

A Causal Theory of Suppositional Reasoning*

Alexander Gebharter

Christian J. Feldbacher-Escamilla

Michał Sikorski

Abstract: Suppositions can be classified as indicative vs. subjunctive and full vs. partial. We propose a causal account of suppositional reasoning that naturally unifies all four types of reasoning based on this classification, provides a justification of the rather heterogenous canonical update rules, and gives rise to a new update rule for the partial subjunctive case in terms of generalized imaging.

Following orthodoxy, we can distinguish between supposing in the indicative and in the subjunctive mood (see, e.g., [Eva & Hartmann, 2021](#)). When an agent supposes a in the *indicative* mood, she revises her epistemic state just as if she had learned that a . She might update her credences in a 's consequences as well as parts of a 's genesis. When she supposes a in the *subjunctive* mood, on the other hand, the agent assumes that a while keeping a 's history fixed. Thus, supposing a might have an influence on what the agent believes is affected by a , but not on what the agent believes happened before or independently of a . She revises her epistemic state as if she had learned that a had been brought about by a “local miracle”.¹

*This is a draft paper. Please do not cite or quote without permission. The final version of this paper is forthcoming under the following bibliographical data: Gebharter, A., Feldbacher-Escamilla, C. J., & Sikorski, M. (forthcoming). A causal theory of suppositional reasoning. *Philosophical Studies*. doi:10.1007/s10670-023-00684-5

¹It has been argued that the indicative and the subjunctive mood reflect merely grammatical surface features and do not necessarily correspond to distinct update rules. For example, the sentence “If the flag had been up, the

We use a simple toy example throughout the paper: Alice works in an office. Assume that Alice supposes that the fire alarm goes off (a). If she does so in the indicative mood, then she will also come to a couple of other beliefs. She will, for example, believe that there might be a fire (f), that the building soon will get very hot (h), that her colleague Bob, who she knows is in the building at the time, heard the alarm too (b), and that the building will be evacuated and employees will meet at the meeting point in front of the main entrance (e). If Alice supposes the alarm to go off in the subjunctive mood, on the other hand, then she will only come to believe the claim about Bob’s hearing the alarm too (b) and about the evacuation (e). The reason is that she epistemically behaves as if the alarm went off independently of what had happened before. So she would only have to come to the beliefs regarding the consequences of the alarm going off, but not to any other beliefs such as the one about the fire (f) or about increased heat (h) that have to do only with the event’s cause.

Suppositions can also be classified as full and partial suppositions. In the case of a *full* supposition, an agent treats their belief like certain knowledge. In the case of a *partial* supposition, on the other hand, the agent only increases her credence to a degree less than certainty. Accordingly, if Alice’s supposition that the alarm went off (a) is a full supposition, then Alice treats it like certain knowledge. If her supposition is a partial supposition, on the other hand, she only increases her credence in the alarm’s activation (a) to some degree below certainty.

The distinctions introduced give rise to the following classification:

full indicative supposition (FIS)	partial indicative supposition (PIS)
full subjunctive supposition (FSS)	partial subjunctive supposition (PSS)

King would have been in the castle” seems to be a subjunctive conditional that needs to be interpreted as expressing an evidential relation (cf. Günther, 2022; Kratzer, 1989). Regardless of whether the indicative vs. subjunctive distinction is substantive, the syntactic markings of these two moods do not have to always align with the semantic distinction captured by distinctive updates rules (cf. Bennett, 2003; Khoo, 2015a). In this paper, we allow the canonical distinction used throughout the paper to depart from the grammatical mood used. We would also like to acknowledge that the semantic distinction is discussed under different names in the literature (see, e.g., Dudman, 1984; Khoo, 2015b; Lindström & Rabinowicz, 1992; Rott, 1999).

All of these combinations have their role in everyday life. Alice might reason according to FIS if, for example, her friend Jill just told her that the alarm went off and Alice fully trusts Jill in this matter. She would reason along the lines of PIS if Jill would give her the same piece of information, but Alice does not fully trust Jill's verdict (for whatever reason), thus allowing for some uncertainty. Likewise for FSS and PSS: Alice reasons according to FSS if, for example, Jill told her that the alarm was activated due to a malfunction and she fully trusts Jill about this. As an effect, Alice would only update the consequences (b and e) of the alarm going off (a) while keeping its history (f and h) fixed as if the alarm was activated due to a local miracle. Finally, Alice would reason in accordance with PSS if Jill told her that the alarm went off by accident, but she does not fully trust Jill in that matter. Since the alarm having been activated due to a malfunction would not be informative about whether there was a fire (because whether there was a fire would not cause the alarm to go off in that case), Alice's credences about the history (f and h) of the alarm going off (a) will not change, but she might update her credences in the alarm's activation's consequences (b and e) depending on how much she trusts Jill's verdict.

