

Imaging Uncertainty

Benjamin Eva* Stephan Hartmann†

August 4, 2017

Abstract

The technique of imaging was first introduced by Lewis (1976), in order to provide a novel account of the probability of conditional propositions. In the intervening years, imaging has been the object of significant interest in both AI and philosophy, and has come to be seen as a philosophically important approach to probabilistic updating and belief revision. In this paper, we consider the possibility of generalising imaging to deal with uncertain evidence and partial belief revision. In particular, we introduce a new logical criterion that any update rule should satisfy, and use it to evaluate a range of different approaches to generalising imaging to situations involving uncertain evidence. We show that none of the currently prevalent approaches to imaging allow for such a generalisation, although a lesser known version of imaging, introduced by Joyce (2010), can be generalised in a way that mitigates these problems.

1 Introduction

Standard Bayesian epistemology tells us that upon learning some new piece of evidence, E , a rational agent should update their beliefs via Bayesian conditionalisation, i.e. $P_1(A) = P_0(A|E)$ should hold for any other proposition A (where P_0 is the probability distribution representing the agent's prior belief state, and P_1 represents their belief state after learning E).¹ The alternative technique of Imaging was first introduced by Lewis in his seminal *Probabilities of Conditionals and Conditional Probabilities* (Lewis, 1976). Its introduction was motivated mainly by the need for a

*Munich Center for Mathematical Philosophy, LMU Munich, 80539 Munich (Germany) – <http://be0367.wixsite.com/benevaphilosophy> – benedgareva@icloud.com.

†Munich Center for Mathematical Philosophy, LMU Munich, 80539 Munich (Germany) – <http://www.stephanhartmann.org> – s.hartmann@lmu.de.

¹ $P(A|E)$ denotes the conditional probability of A given E , defined by $P(A|E) := P(A \wedge E)/P(E)$ (where $P(E) > 0$).

new account of the probability of conditionals. In particular, Lewis’s famous triviality results are often taken to show the untenability of ‘Adam’s thesis’, which states that for any conditional proposition $A \rightarrow B$, we should set $P(A \rightarrow B) = P(B|A)$, i.e. the probability of the conditional is the conditional probability of the consequent given the antecedent. Lewis considered the possibility that when assessing the probability of a conditional, we should not conditionalise on the antecedent, but rather we should ‘image’ on it. Before we see the definition of ‘imaging on A ’, we briefly introduce some formal notation.

W will denote the set of possible worlds, and W_A will denote the subset of W that satisfies A (for any sentence A). μ will denote a ‘measure’ on the powerset $\mathbb{P}(W)$ that assigns a weight $\mu(X) \in [0, 1]$ to each subset $X \subseteq W$. We require that $\sum_{w \in W} \mu(\{w\}) = 1$ and $\mu(X) + \mu(Y) = \mu(X \cup Y)$ whenever $X \cap Y = \emptyset$. We then set $P(A) = \sum_{w \in W_A} \mu(w)$ (it is easy to see that P will be probabilistic). Finally, for any sentence A and any world w , $\sigma(A, w)$ is defined to be the closest possible world to w at which A holds (where ‘closest’ is defined by a *similarity relation* $S : W \times W \rightarrow [0, 1]$).² For now, we assume that $\sigma(A, w)$ picks out a single closest A -world, for each w . We define the result of ‘imaging on A ’ as follows,

$$\mu_A(w) = \begin{cases} \mu(w) + \sum_{w' \in W_{\neg A} | w = \sigma(A, w')} \mu(w'), & \text{if } w \in W_A \\ 0, & \text{if } w \in W_{\neg A}. \end{cases}$$

The idea is that, when we learn A , we should update our probability assignment by transferring the probability of each $\neg A$ -world to the closest A -world. As Lewis puts it:³

“Intuitively, the image on A of a probability function is formed by shifting the original probability of each world w over to $\sigma(A, w)$, the closest A -world to w . Probability is moved around, but not created or destroyed, so the probabilities of worlds still sum to 1. Each A -world keeps whatever probability it had originally, since if w is an A -world then $\sigma(A, w)$ is w itself, and it may also gain additional shares of probability that have been shifted away from $\neg A$ worlds. The $\neg A$ worlds retain none of their original probability and gain none. All the probability has been concentrated on the A -worlds. And this has been accomplished with no gratuitous movement of probability. Every share stays as close as it can to the world where it was originally located.” (Lewis (1976): 310–311)

It’s easy to see⁴ that imaging and Bayesian conditionalisation can give very different answers as to how an agent should update their epistemic state in the face of new

²Intuitively, the similarity relation measures the similarity between pairs of worlds in W .

³We have edited the formal notation in the quote so that it accords with our own.

⁴Lewis gives the following example. Let $\mu_0(w_1) = \mu_0(w_2) = \mu_0(w_3) = 1/3$. Let A hold at w_1 and w_2 , but not at w_3 , and let w_2 be the closest A world to w_3 . Then, while conditioning on A would give

evidence. However, both approaches can claim to provide the ‘minimal’ revisions to the original probability functions that satisfy the new constraints implied by learning A . Thus, we read:

“Imaging P on A gives a minimal revision in this sense: unlike all other revisions of P to make A certain, it involves no gratuitous movement of probability from worlds to dissimilar worlds. Conditioning on A gives a minimal revision in this different sense: unlike all other revisions of P to make A certain, it does not distort the profile of probability ratios, equalities and inequalities among sentences that imply A .” (Lewis (1976): 311)

