The Logic of Partial Supposition

Benjamin Eva[1] and Stephan Hartmann[2]

1: Department of Philosophy, Duke University, benjamin.eva@duke.edu
2: Munich Center for Mathematical Philosophy, LMU Munich, s.hartmann@lmu.de

Forthcoming in *Analysis*

Abstract

According to orthodoxy, there are two basic moods of supposition: indicative and subjunctive. The most popular formalisations of the corresponding norms of suppositional judgement are given by Bayesian conditionalisation and Lewisian imaging, respectively. It is well known that Bayesian conditionalisation can be generalised (via Jeffrey conditionalisation) to provide a model for the norms of *partial* indicative supposition. This raises the question of whether imaging can likewise be generalised to model the norms of 'partial subjunctive supposition'. The present article casts doubt on whether the most natural generalisations of imaging are able to provide a plausible account of the norms of partial subjunctive supposition.

1 Introduction

According to orthodoxy, there are two basic moods of supposition: indicative and subjunctive. They are typically characterised roughly as follows.

Indicative Supposition: When a rational agent A supposes the truth of a proposition p in the indicative mood, they revise their epistemic state in exactly the way they would if they were to learn the truth of p.

Subjunctive Supposition: When a rational agent A supposes the truth of a proposition p in the subjunctive mood, they revise their epistemic state in exactly the way they would if they were to learn that p had been made true by some 'local miracle' or 'ideal intervention'.

When Alice supposes that the barometer reading is low in the indicative mood, she revises her epistemic state just as she would if she were to learn that the barometer reading is in fact low. Since seeing a low barometer reading would lead Alice to believe that a storm is imminent, Alice comes to believe that there will be a storm after indicatively supposing that the barometer reading is low. In contrast, when Alice supposes that the barometer reading is low in the subjunctive mood, she revises her epistemic state just as she would if she were to learn that the barometer reading had been *made to be low* by some local miracle.[1] Clearly, this will not lead Alice to believe that a storm is imminent, since the barometer has been tampered with and is therefore no longer indicative of future weather. However, if Alice were to subjunctively suppose that the atmospheric

[1]For example, that God had reached down to set the barometer reading by hand.

pressure is low, that would lead her to believe that there will be a storm, since atmospheric pressure causally influences the weather, regardless of how the atmospheric pressure has reached its actual value.

In Bayesian epistemology, the norms of indicative and subjunctive supposition are typically modelled by means of Bayesian conditionalisation and (variations of) Lewis's (1976) imaging rule, respectively.[2] In the indicative case, the motivation for positing conditionalisation as the primary norm of suppositional judgement is straightforward. Since the rationality norms for indicative supposition are phrased in terms of learning dynamics, and Bayesianism posits conditionalisation as the unique rational learning rule, Bayesians are compelled to posit conditionalisation as the fundamental norm of indicative supposition. The subjunctive case is less straightforward, and there has been substantial debate regarding the correct Bayesian formalisation of the norms of subjunctive supposition (see e.g. Lewis (1981) and Joyce (2008)). But most authors endorse some variation of Lewis's (1976) probabilistic imaging rule.

One challenge for advocates of specific normative models of suppositional judgement is to show how those models can be generalised to give an account of the norms of *partial supposition.* To illustrate, note that the process of standard suppositional reasoning is intuitively characterised by the following three step procedure.

1: The agent temporarily treats the supposed proposition p as certain knowledge.

2: The agent makes some suitable adjustments to their epistemic state, in order to coherently integrate the new supposition.

3: The agent deliberates/infers/plans/decides/evaluates on the basis of their newly adjusted post-supposition epistemic state.

The difference between indicative and subjunctive supposition can be characterised as a difference in the second step of the process – where the agent needs to adjust their epistemic state in order to coherently integrate the new supposition. When Alice indicatively supposes that the barometer reading is low she adjusts her epistemic judgements differently to when she makes the same supposition in the subjunctive mood. The difference between partial supposition and standard (full) supposition is a difference in the first step. Whereas instances of full supposition involve the agent temporarily treating the supposed proposition as certain knowledge (as in 1), instances of partial supposition involve the agent temporarily treating the supposed proposition as having *an increased degree of plausibility.* So when Alice partially supposes that the barometer reading is low, she does not temporarily treat the low reading as certainly true, but only as being more probable than she actually judges it to be. She then makes some suitable adjustments to ensure that this change to her epistemic judgements is integrated in a coherent state, and then goes on to use the newly adjusted state in her future deliberations.

