Jeffrey Imaging Revisited*

Melissa Fusco

Abstract

In "The Logic of Partial Supposition", Ben Eva and Stephan Hartmann investigate *partial imaging*, a credence-revision method which combines the partiality familiar from Jeffrey Conditioning (Jeffrey, 1983) with the formal notion of imaging familiar from Lewis (1981). They argue that because partial imaging is non-monotonic, it "fail[s] to provide a plausible account of the norms of partial subjunctive suppositions" (1).

In this note, I present a notion of partial imaging that *does* satisfy monotonicity, and discuss some of the applications and ramifications. The account frames conditioning as a form of imaging, and rejects Gärdenfors (1982)'s principle of linearity.

1 Introduction

In "The Logic of Partial Supposition" (2021), Ben Eva and Stephan Hartmann investigate the notion of *partial imaging*, a credence-revision method which combines the partiality familiar from Jeffrey Conditioning (Jeffrey, 1983) with the formal notion of imaging familiar from Lewis (1981) and Gärdenfors (1982). As Eva and Hartmann (henceforth E&H) note, partial imaging occupies a natural "fourth quadrant" in a space of possible credence-revision methods (Table 1, highlighted):

	conditioning	imaging
total	$P_X(\cdot)$	$P^X(\cdot)$
partial	$P_{\{X\}}(\cdot)$	$P^{\{X\}}(\cdot)$

Table 1: Four ways of updating a prior $P(\cdot)$.

Whereas X is a proposition—an element of the algebra taken to be learned or supposed with certainty—$\{X\}$ is a partition of W, which provides new *input probabilities* for its members.[1] These new input probabilities may lie strictly between 0 and 1. The

*To appear. I am grateful to David Boylan, Simon Huttegger, Ginger Schultheis, Snow Zhang, and an anonymous *Analysis* referee for generous commentary and discussion.

[1]Everything I will talk about will be finite.

task of partial update—of either the “indicative” (conditionalising) or “counterfactual” (imaging) kind—is to redistribute credence over the space of possible worlds in accord with the input provided by $\{X\}$.

E&H argue for a pessimistic moral concerning partial imaging. “[T]he most obvious formal generalisations” of imaging to the partial case, they write, “fail to provide a plausible account of the norms of partial subjunctive suppositions” (1). In particular, E&H are interested in a *monotonicity condition*, according to which (directly) increasing one’s credence in a proposition p should never decrease the probability of any of p’s logical consequences. Montonicity is violated by the accounts of partial imaging E&H investigate.

The purpose of this note is twofold. The first is to present a notion of partial imaging that *does* satisfy monotonicity. The second is to evaluate its straightforwardness and intuitive appeal. The account I advocate for rejects Gärdenfors (1982)’s principle of linearity. In wrapping up, I note some helpful connections with recent work in philosophy of language.

2 Background

Table 1 follows Joyce (1999) in using subscripting for conditioning and superscripting for imaging. Standard conditioning is thus defined by the Ratio Formula. For any $w \in W$ and X s.t. $P(X) > 0$:

$$P_X(w) := \frac{P(w \wedge X)}{P(X)} \tag{1}$$

Partial, or Jeffrey, Conditioning for the simple partition $\{X, \neg X\}$ is:

$$P_{\{X,\neg X\}}(w) = (P_X(w) \cdot x) + (P_{\neg X}(w) \cdot (1 - x)) \tag{2}$$

where x and $(1 - x)$ are new input probabilities for X and $\neg X$, respectively.

Imaging comes in two flavours: *sharp* and *blurred* (or *general*). Both require a *selection function* σ, which takes a world and a proposition as arguments.

When imaging is sharp, $\sigma(w', X)$ is the unique world w to which w' wills its mass when P is imaged on the proposition X.

$$P^X(w) := \begin{cases} 0 & \text{if } w \in \overline{X} \\ P(w) + \sum_{w' \in \overline{X} | w = \sigma(w', X)} P(w') & \text{if } w \in X \end{cases} \tag{3}$$

A common way of understanding the selection function σ is that $\sigma(w', X)$ is the *closest* or *most similar* world to w where X is true.[2]

Similarity, however, admits of ties, as Gärdenfors (1982, §1) notes. He thus defines $\sigma(w, X)$ more generally as a *set* of worlds $Y \subseteq W$. The definition of general imaging additionally has recourse to a *transfer function* $T_{w,X} : \{v \in \sigma(w, X)\} \rightarrow [0, 1]$. For example, when $T_{u,X}(v) = .25$, then u sends exactly 25% of its probability mass to v when the probability space is imaged on X. $P^X(w)$ is defined with the aid of $\sigma(\cdot)$ and $T_{(\cdot)}$, as follows:

$$P^X(w) := \begin{cases} 0 & \text{if } w \in \overline{X} \\ P(w) + \sum_{w' \in \overline{X} | w \in \sigma(w', X)} P(w') \cdot T_{w',X}(w) & \text{if } w \in X \end{cases} \tag{4}$$

For any world w' and proposition X, we assume:

- **Success:** $\sigma(w', X) \subseteq X$
- **Strong Centring:** if $w' \in X$, then $\sigma(w', X) = \{w'\}$

On either the sharp (equation (3)) or general (equation (4)) version, imaging thus taps genuine extra-Bayesian structure. When a world w "dies" under imaging, $\sigma(\cdot)$ and $T_{(\cdot)}$ record how it bequeaths its probability mass to its survivors. w may dole out this mass unequally; the only requirement is that "it all goes somewhere": $\sum_{w' \in \sigma(w,X)} T_{w,X}(w') = 1$. For this reason, Lewis influentially described imaging as a process according to which

[2]My notation here closely follows that of Eva and Hartmann, though I adjust slightly in order not to beg questions about the relata of the semantic relation $\vDash$. For readability, I toggle freely between $\overline{X}$ and $\neg X$ throughout.

probability "is moved around" though it is "neither created nor destroyed" (1976, pg. 310); the presumptive contrast is that when a world "dies" under conditionalisation, probability mass *is* destroyed. (As we will see below, however, this framing is optional.)