Since its introduction, imaging has gone on to be used in a variety of contexts, both technical and philosophical. For example, Joyce (2008) and Lewis (1981) use general imaging (a variation on the ‘standard’ form of imaging that will be introduced later) to formulate key aspects of causal decision theory, Cozic (2011) proposes a solution to the Sleeping Beauty problem by replacing conditionalisation with imaging, artificial intelligence researchers have used imaging to address fundamental problems in the field of information retrieval (Crestani, 1998), Leitgeb (2016) recently applied imaging to probabilistic judgement aggregation, and there are even empirical studies suggesting that (general) imaging is the best current formal model of the way that people actually go about updating their belief states in the face of new evidence (Baratgin & Politzer, 2011). Cumulatively, this looks like sufficient grounds for taking imaging seriously as an approach to probabilistic updating. However, it has been argued (see e.g. chapter 6 of Joyce (2008)) that imaging rules are fundamentally different to the updating mechanisms of Bayesian epistemology, in the sense that they are not genuinely ‘evidential’, i.e. they do not preserve the antecedently known evidence. Proponents of this kind of view typically construe imaging as a method for counterfactual belief revision. While it is impossible to conditionalise on a proposition that we know to be false (because the corresponding conditional probability is undefined), it is possible to image on an event that we know to be false. So imaging allows us to determine what we would believe under the counterfactual supposition that a false proposition is true. Here, we remain agnostic about whether or not imaging rules can only be used to model counterfactual belief revision. Certainly, the extant philosophical applications of imaging have gone far beyond counterfactual belief revision, and Lewis and Stalnaker both treated imaging as a potential approach to evidential updating. But the considerations discussed here will be of significant interest regardless of whether one restricts the use of imaging rules to counterfactual belief revision or allows for their use as evidential updating rules.

us the posterior distribution $\mu_1(w_1) = \mu_1(w_2) = 1/2$, imaging on A would give us $\mu_1(w_1) = 1/3$ and $\mu_1(w_2) = 2/3$.

In what follows, we will study how imaging deals with one of the basic problems facing any probabilistic updating or belief revision rule, namely, the problem of generalising the rule to deal with uncertain evidence. In Section 2, we introduce a new criterion that any such generalisation should satisfy, and show that some well known rules (Jeffrey conditionalisation and Leitgeb-Pettigrew updating) from the literature satisfy this criterion. In Section 3, we present a range of new formal results concerning the possibility of extending imaging to uncertain evidential contexts in a way that satisfies the criterion. Section 4 considers the implications of these results and concludes.

2 Updating on Valid Arguments

Let $\Gamma \vdash \phi$ be a valid argument scheme. Usually, we describe this as meaning that there is no possible world where Γ is true and ϕ is false, i.e. every Γ -world is a ϕ -world. Now let's suppose that I have some prior belief state represented by the probability distribution P_0 . It seems very natural to require that, upon learning the evidence Γ , my belief in ϕ should go to 1. And indeed, it is a standard result (see chapter 3 of Adams (1996)) that this is the case if one updates by conditionalisation. However, we are more interested in the general case where our belief in Γ increases without going to 1, i.e. where we only have the constraint that $P_1(\Gamma) \geq P_0(\Gamma)$. The idea is that we should require, for any update rule, that in this situation $P_1(\phi) \geq P_0(\phi)$ is guaranteed. The motivation is fundamentally the same as in the special case where $P_1(\Gamma) = 1$, i.e. since the truth of Γ guarantees the truth of ϕ , learning that Γ is more likely to be true than we previously thought should guarantee a corresponding increase in the probability of ϕ being true. This gives us the following criterion:

Definition 1 (Validity Criterion (VC)) *Let $\Gamma \rightarrow \phi$ be an instance of a valid argument scheme, and let P_0 be our prior probability distribution. If R is our update rule and we learn the evidence Γ with probability $\gamma \geq P_0(\Gamma)$ ⁵, then the result, P_1 , of updating P_0 according to R should guarantee that $P_1(\phi) \geq P_0(\phi)$.*

We know that valid arguments lead us from true premises to true conclusions, which means that certainty about the premises should guarantee certainty about the conclusion. But it seems clear that we also want to be able to use valid arguments to support conclusions based on uncertain premises. Indeed, if this were not the case, one could reasonably ask “What is the value of logical validity if we are in the business of making inferences with uncertain premises?”. Given the ubiquity of arguments based on uncertain premises in both scientific and everyday reasoning,

⁵Whenever Γ is a set of more than one formulae, we use $P(\Gamma)$ to denote the probability of the conjunction of all the formulae in Γ .

this is an important question, and **VC** ensures that we are able to answer it satisfactorily. We now show that **VC** is satisfied by two of the most important update rules from the current literature.

2.1 VC and Jeffrey Conditionalisation

Jeffrey conditionalisation is the natural extension of standard Bayesian conditionalisation to situations involving uncertain evidence. Specifically, if we learn some new evidence E with probability e , then the posterior probability of any proposition A after Jeffrey conditionalisation (henceforth **JC**) on E is defined to be

$$\begin{aligned} P_1(A) &= P_0(A|E) P_1(E) + P_0(A|\neg E) P_1(\neg E) \\ &= P_0(A|E) e + P_0(A|\neg E) (1 - e). \end{aligned}$$

It is easy to see that, in the case where $e = 1$, this reduces to standard conditionalisation. **JC** is by far the most popular rule for updating in the face of uncertain evidence, and we take it to be one of the strongest arguments in favour of Bayesian conditionalisation that it generalises so naturally to uncertain contexts. The following result adds to the already substantial body of considerations supporting the adoption of **JC** as one's updating rule for uncertain propositional evidence. (All absent proofs are in the Appendix.)

Proposition 1 *JC satisfies VC.*

2.2 VC and Leitgeb-Pettigrew Updating

Leitgeb-Pettigrew updating (**LP**) was first introduced in Leitgeb & Pettigrew (2010), as the unique update rule that maximises diachronic accuracy (as measured by the Brier score). Specifically, if we learn some new evidence E with probability e , **LP** requires us to update our epistemic state in the following way, where μ_1 is the new measure obtained by applying **LP** to μ :

$$\mu_1(w) = \begin{cases} \mu(w) + \alpha_E, & \text{if } w \in W_E \\ \max\{\mu(w) - \alpha_{\neg E}, 0\}, & \text{if } w \in W_{\neg E}, \end{cases}$$

where α_E and $\alpha_{\neg E}$ are the unique real numbers such that

$$\begin{aligned} \sum_{w \in W_E} (\mu(w) + \alpha_E) &= e \\ \sum_{w \in W_{\neg E} | \mu(w) - \alpha_{\neg E} > 0} (\mu(w) - \alpha_{\neg E}) &= 1 - e. \end{aligned}$$

LP has a much shorter history than **JC**, and is correspondingly less well known. Its independent justification by arguments from accuracy certainly count in its

favour, but it has already been subjected to criticism concerning the effect it has on the likelihood ratios between propositions (see Levinstein (2012)). However, the following result adds to the rule’s plausibility.