As it turns out, the Bayesian model of indicative supposition already has a well known and widely employed generalisation that can be used to model the norms of partial indicative supposition, namely Jeffrey conditionalisation.

[2] We focus here on quantitive, credence based normative theories of supposition. For an overview of the relationship between these quantitative models and extant qualitative models articulated in terms of categorical belief change, see Eva, Shear and Fitelson (forthcoming).

Full Indicative (Conditionalisation)	Partial Indicative (Jeffrey Conditionalisation)
Full Subjunctive (Imaging)	Partial Subjunctive (?)

Figure 1: Theories of full and partial supposition.

Jeffrey Conditionalisation (JC): When an agent partially indicatively supposes a proposition p in the sense of temporarily increasing their credence in p to some new value x, they should adjust their epistemic state by Jeffrey conditionalising to set p's post supposition credence to x, i.e they should adopt the post supposition credence function c^* defined below, where c denotes the agent's prior credence function

$$c^*(-) = c(-|p) \cdot x + c(-|\neg p) \cdot (1 - x).$$

JC provides a natural generalisation of the standard Bayesian model of indicative supposition, and is widely accepted as the most plausible existing model of the norms of *partial* indicative supposition. But what of *partial subjunctive supposition*? Surprisingly, the literature appears to lack any prospective normative model of partial subjunctive supposition whatsoever.[3] However, there is a natural way that one might try to fill in the empty cell of the table in Figure 1 corresponding to the normative theory of partial subjunctive supposition, namely

Jeffrey Imaging (JI): When an agent partially subjunctively supposes a proposition p in the sense of temporarily increasing their credence in p to some new value x, they should adjust their epistemic state by *Jeffrey imaging* to set p's post supposition credence to x, i.e they should adopt the post supposition credence function c^* defined below, where c denotes the agent's prior credence function and c_p denotes the result of imaging c on p.

$$c^*(-) = c_p(-) \cdot x + c_{\neg p}(-) \cdot (1 - x)$$

Before evaluating the extent to which JI provides a satisfactory explication of the norms of partial subjunctive supposition, it is worth pausing to consider what role the notion of partial subjunctive supposition is supposed to play in our understanding of rational inference and decision making. In chapter 11 of *The Logic of Decision*, Richard Jeffrey considers the problem of generalising his theory of rational decision making to deal with cases in which intending (or deciding) to do some action x does not guarantee that x is actually performed.[4] Standard theories of rational decision making (including Jeffrey's) evaluate the rationality of deciding to x by calculating expected utility under the full supposition that x is successfully performed *with certainty.* In evidential decision theories (like Jeffrey's), this supposition is indicative. In causal decision theories, it is

[3]The Jeffrey imaging rule described below has appeared in the literature (see e.g. Crestani (1998) and Günther (2018)), but not in the context of discussions surrounding the norms of suppositional reasoning.

[4]By way of illustration, I might decide to play football on Sunday, whilst readily acknowledging that any number of factors (bad weather, illness etc) could prevent this decision from being actualised. For more discussion of cases like these and their significance for theories of rational decision making, see e.g. Bratman (1987, 2009) and Stern (forthcoming).

subjunctive. Jeffrey proposes that we can generalise standard decision theory to deal with cases where intending/deciding to x does not guarantee the successful performance of x by calculating expected utility under the partial supposition that x is successfully performed with some non maximal probability. Since his decision theory is evidential and the relevant species of supposition is therefore indicative, he uses JC to model this formally. But causal decision theorists will find this unsatisfactory for the same reasons that they disapprove of standard versions of evidential decision theory: namely that it renders non-causal correlations between acts (or intentions to act) and outcomes relevant to an agent's deliberations. Just as causal decision theorists think that we should evaluate the decision to x by calculating expected utility under the full subjunctive supposition that x is performed in the standard setting where deciding to x guarantees the performance of x, they will presumably also think that we should evaluate the decision to x by calculating expected utility under the partial subjunctive supposition that x is performed in the generalised setting where deciding to x does not guarantee the performance of x. So if causal decision theorists want to generalise their decision theories to allow for the evaluation of intentions whose fulfilment isn't guaranteed in something like the way proposed by Jeffrey, then it seems that they will first need to properly explicate the notion of partial subjunctive supposition. The most obvious way to do this is via JI.

