"

[2] In my (2018) paper, "Diagnosing Newcomb's Problem with Causal Graphs," I argue that we can accomplish a related dialectic feat, but my approach there requires that we abandon evidential decision theory for a distinct decision theory—namely, what I call "Generalized Interventionist

But before I proceed to argue that this commonsensical view can rescue MG's idea, I need to explain why it is in need of rescuing. Toward this end, I will reconstruct an argument for the view (§2), before explaining where it falters (§3). Then, I will argue that that the holes in this reconstructed argument can be patched if we additionally posit constraints on the chances of choices (§4 and §5). Here, we will find that the generality of the rescue attempt's success will depend on the operative chance constraint, but that we can deliver a full-blown vindication of causal-decision-theoretic recommendations from within evidential decision theory if we posit something like the commonsensical view sketched above. Does this mean that we should accept the commonsensical view? I won't definitively answer this question (largely because it's too hard), but I will argue that the view is attractive insofar as it solves all of MG's problems in one fell swoop. Finally, I will conclude by taking stock of what we're licensed to believe about the chances of choices, given the arguments of this paper (§6).

2 Causal Expected Utility as Conditional Expected Utility

Evidential decision theorists follow Jeffrey (1983) in maintaining that agents should opt for whatever option, x, maximizes conditional expected utility when defined as follows, where $P(y|x)$ corresponds to the conditional probability that state y will obtain given that x obtains and where $V(x, y)$ corresponds to the value of the outcomes associated with taking action x in state y.

$$CEU(x) = \sum_{y} P(y|x)V(x, y)$$

For causal decision theorists, the problem with maximizing CEU is that it sometimes recommends opting for an action on the grounds that doing so makes some state more (or less) probable, even though the agent knows full well that their

Decision Theory." Here, my aim is to use constraints on genuine choice to vindicate causal-decision-theoretic verdicts from within evidential decision theory—i.e., without breaking from the Jeffrey's (1983) insight that we should maximize conditional expected utility.

action exerts no causal influence over the state. Newcomb's Problem notoriously exemplifies one such decision-making context.[3]

> **Newcomb's Problem (NP)**: You stand before two boxes. One is transparent and contains \$1,000. The other is opaque. You have a choice. You can one-box (i.e., take the contents of the opaque box) or two-box (i.e., take the contents of both boxes). The game is set up such that the contents of the opaque box always depend on the earlier prediction of a remarkably successful predictor. If the predictor predicts that you will one-box, she places \$1,000,000 inside the opaque box. If she predicts that you will two-box, the opaque box is empty. Should you one-box or two-box?

When confronted with this decision, your should consider it far more likely that the predictor has predicted that you will one-box in the event that you opt to one-box than in the event that you opt to two-box (because the predictor is really good at predicting the actions of Newcomb subjects), but you also know that the predictor's prediction is not causally influenced by whether you one-box or two-box (since the predictor has already made their prediction). Causal decision theorists say that the probabilistic dependence that obtains between your action and the predictor's prediction is irrelevant to rational choice, and thus maintain that you should two-box (because two-boxing is \$1,000 better than one-boxing both when the predictor has predicted that you'll one-box and when the predictor has predicted that you'll two-box). But if you maximize CEU, you'll end up one-boxing because the probability that the predictor predicts that you will one-box given that you one-box is so much greater than the probability that the predictor predicts that you will one-box given that you two-box.

In order to block this result without abandoning the view that we should maximize CEU, we need some reason to think that we should not regard the relevant causally independent state (e.g., the predictor's prediction) as correlated with the decision-maker's choice.[4] Along these lines, a number of authors have argued that

[3] Newcomb's Problem was first intoduced to philosophers by Nozick (1969).

[4] Throughout this manuscript, when I say that X is "correlated" with Y. I mean that X and Y probabilistically depend on each other. In this particular case, the relevant probabilistic dependence is defined in terms of the decision-maker's subjective probability function.

the decision-maker should cease to regard their choice as correlated with the relevant causally independent state when the decision-maker conditions on all of their available evidence (because the available evidence screens off the correlation between their action and the relevant state).[5] Applied to NP, the idea is that once you condition on all of your available evidence (including, e.g., whether you feel inclined to reach for one or both boxes), you should no longer regard whether you one-box or two-box as correlated with the predictor's prediction (e.g., because the predictor's prediction is based on the inclination that you've already registered). But while this strategy may preserve causal-decision-theoretic intuitions in a number of important cases, it is now widely acknowledged that there are scenarios in which decision-makers should not regard the available evidence as screening off correlations between acts and causally independent states (e.g., when the Newcomb subject is told that the predictor not only bases their prediction on the Newcomb subject's inclination, but also on whether the Newcomb subject follows through with their inclination). Thus we are left in need of some other reason that agents should regard their choices as probabilistically independent from any states that they regard as causally independent.

Enter the interventionist approach. Armed with the graphical approach to causal modeling,[6] MG argue that we can generally secure causal-decision-theoretic recommendations while maximizing CEU by treating decision-makers' options as interventions in a causal graph. Their key insight is that one of the axioms of the graphical approach to causal modeling—namely, the Causal Markov Condition (CMC)—implies constraints on the probability distributions that are compatible with a causal graph, and, more specifically, implies that the intervention to x must be unconditionally probabilistically independent from any variable that is not causally downstream from X itself.[7]

How does this follow from the CMC? Allow $\mathbf{V}$ to denote the set of variables over which the relevant probability distribution and causal graph are defined. According

[5] Eells (1982) and Price (1986) are prominent examples of authors who adopt this strategy.

[6] There are many resources that one can consult in order to learn about graphical causal models and their many uses—e.g., Pearl (2009), Pearl and Mackenzie (2020), Peters, Janzing, and Schölkopf (2017), and Spirtes, Glymour, and Scheines (2000).

[7] Here, and throughout the body of this paper (but not the appendix), I use capital italicized letters to denote variables and lowercase italicized letters to denote values of variables.

to the CMC, if two variables, X and Y, are *d-separated* by a (possibly empty) set of variables, Z, in some causal graph over $\mathbf{V}$, then X and Y must be probabilistically independent of each other conditional on any assignment of values over Z in any probability distribution that is compatible with that graph. X and Y are *d-separated* by Z exactly when every path between X and Y is *blocked* by Z, where a path between X and Y is *blocked* by Z exactly when:

1. the path between X and Y contains a non-collider that is in Z, or,
2. the path contains a collider, and neither the collider nor any descendant of the collider is in Z.

If you are unfamiliar with graphical causal models, then this language of "d-separation" and "colliders" probably is not very helpful on first read. But since a collider is just a common effect of two variables along an undirected path[8]—e.g., C along the path, $A \leftarrow B \rightarrow C \leftarrow D$—the CMC can be parsed as saying that any two variables represented in some causal graph must be probabilistically independent of each other unless (i) they share a (direct or indirect) common cause, (ii) one is a (direct or indirect) cause of the other, or (iii) they are both (direct or indirect) causes of some common effect that has been conditioned on.[9] This means that if we define the intervention on X as an exogenous cause of X that we can use to deterministically set X to any of X's values,[10] then provided that we have not conditioned on any variables that are causally downstream from X, the intervention on X (unlike X itself), must be probabilistically independent from any variables that are not causally downstream from X.[11] So it seems that even

[8]An undirected path is just a a sequence of variables such that there is an arrow (going in either direction) between each variable and the next. Intuitively, causal arrows *collide* along paths at *colliders*.

[9]Since the CMC implies the Common Cause Principle (i.e., that if two variables are unconditionally probabilistically dependent, then either one is causally downstream from the other or they are joint effects of some common cause), it is only appropriate to assume the CMC for variable sets that are *causally sufficient* in the sense that they include the common causes (or some important subset of the common causes) of the variables included therein.