Proposition 2 *LP satisfies VC.*

3 Imaging Uncertainty

3.1 Standard Imaging

We’ve seen that the generalisation of conditionalisation to uncertain evidence satisfies our new criterion. The challenge for the proponent of imaging as an evidential update rule is to find a similar generalisation of imaging that also satisfies **VC**. First, we note that there is no problem in the special case where we learn the evidence with certainty. Standard imaging is able to distinguish between valid and invalid arguments in the special case where we learn the premises with certainty. But of course we are more interested in the general case where, instead of learning the premises Γ with certainty, we get some non-maximal increase in the probability of Γ . Before we can consider this more general case, we need a new form of imaging that allows us to have $1 > P_1(\Gamma) \geq P_0(\Gamma)$. This is something that will require a fairly different approach, since standard imaging is only defined for the case where we learn the premises with certainty. Luckily, such an approach already exists (it was first defined in Sebastiani (1998)). In particular, there is a form of imaging known as ‘Jeffrey Imaging’ (inspired by Jeffrey conditionalisation) that allows us to update on uncertain evidence. It’s defined below, where P_Γ^* is the new probability distribution obtained by Jeffrey imaging (**JI**) on Γ , $P^*(\Gamma)$ is the new probability constraint that reflects our new level of confidence in Γ and μ_Γ is the measure obtained by standard imaging on Γ .

$$P_\Gamma^*(\phi) = \sum_{w \in W_\phi} (\mu_\Gamma(w) P^*(\Gamma) + \mu_{\neg\Gamma}(w) P^*(\neg\Gamma))$$

Phrased in terms of the measure on the possible worlds, this gives

$$\mu_\Gamma^*(w) = \mu_\Gamma(w) P^*(\Gamma) + \mu_{\neg\Gamma}(w) P^*(\neg\Gamma).$$

Clearly, **JI** is to imaging what **JC** is to conditionalisation. In the special case where $P^*(\Gamma) = 1$, **JI** reduces back to standard imaging, which makes **JI** the obvious generalisation of imaging to uncertain contexts. Furthermore, **JI** appears to be the *only* such generalisation currently on the market, which makes its plausibility all the more crucial for those who advocate imaging as an evidential updating rule.

Note that **JI** can also be used to generalise the counterfactual belief revision method given by standard imaging. For, under the standard interpretation, imaging

on a proposition A that we are certain is false tells us what it would be rational to believe under the counterfactual supposition that A is true. Thus, imaging tells us what a rational agent would believe if their degree of belief in A were different than it actually is (i.e. if it were 1 rather than 0). But this can be generalised in an obvious way. Rather than simply asking what the agent would believe if their degree of belief in A were 1 rather than 0, we could likewise ask what they would believe if their degree of belief in A were a rather than 0, where $a \in (0, 1)$. This is a philosophically rich and interesting question, which has already been the subject of debate in the belief revision literature (see e.g. Lin & Kelly (2013)). And **J1** provides a natural answer that is obtained by generalising the full counterfactual belief revision given by standard imaging. Furthermore, it seems that **VC** should still be taken as a strong normative restriction on generalised counterfactual belief revision of this kind. For, it seems natural to require of any rational agent that, were they to be more confident of the premises of a valid argument than they actually are, they would also be more confident in the conclusion of that argument. Again, the motivation is the same as before. The premises guarantee the truth of the conclusion, and so it would seem strange to deny that, were a rational agent to become more confident of the premises, they would also become more confident of the conclusion. So regardless of whether one views imaging as a counterfactual belief revision method or a full blown evidential updating rule, it seems that **VC** should act as a normative constraint on any generalisation of imaging to uncertain evidence/generalised belief revision. Unfortunately, we have the following result:

Proposition 3 ***J1**, as defined above, doesn't generally satisfy **VC**.*

Proof To see this, let's consider the example of modus ponens, i.e. $\Gamma = \{A \rightarrow B, A\}$, $\phi = B$, $W = \{w_1, w_2, w_3, w_4\}$ where $w_1 = \{A, B\}$, $w_2 = \{A, \neg B\}$, $w_3 = \{\neg A, B\}$, $w_4 = \{\neg A, \neg B\}$. Here, it is clear that $\sigma(\Gamma, w_2) = \sigma(\Gamma, w_3) = \sigma(\Gamma, w_4) = w_1$ (since $W_\Gamma = \{w_1\}$). However, there are two possibilities for $\sigma(\neg\Gamma, w_1)$, since both w_2 and w_3 satisfy $\neg\Gamma$, and they both differ from w_1 by one truth value. Now, suppose we choose the case where $\sigma(\neg\Gamma, w_1) = w_2$. In this case (setting $\alpha = P^*(\Gamma) - P_0(\Gamma)$), we have that

$$\begin{aligned}
P_\Gamma^*(\phi) &= \mu_\Gamma^*(w_1) + \mu_\Gamma^*(w_3) \\
&= \mu_\Gamma(w_1) P^*(\Gamma) + \mu_{\neg\Gamma}(w_3) P^*(\neg\Gamma) \\
&= P^*(\Gamma) + \mu_{\neg\Gamma}(w_3) P^*(\neg\Gamma) \\
&= P^*(\Gamma) + \mu(w_3) P^*(\neg\Gamma) \quad (\text{since } \sigma(\neg\Gamma, w_1) = w_2) \\
&= \mu(w_1) + \alpha + \mu(w_3) (1 - \mu(w_1) - \alpha) \\
&= \mu(w_1) + \alpha + \mu(w_3) - \mu(w_1) \mu(w_3) - \alpha \mu(w_3).
\end{aligned}$$

So, since $P_0(\phi) = \mu(w_1) + \mu(w_3)$, we know that

$$P_\Gamma^*(\phi) - P_0(\phi) = \alpha - \mu(w_1) \mu(w_3) - \alpha \mu(w_3),$$

which is clearly negative for large values of $\mu(w_3)$ and small values of α . ■