Against this background, E&H approach partial imaging on the two-member partition $\{X, \neg X\}$ via the equation (EH):

$$P^{\{X,\neg X\}}(w) = (P^X(w) \cdot x) + (P^{\neg X}(w) \cdot (1 - x)) \tag{EH}$$

where x is the new input probability of X. As the reader can verify, (EH) presents partial imaging as a simple linear combination of $P^X(A)$ and $P^{\neg X}(A)$ (for arbitrary A), in the same way Jeffrey conditionalisation states the posterior as a linear combination of $P_X(A)$ and $P_{\neg X}(A)$.

2.1 Monotonicity

(EH) does seem natural. But as advertised, Eva & Hartmann find it unsatisfactory. To see why, we need a definition:

> A belief-revision method that changes $\{P(p), P(\neg p)\}$ to $\{P'(p), P'(\neg p)\}$ (where $P'(p) > P(p)$) is **monotonic** iff $p \vdash q$ entails $P'(q) \geq P(q)$.

E&H endorse a *Monotonicity principle* (MC), according to which any belief revision method—a fortiori, partial imaging—should be monotonic.

> 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. (Eva and Hartmann, §1)

Jeffrey conditionalisation is monotonic in the following sense: if one Jeffrey conditions in a way which raises the probability of p, *and* $\{p, \neg p\}$ *is the input partition*, then for any logical consequence q of p, the probability of q is also increased. (Without the italicized restriction, Jeffrey Conditionalisation is *not* monotonic. Example: begin with a uniform distribution on $\{pq, p\bar{q}, \bar{p}q, \bar{p}\bar{q}\}$ (as in Figure 1, below) and Jeffrey Condition to reduce $P(pq)$ from .25 to .1. This operation raises $P(p\bar{q})$ from .25 to .3 but reduces $P(p)$ from .5 to .4, where p is a logical consequence of $p\bar{q}$.[3])

Why, intuitively, does the partial imaging method described in (EH) violate (MC)? According to (EH), partial imaging in a way that increases the probability of X linearly (but unequally) combines two revision methods: fully imaging $P(\cdot)$ on X—call this the *major* component—and fully imaging $P(\cdot)$ on $\neg X$—call this the *minor* component. Hence even if a world in X gains probability mass overall, it will wind up transferring some of its probability mass away according to the minor component of the partial imaging operation. When this happens—assuming the X-world images unequally—it benefits some of X's logical consequences at the expense of others.

The following simple example (a more concrete version of the counterexample to (MC) discussed in E&H's appendix) provides a demonstration. Suppose we have a uniform probability measure $P(\cdot)$ over four worlds comprising two independent propositions p and q (top of Figure 1), and that we wish to partial-image in a way that increases the probability of $(p \wedge q)$—viz., $\{w_1\}$—from .25 to .3. According to (EH), the way to do this is to average $P^{(p\wedge q)}(\cdot)$, weighted to degree 0.3, with $P^{(\neg(p\wedge q))}(\cdot)$, weighted to degree 0.7. Even though the probability of w_1 strictly grows during the overall process, it will matter for posterior probabilities overall what $\sigma(w_1, (\neg(p \wedge q)))$ is—that is, which world w_1 sees as the closest *not*-$(p \wedge q)$ world.

[3] Thanks to Michael Nielsen for this simple example.

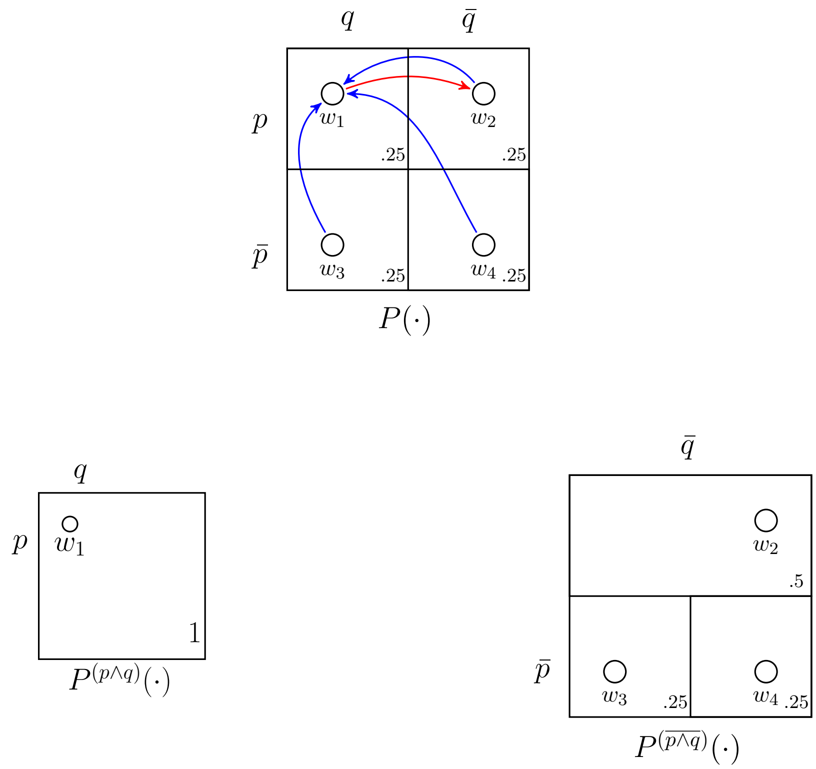

Figure 1: In a prior probability space $P(\cdot)$ (top), blue arrows show how $\neg(p \wedge q)$-worlds image on $(p \wedge q)$; red arrows show how $(p \wedge q)$-world(s) image on $\neg(p \wedge q)$. When $P'(p \wedge q) = .3$, (EH) calculates the posterior as a weighted linear combination (.3, .7) of $P^{(p \wedge q)}(\cdot)$ (bottom left) and $P^{\neg(p \wedge q)}(\cdot)$ (bottom right).