In this paper, we discuss all four kinds of suppositions from a causal perspective. We show how reasoning based on these types can be captured by simple updating and/or transformations of a single causal model which the agent believes to adequately represent the causal structure underlying the specific situation. We then demonstrate how each kind of suppositional reasoning in the causal account relates to one of the canonical update rules for suppositional reasoning: Bayesian update (FIS), Jeffrey conditionalization (PIS) (Jeffrey, 1983), and generalised imaging (FSS) (Lewis, 1976, 1981). There is no generally accepted update rule for PSS in the philosophical literature (cf. Günther, 2018; Eva & Hartmann, 2021; Fusco, 2023 on Jeffrey imaging). We formulate a new update rule that naturally follows from our analysis. We take the fact that such a heterogeneous set of update rules falls out from different transformations of a single causal model as an indicator of the unificatory power of the causal account. In addition to unifying the different types of suppositional reasoning, the analysis shows that suppositional reasoning and especially update rules for the subjunctive mood in terms of generalized imaging require guidance by a fully-fledged account of causation in order to avoid possible causal pitfalls.

1 A causal account of suppositional reasoning

Our causal account is based on the causal interpretation of Bayesian networks (BNs) (Pearl, 2000; Spirtes, Glymour, & Scheines, 1993). Nodes are random variables whose set we denote by \mathbf{V} . Random variables X_i are represented by upper case letters; their values by lower case letters x_i , where \bar{x}_i stands for x_i 's negation. Directed edges (\longrightarrow) between nodes are interpreted as direct causal dependence. BNs (regardless of whether interpreted causally or not) are defined as pairs of directed acyclic graphs and probability distributions that factor as

$$Pr(X_1, \dots, X_n) = \prod_{i=1}^n Pr(X_i | \mathbf{Par}(X_i)), \quad (1)$$

where $\mathbf{Par}(X_i)$ stands for the set of a variable X_i 's parents, i.e., the variables X_j connected to X_i via a directed edge $X_j \longrightarrow X_i$. The probabilities at the right hand side of Equation 1 are called X_i 's *parameters*.

We use the fire alarm example introduced earlier for illustrating the causal account of suppositional reasoning. We assume that the causal BN in Figure 1(a) represents the causal structure Alice believes to underlie the situation at hand. The variables involved are all binary:

- A : whether the alarm goes off
- F : whether there is fire
- H : whether it is hot in the building
- B : whether Bob hears the alarm
- E : whether the building is evacuated

The task of unifying all four types of suppositional reasoning amounts to showing that each type can be squeezed out of the causal structure in Figure 1(a). We use the fire alarm example as a proxy for arbitrarily complex causal settings; all methods and transformations we use can be generalized.

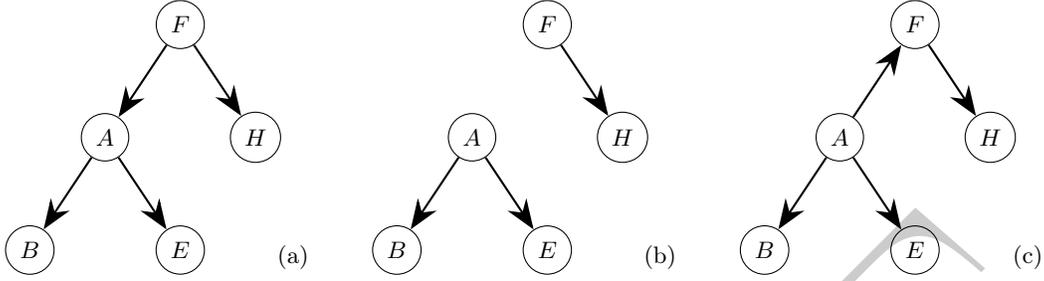


Figure 1

FIG. FIS can be captured by ordinary conditionalization on the basis of the causal structure in Figure 1(a), which amounts to computing

$$Pr_a^{FIS}(\cdot) = Pr(\cdot|a), \quad (2)$$

where $Pr(\cdot)$ is the original probability distribution associated with the structure in Figure 1(a) and $Pr_a^{FIS}(\cdot)$ represents Alice’s updated distribution after full indicative supposition of a . Since F, H, B, E are each d -connected² to A , full indicative supposition of a might lead to a change in credences in each of the events represented by these variables. This fits our intuitions: If Alice supposes a in the full indicative mood, her credences for f, h, b , and e can be expected to go up as well.