The intuitive motivation for JI is that it stands in the same relation to standard imaging that JC does to standard conditionalisation. So assuming that the relationship between the norms of partial and full supposition is the same on both the indicative and subjunctive side, and that we are happy to accept JC as a normative model of partial indicative supposition, we can conclude that JI provides a compelling normative model of partial subjunctive supposition. The basic aim of this paper, though, will be to cast doubt on whether JI is capable of providing a plausible account of the norms of partial subjunctive supposition. Our starting point will be the following criterion for theories of (indicative or subjunctive) partial supposition:

Monotonicity Condition (MC): If $p \vdash q$ then partially supposing the truth of p (in the indicative or subjunctive moods) should never lead one to decrease one's credence in q.

The basic intuition behind MC is simple. If q is a logical consequence of p, then p constitutes a perfect guarantee of q's truth. So whenever I temporarily treat p as certain knowledge (as in full supposition) or increase my credence in p (as in partial supposition), this should not lead me to decrease my credence in q. Otherwise, I would have increased my credence in a perfect guarantee of q's truth whilst decreasing my credence in q itself. In slogan form: it's impossible to truly suppose p (either fully or partially) without also supposing (either fully or partially) p's logical consequences. It seems to us that this criterion is intuitively compelling for both subjunctive and indicative supposition. For example, in the generalised decision theoretic setting described above, agents evaluate potential decisions by calculating expected utility under the partial supposition that they perform the act to which the relevant decision corresponds. This makes sense because while deciding to act may not render the action certain, it certainly does not decrease the probability of its performance, and typically raises it. Likewise, deciding to act surely does not render the logical consequences of the relevant action less probable than they previously were. If I decide to play football, that both makes it more likely that I play football, and more likely that I either play football or basketball. This is exactly what is guaranteed by MC, in the case where we (following Jeffrey) use partial supposition to evaluate decisions in settings where deciding to act does not guarantee that the act if performed. If MC were violated, then in the course of evaluating a potential decision, we would use a probability distribution that treated the logical consequences of the relevant action as being less likely than we currently judge them to be, which seems implausible.

It is well known that that MC is satisfied by JC in full generality (see e.g. Eva and Hartmann (2018)),[5] and hence that MC makes no problems for partial indicative supposition. In the next section, we'll show that MC is systematically violated by JI, under all extant variations of the imaging rule, and hence that JI does not seem to adequately encode the norms of partial subjunctive supposition. One immediate corollary is that the relationship between the norms of full and partial subjunctive supposition seems to be fundamentally different to the relationship between the norms of full and partial indicative supposition.

2 JI and MC

We begin by recalling Lewis's definition of the standard imaging rule, which is widely employed as a normative Bayesian model of full subjunctive supposition (see e.g. Lewis (1981) and Joyce (2008)). W will denote the set of possible worlds, and W_p will denote the subset of W that satisfies p (for any sentence p). For any sentence p and any world w, $\sigma(p, w)$ denotes the 'closest' possible world to w at which p holds. For now, we assume that (i) $\sigma(p, w)$ picks out a single closest p-world, for each w, and (ii) $\sigma(p, w) = w$ whenever $w \models p$.[6] Lewis defines the result of 'imaging' a prior credence function c on p as follows,

$$c_p(w) = \begin{cases} c(w) + \sum\limits_{w' \in W_{\neg p} | w = \sigma(p, w')} c(w'), & \text{if } w \in W_p \\ 0, & \text{if } w \in W_{\neg p}. \end{cases}$$

The idea is that, when we subjunctively suppose p, we should adjust our credences by transferring the probability of each $\neg p$ -world to the closest p -world. We can now combine the above definition of the standard imaging rule with JI to check whether the resulting model of partial subjunctive supposition satisfies the MC criterion forwarded in the previous section.