[10]A variable is "exogenous" relative to $\mathbf{V}$ when it is not causally downstream from any other variable in $\mathbf{V}$.

[11]MG make some additional assumptions about interventions that we'll see are arguably problematic in the context of using evidential-decision-theoretic machinery to vindicate causal-decision-theoretic intuitions, but the characterization of interventions provided here is sufficient

when a decision-maker has reason to regard whether they x as evidentially relevant to whether some causally independent state, y, obtains, the CMC says that they should *not* regard their *intervention* as evidentially relevant to whether y obtains.

Now, in decision-making contexts like NP, it can be unreasonable to have the conviction that you're intervening as you make your choice. After all, representing yourself as intervening when confronted with NP involves representing the predictor as not being able to predict your choice (since it involves representing the predictor's prediction as uncorrelated with your intervention), but you know that the predictor reliably predicts the choices of Newcomb subjects, and you have no reason to regard yourself as an exception to this rule. So you have no reason to believe that you're intervening in NP, and it thus may seem moot that the intervention to one-box (or two-box) is not correlated with the predictor's prediction. But interventionist two-boxers have a retort. Specifically, they can maintain that a choice is genuine only when it is causally autonomous from any other factors under consideration (which effectively means that a choice is genuine only if it is aptly respresented as an intervention),[12] and that the standards of rational choice apply only to genuine choices. According to this line of reasoning, then, it's true that you should be confident that you're not intervening as you choose whether to one-box or two-box, but this just means that you should be confident that you're not making a genuine choice.[13] And since the norms of rational choice apply only to genuine choices, you should bracket the possibilities in which you're not intervening, and should choose only for your intervening self.

The interventionist two-boxer can thus agree with the evidentialist one-boxer that we should maximize conditional expected utility. Where the interventionist two-boxer disagrees with the one-boxer is in the specification of the options on which we should condition as we calculate conditional expected utility.[14] The

for our purposes since the relevant probabilistic independencies follow just from the location of the intervention in the graph under consideration.

[12]See Stern (2018) for more discussion of the relationship between causal autonomy and interventions.

[13]The idea that there is reason to think that the Newcomb subject is not making a genuine choice is not unique to interventionists. See, e.g., Jeffrey (2004) and Slezak (2006; 2023).

[14]This echoes MG (1994, p. 1015): "Our suggestion is that the differences in recommendations offered by causal decision theorists and most of their critics do not result from whatever differences they may have about the principles of rational choice... Where they recommend different decisions in particular cases causal decision theorists have discussed it is because they differ

interventionist two-boxer maintains that we should evaluate the conditional expected utility of *genuinely deciding* to one-box and *genuinely deciding* to two-box (which by their lights amounts to evaluating the conditional expected utility of intervening to one-box and intervening to two-box), while the classical evidentialist one-boxer maintains that we should evaluate the conditional expected utility of one-boxing and two-boxing, simpliciter (perhaps because they adopt considerably weaker constraints on what constitutes a genuine choice).

If we abstract away from the interventionist details, then, the rationale for two-boxing is that we should maximize conditional expected utility when defined as follows.

$$CEU(decide(x)) = \sum_y P(y|decide(x))V(decide(x), y)$$

The calculation is mathematically the same as the standard calculation of conditional expected utility, but it is now explicit that the decision-maker's options should be construed as (genuine) decisions. MG maintain that we can secure causal-decision-theoretic recommendations by adopting this decision rule and maintaining that every genuine decision is an intervention, but the general schema is flexible enough to vindicate recommendations associated with evidential decision theory (since once could part ways with the interventionist conception of choice and instead maintain some thinner conception of choice that is compatible with representing your decision as causally endogenous to the model at hand). So it is by adopting this general decision rule that MG are able to deliver the result that we can fruitfully view the dispute between one-boxers and two-boxers as stemming from a dispute about what it means to make a genuine choice, rather on which of two irreconcilable decision rules is more intuitive.

If MG are right that the view that all genuine choices are interventions implies that decision-makers' options are uncorrelated with their non-effects in their subjective probability distributions, then MG successfully show that we can use this general schema to secure causal-decision-theoretic recommendations without breaking from the evidentialist's insight that we should maximize CEU. But the

about whether an action is an intervention... If so, then a different event must be conditioned on than if not, and a different calculation results."

the view that all genuine choices are interventions does not by itself imply that decision-makers' options are uncorrelated with their non-effects in their subjective probability distributions. Or so I argue in the next section.

3 Meek and Glymour's Mistakes

There are two holes in the case for the view that we can secure causal-decision-theoretic recommendations by maintaining that agents should maximize CEU while treating their options as exogenous interventions. The first is that the CMC doesn't imply that an agent should regard their intervention as probabilistically independent from its non-effects when the agent has "exotic evidence"—i.e., evidence about some variable that they regard as causally downstream from their intervention.[15] The second is that while the CMC is arguably plausible when interpreted as implying constraints on the objective (statistical) probabilistic dependencies that are compatible with a given causal graph,[16] the CMC is not plausible when we interpret it in terms of implying constraints on a decision-maker's subjective (credal) probability function given the causal graph(s) that they accept, because there are cases where it is rational for an agent to regard variables as causally independent but subjectively probabilistically dependent.[17]

[15]See Stern (2021).

[16]There are two prominent objections to the CMC's application to objective probability distributions. The first, discussed by Sober (2001), is that there are empirical counter-examples to the CMC involving coincidental correlations. But there are responses to Sober (e.g., Hoover 2003) to the effect that the "correlations" that figure in Sober's examples are not actually correlations at all (once we are careful to distinguish between association and correlation). No matter whether such responses are successful, it is clear that there is a version of the CMC that survives Sober's objection unscathed—namely, one that applies only to non-coincidental (or nomic) correlations. The second objection comes from quantum correlations and is somewhat more pressing because no one doubts that the correlations at issue are nomic. One response to these cases—explicitly adopted by Hausman (1999)—is to further limit the domain of the CMC so that it doesn't apply to these correlations. But there also may be possible responses to the tune that the relevant quantum correlations actually do not violate the CMC, but rather violate what's known as the Causal Faithfulness Condition, since the formal results in this area are that the quantum correlations cannot satisfy both the CMC and the Causal Faithfulness Condition. See Näger (2016).

[17]A variant of this objection is mentioned by Meek and Glymour (1994) themselves, though they arguably do not fully appreciate its importance, given the supposed lesson of their paper. Meanwhile, both Zhang, Seidenfeld, and Liu (2019) and I have developed this point against standard attempts to secure causal-decision-theoretic recommendations while maximizing CEU.

I make the case for the first point in great detail in Stern (2021), but it is helpful to present an exotic decision-making scenario here that makes the point as simply as possible.

> **Birthweight Problem (BP)**:[18] Suppose that you're pregnant and that an oracle has just informed you that your baby will unfortunately be born underweight. You recently learned that when a baby is born underweight, their prognosis is significantly better if their mother smoked during gestation than if she did not. This is because the odds of survival are better when the newborn is underweight because its mother was a smoker than because the newborn suffers from some genetic condition. If you would find smoking to be a bit unpleasant, should you take up the habit?

Figure 1: The Causal Structure of BP

(See my 2017 paper, "Interventionist Decision Theory.") I present a similar line of reasoning here in a slightly different light largely in order to bring more attention to these extant objections. But it's also worth noting that none of this previous work coheres with my current thinking about the upshot of this objection. Zhang, Seidenfeld, and Liu ultimately champion a response to the kind of uncertainty at issue that is odds with what I argue here is rational, while I previously argued that this objection spells doom for interventionist attempts to secure causal-decision-theoretic recommendations while maintaining that decision-makers should maximize CEU. Here and now, I am effectively arguing that my previous conclusion was somewhat hasty, since I am arguing that we *can* secure causal-decision-theoretic recommendations while maintaining that decision-makers should maximize CEU, provided that we sign up for significant constraints on the objective probabilities for decision-theoretic options. But interestingly, my current line of reasoning does not undercut the interventionist decision theory (IDT) that I developed in order to address this objection in Stern (2017). Instead, against the backdrop of my current argument, IDT can be seen as a decision theory that manages to secure causal-decision-theoretic recommendations (provided that choice is not exotic), regardless of whether there are any constraints on the chances of choices. Of course, it accomplishes this feat at the expense of capturing Jeffrey's insight that we should maximize CEU. But to some, this may be a price worth paying.