So, we've seen that, in its current form, **J1** fails to satisfy our new criterion. Now, the obvious way to fix the problem is to require, in the above example, that $\sigma(\neg\Gamma, w_1) = w_3$, since this implies that

$$\begin{aligned}
P_\Gamma^*(\neg\phi) &= \mu_\Gamma^*(w_2) + \mu_\Gamma^*(w_4) \\
&= \mu_{\neg\Gamma}(w_2) P^*(\neg\Gamma) + \mu_{\neg\Gamma}(w_4) P^*(\neg\Gamma) \\
&= \mu(w_2) P^*(\neg\Gamma) + \mu(w_4) P^*(\neg\Gamma) \quad (\text{since } \sigma(\neg\Gamma, w_1) = w_3) \\
&\leq \mu(w_2) + \mu(w_4) \\
&= P(\neg\phi).
\end{aligned}$$

But of course, the key assumption that $\sigma(\neg\Gamma, w_1) = w_3$ looks ad-hoc and unjustified. What's more, this assumption guarantees that the argument works for *any value of* α . Indeed, α could even have been negative and the probability of ϕ would still have been guaranteed not to decrease. So as well as being ad-hoc, this assumption seems to trivialise the evidential relationship between the conclusion and the premises of valid arguments. This looks like a fundamental problem for **J1**.⁶

3.2 General Imaging

We can now consider another possible amendment to **J1** that one might hope would allow it to satisfy the criterion. The technique of standard imaging has been generalised (Gärdenfors, 1982) to allow for cases where $\sigma(\Gamma, w)$ can be a set of worlds, instead of just a single world. Specifically, the technique of 'general imaging' is defined as follows. First, for any pair of worlds w, w' and any proposition E , we define the value $T_E(w', w) \in [0, 1]$ (satisfying $\sum_{w' \in W} T_E(w, w') = 1$). Intuitively, $T_E(w, w')$ tells us what percentage of $\mu(w)$ will be transferred to w' after we image on E . In the case where $w \models E$, we will have $T_E(w, w') = 0$ for any $w' \neq w$, and $T_E(w, w) = 1$. Similarly, we can view standard imaging as the special case where, for any E and any w , there exists exactly one w' such that $T_E(w, w') = 1$, and $T_E(w, w'') = 0$ for any other w'' . But we are more interested in the general case where there can exist $w' \neq w''$ such that $T_E(w, w') > 0$, $T_E(w, w'') > 0$. Then, we can set

$$\mu_E(w) = \begin{cases} \mu(w) + \sum_{w' \in W_{\neg E} | w \in \sigma(E, w')} \mu(w') T_E(w', w), & \text{if } w \models E \\ 0, & \text{if } w \models \neg E \end{cases}$$

⁶It should also be noted that the proof of Proposition 3 assumed very little about the semantics of the conditional. For, it seems entirely uncontroversial, under any account of the semantics of conditionals, that w_1 is the only one of the four worlds according to which both $A \rightarrow B$ and A hold. We also take it to be uncontroversial that w_2 violates $A \rightarrow B$. And this is all that is assumed in the proof.

This is the formal definition of ‘general imaging’. Now, as with standard imaging, we can use this definition to define a direct analogue of **JC**. In particular, we define **JJ** in exactly the same way as before,

$$\mu_{\Gamma}^*(w) = \mu_{\Gamma}(w) P^*(\Gamma) + \mu_{\neg\Gamma}(w) P^*(\neg\Gamma).$$

But here, it is understood that μ_{Γ} denotes the result of general imaging on Γ , rather than standard imaging. Now, we can look again at the question of whether **JJ** satisfies **VC**. Let’s start by recalling the case of modus ponens. The situation looks more hopeful this time since we don’t need to transfer all of the probability from w_1 to exactly one of w_2, w_3 (we’ve already seen that either of these choices is unsatisfactory) when we image on $\neg\Gamma$. Rather, we can now spread the probability between w_2 and w_3 however we like. One natural idea here is to adopt the following principle:

Definition 2 *Worldly Indifference (WI)* Let $w \in W$, and let Γ be some non-empty set of propositions such that $w \models \Gamma$. Then,

$$\mu_{\Gamma}(w) = \mu(w) + \sum_{w' \in W | w \in \sigma(\Gamma, w')} \frac{\mu(w')}{|\sigma(\Gamma, w')|}.$$

That is, we require that the weight of any non- Γ world gets distributed evenly across all of the maximally similar Γ worlds when we image on Γ . In the current case of modus ponens, this requires that $\mu_{\neg\Gamma}(w_3) = \mu(w_1)/2 + \mu(w_3)$ and that $\mu_{\neg\Gamma}(w_2) = \mu(w_1)/2 + \mu(w_2)$, i.e. $T_E(w_1, w_2) = T_E(w_1, w_3) = 0.5$. Before we go on to see what effect this assumption has on **JJ**’s relation to **VC**, it’s worth noting that it looks very intuitive from a philosophical perspective. At first blush, there is no obvious reason why either w_2 or w_3 should gain the lion’s share of w_1 ’s probability. They are both equally similar worlds (from w_1 ’s perspective) that are compatible with the proposition that we’re imaging on ($\neg\Gamma$). However, we have the following result:

Proposition 4 *Assuming **WI**, **JJ** does not satisfy **VC**.*

Now, one might want to play around with variations of **WI** to see if there is some other way of distributing the probability of w_1 so that **VC** isn’t violated. We now present a general result about the possibility of such variations.

Theorem 1 *There is no general imaging rule that allows for the satisfaction of **VC** by **JJ**.*

So **JJ** violates **VC** for both standard and general imaging. At this stage, **JJ** looks dead in the water as a strategy for extending imaging to updating in uncertain evidential contexts or performing generalised counterfactual belief revision of the

type described in Section 3.1. However, there is another alternative that we’ve not yet considered. In particular, recall the assumption **WI**. The form of imaging defined by this assumption has actually already been considered in the literature (Joyce (2010) refers to it as ‘Laplacian Imaging’). Now, Laplacian imaging actually bears a very close resemblance to **LP**.⁷ Indeed, it is easy to see that Laplacian imaging coincides with **LP** when the evidence is certain and we assume that the similarity relation is the trivial maximal relation,⁸ i.e. $\sigma(\Gamma, w) = W_\Gamma$ for any $w \in W_{-\Gamma}$. Letting $w \in W_\Gamma$, we get

$$\begin{aligned} \mu_\Gamma(w) &= \mu(w) + \sum_{w' \in W | w \in \sigma(\Gamma, w')} \frac{\mu(w')}{|\sigma(\Gamma, w')|} \\ &= \mu(w) + \sum_{w' \in W_{-\Gamma}} \frac{\mu(w')}{|W_\Gamma|} \\ &= \mu(w) + \frac{P_0(-\Gamma)}{|W_\Gamma|} \\ &= \mu(w) + \frac{\alpha}{|W_\Gamma|}. \end{aligned}$$