Here is a numerical example: The probability distribution of the BN depicted in Figure 1(a) is fully determined by the parameters in Table 1. The base rate of fires in office buildings is 0.5%, the chance that fire detectors indicate fire is 98%, the chance of a false alarm is 15%, etc. By applying Equation 1 to the structure in Figure 1(a), we get the joint distribution $Pr(F, H, A, B, E)$. By marginalization we can then compute the following probabilities:³

² X_i, X_j are d -connected iff they are connected by a chain of arrows not featuring a collider $\rightarrow X_k \leftarrow$. Variables that are d -connected can be correlated. Variables that are d -separated (i.e., not d -connected) are always independent. For the full definition also covering the conditional versions, see (Pearl, 2000, sec. 1.2.3).

³All numbers are rounded to three digits.

$$\begin{array}{ccccc}
Pr(f) = 0.005 & Pr(h|f) = 1.00 & Pr(a|f) = 0.98 & Pr(b|a) = 0.95 & Pr(e|a) = 0.98 \\
Pr(h|\bar{f}) = 0.05 & Pr(a|\bar{f}) = 0.15 & Pr(b|\bar{a}) = 0.01 & Pr(e|\bar{a}) = 0.01 &
\end{array}$$

Table 1

$Pr(f)$	$Pr(h)$	$Pr(a)$	$Pr(b)$	$Pr(e)$
0.005	0.055	0.154	0.155	0.16

These probabilities represent Alice’s credences in these propositions before supposing a . If Alice supposes that the alarm goes off in the full indicative mood, we end up with the following updated probabilities:

$Pr_a^{FIS}(f)$	$Pr_a^{FIS}(h)$	$Pr_a^{FIS}(a)$	$Pr_a^{FIS}(b)$	$Pr_a^{FIS}(e)$
0.032	0.08	1	0.95	0.98

As expected, all of Alice’s credences go up.

PIS. This case requires that we increase a ’s probability to a degree less than 1 before updating. Equation 1 allows us to store any joint probability distribution in terms of the model’s variables’ probabilities given their parents. Thus, the probability distributions over exogenous variables – variables with no incoming arrows – can be manipulated individually without having to change any of the BN’s other parameters. But since A is not exogenous in Figure 1(a), its distribution cannot be changed in isolation. Here is one strategy to overcome this problem: For any endogenous variable – a variable with at least one incoming arrow – we can construct a new BN in which this variable is exogenous. Note that this new BN does not fully preserve the causal structure. However, it preserves how probabilistic information spreads amongst variables according to the original BN’s causal structure. Since we only need to calculate the impact of changing A ’s value with respect to all variables belonging to A ’s history and consequences, this is all we need. For constructing this new BN, we first choose an ordering of variables agreeing with the direction of the arrows in the original structure. It does not matter which such ordering we choose, though choosing one that keeps the overall number of arrows to draw sparse is preferable in order to keep computational complexity low. One such ordering is $\langle F, H, A, B, E \rangle$. Next, we move A to the beginning of the

ordering. This results in the ordering $\langle A, F, H, B, E \rangle$. If we then apply the chain rule formula

$$Pr(X_1, \dots, X_n) = \prod_{i=1}^n Pr(X_i | X_1, \dots, X_{i-1}) \quad (3)$$

to $\langle A, F, H, B, E \rangle$, we get

$$Pr(A, F, H, B, E) = Pr(A)Pr(F|A)Pr(H|A, F)Pr(B|A, F, H)Pr(E|A, F, H, B). \quad (4)$$

From Equation 1 applied to the original structure in Figure 1(a) we know that $Pr(H|A, F) = Pr(H|F)$, $Pr(B|A, F, H) = Pr(B|A)$, and $Pr(E|A, F, H, B) = Pr(E|A)$. Hence, Equation 4 can be simplified to

$$Pr(A, F, H, B, E) = Pr(A)Pr(F|A)Pr(H|F)Pr(B|A)Pr(E|A). \quad (5)$$

With the definition of a BN and Equation 1 it follows that the structure in Figure 1(c) forms another BN satisfying Equation 1. This new BN has A as an exogenous variable and A 's probability distribution can be changed without changing any of the new BN's other parameters. We can now compute the epistemic effects of Alice supposing a in the partial indicative mood as

$$Pr_a^{PIS}(\cdot) = Pr^*(\cdot), \quad (6)$$

where $Pr^*(\cdot)$ is the probability distribution associated with the structure in Figure 1(c) after a 's probability has been increased to some extent. Since A is d -connected to all other variables in Figure 1(c), we can expect an influence of increasing a 's probability on F, H, B, E , which fits our intuitions: If Alice supposes that the alarm goes off in the partial indicative mood, then her credences in f, h, b , and e can be expected to go up as well. If the probability of a were raised to 0.8 in the numerical example introduced earlier (see Table 1), Alice's updated credences would be:

$$\begin{array}{ccccc}
Pr_a^{PIS}(f) & Pr_a^{PIS}(h) & Pr_a^{PIS}(a) & Pr_a^{PIS}(b) & Pr_a^{PIS}(e) \\
0.025 & 0.074 & 0.8 & 0.762 & 0.786
\end{array}$$

If we compare these probabilities with the ones in the full indicative supposition case, we see that PIS has lower impact than FIS, which is expected because the degree of belief in a after the supposition is lower in the former than in the latter case.