[18] I briefly discuss a version of this case in fn. 25 of Stern (2021). Pearl (2016) explains the closely related "birthweight paradox" in terms of conditioning on a collider.

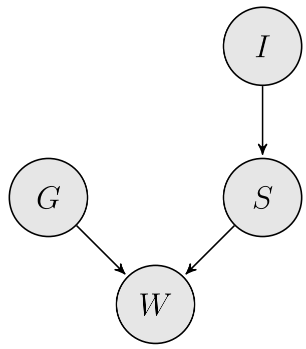

When confronted with BP, it's clear that causal-decision-theoretic reasoning favors abstaining from smoking (if you mind smoking at all) since your choice whether to smoke exerts no causal influence over whether your newborn will suffer from a genetic condition—or put differently, since there is nothing that you can do now by deciding whether to smoke to change your newborn's genetic makeup. But given the Figure 1 representation of BP, the CMC allows for you to regard whether you intervene to smoke or not smoke (I) as correlated with whether your newborn will suffer from the genetic condition (G), given the oracle's foreknowledge that your baby will be underweight ($W = \mathit{underweight}$).[19] This is because W is a collider on the undirected path that links I and G.[20] So if we maintain that decision-makers should maximize CEU while treating their options as interventions, then there will be reason to intervene to smoke that stems from the correlation that you rightly take to obtain between I and G. Ergo, when choice is exotic (and the possibility of conditioning on a collider that is downstream from the decision-maker's intervention is correspondingly introduced), we cannot secure causal-decision-theoretic recommendations simply by maintaining that decision-makers should maximize CEU while treating their options as interventions.

This makes it clear that MG's preferred interventionist method for vindicating causal-decision-theoretic recommendations does not generally deliver the desired results when choice is exotic, but the severity of this problem can perhaps be con-

[19] To be clear, the CMC is right to allow for this correlation. After all, it's a live possibility that your newborn will be born underweight *because* of your choice to smoke, and this would be *good news* insofar as it'd decrease the odds that your newborn would be born underweight because of some genetic condition, even if this good news doesn't reflect your causal influence over G.

[20] This means that I and G are *not* d-separated by W.

tested on the grounds that we should flag exotic decisions involving foreknowledge as insufficiently unrealistic to carry substantial normative weight.[21] The same is not true of the second problem. For if the causal graphs that a decision-maker accepts do not constrain the decision-maker's subjective probability judgments, it is unclear how to even pursue the project of using interventionist machinery to deliver causal-decision-recommendations while maximizing CEU, given the foundational role that the CMC plays in both MG's argument and the graphical approach to causal modeling, more generally.

Like the first problem, the second problem can be illuminated by a simple example.

> **Dueling Meteorologists (DM)**: Suppose that you're certain that whether it rains on July 31st in New Orleans is completely causally unconnected from whether there is bad air quality in Shanghai on July 31st (in the sense that neither variable causes the other, and they share no common cause). Now suppose that you've just consulted two meteorologists' forecasts for July 1st, and that you know that one of them is right about the chance of rain in New Orleans and the chance of bad air quality in Shanghai, but that you are indifferent as to which is right. Moreover, suppose that both meteorologists treat the chance of rain in New Orleans as independent from the chance of bad air quality in Shanghai, but that PESSIMIST regards both chances as relatively high (say, 0.8) while OPTIMIST regards both chances as relatively low (say, 0.4). Suppose now that the 31st comes around and that you look out your New Orleans window to discover that it's raining. Should your glance out the window increase your confidence that there is bad air quality in Shanghai?

When confronted with DM, it is intuitive that you *should* increase your confidence that there is bad air quality in Shanghai. After all, you know that either PESSIMIST or OPTIMIST has correctly identified the chances, and your discovery

[21] See Lewis (1982), Rabinowicz (2009) and Stern (2021) for various arguments to the contrary. See Gallow (forthcoming) for extensive discussion of the role that foreknowledge plays in the dispute between causal decision theorists and evidential decision theorists.

of rain in New Orleans gives you reason to trust PESSIMIST's chance estimates over OPTIMIST's, since PESSIMIST's estimate for the chance of rain in New Orleans was substantially greater than OPTIMIST's. More formally, while it initially would have been rational to assign a subjective probability of 0.6 to bad air quality in Shanghai (since you were indifferent about which meteorologist is right), it is clear that your subjective probability should increase to some value greater than 0.6 upon learning that it's raining since you now should be at least somewhat more confident that PESSIMIST was right about the chances than that OPTIMIST was right. But if it's rational to update your subjective probability estimate in this way while accepting a causal graph according to which whether it rains on July 31st in New Orleans is completely causally unconnected from whether there is bad air quality in Shanghai on July 31st, then MG's application of the CMC to subjective probability estimates is in trouble (even though it still constrains the objective probability distributions that you regard as live possibilities). If we allow R to represent whether it rains in New Orleans on July 31st and A to represent whether there is bad air quality in Shanghai on July 31st, the problem is that it is reasonable to accept the Figure 2 graph (according to which R and A are trivially d-separated by the empty set), but to treat R and A as unconditionally probabilistically dependent in your subjective probability distribution. This means that CMC cuts no ice when interpreted in terms of implying constraints on a decision-maker's subjective probability distribution given the causal graph(s) that they accept.

Figure 2: The Causal Structure of DM

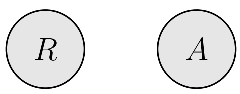

There are at least two ways that one might try to rescue the CMC's application to subjective probabilities from this objection.[22]

First, one might try to argue that even though it is stipulated that you accept the Figure 2 graph according to which R and A are completely causally unconnected, this acceptance is unjustified because there is good reason to believe that

[22]The reader who is already convinced by the case against the CMC's application to subjective probabilities can skip to (§4) without any loss of comprehension.

the Figure 2 graph excludes some path between R and A that renders R and A possibly correlated as far as the CMC is concerned.[23] This reply grants that you should regard R and A as correlated in your subjective probability distribution, but denies that it's rationally permissible to do so while accepting the Figure 2 graph on the grounds that it was never rationally permissible to accept the Figure 2 graph in the first place. What path might Figure 2 unjustifiably exclude? While it's clear that R doesn't causally influence A and that A doesn't causally influence R (since they describe roughly simultaneous weather events on opposite sides of the world), there may be room to argue that there is some common cause of R and A, or some collider between R and A on which the DM subject has implicitly conditioned. Let us address these possibilities in turn.

The first thing to say in response to the idea that there is a common cause of R and A is that the independence of R and A in both candidate chance distributions seems to reflect the received wisdom that R and A share no common cause. Strictly speaking, it's compatible with CMC's application to objective probabilities that these chances be independent even though R and A share a common cause (since the CMC is in the business of implying probabilistic independencies, not dependencies), but other conditions that play an essential role in the graphical approach to causal modeling imply that the chances of R and A be dependent if they share a common cause.[24] This means that we can arguably justifiably infer that R and A do not share a common cause from their independence in the live objective probability distributions under consideration. Moreover, while it may not be completely unreasonable to think that there could be some distal common

[23]It is important to remember that the graphs at issue are *causal* in the sense that the directed edges in the graph represent causal relations. As an anonymous referee helpfully points out, one might be tempted to account for the correlation between R and A by positing some non-causal relationship (e.g., like those that appear in Bovens and Hartmann's (2003) epistemic Bayes nets), but this would not save the application of the *Causal* Markov Condition to subjective probabilities.