This is exactly the same result we obtain with **LP** (and by definition, any $w \in W_{-\Gamma}$ is given a weight of 0 by both μ_Γ and **LP**). So, under these restrictive assumptions, **LP** can be seen as a generalisation of Laplace imaging to uncertain evidential contexts. Furthermore, since we’ve already seen (Proposition 2) that **LP** satisfies **VC**, we can claim to have succeeded in finding a generalisation of imaging to uncertain situations that satisfies our criterion. However, a key part of the original motivation for imaging has been sacrificed here. In particular, the original argument that Lewis used to justify imaging was that it involved ‘no gratuitous movement of probability from worlds to dissimilar worlds’. But this justification is no longer applicable if we adopt the present strategy, since the present strategy essentially requires us to surrender any meaningful notion of similarity between worlds. Of course, this is likely to be unsatisfactory for the advocate of imaging. So we take it that **LP** can be discounted as a viable candidate for generalising imaging in the desired way. However, there is one final alternative form of imaging that we’ve not yet considered.

3.3 Proportional Imaging

Joyce (2010) introduces yet another form of imaging. In particular, working in the framework of general imaging where $\sigma(\Gamma, w)$ can have multiple elements, he

⁷This is noted in chapter 15 of Pettigrew (2016).

⁸To see that this assumption is necessary, consider the following example. Let $W_\Gamma = \{w_1, w_2, w_3\}$ and $W_{-\Gamma} = \{w_4\}$. Then, if $\mu(w_1) = \mu(w_2) = \mu(w_3) = 0.1, \mu(w_4) = 0.7, \sigma(\Gamma, w_4) = \{w_1, w_2\}$ and we learn Γ with certainty, then it’s easy to see that Laplace imaging gives us $\mu_\Gamma(w_1) = \mu_\Gamma(w_2) = 0.45, \mu_\Gamma(w_3) = 0.1, \text{ and } \mu_\Gamma(w_4) = 0$, whereas updating with **LP** gives us $\mu_1(w_1) = \mu_1(w_2) = \mu_1(w_3) = 1/3$, which is what Laplacian updating would give if $\sigma(\Gamma, w_4) = W_\Gamma$ held.

suggests that we define (where $\neg w \models E$)

$$T_E(w, w') := \begin{cases} \frac{\mu(w')}{\sum_{w'' \in \sigma(E, w)} \mu(w'')}, & \text{if } w' \in \sigma(E, w) \\ 0 & \text{if } w' \notin \sigma(E, w) \end{cases}$$

i.e. the amount of w 's probability that gets transferred to any of the closest E -worlds after imaging on E will be proportional to the prior probability of those worlds. Note that this approach is fundamentally different to the form of general imaging we defined above, where we assume that $T_E(w, w')$ is always independent of $\mu(w')$. We will call this new form of general imaging 'proportional Imaging'. Joyce gives the following justification for proportional imaging as opposed to the other forms of general imaging,

“... it now seems obvious that imaging *must* involve the *combined* effects of judgments of similarity among worlds and prior probabilities. When the information about similarity runs out and an imager is left with a non-trivial set of E -worlds that are most like w , she still has excellent reasons for treating some worlds in this set differently from others. After all, she began by regarding some worlds in the set as more likely than others, and her evidence has not changed. Imaging should thus be *Bayesianized*, so that probabilities are spread over sets of ‘most similar’ worlds in a way that preserves the imager’s prior beliefs.” (Joyce, 2010)

Consider the following example. Ettie is wondering what's for dinner. She knows that she will either have rice or chips in her meal. She also knows that there are two possible situations if she's going to have chips: either she will have fish and chips, or she will have tofu and chips, but fish with chips is much more likely (since tofu and chips is a bit strange). For simplicity, let's say there is only one possible situation where she gets rice (where she has rice and beans for dinner), and this situation is, in her humble opinion, equally similar to each of the two situations where she gets chips. Now, suppose she sees her dad peeling some potatoes, so she knows that she will be having chips for dinner, not rice. As a staunch advocate of probabilistic imaging, she will now transfer all the prior probability stored in the rice world to the two chip worlds. How should she do this? According to Joyce, she should transfer the majority of the probability of the rice-world to the fish-and-chip-world, since this is by far the more likely of the two chip worlds. Otherwise, she would risk distorting the ratio of the probability of her getting fish and chips compared to the probability of her getting tofu and chips.

To illustrate, suppose she has good reasons for believing that fish and chips is twice as likely as tofu and chips, and suppose that the probability of the rice world is 0.4, so $\mu(TC) = 0.2$, $\mu(FC) = 0.4$, $\mu(R) = 0.4$. Then, according to Laplacian imaging, the new evidence would require her to add 0.2 to each of $\mu(TC)$, $\mu(FC)$ to get the posterior probabilities of those worlds, $\mu_1(TC) = 0.4$ and $\mu_1(FC) = 0.6$.

But then $\mu_1(TC) \neq \mu_1(FC)/2$. So Ettie no longer thinks that fish and chips is twice as likely as tofu and chips, which seems wrong, given that she hasn't learned anything about the comparative probabilities of these two events. Whereas if she uses Joyce imaging, she will get

$$\begin{aligned}\mu_1(TC) &= \mu(TC) + \frac{\mu(R)\mu(TC)}{\mu(TC) + \mu(FC)} = 1/3 \\ \mu_1(FC) &= \mu(FC) + \frac{\mu(R)\mu(FC)}{\mu(TC) + \mu(FC)} = 2/3,\end{aligned}$$

i.e. $\mu_1(TC) = \mu_1(FC)/2$, i.e. Ettie has successfully updated her probabilities without unnecessarily compromising the probability ratios of her prior beliefs. This kind of example makes Joyce's arguments on the previous page look very convincing.