Suppose that $\sigma(w_1, \neg(p \wedge q)) = \{w_2\}$.[4] What results is a transformation on which $q = \{w_1, w_3\}$ loses probability; world w_3 loses more mass than w_1 gains. w_3 loses probability in the major component of the linear combination, of course, because $w_3 \notin (p \wedge q)$; but w_3 *also* loses probability in the minor component of the linear combination—imaging on $\neg(p \wedge q)$—because $w_3 \notin \sigma(w_1, \neg(p \wedge q))$. q, however, is obviously a logical consequence of $(p \wedge q)$. $P'(q) = .3 + .7(.25) < .5$, and so (MC) is violated.

[4] I use general imaging here, so $\sigma(w_1, \neg(p \wedge q))$ is a set (albeit a singleton set, $\{w_2\}$). In the case of sharp imaging, we can simply say instead that $\sigma(w_1, \neg(p \wedge q)) = w_2$. It follows that when imaging is sharp and partial imaging is calculated using (EH), violations of (MC) are not merely possible but *guaranteed*.

3 A Different Rule

Suppose we accept E&H's argument: Monotonicity governs all forms of (total and partial) update. We would then want a new monotonic rule for partial imaging to replace (EH). The rule I introduce below, (6), is also a way of creating a graded notion of imaging on $\{X\}$ from the model-theoretic tools already introduced. However, unlike (EH), it avoids incorporating imaging on $\neg X$. Intuitively, the new rule ensures that no world which grows in probability overall (like w_1) will transfer away any of its mass.

Suppose, as before, that we want to partially update the space from Figure 1 on $\{(p \wedge q), \neg(p \wedge q)\}$ from prior $P(p \wedge q) = .25$ to posterior $P'(p \wedge q) = .3$. Since w_1 is the unique pq-image of each of $\{w_2, w_3, w_4\}$, this can be accomplished by importing $.05/3 = .01\bar{6}$ units of probability mass from each. This means each world in the $\neg(p \wedge q)$-region retains $(.25 - .01\bar{6})/.25 = 14/15$ of its mass.

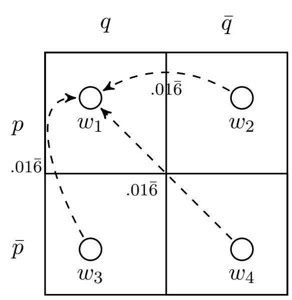

Figure 2: simple example of partial sharp imaging

That's it; we don't need to image the space on $\neg(p \wedge q)$ at all to satisfy the desideratum set by the new input prior.

In the general case, assume $P'(X) > P(X)$. Let $Z = \left(\frac{P'(\neg X)}{P(\neg X)}\right)$ (this was the role played by 14/15 in the foregoing example). In the sharp case, we can write a partial version of (3) as follows:

$$P'(w) = \begin{cases} Z \cdot P(w) & \text{if } w \in \overline{X} \\ P(w) + \sum_{w' \in \overline{X} | w = \sigma(w', X)} (1 - Z) \cdot P(w') & \text{if } w \in X \end{cases} \tag{5}$$

This example showcases partial imaging on the input partition $\{X, \neg X\}$, in a model where imaging is sharp.

Our final step is to upgrade this operation to the case of general, rather than sharp, imaging. Things do not change much: we make recourse to $T_{v,X}(\cdot)$, the function which further specifies what proportion of v's lost mass goes to which worlds. Once again, let $Z = \left(\frac{P'(\neg X)}{P(\neg X)}\right)$.

$$P'(w) = \begin{cases} Z \cdot P(w) & \text{if } w \in \overline{X} \\ P(w) + \sum_{v \in \overline{X} | w \in \sigma(v,X)} (1 - Z) \cdot P(v) \cdot T_{v,X}(w) & \text{if } w \in X \end{cases} \tag{6}$$

Figure 3 shows such a case, where probability mass transfer is split between w_0 and w_1:

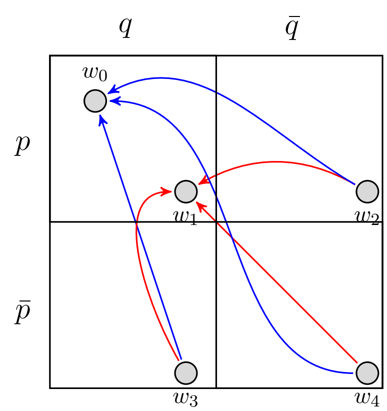

Figure 3: a simple example of general (or blurred) imaging, where $P'(p \wedge q) = .3$.

The $T_{(\cdot)}$ apparatus doesn't make any difference to the calculation in the example above, though of course it will to any further proposition r which partitions w_0 from w_1.

3.1 An n-ary Partition

The general version of imaging on a partition involves an input of the form $\{X_1, \ldots, X_n\}$, rather than the binary $\{X, \bar{X}\}$. Our new input probabilities are $P'(X_1), \ldots, P'(X_n)$.[5] Two natural questions arise: first, how can (6) be extended to such n-ary input partitions? Second, how should Monotonicity be formulated in such cases?

To extend (6) to the n-ary case, we identify which X_n are s.t. $P'(X_n) \leq P(X_n)$. Let us arrange them thus: $\{L_1, \ldots L_k, G_1, \ldots G_m\}$ are such that $P'(L_i) < P(L_i)$ ("L" for "losers") and $P'(G_j) \geq P(G_j)$ ("G" for "gainers"). I will use L for $\bigcup_i L_i$ and G for $\bigcup_j G_j$.