Before going on, we would like to address a possible worry. One might be worried that the newly created BN in which we can directly manipulate A 's probability distribution does not correctly represent the causal structure Alice believes to underly the situation at hand. This is true, but, as indicated above, unproblematic. The new BN should indeed not be causally interpreted. It should rather be seen as a reasoning device generated on the basis of that causal structure in [Figure 1\(a\)](#) Alice assumes. The arrow structure of the newly created BN does not have to exactly match the structure of the original causal BN when it comes to PIS. The reason is that PIS (and the indicative mood in general) requires Alice to treat the new information – i.e., the updated probability distribution over A – in such a way that it may impact Alice's credences in A 's consequences as well as in A 's history. For this kind of evidential update, the specific direction of the arrows does not matter as long as all d -connection relations featured by the original causal BN are preserved in the newly created inference device. Recall that d -connection captures all the variables that might be affected by a change in A or its probability distribution regardless of whether they belong to A 's consequences or history (see [fn. 2](#)). Strictly preserving the direction of arrows is only necessary for reasoning in the subjunctive mood according to which changes in A or its probability distribution are expected to impact A 's consequences or effects only.

FSS. We propose to handle FSS in terms of *surgical interventions* ([Pearl, 2000](#)). We first modify the causal structure in [Figure 1\(a\)](#) by deleting all arrows into A . This results in the structure in [Figure 1\(b\)](#). Next, one assumes that A takes value a . One can then apply [Equation 1](#) to the modified structure (b) for computing Alice's updated distribution after full subjunctive supposition of a . We follow Pearl and mark the procedure of deleting all arrows into A and assigning probability 1 to $A = a$ with a hat symbol as \hat{a} and the post-intervention distribution as $Pr_{\hat{a}}(\cdot) = Pr(\cdot|\hat{a})$. We

can compute the distribution one gets by supposing a in the full subjunctive mood as

$$Pr_a^{FSS}(\cdot) = Pr_a(\cdot). \quad (7)$$

Again, the result fits our intuitions: Assigning probability 1 to $A = a$ amounts to treating a like certain knowledge, and deleting the arrow into A guarantees that supposing a will not influence any of Alice’s credences about events that are not consequences of A . This procedure can be considered to flesh out what Lewis (1976) has called a “minimal revision”. More generally, our treatment of FIS and FSS can be reformulated in terms of causal modeling (cf. Pearl, 2000, sec. 1.3.1) as follows: Full indicative supposition of a amounts to treating a like it would have been observed, while full subjunctive supposition of a amounts to treating a like it would have been brought about by intervention.

If we calculate the probability distribution resulting from the parameters in Table 1 by full subjunctive supposition of a , we arrive at:

$Pr_a^{FSS}(f)$	$Pr_a^{FSS}(h)$	$Pr_a^{FSS}(a)$	$Pr_a^{FSS}(b)$	$Pr_a^{FSS}(e)$
0.005	0.055	1	0.95	0.98

As expected, supposing a in the full subjunctive mood has the same effect on Alice’s credences in b and e (values of consequences) as supposing a in the full indicative mood, but no effect at all on her credences in f and h (values of non-consequences).

PSS. To model PSS on the basis of the causal structure in Figure 1(a), we have to do two things: Firstly, we need to guarantee that supposing a in the partial subjunctive mood does not influence any credences about non-consequences of A . Secondly, we need to increase the probability of a to some extent. The first task can, just like in FSS, be accomplished by deleting all arrows into A , which results in Figure 1(b). In order to accomplish the second task back then when we discussed PIS, we constructed the non-causal BN in (c) on the basis of the causal structure in (a). That BN featured A as an exogenous variable, which allowed us to directly manipulate A ’s probability distribution. A similar strategy is not possible for PSS, since PSS requires that only Alice’s credences about A ’s

consequences are updated. Luckily, such a strategy is also not necessary because cutting the arrows into A already guarantees that A is an exogenous variable. Thus, to achieve the second task for PSS, we can simply increase a 's probability and update on the basis of the structure in (b). We can thus identify the epistemic effects of Alice supposing a in the partial subjunctive mood as

$$Pr_a^{PSS}(\cdot) = Pr_A^*(\cdot), \quad (8)$$

where $Pr_A^*(\cdot)$ is the distribution one gets from $Pr_A(\cdot)$ by increasing a 's probability to a value less than 1 and $Pr_{\bar{A}}(\cdot)$ is the distribution one gets from applying Equation 1 to the structure in Figure 1(b) which one gets from deleting all arrows into A .⁴ If we assume that Alice increases her original credence in a to 0.8 by supposing a in the partial subjunctive mood, we get the following probabilities (based on the parameters in Table 1):