However, one might be tempted to ask whether this 'Bayesianisation' of imaging does not render the approach equivalent to standard conditionalisation. The following example shows that this is not the case. Let $W_A = \{w_1, w_2\}$, $W_{\neg A} = \{w_3\}$ and $W_B = \{w_1, w_3\}$. Then, if we learn the evidence A and update by conditionalisation, our posterior probability for B will be

$$\begin{aligned}P_1(B) &= P_0(B|A) = \frac{P_0(A \wedge B)}{P_0(A)} \\ &= \frac{\mu(w_1)}{\mu(w_1) + \mu(w_2)}.\end{aligned}$$

But now, let's assume that we have a similarity relation such that $\sigma(A, w_3) = \{w_1\}$. Then, according to Joyce imaging, we get

$$\begin{aligned}P_1(B) &= \mu_A(w_1) \\ &= \mu(w_1) + \frac{\mu(w_1)\mu(w_3)}{\mu(w_1)} \\ &= \mu_1 + \mu_3.\end{aligned}$$

Clearly, the two answers are very different here, and this difference is determined completely by the choice of similarity relation. Just as **LP** can be seen as the limit case of Laplacian imaging where the similarity relation is trivial, conditionalisation can be seen as the limit of proportional imaging when we assume an indiscriminate similarity relation. When the similarity relation is non-trivial, we will no longer be guaranteed to preserve probability ratios, because we will be more concerned with ensuring that there is 'no gratuitous movement of probability from worlds to dissimilar worlds'. But proportional imaging (unlike Laplacian imaging, for example) still has the advantage that, when these two desiderata are jointly satisfiable, it will satisfy them both. Furthermore, we will now show that proportional imaging can be generalised to deal with uncertain evidence in a way that avoids at least some of the issues described in previous sections. Recall that both standard and general imaging violated **VC** even for the simple case of modus ponens, and even once one allows

for the world w_1 to have two closest worlds. This is not the case for proportional imaging.

Proposition 2 *Let W , Γ , ϕ be as in the proof of Proposition 3 and suppose that $\sigma(\neg\Gamma, w_1) = \{w_2, w_3\}$. Then assuming proportional imaging, using **JI** to increase the probability of Γ will never lead to a decrease in the probability of ϕ .*

So proportional imaging succeeds where standard and general imaging failed. Given a suitable similarity relation (i.e. one that includes w_3 amongst the closest worlds to w_1), increasing the probability of the premises of a modus ponens argument will never lead to a decrease in the probability of the conclusion. This result is an instance of a more general property of proportional imaging, captured by the following result.

Theorem 3 *Let $\Gamma \rightarrow \phi$ be an instance of a logically valid argument scheme and let W be such that for any $w \in W_\Gamma$, there exists some $w' \in W_{\neg\Gamma \wedge \phi}$ such that $w' \in \sigma(\neg\Gamma, w)$. Then assuming proportional imaging, using **JI** to increase the probability of Γ will never lead to a decrease in the probability of ϕ .*

Theorem 3 guarantees that, assuming proportional imaging, **JI** will always satisfy **VC** except for in cases where there exists a world w that satisfies the premises of a valid argument but is such that all the closest worlds to w at which the premises are false are also worlds at which the conclusion of the argument fails to obtain. Thus, although **JI** does not satisfy the validity criterion **VC** in full generality, we have at least identified the exact conditions under which the criterion can be violated. Whereas other forms of imaging will force **JI** to violate **VC** in an unconstrained and ubiquitous way, proportional Jeffrey imaging will only violate **VC** in special cases where the similarity relation somehow overrides the evidential relationship between the premises and conclusions of valid arguments.

4 Conclusion

In summary, we considered the problem of generalising imaging to deal with uncertain evidential contexts and partial belief revision. We saw that the most natural such generalisation of the best known forms of imaging from the literature violated an intuitive logical criterion in a radical way. We then proved that a lesser known form of imaging (proportional imaging) is better able to satisfy this criterion, although there are still special cases in which the criterion is violated. We hope to investigate these cases in further detail in future work. More generally, the results presented here elucidate a number of important and hitherto neglected structural features of imaging rules and illustrate the obstacles and possibilities facing those

who are interested in generalising imaging to uncertain evidential contexts and/or partial belief revision.

References

- Adams, E. (1996). *A Primer of Probability Logic*. Chicago: The University of Chicago Press.
- Baratgin, J. and G. Politzer (2011). Updating: A Psychologically Basic Situation of Probability Revision. *Thinking and Reasoning* 15: 253–287.
- Cozic, M. (2011). Imaging and Sleeping Beauty: A Case for Double-Halfers. *International Journal of Approximate Reasoning* 52 (2): 137–143.
- 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.
- Gärdenfors, P. (1982). Imaging and Conditionalization. *The Journal of Philosophy* 79(12): 747–760.
- 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.
- Leitgeb, H. and R. Pettigrew (2010). An Objective Justification of Bayesianism II: The Consequences of Minimizing Inaccuracy. *Philosophy of Science* 77: 236–272.
- Lin, H. and Kelly, K.T. (2013). Comments on Leitgeb’s Stability Theory of Belief. *Logic Across the University: Foundations and Applications*, Studies in Logic Vol. 47, London: College Publications.
- Leitgeb, H. (2016). Imaging all the People. To appear in *Episteme*, DOI: 10.1017/epi.2016.14.
- Levinstein, B. (2012). Leitgeb and Pettigrew on Accuracy and Updating. *Philosophy of Science* 79(3): 413–424.
- 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.

Pettigrew, R. (2016). *Accuracy and the Laws of Credence*. Oxford: Oxford University Press.

Sebastiani, F. (1998). Information Retrieval, Imaging and Probabilistic Logic. *Computers and Artificial Intelligence* 17: 1–16.