We want probability mass transfer to be determined by the "smaller" posteriors $P'(L_i)$ and the probability mass transfer information $\sigma(x, G)$ for $x \in L$, as we did in the simple example above. Each L_i in the shrinking part of the partition shrinks by a different amount (determined by the input probabilities); each G_j gains a certain proportion of the mass drawn from L, in the simplest way necessary for each cell to reach the posteriors enumerated in the input probabilities. So we write (6) in a way which is indexed to the partition-cell L_i (or G_i) of which w is a member.

$$P'(w) = \begin{cases} \frac{P'(L_i)}{P(L_i)} \cdot P(w) & \text{if } w \in L_i \\ P(w) + \lambda_i \left[\sum_{v \in L | w \in \sigma(v, G_i)} \left(1 - \frac{P'(L)}{P(L)}\right) \cdot P(v) \cdot T_{v,G_i}(w) \right] & \text{if } w \in G_i \end{cases} \tag{6$'$}$$

where, for arbitrary $w \in G_i$, λ_i is the proportion of the transferred mass that achieves the target posterior. Routine algebra gives the following value for λ_i:

$$\lambda_i := \frac{P'(G_i) - P(G_i)}{P(L) - P'(L)}$$

The analogue of monotonicity for this n-ary case is that any logical consequence of G gains mass:

[5] Here it is worth noting the n-ary form of Jeffrey Conditioning. Where $P'(X)$ is the new input probability for $X \in \{X\}$: $P'(w) = \sum_{X \in \{X\}} P(w \mid X) P'(X)$.

A belief-revision method that changes $\{P(L_1), \ldots P(L_k), P(G_1), \ldots P(G_m)\}$ to $\{P'(L_1), \ldots P'(L_k), P'(G_1), \ldots P'(G_m)\}$ is **n-ary monotonic** iff: $G \vdash q$ entails $P'(q) \geq P(q)$.

Note that no stronger conclusion can be drawn for arbitrary $G_i \subsetneq G$. To see why, suppose that $P(\cdot)$ is an indifferent prior over $\{p_1, p_2, p_3\}$ and that $P'(p_1) = .65, P'(p_2) = .34$, and $P'(p_3) = .01$. Hence $G = p_1 \cup p_2$ and $L = p_3$. Note that $G_i \vdash q$ does *not* entail $P'(q) \geq P(q)$: $G_i = p_2$ and $q = (p_2 \vee p_3)$ is a counterexample. Here we see that what matters about E&H's original statement of monotonicity is not (only) that proposition p gains mass, but that it is the *union* of *all* propositions in the input partition that gain mass (viz., $p = G$). As noted in §2.1, this is also the only sense in which partial *conditioning* is monotonic. So (6′) is as true to E&H's intuitive defense of Monotonicity as it can be, subject to the requirement that conditioning itself also satisfies their desideratum.

3.2 Proving Monotonicity

Recall that (partially) imaging $\{p, \bar{p}\}$ to $\{P'(p), P'(\bar{p})\}$ (where $P'(p) > P(p)$) is **monotonic** iff $p \vdash q$ entails $P'(q) \geq P(q)$. It's now easy to show:

Observation 1 (Monotonicity)**.** *Binary Partial Imaging, as defined by (6), is monotonic.*

Proof. Suppose $p \vdash q$. Then $q = p \cup r$ for some (possibly empty) $r \subseteq \bar{p}$. $P'(q) = P'(p) + P'(r)$. It is clear that $P'(p) \geq P(p)$, since for any $w \in p$, $P'(w) = P(w) + \alpha$, where α is a positive number.[6] If r is empty, then $p = q$, and we are done.

If r is nonempty: let v be an arbitrary world in r. By (6), $P'(v) = Z \cdot P(v)$, where $Z \in [0, 1]$. The remainder of v's probability mass is distributed across worlds $u \in p$. Hence it remains in the union of r and p. Hence it remains in q. □

Observation 2 (n-ary Monotonicity)**.** *n-ary Partial Imaging, as defined by (6′), is n-ary monotonic.*

[6] To wit, α is: $\sum_{v \in \overline{X} | w \in \sigma(v,X)} (1 - Z) \cdot P(v) \cdot T_{v,X}(w)$, by (6).

Proof. Similar to (binary) monotonicity. Suppose $G \vdash q$. Then $q = G \cup r$ for some (possibly empty) $r \subseteq L$. $P'(q) = P'(p) + P'(r)$. By stipulation, $P'(G) \geq P(G)$. If r is empty, then $G = q$, and we are done.

If r is nonempty: let v be an arbitrary world in r and let L_i be the cell of L such that $v \in L_i$. By (6′), $P'(v) = Z \cdot P(v)$, where $Z \in [0, 1]$. The remainder of v's probability mass is distributed across worlds $u \in G$. Hence it remains in the union of r and G. Hence it remains in q. □

3.3 The naturalness of this generalisation

For the remainder of this paper I return to discussing the case of imaging on a binary partition, since this is more directly in dialogue with E&H. We rejected (EH) (repeated below) as a form of partial imaging:

$$P^{\{X,\neg X\}}(w) = (P^X(w) \cdot x) + (P^{\neg X}(w) \cdot (1 - x))$$

in favor of (6) (repeated):

$$P^{\{X,\neg X\}}(w) = \begin{cases} Z \cdot P(w) & \text{if } w \in \overline{X} \\ P(w) + \sum_{v \in \overline{X} | w \in \sigma(v,X)} (1 - Z) \cdot P(v) \cdot T_{v,X}(w) & \text{if } w \in X \end{cases}$$

Is (6) *uglier* and *less natural* than (EH)?

I think it is not, if one looks at things "imaging-first". To see this, let us invert precedent and write (1) total and (2) partial (Jeffrey) conditioning as imaging. We'll do these in order.