Proposition 1 *JI does not generally satisfy MC.*

The proof of Proposition 1 demonstrates that Jeffrey imaging to increase the probability of a conjunction

[5] Important technical clarification: Jeffrey conditionalizing to increase the probability of p *relative to the partition* $\{p, \neg p\}$ always leads to an increase in the probability of p's logical consequences. This corresponds to the case in which the content of the partial supposition is only to regard p as more likely than previously, which is the case we're focusing on here. Jeffrey conditionalizing to increase the probability of p can lead to a decrease in the probability of p's logical consequences if one also updates the distribution of probability over subregions of the $\neg p$ cell of the partition, in which case the content of the partial supposition is not exhausted by p.

[6] Assumption (ii) is often referred to as 'strong centering'. It encodes the idea that each world is uniquely closest to itself.

$p \wedge q$ can lead to a decrease in the probability of one of the conjuncts. To illustrate the reason for this, suppose that the set of worlds W is given by $w_1 \models p, q$, $w_2 \models p, \neg q$, $w_3 \models \neg p, q$, $w_4 \models \neg p, \neg q$. Then any of the three possible choices of $\sigma(\neg(p \wedge q), w_1)$ can lead to MC violations when we use JI to increase the probability of $p \wedge q$. One might hope to escape the problem by relaxing our initial requirement that there is always a unique closest world for every proposition/world pair, i.e. to allow $\sigma(\neg(p \wedge q), w_1)$ to be a *set* of equally closest worlds. Gärdenfors (1982) generalised Lewis's original imaging rule in exactly this way. Specifically, the technique of 'general imaging' is defined as follows (where it is assumed that $\sigma(w, p) = \{w\}$ whenever $w \models p$). For every w, p, define a corresponding 'transfer function' $T_{w,p} : \sigma(w, p) \to [0, 1]$ with the property that $\sum_{w' \in \sigma(w,p)} T_{w,p}(w') = 1$. Intuitively, the transfer function describes how much of w's prior probability will be shifted to each of the (possibly many) closest p worlds when one images on p. Then, we can set

$$c_p(w) = \begin{cases} c(w) + \sum\limits_{w' \in W_{\neg p} | w \in \sigma(p, w')} c(w')\, T_{p,w'}(w), & \text{if } w \models p \\ 0, \quad \text{if } w \models \neg p. \end{cases}$$

This definition avoids the problem of having to choose a single world to represent $\sigma(\neg(p \wedge q), w_1)$, and it is natural to conjecture that setting $\sigma(\neg(p \wedge q), w_1) = \{w_2, w_3\}$ (for instance) may allow us to ensure that MC is satisfied when we plug the above definition of 'general imaging' into JI. This hope is misplaced.

Theorem 1 *There is no general imaging rule that allows for the satisfaction of MC by JI – i.e. however we specify the transfer function in general imaging, it is possible for JI to violate MC.*

There is one last possible escape route for the advocate of JI. Namely, to make the transfer function itself a function of the prior distribution, so that how one distributes post-imaging probability amongst the closest worlds depends on the prior probability of those worlds. This will block the route to the counterexamples used in the proof of Theorem 1. As it happens, Joyce (2010) proposes a variant of imaging, called 'proportional imaging', which makes the transfer function dependent on the prior distribution. Specifically, proportional imaging is the general imaging rule that one obtains by setting

$$T_{p,w} = \frac{c(w')}{\sum\limits_{w'' \in \sigma(p,w)} c(w'')},$$

i.e. proportional imaging sets the share of probability that gets transferred from w to w' (after imaging on p) to be equal with the proportion of the prior probability mass distributed over the worlds in $\sigma(p, w)$ that is stored at w'.

Theorem 2 *Assuming proportional imaging, JI doest not satisfy MC.*

So whichever of the extant variations of imaging one adopts as a normative model of full subjunctive supposition, the corresponding model of partial subjunctive supposition induced by JI will end up violating MC in some circumstances. So we are forced to choose between JI and MC as norms for partial subjunctive supposition. They can't both be right. Given that many authors accept imaging as a normative model of full subjunctive supposition, and that JI is the most natural partial generalisation of imaging, some may be tempted to reevaluate the intuitive normative status of MC. But whichever horn of the dilemma one embraces, it is now clear that the norms of partial subjunctive supposition must differ from their indicative counterparts in unexpected ways. Taxonomising and understanding these differences strikes us as a pressing philosophical project that remains largely untouched.