$Pr_a^{PSS}(f)$	$Pr_a^{PSS}(h)$	$Pr_a^{PSS}(a)$	$Pr_a^{PSS}(b)$	$Pr_a^{PSS}(e)$
0.005	0.055	0.8	0.762	0.786

Again, this fits our intuitions: If Alice supposes that the alarm goes off in the partial subjunctive mood, then her credences in b and e (consequences) go up, while her credences in f and h (non-consequences) remain the same. However, the increase in credences in b and e is less than if Alice would have supposed a in the full subjunctive mood. Finally, Alice's new credences in the consequences of the alarm going off in b and e after increasing her degree of belief in a to 0.8 are identical to her credences in b and e after supposing a in the partial indicative mood. This reflects the intuition that whether Alice supposes a in the indicative or the subjunctive mood does not matter for a 's consequences as long as Alice's supposition is a partial supposition and she increases her degree of belief similarly in both cases.

⁴Note that computing $Pr_{\bar{A}}(\cdot)$ as well as $Pr_{\bar{a}}(\cdot)$ requires that the arrows into A are deleted. The difference between these two distributions is that the probability of a is set to 1 in $Pr_{\bar{a}}(\cdot)$, while the probability of a in $Pr_{\bar{A}}(\cdot)$ is unaltered, i.e., equals $Pr(a)$.

2 Canonical update rules and a new rule for PSS

In the previous section, we showed how the four different kinds of suppositional reasoning can be squeezed out of the causal structure the agent believes to represent the situation of interest. In this section, we show that the different types of update within our causal account give rise to the canonical update rules from the literature as well as to a novel update rule for PSS. This section is more technical and we will not come back to the fire alarm example for illustration purposes. The reason for this decision is that since the specifications of the update rules we give follows from the analysis provided in [section 1](#). Thus, nothing would be gained by repeating the same plausible outcomes in terms of the example again.

FIS and Bayesian update. One standard way to model full indicative suppositional reasoning is *Bayesian updating*. FIS on the basis of the causal structure in [Figure 1\(a\)](#) is closely linked to the orthodox treatment: $Pr_a^{FIS}(\cdot) = Pr(\cdot|a)$ simply amounts to the classical Bayesian strategy to identify posterior probabilities with conditional prior probabilities. The only difference is that the causal treatment operates not only on the probability distribution, but also on the causal structure. Strictly speaking, the causal structure does not matter for FIS if treated in isolation. Conditionalizing in the causal setup is only about reading off conditional probabilities. However, the relevance of causal structure will become apparent when taking a look at the full picture and also considering FSS and generalized imaging. Without it, the framework would not succeed in justifying all canonical update rules. And as we will see soon, causal structure also already plays a role for PIS.

PIS and Jeffrey conditionalization. The traditional way to model PIS is *Jeffrey conditionalization*:

$$Pr^*(x) = Pr(x|y) \cdot Pr^*(y) + Pr(x|\bar{y}) \cdot Pr^*(\bar{y}) \quad (9)$$

Pr is the prior distribution and Pr^* is the posterior probability after learning something about y , i.e., after changing the probability of y . This amounts to computing the distribution describing

Alice’s credences after supposing a in the partial indicative mood as

$$Pr^*(\cdot) = Pr(\cdot|a) \cdot Pr^*(a) + Pr(\cdot|\bar{a}) \cdot Pr^*(\bar{a}). \quad (10)$$

Note that Jeffrey conditionalization requires that the conditional probabilities $Pr(\cdot|a)$ and $Pr(\cdot|\bar{a})$ are rigid (cf. Schwan & Stern, 2017), meaning that they remain invariant when changing a ’s probability. Strictly speaking, the same rigidity is required also for how posterior distributions are computed in ordinary Bayesian update, which is a special case of Jeffrey conditionalization we get by setting $Pr(a) = 1$. Next, we show that Jeffrey conditionalization coincides with our causal treatment of PIS and that the causal treatment provides an independent justification of the rigidity assumption that is crucial for Jeffrey conditionalization.