We begin with total conditioning. We want to preserve the probability ratios between worlds in p when we "image" on p. We proceed by taking it in general that for any proposition X, and world $v \in \overline{X}$, when $P(\cdot)$ is condition-imaged on X, then

$$\sigma(v, X) = \{u \in X : P(u) > 0\} \tag{i}$$

Next, we set values $T_{v,X}(w)$ for $v \in \overline{X}$, $w \in \overline{X}$ as functions of $P(w)$ and $P(X)$, but not v. In particular,

$$T_{v,X}(w) = \frac{P(w)}{P(X)} \tag{ii}$$

According to (i), when a $\neg X$-world v loses probability under partial imaging to raise $P(X)$, it sees *every live* X-world as equally "similar" or "close". According to (ii), how much of v's lost mass gets transferred to an individual X-world w is determined simply by how much of w already contributes to the probability of X in the prior.

Officially, then, we should relativise the selection function $\sigma(\cdot)$ and the transfer function $T_{(\cdot)}$ to a further parameter (the prior, $P(\cdot)$). For any proposition X, and world $v \in \overline{X}$,

$$\sigma_P(v, X) = \{u \in X : P(u) > 0\} \tag{i′}$$

$$T^P_{v,X}(w) = \frac{P(w)}{P(X)} \tag{ii′}$$

(i) and (ii) make it possible to write (non-partial) conditioning as follows:

$$P'(w) = P'_X(w) = \begin{cases} 0 & \text{if } w \in \overline{X} \\ P(w) + \sum_{v\in\overline{X}} P(v) \cdot \frac{P(w)}{P(X)} & \text{if } w \in X \end{cases} \tag{7}$$

So, combining (6) and (7), we can write Jeffrey Conditioning (to directly increase $P(X)$) as:

$$P'(w) = \begin{cases} Z \cdot P(w) & \text{if } w \in \overline{X} \\ P(w) + \sum_{v\in\overline{X}} (1 - Z) \cdot P(v) \cdot \frac{P(w)}{P(X)} & \text{if } w \in X \end{cases} \tag{8}$$

...where, once again, $Z = \left(\frac{P'(\neg X)}{P(\neg X)}\right)$. I submit that (6) is not unnatural at all, when it's seen in comparison to (8). In sum:

	conditioning	imaging
total	(7)	(4)
partial	(8)	(6)

Table 2: Four quandrants of belief revision.

Or, in full glory:[7]

	conditioning	imaging
total (on X)	0 $P(w) + \sum_{v \in \overline{X} \mid w \in \sigma_P(v,X)} \cdot P(v) \cdot \frac{P(w)}{P(X)}$	0 $P(w) + \sum_{v \in \overline{X} \mid w \in \sigma_P(v,X)} P(v) \cdot T^P_{v,X}(w)$
partial (on $\{X, \neg X\}$)	$Z \cdot P(w)$ $P(w) + \sum_{v \in \overline{X} \mid w \in \sigma_P(v,X)} (1 - Z) \cdot P(v) \cdot \frac{P(w)}{P(X)}$	$Z \cdot P(w)$ $P(w) + \sum_{v \in \overline{X} \mid w \in \sigma_P(v,X)} (1 - Z) \cdot P(v) \cdot T^P_{v,X}(w)$

Table 3: Posteriors $P'(w)$ where $Z := \frac{P'(\neg X)}{P(\neg X)}$
(top line: for $w \in \overline{X}$; bottom line: for $w \in X$)

3.4 Application to cases

Imaging is rarely defended as a general method of belief revision. But imaging *is* often taken to reflect how causal beliefs factor into particular instances of rational decision-making. Perhaps the best-known examples are Medical Newcomb problems (Briggs, 2010; Ahmed, 2014, ch. 4), such as Coffee Lesion, below.

> *Coffee Lesion.* Otto has a choice between drinking coffee (*coffee*) or refraining (*no coffee*). He believes that coffee consumption is positively correlated with cancer, but only because there is a common cause—a brain lesion (*lesion*) that tends to cause both. Otto enjoys *coffee* more than *no coffee*.

[7] Here I skate over sharp imaging to present only the blurred (more general) case. With sharp imaging, we would have this:

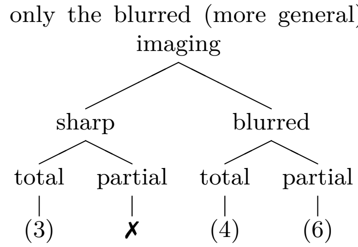

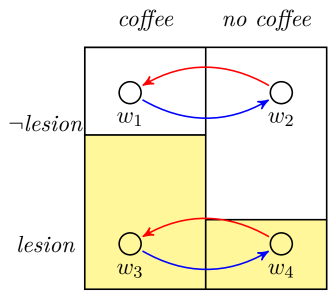

Figure 4: Red arrows image *no coffee* to *coffee* worlds; blue arrows image *coffee* to *no coffee* worlds.

In a decision problem like this, it is routine to use imaging to calculate the Causal Expected Utility (CEU) of the available pure options *coffee* and *no coffee*. The dependency hypotheses *lesion* and ¬*lesion* provide a partition of the sample space within which, as Lewis noted, worlds "image alike" (diagram arrows). Causal Decision Theory thus contrasts with (naïve) Evidential Decision Theory (EDT), which uses conditional, rather than imaged, probabilities to calculate expected utility.[8,9]

$$CEU(A) = \sum_S P^A(S) \cdot U(A \wedge S) \tag{9}$$

Familiarly, in this case, $CEU(\textit{coffee}) > CEU(\textit{no coffee})$, reflecting the popular ("causalist") intuition that Otto should go ahead and drink his coffee. We can also use (9) to estimate the (causal) expected utility of the status quo, $\top$:

$$CEU(\top) = \sum_A P(A) \cdot CEU(A) = \sum_A P(A) \sum_S P^A(S) \cdot U(A \wedge S) \tag{10}$$

In a case like this, $CEU(\top)$ will be sensitive to Otto's act-probabilities $P(A) : A \in \{\textit{coffee}, \textit{no coffee}\}$. These act-probabilities may lie strictly between 0 and 1: for example, Otto might consider drinking coffee to be the rational choice, but be too lazy to procure

[8] In particular, the EDTer uses the following formula for *Evidential Expected Utility* (EEU): $EEU(A) = \sum_S P_A(S) \cdot U(A \wedge S)$. The difference with equation (9) is in the substitution of $P_A(S)$ for $P^A(S)$.