[24]The details are somewhat complicated here. The widely assumed Causal Faithfulness Condition straightforwardly implies that R and A must be unconditionally correlated if they share a common cause, but the Causal Faithfulness Condition is typically regarded as a simplifying assumption rather than an axiom because it falls prey to rare counterexamples. Meanwhile, the strictly weaker and arguably axiomatic Causal Minimality Condition implies that R and A must be unconditionally correlated when they share a common cause, provided (i) that the variables at play are binary, and (ii) that there are no paths between R and A other than that which includes the operative common cause.

cause of two isolated weather events, it is easy to construct examples that make the same point as DM in terms of variables that are even more removed from each other than R and A.[25] This means that even if someone were to successfully argue that there actually is reason to think that there is a common cause of R and A, their argument would not vindicate the CMC's general application to subjective probabilities.

The other possibility is that Figure 2 excludes some collider between R and A on which the DM subject has conditioned. Though it's hard to identify any such collider since no such variable is explicitly mentioned in the statement of DM, one might try to argue that you implicitly update on some variable that is accurately represented as a collider between R and A when you update on the meteorologists' reports. It is implausible that the reports should be represented as a common effect of R and A (since you update on their reports *before* July 31st, and there is no good reason to think there is retro-causation at play), but it is possible for a variable to be a collider between R and A without being causally downstream from R and A—e.g., if there is one common cause of R and the report variable and another common cause of A and the report variable.[26] This possibility arguably has more going for it than the first insofar as it captures the received wisdom that R and A share no common cause. But without some argument that we *always* update on a collider between two variables when we get information about their respective chances, it is hard to see how it could fare any better than the first at vindicating the CMC's general application to subjective probabilities. Indeed, even when we focus on DM itself, and even when we insist on explicitly modeling the meteorologists' reports, it is not at all clear why we should (or must)

[25]The abstract structure of the case can be realized by any plausibly causally independent variables. I chose to make the point with DM because it describes a relatively familiar kind of disagreement between experts, but we can consider a revised version of DM that is not concerned with the chances of two weather events, but rather with the chances of two completely different kinds of events—e.g., R and another variable WC that expresses whether a European team will win the next World Cup. If PESSIMIST treats R and WC as independent and assigns 0.8 both to rain in New Orleans and to a European team winning the next World Cup, while OPTIMIST treats R and WC as independent but assigns 0.4 to both events, then for the same reasons that it was rational to increase your confidence in bad Shanghainese air quality upon discovering New Orleans rain, it is rational to become more confident that a European team will win the next World Cup upon discovering New Orleans rain.

[26]I am grateful to an anonymous referee for raising this possibility.

understand the reports in terms of a collider between R and A.[27]

This leaves us with the second strategy for rescuing the CMC's application to subjective probabilities—namely, championing some alternative response to the uncertainty about the underlying objective probabilities that satisfy the CMC. This effectively is the approach taken by Zhang, Seidenfeld, and Liu (2021), who seem to argue that in circumstances like those we confront in DM, we should adopt an imprecise or indeterminate subjective representor that consists of a non-convex set of probability distributions that includes the candidate objective probability distributions (each of which includes the independencies implied by the CMC), rather than a singleton probability distribution that incorporates the decision-maker's uncertainty over the candidate objective probability distributions.[28] One might take issue with this response in the context of using evidential-decision-theoretic machinery to secure causal-decision-theoretic recommendations since it requires that we adopt an imprecise Bayesian approach to representing uncertainty, which plays no role in Jeffrey's (1983) classical treatment of evidential decision the-

[27] For example, if we update on the meteorologists' reports at different times, then it is natural to represent each report with its own variable. It is hard to identify a plausible causal graph according to which either individual variable is represented as a collider between R and A.

[28] I've included "seem to" here because it is somewhat difficult to interpret Zhang, Seidenfeld, and Liu, given the inclusion of following (fn. 6) passage of their paper.

> "What if the agent has a prior or higher-order probability distribution over the set of probability functions? Should the prior be ignored or should it be used to integrate out parameters to form a single, precise probability function? This is an important question to which we cannot do justice here, in part because we are yet to make up our (group) mind about it. However, it is fair to say that for the reason elaborated earlier, unless we give up the subjective Markov condition, we cannot simply collapse the hierarchical structure into a precise distribution over the substantive variables and treat that single distribution as all there is to the agent's doxastic state (though this single distribution may well be useful for other purposes)."

If DM is countenanced as a case in which you have a "higher-order probability distribution over the set of probability functions" (on the grounds that you have a subjective probability distribution that is defined over the candidate objective probability distributions) then Zhang, Seidenfeld, and Liu can perhaps be interpreted as punting when it comes to the epistemically rational response here. But the passage also makes it clear that their overall project is at odds with the response that I endorse as epistemically rational, since my favored response amounts to accepting a causal graph according to which two variables are causally independent while adopting a subjective probability function according to which they're probabilistically dependent.

ory.[29] But either way, it is still worth considering whether this response can save a subjective interpretation of the CMC. Though I agree with Zhang, Seidenfeld, and Liu that it is sometimes rationally permissible to adopt an imprecise representor (and perhaps even a non-convex imprecise representor),[30] it is likewise sometimes permissible to adopt a singleton probability distribution that incorporates the decision-maker's uncertainty over the objective probability distributions that they regard as live possibilities. The case for this is partially based on intuition—e.g., it seems obviously rationally permissible to adopt a precise subjective probability of 0.75 that a coin's toss will land heads when 50/50 with respect to whether the coin is fair or two-headed—but it is also in keeping with imprecise Bayesian orthodoxy. The standard epistemic motivation for imprecise Bayesianism is that there are evidential contexts in which it is rationally permissible for an agent to be unopinionated toward some space of possibilities in a sense that is ruled out by a precise singleton subjective probability distribution over said possibilities.[31] But even granting that such lack of opinionation is sometimes warranted, there is near consensus among imprecise Bayesians that when an agent does have an opinion about the relative plausibility of some candidate objective probability distributions (in the form of a precise probability distribution that is defined over the candidate objective probability distributions), the agent should integrate this opinion by adopting a linear mixture of the candidate objective distributions, where the weights that are utilized in the construction of the mixture are supplied by the agent's subjective probability judgments in the candidate objective distributions.[32]

[29]This argument is arguably not very strong, since there are evidential-decision-theoretic and causal-decision-theoretic versions of imprecise Bayesian decision theory.

[30]See Levi (2009) for various arguments that imprecise representors should be convex. If Levi is right, then as Zhang, Seindenfeld, and Liu acknowledge, we very clearly cannot rescue a subjective version of the CMC by adopting an imprecise framework, since a convex set of probability distributions that includes the two meteorologists' distributions will contain every linear mixture of these two distributions, and these linear mixtures will not satisfy the CMC.

[31]When an agent has a precise subjective probability distribution toward some space of possibilities, their comparative confidence judgments are completely ordered. But it seems to many that there are evidential contexts in which an agent can justifiably be unordered toward some of the possibilities under consideration. Imprecise Bayesianism allows for this latter possibility, since we can represent an agent as unordered toward some space of possibilities when the probability judgments of the probability functions that comprise the imprecise representor are in disagreement about which possibilities are more probable than which others. See Eva (2019), Joyce (2010), and Levi (2009) for further discussion of this point.

[32]See Joyce (2010) and Levi (2009) for recent discussion of this point by two influential ad-

When confronted with DM, this amounts to adopting a subjective probability distribution that incorporates your indifference between the two candidate objective distributions. In the resulting mixture, R and A are probabilistically dependent even though they are probabilistically independent in both candidate objective distributions.