We start with the structure in Figure 1(a). For causal settings it is essential “that each parent-child relationship in the network represents a stable and autonomous physical mechanism – in other words, that it is conceivable to change one such relationship *without* changing the others” (Pearl, 2000, p. 22). This feature of causal structures is often referred to as *modularity* (ibid.). It guarantees, among other things, that changing a variable’s value does not affect the model’s parameters. In fact, modularity does not only hold for causal BNs, but for BNs in general since it directly follows from the fact that a BN’s probability distribution factors according to Equation 1. Since the structure in (c) we constructed as a device for computing the effects of changing A ’s probability distribution preserves the d -connection relations in the original causally interpreted BN in (a), it will imply the same consequences when supposing a in the indicative mood (see also our discussion in section 1). Now, since raising a ’s probability from $Pr(a)$ to $Pr^*(a)$ – which corresponds to supposing a in the partial indicative mood – will not change the parameters of the model in (c) and since the probability of a is irrelevant for computing $Pr(\cdot|a)$ and $Pr(\cdot|\bar{a})$, it follows from our analysis of PIS from section 1 that $Pr^*(\cdot|a) = Pr(\cdot|a)$ and that $Pr^*(\cdot|\bar{a}) = Pr(\cdot|\bar{a})$. Thus, we get

$$Pr_a^{PIS}(\cdot) = Pr(\cdot|a) \cdot Pr^*(a) + Pr(\cdot|\bar{a}) \cdot Pr^*(\bar{a}) = Pr^*(\cdot). \quad (11)$$

FSS and generalized imaging. The usual way to tackle FSS is *generalized imaging*. The idea is that if one supposes a in the full subjunctive mood, the full probabilistic mass of possible non- a -worlds is moved over to possible a -worlds. Here is the corresponding update rule (cf. Leitgeb, 2017, p. 4):⁵

$$Pr(w|_T a) = \sum_{w' \in W} Pr(w') \cdot T_a(w', w) \quad (12)$$

w and w' are possible worlds, W is the set of all possible worlds, $Pr(w|_T a)$ is the probability of possible world w when supposing a , and $T_a(w', w)$ is a function indicating how much of the initial probability of w' is transferred over to w when supposing a . $T_a(w', w)$ outputs values within the interval $[0, 1]$, where 0 means that no probability mass, 1 that all probability mass, and any value r between these extremes that a proportion of r of the probability mass is transferred from w' to w .

For all propositions a and all $w' \in W$ it is assumed that

$$\sum_{w \in W} T_a(w', w) = 1, \quad (13)$$

which renders $T_a(w', w) = 1$ to express weights. It guarantees that the result of the probability transfer is again a probability function. Differently from Bayesian updating (for FIS) and Jeffrey conditionalization (for FSS), generalized imaging, as characterized here, does not determine a single posterior probability distribution but, depending on the exact characterization of the transfer function T_a , a variety of such possible probability distributions. Strictly speaking, it makes sense, therefore, to consider generalized imaging as a *family* of update rules. As we will see, different specifications of T_a will result in different specific update rules.

To link generalized imaging to our causal analysis, let us first agree that a possible world w can be understood as an instantiation of all causal variables in \mathbf{V} to particular values (cf. Briggs, 2012). Thus, $\langle f, h, a, b, e \rangle$ is a possible world, $\langle f, h, a, b, \bar{e} \rangle$ is a possible world, and so on. The proposition a

⁵Lewis (1976) originally suggested to analyse supposing a in the full subjunctive mood by shifting the entire probabilistic mass from non- a -worlds to the most similar a -world. This presupposes that each world has a unique most similar a -world, an assumption which he lifted in (1981).

then corresponds to the set of all possible worlds in which a is true, the variable A to the partition between a -worlds and non- a -worlds, etc.

According to our analysis, $Pr_a^{FSS}(\cdot) = Pr_{\hat{a}}(\cdot)$ when reasoning in the full subjunctive mood. Thus, we need to utilize the post-intervention distribution in order to specify the transfer function $T_a(w', w)$. There are two cases we have to distinguish: If a is true in world w to which the function transfers probability mass, then $T_a(w', w)$ needs to transfer a proportion of $r = Pr(w|\hat{a})$ of w' 's probability mass over to w . Thus, $T_a(w', w)$ will be $Pr(w|\hat{a})$. If a is not true in w , on the other side, no probability mass should be moved from w' to w . Hence, $T_a(w', w)$ will be 0. This amounts to

$$T_a(w', w) = Pr(w|\hat{a}), \quad (14)$$

where $Pr(w|\hat{a})$ can be computed as

$$Pr(w|\hat{a}) = \prod_{X \in \mathbf{V} \setminus \{A\}} Pr(A, F, H, B, E|\hat{a}). \quad (15)$$