[9] I say *naïve* evidential decision theory because some evidential decision theorists have argued that $P_A(S) = P^A(S)$ when *all* a reflective agent's available evidence is taken into account. For this move, see e.g. Eells (1982)'s "Tickle Defense".

any.

It is here that we can bring partial imaging onstage. Suppose we are in a case where partial supposition alters Otto's probabilities in *coffee* and *no coffee* without taking him all the way to certainty in either. We should be able to apply (6) to calculate the CEU of the new status quo. If partial imaging gives us a new posterior $P'(\cdot) = P^{\{A\}}(\cdot)$, it is routine to extend (10) as follows:

$$CEU(\top) = \sum_A P'(A) \cdot CEU(A) = \sum_A P(A) \sum_S P'(S) \cdot U(A \wedge S) \qquad (11)$$

If, for example, Otto uses partial imaging to move from indifference over {*no coffee, coffee*} to $\{.2, .8\}$, he will see $CEU(\top)$ increase.[10] From a CDT point of view, this seems like the right result.

To see how the the monotonicity of (6)—reasoning through E&H's "perfect guarantee"—works in a Medical Newcomb context, though, it will be helpful add a third option.

> *Soda Lesion.* Serafina is debating whether to drink water, to drink Mountain Dew, or to drink coconut LaCroix. Serafina presently prefers Mountain Dew to LaCroix to water.

	$\neg$*lesion*	*lesion*
LaCroix	+9 (w_1)	-11 (w_2)
Mountain Dew	+10 (w_3)	-10 (w_4)
no soda	0 (w_5)	-20 (w_6)

> Serafina believes that Mountain Dew consumption causally contributes to lesion formation, and hence to cancer. She believes LaCroix consumption, on the other hand, is merely an unfortunate sign of an underlying tendency to develop the lesion.

[10] Suppose we use the following payoff table:

	$\neg$*lesion*	*lesion*
no coffee	0	-20
coffee	+10	-10

and that initially, $P(\textit{coffee}) = P(\textit{no coffee}) = P(\textit{lesion}) = P(\neg\textit{lesion}) = 0.5$, and let the posterior $P'(\cdot)$ partially imaged on {*no coffee, coffee*} be $P'(\textit{no coffee}) = .2$ and $P'(\textit{coffee}) = .8$, as described. A (trivial) application of (6) gives us posteriors $P'(\neg\textit{lesion}) = P'(\textit{lesion}) = .5$. Hence while $CEU(\top) = .5[.5(0) + .5(-20)] + .5[.5(10) + .5(-10)] = -5$ by Equation (11), the posterior $CEU(\top)$ is $.2[.5(0) + .5(-20)] + .8[.5(10) + .5(-10)] = -2$ by Equation (11).

An agent committed to naive EDT will avoid *LaCroix* in Soda Lesion, via a line of reasoning similar to a dominance argument: LaCroix brings just as much bad news as Mountain Dew, but less pleasure! A CDTer rejects this line of reasoning, because her subjective outlook distinguishes between bad news and bad causal effects.[11] More to the point, from the CDTer's perspective, LaCroix has a lot going for it: it brings *almost* as much pleasure as Mountain Dew, with none of the bad causal consequences. The case for *LaCroix* is much like the case for *coffee* in Coffee Lesion.

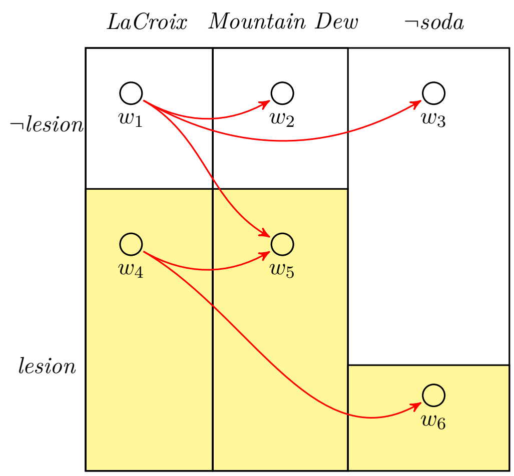

Figure 5: Red arrows image *LaCroix* worlds to ¬*LaCroix* worlds.

In terms of imaging, what is important here can be illustrated in a model where (i) $\sigma(w_1, \text{Mountain Dew}) \ni w_5$, and (ii) $\sigma(w_3, \text{LaCroix}) \not\ni w_4$. (i) reflects the fact that, in a world where Serafina drinks coconut LaCroix and is lesion-free, she *might have* gotten the lesion if she had not had LaCroix: for if she had not had LaCroix, she might have had Mountain Dew, and *that* might have been enough to cause a lesion. The selection

[11]Indeed, a non-akratic CDTer will have soda of some kind, for *LaCroix* dominates *no soda* relative to a partition of dependency hypotheses (Lewis, 1981):

	¬*lesion*	*lesion*	$(lesion \leftarrow\Box\Box\rightarrow \textit{Mountain Dew})$
LaCroix	+9	-11	+9
Mountain Dew	+10	-10	-10
no soda	0	-20	0

The dependency hypotheses in these columns capture the same information stated by the selection function in Figures 5-6.