4 The Chances of Choices and the CMC

Given the case against the CMC's application to subjective probabilities, it is clear that if we're going to use the CMC to vindicate MG's idea, we'll have to do so indirectly—i.e., by way of its application to objective probabilities. Now that we know that the causal graphs that an agent accepts do not directly constrain their subjective probability judgements, the only remaining possibility is that the causal graphs that an agent accepts constrain the objective probability distributions that the agent regards as live options, and that the relevant independencies in the agent's subjective probability distribution can somehow be derived from the objective independencies that are implied by the CMC. For example, were it the case that any rational subjective probability distribution over $\mathbf{V}$ must integrate the independencies that are shared by every candidate objective probability distribution over $\mathbf{V}$, then we could straightforwardly rescue MG's insight (since the CMC's application to objective probabilities would straightforwardly give us everything that we need). But DM already reveals that this is not so since DM exemplifes a scenario in which it is rational for an agent to regard two variables as probabilistically independent in every candidate objective probability distribution, but as probabilistically dependent in their subjective probability distribution. Does this spell doom for MG's idea?

Mathematically, the problem revealed by DM is that (i) when an agent has subjective probabilities toward some partition of candidate objective probability distributions, the agent should adopt the linear mixture of the objective proba-

vocates of imprecise Bayesianism. Bradley (2017) goes even further and maintains that when an imprecise representor is permissible, further opinionation is always rationally permissible. Bradley's idea is that if it's rationally permissible to adopt a non-singleton set of probability functions as one's representor in some evidential context, it is likewise rationally permissible to adopt any member of this set as one's representor in said evidential context.

bility distributions that incorporates their uncertainty about the underlying objective probabilities, but (ii) probabilistic independence isn't generally preserved under linear mixtures. However, the fact that probabilistic independence isn't *generally* preserved under linear mixtures does not by itself imply that the particular probabilistic independencies between the decision-maker's intervention to x and X's non-effects are not preserved. Might there be some reason to think that these particular independencies always manifest some property that renders them special in a sense that guarantees their preservation in linear mixtures?

It is easy to prove that the unconditional probabilistic independence of X and Y in a set of probability distributions is preserved under their linear mixtures when the unconditional probabilities of either X or Y are the same in every initial candidate distribution (see Appendix).[33] This means that if there were some constraint on the chances of genuine choices that renders them the same in every candidate objective distribution, then the probabilistic independencies that obtain (by the CMC) between any genuine choice and its non-effects in every candidate objective distribution would be preserved in any linear mixture of the candidate distributions.[34] So if we supplement a constraint like this one with the now familiar idea

[33]Though the appendix contains a proof of this result, it is perhaps worth providing an intuitive explanation here. If the unconditional probabilities for some variable, X, are the same in every objective distribution under consideration, then conditioning on $X = x$ does not give us any information about which underlying objective distribution is true (since the candidate distributions all agree about the probability that $X = x$). Thus even if some of the candidate distributions assign different probabilities to $Y = y$, learning that $X = x$ does not confirm or disconfirm any of these distributions, and our subjective probability for $Y = y$ should thus stand pat. Meanwhile, if we learn that $Y = y$, this *can* confirm or disconfirm which underlying objective distribution is true (since the distributions are allowed to disagree about their probabilities for Y), but this doesn't mean that the subjective probabilities for X should change since every candidate distribution says the same exact thing about the probabilities for X. This allows us to see why these independencies are preserved, unlike those in DM.

[34]One might argue that MG's characterization of interventions already accomplishes this feat, but the argument for this point would undercut MG's ability to use evidential decision theory to vindicate causal-decision-theoretic recommendations. Though we have characterized interventions largely in terms of their location in causal graphs, MG have more to say about how we should construe intervention variables. Specifically, they maintain that we should partition the intervention variable for X so that it contains (i) a value on which one can condition to deterministically set X to x for every x in X and (ii) a value that corresponds to *not* intervening on X (i.e., to allowing the probability distribution over X to be determined by X's causes in $\mathbf{V}$). Since the observational or "unmanipulated" distribution is that which results from conditioning on the value of the intervention variable that corresponds to *not intervening*, MG maintain that we should think of the intervention variable as actually conditioned on this value, and should think

that genuine choices are interventions, then when choice is ordinary (non-exotic), we secure the result that if a decision-maker accepts a causal graph according to which Y is a non-effect of X, they should regard the genuine choice to x as probabilistically independent of Y in their subjective probability distribution (since the CMC and the accepted graph jointly imply that the decision-maker should regard the intervention to x as probabilistically independent of Y in every candidate objective distribution). But are there any such constraints on the chances of choices that are at all plausible? And what about when choice is exotic?

Before taking up these questions, I should say a bit more about what I mean by "chances"—especially since my usage of the term deviates in some important respects from other philosophers'. When I speak of the chance of anything—as I do in the title of this paper—my intention is simply to denote an objective probability distribution that is constrained by the CMC (or some aspect of an objective probability distribution that is constrained by the CMC), rather than a subjective (credal) probability distribution. These distributions are objective because they can be true or false of the world and provide grounds for deference,[35] but there are also some features of these distributions that disqualify them from counting as examples of the Lewisian chance distributions that sometimes dominate philosophical discussion. First, while Lewis maintains that the chances of past events must be 0 or 1, this cannot be true of the distributions at play since we need to permit

of the values that correspond to intervening as merely hypothetical. Thus one might maintain that MG do assign interventions the same antecedent chances in every distribution—namely, certainty to *not intervening* and zero to every value that corresponds to intervening. But in so doing, MG break from evidential decision theory in a rather severe way, and thus rule out the possibility of using evidential decision theory to vindicate causal-decision-theoretic recommendations. Formally, they introduce a species of counterfactual supposition that can not be spelled out in terms of the standard definition of conditional probability that evidential decision theorists like Jeffrey deploy (since it involves conditioning on events that are not assigned positive probability). Informally, the decision theory no longer asks us to consider what's likely given our actual action, but rather asks us to consider what *would* be likely given a *hypothetical* action. This latter kind of supposition has more in common with extant versions of causal decision theory that deploy distinctly counterfactual kinds of supposition (e.g., imaging) than with evidential decision theory. The same trouble arises if we seek to vindicate causal-decision-theoretic recommendations by deploying Pearl's (2009) do-operator, since $P(y|do(x))$ is not a true conditional probability, despite appearances.

[35]For those familiar with the literature on structural equation models, we can regard an objective probability distribution over $\mathbf{V}$ as true when it is generated by a true structural equation model over $\mathbf{V}$. See Stern (2017) for further discussion.

correlations between variables about the past and variables about the present and future in order to adequately represent causal relationships between them, and variables whose values have been conditioned on are trivially uncorrelated with any other variables.[36] Second, Lewis maintains that it's not always rational to defer to an objective chance distribution even when you're certain that it's true, because you might have "inadmissible" evidence about the future that provides you with reason to deviate from the known chance distribution. For example, according to Lewis, even if I am certain that the chance that a coin will land heads on the next toss is 0.5, I should assign the toss a subjective credence of 1 if an oracle has told me that it will, in fact, land heads (since the oracle's proclamation is constitutes "inadmissible" evidence). The reason that Lewis needs to include the proviso about inadmissible evidence is actually the same as the reason that he treats the chances of past events as 1 or 0—namely, Lewis seeks to understand *the* chance distribution that obtains at a world at a time, and thereby must condition the distribution on the relevant world's history, rather than on some specific body of evidence that is particular to an individual. But since I don't share this metaphysical ambition of Lewis's, I am happy to relativize chance distributions to evidential context (rather than to world history) in a manner that allows us to maintain that it's always rational to defer to an objective chance distribution when certain that it's true—e.g., so that the chance of the fair coin landing heads is itself 1 given an evidential context in which the oracle's proclamation is included as evidence. The idea here is that because the chance distribution in which you are certain says that $P(heads|oracle_{heads}) = 1$ (prior to updating on your evidence), you should be certain that it will land heads upon becoming certain of the oracle's proclamation. So according to my treatment of the chances, you should defer to an objective chance distribution when you are certain that it's true. It's just that you should defer to the chance distribution that is itself updated on your ("inadmissible") evidence.