Acknowledgements

Many thanks to Jessica Collins, Melissa Fusco, Reuben Stern and an anonymous referee for their many helpful and insightful comments on earlier drafts of this paper. We also thank the Alexander von Humboldt Foundation and LMU Munich's Center for Advanced Studies (CAS) for financial support.

A Appendix: Proofs

A.1 Proposition 1

We show that Jeffrey imaging to increase the probability of a conjunction can lead the probability of the conjuncts to decrease. Let $W = \{w_1, w_2, w_3, w_4\}$ where $w_1 \models p, q,\ w_2 \models p, \neg q,\ w_3 \models \neg p, q,\ w_4 \models \neg p, \neg q$. Suppose we partially subjunctively suppose the truth of $p \wedge q$ and model our post supposition credences via JI. It is clear that $\sigma(p \wedge q, w_2) = \sigma(p \wedge q, w_3) = \sigma(p \wedge q, w_4) = w_1$ (since $W_{p\wedge q} = \{w_1\}$). However, there are two possibilities for $\sigma(\neg(p \wedge q), w_1)$, since both w_2 and w_3 satisfy $\neg(p \wedge q)$, and they both differ from w_1 by one truth value (with respect to the atomic propositions p and q). Now, suppose we choose the case where $\sigma(\neg(p \wedge q), w_1) = w_2$. In this case (setting $\alpha = c^*(p \wedge q) - c(p \wedge q)$), we have that

$$\begin{aligned}
c^*(q) &= c^*(w_1) + c^*(w_3) \\
&= c_{p\wedge q}(w_1)\, c^*(p \wedge q) + c_{\neg(p\wedge q)}(w_3)\, c^*(\neg(p \wedge q)) \\
&= c^*(p \wedge q) + c_{\neg(p\wedge q)}(w_3)\, c^*(\neg(p \wedge q)) \\
&= c^*(p \wedge q) + c(w_3)\, c^*(\neg(p \wedge q)) \quad \text{(since } \sigma(\neg\Gamma, w_1) = w_2) \\
&= c(w_1) + \alpha + c(w_3)\,(1 - c(w_1) - \alpha) \\
&= c(w_1) + \alpha + c(w_3) - c(w_1)\, c(w_3) - \alpha\, c(w_3).
\end{aligned}$$

So, since $c(q) = c(w_1) + c(w_3)$, we know that

$$c^*(q) - c(q) = \alpha - c(w_1)\, c(w_3) - \alpha\, c(w_3),$$

which is clearly negative for large values of $c(w_3)$ and small values of α. ∎

A.2 Theorem 1

Again, we use the same set W as before and let $\sigma(\neg(p \wedge q), w_1) = \{w_2, w_3\}$. First off, we let $c_1 := c(w_1)$, $c_3 := c(w_3)$, $\alpha := c^*(p \wedge q) - c(p \wedge q) = c^*(p \wedge q) - c_1$ (we know this is positive, by the assumption that $c^*(p \wedge q) \geq c(p \wedge q)$), $\beta := c_{\neg(p\wedge q)}(w_3) - c_3$ (again, we know this is positive, by the definition of imaging). Then,

$$\begin{aligned}
c^*(q) - c(q) &= c^*(w_1) + c^*(w_3) - c_1 - c_3 \\
&= c_{(p\wedge q)}(w_1)\, c^*(p \wedge q) + c_{\neg(p\wedge q)}(w_3)\, c^*(\neg(p \wedge q)) - c_1 - c_3 \\
&= c^*(p \wedge q) + c_{\neg(p\wedge q)}(w_3)\, c^*(\neg(p \wedge q)) - c_1 - c_3 \\
&= c_1 + \alpha + (c_3 + \beta)\,(1 - c_1 - \alpha) - c_1 - c_3 \\
&= \alpha + c_3 - c_1\, c_3 - \alpha c_3 + \beta - \beta\, c_1 - \alpha\,\beta - c_3 \\
&= \alpha\,(1 - c_3) + \beta\,(1 - c_1) - \alpha\,\beta - c_1\, c_3.
\end{aligned}$$

Let's define the function

$$f(\alpha, \beta, c_1, c_3) := \alpha\,(1 - c_3) + \beta\,(1 - c_1) - \alpha\,\beta - c_1\, c_3\,.$$