Since $\sum_{w \in W} Pr(w|\hat{a}) = 1$, we get [Equation 13](#) relevant for generalized imaging for free, i.e., without the need to assume it in addition to our analysis. Since $\sum_{w' \in W} Pr(w') = 1$ and $Pr(w|\hat{a}) = 0$ for non- a -worlds w (this follows from [Equation 15](#)), the whole probability mass of all the possible worlds is distributed exclusively towards the worlds in which a is true. Furthermore, since we transfer from each world w' the exact same proportion of $r = Pr(w|\hat{a})$ of its probability mass to each world w in which a is true, each such world w will end up with a probability which is exactly the post intervention probability $Pr(w|\hat{a})$.⁶

Since we can compute the probability of any proposition or variable instantiation x as

$$Pr(x|_T a) = \sum_{w \in x} Pr(w|_T a), \quad (16)$$

⁶Note that our analysis is based on causal models and transformations of these models only. As such, it does not aim at encoding information about a similarity ordering into T_a . Though transformations of causal models could be used to specify a similarity ordering, this route is not further pursued in this paper.

we get

$$Pr_a^{FSS}(x) = Pr(x|Ta) = \sum_{w \in x} Pr(w|Ta) = Pr_{\hat{a}}(x) \quad (17)$$

if T_a is specified according to [Equation 14](#).

PSS and generalized imaging. PSS is the only one of the four types of suppositional reasoning for which there is no canonical update rule. Thus, in this paper we content ourselves by formulating an update rule in terms of generalized imaging that naturally follows from our account along the lines of our treatment of FSS. As before, when supposing a in the partial subjunctive mood, which implies increasing one's probability for a to a value less than 1, we want to transfer probability mass of non- a -worlds over to a -worlds. The difference is that we do not want to transfer all of a non- a -world's probability mass over to a worlds. Again, the main challenge consists in how to determine the transfer function. Our causal analysis suggests

$$T_a(w', w) = Pr_{\hat{A}}^*(w), \quad (18)$$

where $Pr_{\hat{A}}^*(\cdot)$ is defined as in [section 1](#) and $Pr_{\hat{A}}^*(w)$ can be computed as

$$Pr_{\hat{A}}^*(w) = \prod_{X \in \mathbf{V} \setminus \{A\}} Pr_{\hat{A}}^*(A, F, H, B, E). \quad (19)$$

Because we transfer from each world w' the exact same proportion of $r = Pr_{\hat{A}}^*(w)$ of its probability mass to each world w , each such world w will end up with a probability which is exactly the post intervention probability $Pr_{\hat{A}}^*(w)$.

Thus, with [Equation 16](#) we get

$$Pr_a^{PSS}(x) = Pr(x|Ta) = \sum_{w \in x} Pr(w|Ta) = Pr_{\hat{A}}^*(x) \quad (20)$$

if $T_a(w', w)$ is specified as $Pr_{\hat{A}}^*(w)$. We leave it to future research how this rule relates to other proposals in the contemporary literature such as Jeffrey imaging ([Günther, 2018](#); [Fusco, 2023](#)).

Comments on specifying generalized imaging. In a sense our specifications of the transfer function for FSS and PSS can be understood as answers to the question of how different kinds of intervention can be rewritten in terms of a transfer function of probability mass from worlds to worlds. One might thus worry that our treatments of FSS and PSS are too restrictive, which is not the case in the canonical treatment of FSS in terms of generalized imaging. However, we believe that such a restriction is necessary for causal settings. Let us briefly explain.

Without further restricting the transfer function, generalized imaging can easily get FSS and PSS wrong. Consider a simple causal model with $\mathbf{V} = \{X, Y\}$ and causal structure $X \rightarrow Y$. Let us specify $T_y(w', w) = Pr(w|y)$. Now assuming that x and y are correlated can lead to a change in x 's probability after supposing y in the full or partial subjunctive mood though x belongs to y 's history. The problem is that generalized imaging is such a wide concept that it can capture the transformation of any probability distribution to any other. Without further causal constraints, generalized imaging is blind to the restrictions coming with the specific causal structure underlying the situation. We are not saying that the specifications of the transfer function we propose are the only ones that would yield plausible results in causal settings, but it seems clear that one's specification needs to track some kind of intervention, which is necessary to guarantee that supposing a does not impact A 's history. A causal take on subjunctive supposition seems thus indispensable. We conjecture that it might thus also be possible to construct counter examples to other accounts relying on generalized imaging (e.g., Günther, 2018; Fusco, 2023). We postpone a proper analysis that can give these accounts the attention and care they deserve to another paper.