[36] Were Lewis's chance distributions paired with causal graphs that include variables for the past, then there would be rampant violations of the Causal Faithfulness Condition and Causal Minimality Condition. An anonymous referee points out that Lewisians may be able to interpret the objective probability distributions in causal Bayes nets as Lewis's (1980) ur-chance functions, or as Lewis's objective chance functions at a time before any of the variables in the graph take on their values.

With this throat-clearing about the meaning of "chance" in place, we are now in a position to propose and consider constraints on chances of choices. One option here is to take up the commonsensical view introduced at the outset of this paper, according to which the chances of genuine choices must not distinguish between the decision-maker's options. Let us call this view "Chance Indifference" (CI).

> **Chance Indifference (CI)**: A choice from some menu of options is genuine only if for any two of the options, x and y, $Ch(x) = Ch(y)$.

According to CI, genuine choice requires the chance of every option (or intervention) to be the same as every other on the grounds that the chances of choices cannot compel us to choose in any particular way. Since this means that the chances of every option must take on a specific value in every distribution that is compatible with genuine choice, it immediately follows that the chance of every option is the same in every distribution that is compatible with genuine choice—e.g., when a choice is between two options, it follows from CI that each option is assigned an objective probability of 0.5 in every distribution that is compatible with genuine choice, from which it trivially follows that all of the distributions that are compatible with genuine choice agree about the objective probability assigned to each option.

But before we consider the case for CI in greater detail, it is worth noting that there is another constraint on the chances of choices that more directly requires that these chances be the same in every live objective probability distribution. Specifically, one can maintain that the chances of genuine choices must be *transparent* to the decision-maker in the sense that the decision-maker must be certain of the chances that they will choose a particular way as they settle for themselves what to do. This constraint does not license a constraint on the numeric values of the chances of genuine choices in the same way that CI does, but rather just necessitates that they be the same in every candidate chance distribution that is compatible with genuine choice—which, again, is exactly what we need in order to ensure that decision-makers should regard their choices (interventions) as (subjectively) probabilistically independent of their accepted non-effects when choice is not exotic. This means that we can potentially rescue MG's insight (at least when choice is not exotic) by incorporating lessons from the philosophy of action about the "self-presentational" immediacy of intentional action.[37] If we have so-called "non-observational" knowledge of how we're likely to choose as we settle for ourselves what to do when we make a genuine choice, then we're licensed to be certain of the chances that we'll genuinely choose in any particular way, regardless of what these chances are. Let us call the view that the chances of genuine choices are transparent in this way "Chance Transparency" (CT).

> **Chance Transparency (CT)**: A choice from some menu of options is genuine only if for each individual option, x, there is a number, n, such that $P(Ch(x) = n) = 1$.

Strictly speaking, CT doesn't follow from CI (since CI can be satisfied without the decision-maker knowing it),[38] but any decision-maker who represents their own choice as satisfying CI will likewise represent their choice as satisfying CT (since every objective probability distribution under consideration will agree that the chance of every option is the same). The converse does not hold. That is, when a decision-maker represents their choice as satisfying CT, it does not follow that they will represent their choice as satisfying CI (since the decision-maker could be sure about the chances of their options without being sure that these chances are equivalent). This means that CI implies strictly more substantive constraints than CT on how agents must represent their own choices when computing the CEU of genuinely deciding to act. Thus we might be tempted to side with CT over CI on the grounds that we shouldn't saddle our view of genuine choice with CI's extra commitments about self-representation. But there is a cost to siding with CT over CI. Specifically, while CT deals elegantly with the issues that stem from the CMC's implausibility as a constraint on rational subjective probabilities, it does nothing to secure causal-decision-theoretic recommendations when choice is exotic. The same is not true of CI. This is the focus of the next section.

[37] See Anscombe (1963), Moran (2001), and Paul (2009) for prominent discussion of this aspect of intentional action in the philosophy of action literature.

[38] I am grateful to an anonymous referee for bringing this possibility to my attention.

5 The Chances of Exotic Choices

Since the CMC admits spurious correlations between interventions and their non-effects when choice is exotic—or more specifically, when some value of a variable that is causally downstream from both the intervention and the relevant non-effect(s) has been conditioned on—the most straightforward way for the interventionist to vindicate causal-decision-theoretic recommendations is to maintain that the decision-maker has no genuine choice if their choice is genuinely exotic. Though this may seem like punting when it comes to exotic choice (insofar as it involves maintaining that there is no rational decision when choice is exotic), I have argued in the past that such a conception of choice can legitimate an alternative updating procedure that secures causal-decision-theoretic recommendations when a decision-maker has exotic evidence, wherein we update on the exogneous intervention to bring about the exotic evidence, rather than the exotic evidence itself (because the former update is compatible with the genuineness of choice, while the latter is not).[39] So, according to this line of reasoning, just as the Newcomb subject should be confident that they're not intervening, but should be concerned solely

[39]See Stern (2021). In the context of explaining why my approach to updating on exotic evidence is consistent with the view that it's conceptually impossible to make an exotic choice—i.e., the view that "no agent can have evidence about the outcome of her choice as she makes it (because the nature of choice precludes this possibility)," I said the following (pp. 561-562):

> "Since the intervention to bring about the exotic evidence is not exotic itself (because the intervention is not causally downstream from the agent's choice), there is a sense in which the agent must treat her evidence as non-exotic (and therefore a sense in which her choice should be regarded as non-exotic) in order for the evidential autonomy of her choice to be preserved—that is, she must update on the non-exotic intervention rather than the exotic evidence itself.

This aspect of my view may have gone unnoticed by Gallow (forthcoming), who applies my approach to updating on exotic evidence to examples where the decision-maker is certain that their evidence is genuinely exotic (in the sense that the decision-maker is certain that the evidence obtains because of the choice that they're in the process of making, rather than because of some other cause(s)). But this kind of case exemplifies a scenario where the decision-maker should rule out the possibility that their evidence is actually non-exotic (in the sense required for modeling it as an intervention) and should thus be certain that they are not making a genuine choice and that the standards of practical rationality therefore do not apply. This means that Gallow's examples are better construed as cases where my approach would say it doesn't matter what you do (because the decision-maker can be sure that their choice isn't genuine) rather than cases where my approach delivers the results that Gallow says it does.

with what's likely in the event that they're intervening (given the interventionist's conception of what genuine choice requires), the BP subject should be concerned solely with what is likely in the event that their newborn will be underweight for causally exogenous reasons, even though they have epistemic reason to entertain possibilities in which their newborn will be underweight because of their current choice (since genuine choice is compatible with the former kind of evidence, but not the latter).[40]

But what plausible conception of choice delivers the result that there is no genuine choice when the decision-maker's evidence is genuinely causally downstream from their choice? I have previously argued that we can secure this result by adopting a Ramseyan conception of choice according to which genuine choices are *evidentially autonomous* (or subjectively probabilistically independent) from anything that the decision-maker takes to be settled.[41] When choice is ordinary (non-exotic), this evidential autonomy is secured by modeling the decision-maker's options as interventions (since the CMC implies that the decision-maker's interventions are probabilistically independent from their non-effects). But when a decision-maker takes some consequence of their putative choice to be settled, this evidential autonomy is generally compromised (since the CMC permits correlations between the decision-maker's interventions and its effects).[42]

There are two potential issues with my previous line of reasoning. First, absent substantial constraints on the chances of choices like CI and CT, evidential autonomy is not secured by the assumption that genuine choices are interventions even when choice is ordinary. This is the lesson of DM in a nutshell.[43] Second,

[40]The BP subject has epistemic reason to entertain these possibilities because their evidence comes from a clairvoyant oracle.