A Appendix: Proofs

A.1 Proposition 1

Let $\Gamma \vdash \phi$ be a valid argument scheme, and suppose we learn Γ with probability $P'(\Gamma) \geq P(\Gamma)$, and let $\alpha = P'(\Gamma) - P(\Gamma)$. Then,

$$\begin{aligned}
P'(\phi) &= P(\phi|\Gamma) P'(\Gamma) + P(\phi|\neg\Gamma) P'(\neg\Gamma) \\
&= \frac{P(\phi \wedge \Gamma)}{P(\Gamma)} P'(\Gamma) + \frac{P(\phi \wedge \neg\Gamma)}{P(\neg\Gamma)} P'(\neg\Gamma) \\
&= P'(\Gamma) + \frac{P(\phi \wedge \neg\Gamma)}{P(\neg\Gamma)} P'(\neg\Gamma) \\
&= P(\Gamma) + \alpha + \left(\frac{P(\phi \wedge \neg\Gamma)}{P(\neg\Gamma)} \right) (P(\neg\Gamma) - \alpha) \\
&= P(\Gamma) + \alpha + P(\phi \wedge \neg\Gamma) - \frac{\alpha P(\phi \wedge \neg\Gamma)}{P(\neg\Gamma)}.
\end{aligned}$$

So,

$$\begin{aligned}
P'(\phi) - P(\phi) &= P'(\phi) - P(\Gamma) - P(\phi \wedge \neg\Gamma) \\
&= P(\Gamma) + \alpha + P(\phi \wedge \neg\Gamma) - \frac{\alpha P(\phi \wedge \neg\Gamma)}{P(\neg\Gamma)} - P(\Gamma) - P(\phi \wedge \neg\Gamma) \\
&= \alpha - \frac{\alpha P(\phi \wedge \neg\Gamma)}{P(\neg\Gamma)} \\
&\geq \alpha - \frac{\alpha P(\neg\Gamma)}{P(\neg\Gamma)} \\
&= 0. \quad \blacksquare
\end{aligned}$$

A.2 Proposition 2

Let $\Gamma \vdash \phi$ be a valid argument scheme, and suppose we learn Γ with probability $P_1(\Gamma) \geq P_0(\Gamma)$. Again, we set $\alpha = P_1(\Gamma) - P_0(\Gamma)$. Then, we have that

$$\begin{aligned}
P_1(\phi) &= P_1(\Gamma) + P_1(\neg\Gamma \wedge \phi) \\
&= P_0(\Gamma) + \alpha + P_1(\neg\Gamma \wedge \phi).
\end{aligned}$$

Note that $\alpha = \alpha_\Gamma |W_\Gamma|$, and $\alpha \geq \alpha_{-\Gamma} |\{w \in W_{-\Gamma} | \mu(w) - \alpha_{-\Gamma} > 0\}|$. Thus, we have that

$$\begin{aligned}
P_1(\phi) - P_0(\phi) &= P_0(\Gamma) + \alpha + P_1(\neg\Gamma \wedge \phi) - P_0(\Gamma) - P_0(\neg\Gamma \wedge \phi) \\
&= \alpha + P_1(\neg\Gamma \wedge \phi) - P_0(\neg\Gamma \wedge \phi) \\
&= \alpha + \sum_{w \in W_{-\Gamma} | \mu(w) - \alpha_{-\Gamma} > 0} (\mu(w) - \alpha_{-\Gamma} - \mu(w)) \\
&= \alpha - \alpha_{-\Gamma} |\{w \in W_{-\Gamma} | \mu(w) - \alpha_{-\Gamma} > 0\}| \\
&\geq \alpha - \alpha_{-\Gamma} |\{w \in W_{-\Gamma} | \mu(w) - \alpha_{-\Gamma} > 0\}| \\
&\geq 0. \quad \blacksquare
\end{aligned}$$

A.3 Proposition 4

In particular, sticking with the modus ponens example, let's assume the following initial setup: $\mu(w_1) = 0.1$, $\mu(w_2) = 0.05$, $\mu(w_3) = 0.8$ and $\mu(w_4) = 0.05$ (and hence $P_0(\Gamma) = 0.1$) with the following constraints: $P^*(\Gamma) = 0.11$ and $P^*(\neg\Gamma) = 0.89$. Then, we have that $P_0(\phi) = \mu(w_1) + \mu(w_3) = 0.9$ and

$$\begin{aligned}
P_\Gamma^*(\phi) &= \mu_\Gamma^*(w_1) + \mu_\Gamma^*(w_3) \\
&= \mu_\Gamma(w_1) P^*(\Gamma) + \mu_{-\Gamma}(w_3) P^*(\neg\Gamma) \\
&= P^*(\Gamma) + \mu_{-\Gamma}(w_3) P^*(\neg\Gamma) \\
&= P^*(\Gamma) + (\mu(w_3) + \mu(w_1)/2) P^*(\neg\Gamma) \\
&= 0.11 + (0.85)(0.89) = 0.8665.
\end{aligned}$$

Hence, $P_0(\phi) > P_\Gamma^*(\phi)$. \blacksquare

A.4 Theorem 1

Again, we use the example of modus ponens. First off, we let $\mu_1 := \mu(w_1)$, $\mu_3 := \mu(w_3)$, $\alpha := P^*(\Gamma) - P_0(\Gamma) = P^*(\Gamma) - \mu_1$ (we know this is positive, by the assumption that $P^*(\Gamma) \geq P_0(\Gamma)$), $\beta := \mu_{-\Gamma}(w_3) - \mu_3$ (again, we know this is positive, by the definition of imaging). Then, we want to show that

$$P^*(\phi) - P_0(\phi) \geq 0.$$

Now,

$$\begin{aligned}
P_\Gamma^*(\phi) - P_0(\phi) &= \mu_\Gamma^*(w_1) + \mu_\Gamma^*(w_3) - \mu_1 - \mu_3 \\
&= \mu_\Gamma(w_1) P^*(\Gamma) + \mu_{-\Gamma}(w_3) P^*(\neg\Gamma) - \mu_1 - \mu_3 \\
&= P^*(\Gamma) + \mu_{-\Gamma}(w_3) P^*(\neg\Gamma) - \mu_1 - \mu_3 \\
&= \mu_1 + \alpha + (\mu_3 + \beta)(1 - \mu_1 - \alpha) - \mu_1 - \mu_3 \\
&= \alpha + \mu_3 - \mu_1 \mu_3 - \alpha \mu_3 + \beta - \beta \mu_1 - \alpha \beta - \mu_3 \\
&= \alpha(1 - \mu_3) + \beta(1 - \mu_1) - \alpha \beta - \mu_1 \mu_3.
\end{aligned}$$

Let's define the function

$$f(\alpha, \beta, \mu_1, \mu_3) := \alpha(1 - \mu_3) + \beta(1 - \mu_1) - \alpha\beta - \mu_1\mu_3.$$