As a consequence of what we just said, generalized imaging, as a family of update rules, can cover not only FSS and PSS, but also FIS and PIS. In particular, we can force generalized imaging to capture FIS by specifying

$$T_a(w', w) = Pr(w|a), \tag{21}$$

and PIS by specifying

$$T_a(w', w) = Pr^*(w), \tag{22}$$

where $Pr^*(w)$ is specified as in [section 1](#). From this we get, again with [Equation 16](#),

$$Pr_a^{FIS}(x) = Pr(x|_T a) = \sum_{w \in x} Pr(w|_T a) = Pr(x|a) \quad (23)$$

for FIS and

$$Pr_a^{PIS}(x) = Pr(x|_T a) = \sum_{w \in x} Pr(w|_T a) = Pr^*(x) \quad (24)$$

for PIS.

3 Conclusion

We proposed a unifying framework for reasoning on the basis of the canonical types of suppositions: full indicative (FIS), partial indicative (PIS), full subjunctive (FSS), and partial subjunctive (PSS). It turned out that reasoning on the basis of all four types of suppositions can be captured by a single causal model held by the agent. We first showed how this can be done by different transformations of the original causal model. Then, we explored the connection of our causal account to the orthodox update rules for FIS, PIS, and FSS: Bayesian updating, Jeffrey conditionalization, and generalized imaging, respectively. Finally, we introduced a new update rule based on generalized imaging for the until now understudied PSS. The analysis we provided does not only provide a unified treatment of all four types of suppositional reasoning, but also shows that such an account is required in order to avoid possible causal pitfalls especially in cases of the subjunctive mood. We hope that the framework we developed will facilitate future comparative studies of different update rules and suppositional reasoning. In particular, it would be interesting to compare the new update rule for PSS with recent proposals from the literature such as the ones developed by [Günther \(2018\)](#) and [Fusco \(2023\)](#).

Acknowledgments: This research was funded by the European Union – Next Generation EU, Mission 4, Component 1, CUP I53D23006890001 (funds of Alexander Gebharter and Michał Sikorski) and the German Research Foundation (DFG), project *Rudolf Carnap, the Problem of In-*

duction, and the Choice of Scientific Frameworks, 545054032 (fund of Christian J. Feldbacher-Escamilla).

References

- Bennett, J. (2003). *A philosophical guide to conditionals*. Oxford: Oxford University Press.
- Briggs, R. (2012). Interventionist counterfactuals. *Philosophical Studies*, 160(1), 139–166.
- Dudman, V. H. (1984). Parsing ‘if’ sentences. *Analysis*, 44(4), 145–153.
- Eva, B., & Hartmann, S. (2021). The logic of partial supposition. *Analysis*, 81(2), 215–224.
- Fusco, M. (2023). Jeffrey imaging revisited. *Analysis*, 83(3), 447–464.
- Günther, M. (2018). Learning conditional information by Jeffrey imaging on Stalnaker conditionals. *Journal of Philosophical Logic*, 47(5), 851–876.
- Günther, M. (2022). Causal and evidential conditionals. *Minds and Machines*, 32, 613 - 626.
- Jeffrey, R. C. (1983). *The logic of decision*. Chicago: University of Chicago Press.
- Khoo, J. (2015a). On indicative and subjunctive conditionals. *Philosophers’ Imprint*, 15.
- Khoo, J. (2015b). On indicative and subjunctive conditionals. *Philosopher’s Imprint*, 15(32), 1–40.
- Kratzer, A. (1989). An investigation of the lumps of thought. *Linguistics and Philosophy*, 12, 607-653.
- Leitgeb, H. (2017). Imaging all the people. *Episteme*, 14(4), 463–479.
- Lewis, D. (1976). Probabilities of conditionals and conditional probabilities. *The Philosophical Review*, 85(3), 297–315.
- Lewis, D. (1981). Causal decision theory. *Australasian Journal of Philosophy*, 59(1), 5–30.
- Lindström, S., & Rabinowicz, W. (1992). Belief revision, epistemic conditionals and the Ramsey test. *Synthese*, 91(2), 195–237.
- Pearl, J. (2000). *Causality*. Cambridge: Cambridge University Press.

- Rott, H. (1999). Moody conditionals: Hamburgers, switches, and the tragic death of an American president. In J. Gerbrandy, M. Marx, M. de Rijke, & Y. Venema (Eds.), *Essays dedicated to Johan van Benthem on the occasion of his 50th birthday* (pp. 98–112). Amsterdam University Press.
- Schwan, B., & Stern, R. (2017). A causal understanding of when and when not to Jeffrey conditionalize. *Philosophers' Imprint*, 17(8), 1–21.
- Spirtes, P., Glymour, C., & Scheines, R. (1993). *Causation, prediction, and search*. Dordrecht: Springer.

DRAFT