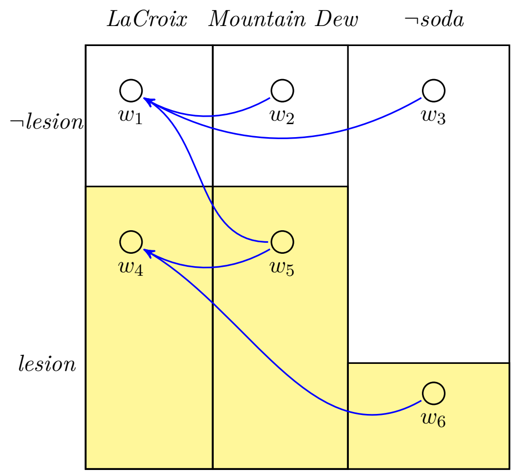

Figure 6: Blue arrows image ¬*LaCroix* worlds to *LaCroix* worlds.

function on (w_1, Mountain Dew) thus "crosses the border" between ¬*lesion* and *lesion* (viz., into the colored region from the white region.) (ii) (in Figure 6) reflects the fact that, in a world where Serafina drinks no soda and is lesion-free, it's *not* possible that she might have gotten the lesion if she had switched to coconut LaCroix. For by hypothesis, LaCroix does not causally contribute to lesions.

Here, *soda* (= {*LaCroix, Mountain Dew*}) is a logical consequence of *LaCroix*. So Monotonicity entails that partially imaging on {*LaCroix*, ¬(*LaCroix*)}—in a way that increases the former—should (weakly) increase the probability of *soda*. A rational agent who "simply" becomes more confident in *LaCroix* via a process of partial imaging thus becomes more confident in *soda*.

This seems correct. For example, in light of the "a lot going for it" argument in the latter case, Serafina might increase her confidence that she will drink LaCroix (say, from indifference)—but not all the way to certainty. (Like Otto, she could be too lazy to go to the store). As we saw, Monotonicity thus entails that Serafina's confidence that she will have soda (weakly) increases. (6) also entails that the causal expected utility of the status quo also will weakly increase, via a dynamic like the one we saw in *Coffee Lesion*.

As we would expect, this contrasts with the (naive) Evidential Decision Theorist's point of view. For the EDTer, partial update $\{.8, .2\}$ on $\{LaCroix, \neg(LaCroix)\}$ will induce a drop in expected utility. After all, for the EDTer, a direct *increase* in the probability of *LaCroix* is (via the Montonicity of Jeffrey conditioning) a strict increase in the probability of bad news.

4 Conclusion

I am not the first person to consider whether imaging is a form of conditioning, or vice-versa.

Thinking of imaging as a form of conditioning is attractive because we might want to reduce, or ground, the ontological commitments of the extra-Bayesian structure in σ and T. The reverse point of view—thinking of conditioning as a special case of imaging—is attractive for the obverse reason: whereas imaging requires extra structure, a probability space alone fixes conditioning kinematics.

The suggestion to frame conditioning as a degenerate case of extra imaging structure has surfaced multiple times in the literature, usually in the context of unifications of different forms of causal decision theory (Pearl, 2000, pg. 73; Joyce *op. cit.*, pg. 149; Gärdenfors (1982) credits Isaac Levi with a similar suggestion, discussed below.) To this consideration, we can now add the one advanced in the present paper: by looking at conditioning as a special case of imaging, we arrive at a better general characterization of the relationship between *total* and *partial* supposition—a distinction which crosscuts the imaging-conditioning divide. I have tried (briefly) to make the case that the (monotonic) notion of partial imaging this makes available (in (6)) is *useful* and *appropriate* in natural extensions of familiar decision problems.

It is notable that, in order to frame conditioning as imaging, we had to relativise both imaging's characteristic functions $\sigma(\cdot)$ and $T_{(\cdot)}$, as well as the weighting coefficient Z in partial cases, to the prior $P(\cdot)$—to relativise it, roughly speaking, to the agent's evidence. From a formal perspective, this is an obvious move—at least, it appears to have been obvious to Levi, Pearl, and Joyce. In closing, though, I want to flag two reasons that such relativisation has been rejected as a mistaken.

The first worry is that relativising the apparatus of imaging to one's evidence conflates objective and subjective components that are, for decision-theoretic applications, best kept separate.[12] The apparatus of imaging, the thought goes, should track the *objective* closeness of possible worlds—worlds which e.g. share our laws of nature, or our fixed past (Lewis, 1979). Objective closeness is independent of subjective information gain. My information might, for example, rule out some nearby possibility in which I am $1,000 poorer than I actually am. This should not change my opinion that in the *objectively* closest possible world where some counterfactual antecedent A is true, I *am* $1,000 poorer than I actually am. This fact is essential to explaining the rationality of my choice in Newcomb's problem—at least, from a CDTer's point of view.

I think this is broadly correct. But it is worth noting that this tidy dichotomy is under threat anyway—for example, from a number of deterministic counterexamples to CDT (Ahmed, 2013; Ahmed, 2014, Ch. 5). Recent CDT-oriented responses to these counterexamples push back on the neat conceptual divide between the objective and the subjective (Solomon, 2019; Williamson & Sandgren, forthcoming, §5). This includes treating worlds as "ruled out" by information in two ways—one which is consistent with such worlds remaining "live" to the selection function, and one which is not. This is to say that, at least on one interpretation, modal closeness *is* subjective.

The second objection to relativising imaging to evidence appears to be more technical. According to Gärdenfors (1982)'s influential discussion, imaging is characterised by a property called *Linearity*, while conditionalisation is characterised by a property called *Conservativity*. A belief-revision method is **linear** iff, if $\exists P_1, P_2$ in the model s.t.

$$P(B) = \alpha P_1(B) + (1 - \alpha)P_2(B),$$

then [where P^{+A} is P belief-revised on A]:

$$P^{+A}(B) = \alpha P_1^{+A}(B) + (1 - \alpha)P_2^{+A}(B).$$

[12] See Bacon (2022). Such a sentiment is also articulated in unpublished work by Rush Stewart and Michael Nielsen.