Now, clearly, this function is decreasing in μ_1 and μ_3 , so to violate the inequality, let's set $\mu_1 + \mu_3 = 1$. Then, we have

$$\begin{aligned} f(\alpha, \beta, \mu_1, \mu_3) &= f(\alpha, \beta, \mu_1, 1 - \mu_1) \\ &= \alpha(1 - (1 - \mu_1)) + \beta(1 - \mu_1) - \alpha\beta - \mu_1(1 - \mu_1) \\ &= \alpha\mu_1 + \beta - \beta\mu_1 - \alpha\beta - \mu_1 + \mu_1^2 \\ &= (\beta - \mu_1)(1 - \alpha - \mu_1). \end{aligned}$$

Clearly, the term on the right is always non-negative, since $\alpha + \mu_1 \leq 1$ (by definition of α). But the term on the left will always be negative except for the case in which $\mu_1 = \beta$ (we know that $\beta \leq \mu_1$ always holds by definition of β). So the only case in which the probability of ϕ doesn't decrease here is where $\mu_1 = \beta$, i.e. where we transfer all of μ_1 's probability to μ_3 . For any other alternative to **WI**, we can find values for μ_1, μ_3, α that violate **VC**. ■

A.5 Theorem 3

Setting $e_w = \{w' \in W | w \in \sigma(\neg\Gamma, w')\}$, where $\Gamma \vdash \phi$ is valid, the assumption (which we will refer to as **VA**) says that for any $w' \in W_\Gamma$, there exists some $w \in W_{\neg\Gamma \wedge \phi}$ such that $w \in e_{w'}$ (of course, we don't need to worry about the case in which every ϕ world is a Γ world, since in this case the simple assumption that $P_1(\Gamma) \geq P_0(\Gamma)$ is sufficient). Next we set

$$S = \sum_{w \in W_{\neg\Gamma \wedge \phi}} \sum_{w' \in e_w} \frac{\mu(w)\mu(w')}{\sum_{w'' \in \sigma(\neg\Gamma, w')} \mu(w'')}.$$

We have

$$\begin{aligned} P_1(\phi) &= P_1(\phi \wedge \Gamma) + P_1(\phi \wedge \neg\Gamma) \\ &= P_1(\Gamma) + P_1(\phi \wedge \neg\Gamma) \\ &= P_0(\Gamma) + \alpha + \left(\sum_{w \in W_{\neg\Gamma \wedge \phi}} \mu_{\neg\Gamma}(w) \right) (1 - P_0(\Gamma) - \alpha) \\ &= P_0(\Gamma) + \alpha + \left(\sum_{w \in W_{\neg\Gamma \wedge \phi}} \mu(w) + S \right) (1 - P_0(\Gamma) - \alpha) \\ &= P_0(\Gamma) + \alpha + (P_0(\neg\Gamma \wedge \phi) + S) (1 - P_0(\Gamma) - \alpha) \\ &= P_0(\Gamma) + \alpha + P_0(\neg\Gamma \wedge \phi) - P_0(\neg\Gamma \wedge \phi)P_0(\Gamma) - P_0(\neg\Gamma \wedge \phi)\alpha + S(1 - P_0(\Gamma) - \alpha). \end{aligned}$$

So,

$$\begin{aligned} P_1(\phi) - P_0(\phi) &= \alpha - P_0(\neg\Gamma \wedge \phi) P_0(\Gamma) - P_0(\neg\Gamma \wedge \phi)\alpha + S(1 - P_0(\Gamma) - \alpha) \\ &= \alpha(1 - P_0(\neg\Gamma \wedge \phi)) - P_0(\neg\Gamma \wedge \phi) P_0(\Gamma) + S(1 - P_0(\Gamma) - \alpha). \end{aligned}$$

Clearly, this function is decreasing in $P_0(\neg\Gamma \wedge \phi)$ and $P_0(\Gamma)$, so setting $P_0(\Gamma) + P_0(\neg\Gamma \wedge \phi) = 1$, we get,

$$P_1(\phi) - P_0(\phi) = \alpha P_0(\Gamma) - P_0(\neg\Gamma \wedge \phi) P_0(\Gamma) + S(P_0(\neg\Gamma \wedge \phi) - \alpha).$$

We now show that $S \geq P_0(\Gamma)$. First, recall that

$$S = \sum_{w \in W_{\neg\Gamma \wedge \phi}} \sum_{w' \in e_w} \frac{\mu(w)\mu(w')}{\sum_{w'' \in \sigma(\neg\Gamma, w')} \mu(w'')}.$$

Let $w' \in W_\gamma$. Then, by **VA**, we know that there exists $w \in W_{\neg\Gamma \wedge \phi}$ such that $w' \in e_w$, i.e. $w \in \sigma(\neg\Gamma, w')$. Furthermore, by the fact that $w' \in e_w$ and the assumption that $P_0(\Gamma) + P_0(\neg\Gamma \wedge \phi) = 1$, we get

$$\sum_{w'' \in \sigma(\neg\Gamma, w')} \mu(w'') = \sum_{w'' \in \sigma(\neg\Gamma \wedge \phi, w')} \mu(w'').$$

Now, note that

$$\mu(w') = \frac{\sum_{w'' \in \sigma(\neg\Gamma \wedge \phi, w')} \mu(w'')\mu(w')}{\sum_{w'' \in \sigma(\neg\Gamma \wedge \phi, w')} \mu(w'')}.$$

So,

$$P_0(\Gamma) = \frac{\sum_{w' \in W_\Gamma} \sum_{w \in \sigma(\neg\Gamma \wedge \phi, w')} \mu(w)\mu(w')}{\sum_{w'' \in \sigma(\neg\Gamma, w')} \mu(w'')}.$$

And, by another application of **VA**,

$$\sum_{w' \in W_\Gamma} \sum_{w \in \sigma(\neg\Gamma \wedge \phi, w')} \mu(w)\mu(w') \leq \sum_{w \in W_{\neg\Gamma \wedge \phi}} \sum_{w' \in e_w} \mu(w)\mu(w').$$

This gives us $S \geq P_0(\Gamma)$, as desired. Thus,

$$\begin{aligned} P_1(\phi) - P_0(\phi) &\geq \alpha P_0(\Gamma) - P_0(\neg\Gamma \wedge \phi) P_0(\Gamma) + P_0(\Gamma) (P_0(\neg\Gamma \wedge \phi) - \alpha) \\ &= 0. \quad \blacksquare \end{aligned}$$