[41]Both here and in my previous work, I follow Ramsey (1929) in using "settled" in a subjective evidential sense according to which the value of a variable is settled for a decision-maker only when the decision-maker treats its value as known. When choice is ordinary (not exotic), this means that only past events are treated as settled, but it does not mean that all past events are treated as settled. When choice is exotic, the possibility of treating future events as settled is introduced.

[42]If we grant the Causal Faithfulness Condition (according to which the only independencies compatible with a graph are those that follow from the CMC), then this evidential autonomy is *always* compromised.

[43]In my previous work, I "abstract away" from this difficulty by limiting my focus to decision-making contexts wherein the decision-maker does not entertain multiple causal hypotheses. See Stern (2021, fn. 13).

even though multiple authors have argued for the intuitive plausibility of the claim that genuine choices are evidentially autonomous in roughly the sense described here,[44] it is unclear what is added (if anything) by maintaining that choices are evidentially autonomous if we are already committed to the view that genuine choices are *causally autonomous* in the sense that is required for them to be accurately modeled as interventions. That is, since the view that genuine choices are interventions already captures a sense in which our choices must be *up to us*—i.e., by requiring that our choices must be uncaused by the factors under consideration in $\mathbf{V}$—it is unclear what is gained by articulating a further (evidential) sense in which our choices must be up to us. Why think that causal autonomy is not enough autonomy for genuine choice? When choice is putatively exotic, this issue is of great importance since genuinely exotic choice is plausibly compatible with the assumption of causal autonomy,[45] but not with the assumption of evidential autonomy.

We have already seen that the first of these problems can be solved by maintaining that in order for a choice to be genuine, it not only must be an exogenous intervention, but also must conform to some significant constraint regarding the chances of choices (e.g., CI or CT). Might the same be true of the second problem? That is, might either CI or CT provide us with reason to think that genuine choices must be evidentially autonomous in addition to causally autonomous?

Let us consider CT first. The view that genuine choices are causally exogenous interventions whose chances are transparent to the decision-maker does not accomplish this feat since there is nothing that prevents the chances of choices from being transparent to the decision-maker when the decision-maker gets evidence that they take to be correlated with their choice. This is perhaps easiest to see in a case where a decision-maker is certain of a single objective probability distribution. Here, CT is trivially satisfied no matter what evidence the decision-maker has since there is only objective probability distribution at play, and disagreement between candidate objective probability distributions is thereby impossible. That is, even when we condition on some exotic evidence that is cor-

[44]See Joyce (2007), Liu and Price (2020), Ramsey (1929), and Velleman (1989).

[45]The plausibility of this compatibility arguably depends on whether exotic choice involves causal loops between an agent's choice and their exotic evidence.

related with the decision-maker's intervention in the single objective probability distribution under consideration,[46] CT will be satisfied since the result of conditioning on some evidence in a single probability distribution is another (updated) single probability distribution, and the chances of choices are trivially transparent to the decision-maker whenever there is just one objective probability distribution under consideration. This means that there is no real tension between an absence of evidential autonomy and CT. So it is hard to see how CT could provide us with any explanation of why genuine choice should be considered evidentially autonomous in addition to causally autonomous.

What about CI? Since CI constrains the decision-maker's representation of their choice more than CT, it is possible that the extra commitments contained within CI endow it with the ability to succeed where CT fails. Indeed, when we posit additional numerical constraints on the chances of genuine choices (as we do when we move from CT to CI), we acquire reason to maintain that decision-makers should regard their own decisions as evidentially autonomous because we introduce the threat of evidence compromising the genuineness of choice by forcing the chances of choices to depart from their enforced (indifferent) values.[47]

Suppose, for example, that you're in a decision-making context that is exactly like BP except that you have no foreknowledge of your future newborn's weight. That is, you're in a decision-making context that many actual pregnant women confront; you simply have a choice whether to smoke or not during pregnancy and have every reason to believe that smoking would causally promote states of affairs in which your future newborn is born underweight. If we assume both CI and that genuine choices are interventions, then so long as you can be sure

[46]Here, if the decision-maker is rational, their subjective probability distribution should be identical to this objective probability distribution since they are certain of the objective probability distribution.

[47]Of course, there are infinitely many logically possible views according to which the chances of choices could be numerically constrained—each with its own specification of the numeric values that the chances of choices must take on—but CI is arguably the only constraint of this stronger variety that has any plausible rationale since it's somewhat plausible that the chances don't nudge the decision-maker in any direction when choice is genuine, but not at all plausible that the chances must nudge the decision-maker towards some particular option(s) to some particular extent when choice is genuine. For example, it would be plainly bizarre to maintain that when a decision-maker is confronted with a choice between two options, the chance of one option must be twice the chance of the other.

that any evidence you will obtain will be ordinary (non-exotic), it follows from the CMC that you can rest assured that any evidence that you might obtain will be (subjectively and objectively) probabilistically independent from your choice (intervention), and therefore cannot compromise the genuineness of your choice by affecting the chances of how you will choose. However, if we acknowledge the possibility that you will receive evidence that is correlated with your choice in the objective probability distributions that are compatible with your genuine choice (e.g., the genuinely exotic evidence that your newborn will be born underweight), then there exists a possibility that the chances of your choices will be forced to deviate from their indifferent values in correspondence with the prior conditional chances of your choices given this exotic evidence. For example, in any reasonable objective distribution that you might consider, receiving the evidence that your newborn is underweight would increase the chance that you'll intervene to smoke to some value above 0.5 and would thereby compromise the genuineness of your choice. To the extent that there is good reason to block this threat (so that the genuineness of choice does not depend on one's evidence), then, there is good reason for the defender of CI to maintain that genuine choice is evidentially autonomous. Thus CI (unlike CT) furnishes us with an explanation of the view that an agent has no genuine choice to make when their evidence is genuinely causally downstream from their intervention on the grounds that such choice is not evidentially autonomous. Were genuine choices not evidentially autonomous, then the chances of one's choices could be infringed upon by incoming evidence.

One might object that this line of reasoning involves justifying one intuitively justified autonomy constraint (evidential autonomy) with another (CI), and therefore does not strengthen the footing of my previous interventionist treatment of exotic choice. But there is an asymmetry between evidential autonomy and CI—specifically, while there is no obvious theoretical reason to posit evidential autonomy on top of causal autonomy, it is clear that we must posit some substantial constraint on the chances of choices (like CI or CT) on top of the view that choices are exogenous interventions in order to obtain causal-decision-theoretic recommendations while maintaining that we should maximize CEU. This makes CI into a skeleton key of sorts—i.e., when the interventionist augments their conception of choice by maintaining that the chances mustn't distinguish between the decision-maker's choices, the problems stemming from exotic choice and the CMC's inapplicability to subjective probabilities are solved in one fell swoop.

6 Conclusion

So what should we believe about the chances of choices? The answer to this question depends on more than the decision-theoretic issues I've surveyed here. But to the extent that it's valuable to secure causal-decision-theoretic recommendations while sticking with the evidentialist's insight that we should maximize CEU, there appears to be reason to adopt some significant constraint on the chances of choices. If there is principled reason to set aside a decision theory's application to exotic choice, then CT arguably provides the interventionist with everything required to rescue MG's idea that we can attain causal-decision-theoretic recommendations by maximizing CEU and treating decision-theoretic options as interventions. But if we should be concerned with a decision theory's application to exotic choice (as many philosophers contend),[48] then CI is attractive—especially since CI (i) capably solves all of MG's problems in one fell swoop and (ii) expresses a commonsensical and pervasive intuition about the autonomy of choice. Of course, one might respond here by arguing that CI is naïve and probably false on action-theoretic grounds. But I have nothing against this response. In fact, to my ear, this sounds like the beginnings of a novel argument against causal-decision-theoretic recommendations from within the interventionist approach to decision theory. And if that's where this paper takes us, then that's fine by me. Thus my aim is not to weigh in on whether we should follow causal decision theorists' advice. Nor is it to argue that we should maximize CEU. It is rather just to show that we can use graphical causal models and evidential-decision-theoretic machinery to attain causal-decision-theoretic recommendations by adopting significant constraints on the chances of genuine choices.