Now, clearly, this function is decreasing in c_1 and c_3, so to minimise f, let's set $c_1 + c_3 = 1$. Then, we have

$$\begin{aligned}
f(\alpha, \beta, c_1, c_3) &= f(\alpha, \beta, c_1, 1 - c_1) \\
&= \alpha\,(1 - (1 - c_1)) + \beta\,(1 - c_1) - \alpha\,\beta - c_1\,(1 - c_1) \\
&= \alpha\, c_1 + \beta - \beta\, c_1 - \alpha\,\beta - c_1 + c_1^2 \\
&= (\beta - c_1) \cdot (1 - \alpha - c_1).
\end{aligned}$$

The term on the right is always non-negative, since $\alpha + c_1 \leq 1$ (by definition of α). But the term on the left will always be negative except for the case in which $c_1 = \beta$ (we know that $\beta \leq c_1$ always holds by definition of β). So the only case in which the probability of q doesn't decrease here is where $c_1 = \beta$, i.e. where we transfer all of w_1's probability to w_3. For any other alternative choice of transfer function, we can find values for c_1, c_3, α that violate MC. ∎

A.3 Theorem 2

Let $W = \{w_1, w_2, w_3, w_4\}$ and let $\sigma(\neg(\{w_1\}), w_1) = \{w_2, w_4\}$, and let $c(w_1) = 0.3$, $c(w_2) = 0.08$, $c(w_3) = 0.22$ and $c(w_4) = 0.4$. Suppose that $c^*(w_1) = 0.35$ (c^* comes from c by Jeffrey imaging to increase the probability of w_1 from 0.3 to 0.35). The result of proportional imaging c on $\neg(\{w_1\})$ is as follows: $c_{\neg(\{w_1\})}(w_1) = 0$, $c_{\neg(\{w_1\})}(w_2) = 0.13$, $c_{\neg(\{w_1\})}(w_3) = 0.22$, and $c_{\neg(\{w_1\})}(w_4) = 0.65$. Hence $c^*(w_2) = 0.65 \cdot 0.13 = 0.0845$, $c^*(w_3) = 0.65 \cdot 0.22 = 0.143$, and $c^*(w_4) = 0.65 \cdot 0.65 = 0.4225$. So

$c^*(\neg\{w_4\}) = c^*(w_1) + c^*(w_2) + c^*(w_3) = 0.5775 < c(\neg\{w_4\})$. Since $w_1 \vdash \neg\{w_4\}$, this constitutes a violation of MC. ∎

References

Bratman, M. (1987). *Intention, Plans, and Practical Reason.* Stanford, CA: CSLI Publications.

Bratman, M. (2009). Intention, Belief, Practical, Theoretical. In S. Robertson (ed.): *Spheres of Reason: New Essays in the Philosophy of Normativity.* Oxford: Oxford University Press, pp. 29–61.

Crestani, F. (1998). Logical Imaging and Probabilistic Information Retrieval. In: C. van Rijsbergen, F. Crestani and M. Lalmas (eds.): *Information Retrieval, Uncertainty and Logics: Advanced Models for the Representation and Retrieval of Information.* Berlin: Springer, pp. 247–279.

Eva, B. and S. Hartmann (2018). Bayesian Argumentation and the Value of Logical Validity. *Psychological Review* 125(5): 806–821.

Eva, B., Shear, T. and B. Fitelson (forthcoming). *Four Approaches to Supposition.*

Gärdenfors, P. (1982). Imaging and Conditionalization. *The Journal of Philosophy* 79(12): 747–760.

Günther, M. (2018). Learning Conditionals by Jeffrey Imaging on Stalnaker Conditionals. *Journal of Philosophical Logic* 47(5): 851–876.

Jeffrey, R. (1965). *The Logic of Decision.*. Chicago, Ill.: The University of Chicago Press.

Joyce, J. (2008). *The Foundations of Causal Decision Theory.* Cambridge: Cambridge University Press.

Joyce, J. (2010). Causal Reasoning and Backtracking. *Philosophical Studies* 147: 139–154.

Lewis, D. (1976). Probabilities of Conditionals and Conditional Probabilities. *Philosophical Review* 85(3): 297–315.

Lewis, D. (1981). Causal Decision Theory. *Australasian Journal of Philosophy* 59(1): 5–30.

Stern, R. (forthcoming). When is it Rational to Make Plans?