A belief-revision method is **conservative** iff, for all sentences A and B, if $P(A) > 0$ and $P(B) = 1$, then $P^{+A}(B) = 1$ Gärdenfors (1982, pgs. 753, 754). Gärdenfors proves that linearity and conservativity impose incompatible constraints on all reasonably rich belief-change models. However, Gärdenfors *also* briefly discusses—with credit to Isaac Levi—what he calls *conservative imaging*, which is essentially the relativisation of $\sigma(w, X)$ to the support of $P(\cdot)$ I advocated above (pg. 756). Conservative imaging sacrifices linearity, so Gärdenfors pronounces it "*not* a form of imaging" (pg. 757, emphasis in original).

Gärdenfors's terminological precedent aside, why must imaging be linear? The chief attraction of linearity is to facilitate a smooth (some would say *naïve*) semantics for an object-language conditional '$\Box\!\!\rightarrow$', as follows:

$$P(A \Box\!\!\rightarrow B) = P^A(B) \tag{12}$$

Seen in this light, Gärdenfors's proof of the incompatibility of Linearity and Conservativity pairs naturally with Lewis (1976)'s Triviality result, which established the *nonexistence* of a (linear) object-language conditional for (conservative) conditionalisation.[13] Both results point to the same forced choice between Gärdenfors's two properties (Table 4):

Revision method	Naive bridge principle to object-language conditional	Conservative	Linear	Triviality Result?
conditioning ($P_A(\cdot)$)	$P_A(B) = P(A \rightarrow B)$	✓	✗	Lewis (1976)
imaging ($P^A(\cdot)$)	$P^A(B) = P(A \Box\!\!\rightarrow B)$	✗	✓	Schultheis (forthcoming)

Table 4: Bridge principles and the
incompatibility of Conservativity and Linearity.

In this paper I have not, of course, considered object-language conditionals at any length. But in implicitly advocating for the rejection of linearity in the characteri-

[13]In particular, the assumption is that $P(A \Box\!\!\rightarrow B) = P^A(B)$ *and* that the semantic value of $\ulcorner A \Box\!\!\rightarrow B \urcorner$ is a proposition, which depends for its semantic value only on w. If this is right, then $P'(A \Box\!\!\rightarrow B) = P'^A(B)$ for *any* probability function P' in the model, not just for the prior. This assumption, in the indicative mode (viz., with $\rightarrow$ for $\Box\!\!\rightarrow$ and P_A for P^A), is the basis for Lewis (1976)'s Triviality Proof.

zation of (partial) imaging, I have also implicitly advocated foreclosing on the naive semantic bridge principle in (12). My position is hence of a piece with two trends in "post-Triviality" work on the semantics of indicative conditionals. The first exploits attractive bridge principles between imaging and conditional chances, due to Skyrms (1981), to derive triviality results for the counterfactual conditional (Schultheis, forthcoming; see also Williams, 2012 and Moss, 2013). The upshot is that the Skyrmsian link to conditional chance restores conservativity—at the cost of linearity. The second trend concerns the probabilities of indicative conditionals and how best to respond to Triviality (Bacon, 2015; Goldstein, 2020; Goldstein & Santorio, 2021). These writers block Lewis's proof by relativising the semantics of the *indicative* conditional to an information-state-like parameter, a move which both sacrifices conservativity and prefigures the relativisation advocated for here.

References

AHMED, ARIF (2013). "Causal Decision Theory: A Counterexample." *Philosophical Review*, 122(2): pp. 289–306.

AHMED, ARIF (2014). *Evidence, Decision and Causality.* Cambridge: Cambridge University Press.

BACON, ANDREW (2015). "Stalnaker's Thesis in Context." *Review of Symbolic Logic*, 8(1): pp. 131–163.

BACON, ANDREW (2022). "Actual Value in Decision Theory." *Analysis.*

BRIGGS, R. A. (2010). "Decision-Theoretic Paradoxes as Voting Paradoxes." *Philosophical Review*, 119(1): pp. 1–30.

EELLS, ELLERY (1982). *Rational decision and causality.* Cambridge: Cambridge University Press.

EVA, BEN, AND STEPHAN HARTMANN (2021). "The logic of partial supposition." *Analysis*, 81(2): pp. 215–224.

GÄRDENFORS, PETER (1982). "Imaging and Conditionalization." *Journal of Philosophy*, 79(12): pp. 747–760.

GOLDSTEIN, SIMON (2020). "Epistemic Modal Credence." *Philosophers' Imprint*, 21(26): pp. 1–24.

Goldstein, Simon, and Paolo Santorio (2021). "Probability for Epistemic Modalities." *Philosophers' Imprint*, 21(33): pp. 1–37.

Jeffrey, Richard (1983). *The Logic of Decision*. Chicago: University of Chicago Press.

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

Lewis, David (1976). "Probabilities of Conditionals and Conditional Probabilities." *Philosophical Review*, 85: pp. 297–315.

Lewis, David (1979). "Counterfactual Dependence and Time's Arrow." *Noûs*, 13(4): pp. 455–476.

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

Moss, Sarah (2013). "Subjunctive Credences and Semantic Humility." *Philosophy and Phenomenological Research*, 87(2): pp. 251–278.

Pearl, Judea (2000). *Causality: Models, reasoning, and inference*. Cambridge: Cambridge University Press.

Schultheis, Ginger (forthcoming). "Counterfactual Probability." *Journal of Philosophy*.

Skyrms, Brian (1981). "The Prior Propensity Account of Subjunctive Conditionals." In W. Harper, R. Stalnaker, and G. Pearce (eds.) *Ifs: Conditionals, Belief, Decision, Chance, and Time*, D. Reidel, Dordrecht, pp. 259–266.

Solomon, Toby (2019). "Causal Decision Theory's Predetermination Problem." *Synthese*, 198(6): pp. 5623–5654.

Williams, J.R.G. (2012). "Counterfactual Triviality: A Lewis-Impossibility Argument for Counterfactuals." *Philosophy and Phenomenological Research*, LXXXV(3): pp. 648–670.

Williamson, Timothy Luke, and Alexander Sandgren (forthcoming). "Law-Abiding Causal Decision Theory." *British Journal for the Philosophy of Science*.