[48]Gallow (forthcoming), Hitchcock (2016), Lewis (1982), Rabinowicz (2009), and Stern (2021) are examples.

7 Appendix

Let Q denote any arbitrary linear mixture of some set of probability distributions that unanimously agree that two arbitrary events, X and Y, are probabilistically independent. The following proof demonstrates that if the set of probability distributions likewise unanimously assign a probability of c to X, then X and Y are guaranteed to be probabilistically independent according to Q.

$$\begin{aligned}
Q(X|Y) &= \frac{Q(X \cap Y)}{Q(Y)} \\
&= \frac{\sum_{i=1}^{n} \alpha_i P_i(X \cap Y)}{\sum_{i=1}^{n} \alpha_i P_i(Y)} \\
&= \frac{\sum_{i=1}^{n} \alpha_i P_i(X) P_i(Y)}{\sum_{i=1}^{n} \alpha_i P_i(Y)} \\
&= \frac{\sum_{i=1}^{n} \alpha_i c P_i(Y)}{\sum_{i=1}^{n} \alpha_i P_i(Y)} \\
&= \frac{c \sum_{i=1}^{n} \alpha_i P_i(Y)}{\sum_{i=1}^{n} \alpha_i P_i(Y)} \\
&= c \\
&= Q(X)
\end{aligned}$$

8 Acknowledgements

For helpful discussion and comments, I am grateful to Benjamin Eva, Shanna Slank, Rush Stewart, two anonymous referees, and the audience at the Workshop on Issues in Dynamic Decision Theory in Konstanz, Germany. Rush Stewart also deserves credit for helping me to see the particular form of the result in the appendix.

9 References

Anscombe, G.E.M. (1963). *Intention (2nd Ed.)*. Oxford: Blackwell Press.

Bovens, L. and Hartmann, S. (2003). *Bayesian Epistemology*. Oxford: Oxford University Press.

Bradley, R. (2018). *Decision Theory with a Human Face*. Cambridge: Cambridge University Press.

Eells, E. (1982). *Rational Decision and Causality*. Cambridge: Cambridge University Press.

Eva, B. (2019). Principles of Indifference. *The Journal of Philosophy* 116: 390–411.

Gallow, J.D. (forthcoming). Decision and Foreknowledge. *Noûs*.

Hausman, D. (1999). Lessons from Quantum Mechanics. *Synthese* 121: 79–92.

Hitchcock, C. (2016). Conditioning, Intervening, and Decision. *Synthese* 193: 1157–1176.

Hoover, K. (2003). Nonstationary Time Series, Cointegration, and the Principle of the Common Cause. *British Journal for the Philosophy of Science* 54: 527–551.

Jeffrey, R. (1983). *The Logic of Decision (2nd ed.)*. University of Chicago Press.

Jeffrey, R. (2004). *Subjective Probability: The Real Thing*. Cambridge: Cambridge University Press.

Joyce, J. (1999). *The Foundations of Causal Decision Theory*. Cambridge: Cambridge University Press.

Joyce, J. (2007). Are Newcomb Problems Really Decisions? *Synthese* 156: 537–562.

Joyce, J. (2010). A Defense of Imprecise Credences in Inference and Decision Making. *Philosophical Perspectives* 24: 281–323.

Levi, I. (2009). Why Indeterminate Probability Is Rational. *Journal of Applied Logic* 7: 364–376.

Lewis, D. (1980). A Subjectivist's Guide to Objective Chance. In *IFS: Conditionals, Belief, Decision, Chance and Time*. Dordrecht: Springer Netherlands, pp. 267–297.

Lewis, D. (1982). Letter to Wlodek Rabinowicz.

Lewis, D. (1986). Causal Decision Theory. In *Philosophical Papers: Volume 2*. Oxford: Oxford University Press, pp. 214–240.

Liu, Y. and Price, H. (2020). Ramsey and Joyce on Deliberation and Prediction. *Synthese* 197: 4365–4386.

Meek, G. and Glymour, C. (1994). Conditioning and Intervening. *The British Journal for the Philosophy of Science* 45: 1001–1021.

Moran, R. (2001). *Authority and Estrangement: An Essay on Self Knowledge*. Princeton: Princeton University Press.

Näger, P. (2016) The Causal Problem of Entanglement. *Synthese* 193: 1127–1155.

Nozick, R. (1969) Nozick, Newcomb's Problem and Two Principles of Choice. In N. Rescher (ed.) *Essays in Honor of Carl G. Hempel*. Dordrecht: D. Reidel: pp. 114–46.

Papineau, D. (2001). Evidentialism Reconsidered. *Noûs* 35: 239–259.

Paul, S. (2009). How We Know What We're Doing. *Philosophers' Imprint* 9: 1–24.

Pearl, J. (1993). Comment: Graphical Models, Causality, and Intervention. *Statistical Science* 8: 266–269.

Pearl, J. (2009). *Causality: Models, Reasoning and Inference.* 2nd edition. Cambridge: Cambridge University Press.

Pearl, J. and Mackenzie, D. (2018). *The Book of Why: the New Science of Cause and Effect.* Basic Books.

Peters, J., Janzing, D., and Schölkopf, B. (2017). *Elements of Causal Inference*, Cambridge, MA: MIT Press.

Price, H. (1986). Against Causal Decision Theory. *Synthese* 67: 195–212.

Rabinowicz, W. (2009). Letter from Long Ago: On Causal Decision Theory and Centered Chances. In J. Lars-Gorän, J. Österberg, and R. Śliwiński (eds.) *Logic, Ethics, and all that Jazz: Essays in Honour of Jordan Howard Sobel.* Uppsala, pp. 247–273.

Ramsey, F. P. (1929). General Propositions and Causality. In D. H. Mellor (ed.) *Foundations: Essays in Philosophy, Logic, Mathematics and Economics.* London: Routledge, pp. 133–151.

Slezak, P. (2006). Demons, Decisions and Liars: Newcomb's Malin Génie. *Theory and Decision* 61: 277–303.

Slezak, P. (2023). *Spectator in the Cartesian Theater: Where Theories of Mind Went Wrong since Descartes.* London: Rowman and Littlefield.

Sober, E. (2001). Venetian Sea Levels, British Bread Prices, and the Principle of the Common Cause. *British Journal for the Philosophy of Science* 52: 331–346.

Spirtes, P., Glymour, C., and Scheines, R. (2000): *Causation, Prediction and Search.* Cambridge, MA: MIT Press.

Stern, R. (2017). Interventionist Decision Theory. *Synthese* 194: 4133–4153.

Stern, R. (2018). Diagnosing Newcomb's Problem with Causal Graphs. In A. Ahmed (ed.) *Newcomb's Problem (Classical Philosophical Arguments).* Cambridge: Cambridge University Press, pp. 201–220.

Stern, R. (2021). An Interventionist's Guide to Exotic Choice. *Mind* 130: 537–566.

Velleman, J. D. (1989). Epistemic Freedom. *Pacific Philosophical Quarterly* 70: 73–97.

Zhang, J., Seidenfeld, T., and Liu, H. (2019). Subjective Causal Networks and Indeterminate Suppositional Credences. *Synthese* 198: 6571–6597.