

Don't Look Now

Arif Ahmed and Bernhard Salow

Draft – Please Cite Published Version

Abstract

Good's Theorem is the apparent platitude that it is always rational to 'look before you leap': to gather (reliable) information before making a decision when doing so is free. We argue that Good's Theorem is not platitudinous and may be false. And we argue that the correct advice is rather to 'make your act depend on the answer to a question'. Looking before you leap is rational when, but only when, it is a way to do this.

1. *Introduction*
2. *Good's Theorem*
3. *Inexact Observation*
4. *Independence*
5. *Conditionality*
 - 5.1 *The principle*
 - 5.2 *Revisiting the counterexamples*
6. *Conclusion*

1 Introduction

Good's Theorem says: look before you leap if looking is free.

'Free' excludes all costs: of money, of time, of opportunity. Good doesn't say that you should observe everything before doing anything; nor that you should pay the monetary cost of checking your tax return with every accountant, the temporal cost of reading everything on a subject before writing a paper on it, or the opportunity cost of constant dithering. He only recommends not choosing until you have learnt what you can learn for free. So qualified, the theorem looks platitudinous. 'Look before you leap' isn't always good advice; 'Look for free before you leap' surely is.

But there are counterexamples to Good's Theorem, described and classified here. Since Good's proof is valid, the counterexamples show that his premises are contentious. First question: which premises and why? Second question: how close we can get to Good's Theorem without them? – is any genuine platitude within the neighbourhood of 'look before you leap'? We answer both.

Section 2 presents Good's proof. Sections 3-4 give two counterexamples and identify which premises cause trouble. Section 5 defends 'Conditionality' (5.1), which is the best we can do whilst avoiding that trouble. Conditionality has two virtues: it follows from Good's uncontroversial assumptions (5.1); and it yields Good's Theorem as a special case in certain environments (5.2).

2 Good's Theorem

Informally, Good's Theorem (GT) says

A rational agent will make any available, free and relevant observation before choosing an act.

Before the proof, we make one comment and stress three points. Comment: ‘relevant’ is interpreted lightly: an observation is relevant unless the agent is certain that it has no bearing on the decision. (The proof also shows that when a free observation isn’t relevant, it is still permissible to make it.)

Three further points: first (as mentioned): calling observation ‘free’ rules out anything the agent considers a cost. Second: making the observation is not supposed to affect any event on which the success of any available act depends. But third: it may affect the agent’s beliefs about these events—that’s the point.

Good’s ([1967]) proof uses standard (i.e. Savage’s) decision theory and standard (i.e. Bayesian) formal epistemology. Specifically, Good assumes that rational agents (a) have preferences that maximize expected utility and choose in accordance with these preferences, and (b) update by conditionalization. Before stating (a) and (b) we outline their background: not because it plays a role in Good’s proof, but because we must relax it if (as we argue) there is reason to dilute GT.

We have a set W of possible worlds (assumed, for simplicity, to be finite) and a set Z of possible outcomes or payoffs (identified, for simplicity, with dollar credits and debits, so that Z is a set of real numbers). The ‘act space’ A then contains all functions from W to Z i.e. $A =_{\text{def.}} Z^W$. An act $a \in A$ specifies a ‘prize’ $a(w)$ in each world w . An ‘event’ or ‘proposition’ is a subset of W . If h is a proposition, and $a(w_1) = a(w_2)$ whenever $w_1, w_2 \in h$, we define $a(h)$ as $a(w)$ for any $w \in h$. So a gamble paying \$1 if this coin lands heads (h_1) and losing \$1 otherwise (h_2) is the act a such that $a(h_1) = 1$ and $a(h_2) = -1$.

Classical decision theory says that for $a_1, a_2 \in A$, a rational agent’s preference satisfies:

$$EU(a) =_{\text{def.}} \sum_{h \in H} Cr(h)a(h) \quad ^1$$

$$a_1 \succ a_2 \text{ iff } EU(a_1) > EU(a_2)$$

‘ $a_1 \succ a_2$ ’ means the agent strictly prefers a_1 to a_2 (so always chooses a_1 when a_2 is the only alternative). Similarly ‘ $a_1 \succeq a_2$ ’ means $\sim(a_2 \succ a_1)$: the agent ‘weakly prefers’ a_1 to a_2 . $EU(a)$ is the expected utility of a . It is a weighted average of the returns to a in the events h_1, h_2, \dots, h_n (chosen so that $a(w_1) = a(w_2)$ whenever $w_1, w_2 \in h_i$), the weight for h_i being $Cr(h_i)$. Cr is the agent’s ‘credence’: a probability function mapping a proposition to the confidence the agent assigns to that proposition. It follows that given a menu M of acts ($M \subseteq A$), her expectation of utility when facing M with belief state Cr is:

$$(1) \quad EU(M, Cr) =_{\text{def.}} \max_{a \in M} \sum_{h \in H} Cr(h)a(h)$$

Thus suppose you are choosing whether to bet \$1 on Red Rum: winning makes a profit of \$1. Let h_1 be the proposition that (that is, the set of possible worlds at which) Red Rum wins and h_2 the proposition that he doesn’t. The available acts are then as follows.

¹ We assume that the agent has increasing linear utility for money. Nothing turns on this.

	h_1 : RR wins	h_2 : RR doesn't win
a_1 : bet	1	-1
a_2 : no bet	0	0

Table 1: *Red Rum*

a_1 and a_2 are functions from $\{h_1, h_2\}$ to $Z = \{-1, 0, 1\}$ s.t. $a_1(h_1) = 1$, $a_1(h_2) = -1$, $a_2(h_1) = a_2(h_2) = 0$.

The expected utility of a_1 is $Cr(h_1) - Cr(h_2)$, which is $2Cr(h_1) - 1$ since $Cr(h_1) + Cr(h_2) = 1$. The expected utility of a_2 is 0. Since you pick whichever act has higher *EU*, the expected utility of choosing from $\{a_1, a_2\}$ is:

$$EU(\{a_1, a_2\}, Cr) = \max\{0, 2Cr(h_1) - 1\}.$$

So much for the background decision theory.

Call a set of propositions $P = \{p_1, p_2, \dots, p_m\}$ a ‘partition’ if those propositions are mutually exclusive ($i \neq j \leftrightarrow p_i \cap p_j = \emptyset$) and jointly exhaustive ($\cup P = W$). Now suppose that before choosing you have the option to learn which element of a partition P is true—that is, you can do what Good calls ‘making an observation’ that settles P . (Since P partitions W , exactly one element of P is true.) As Bayesians, we assume that upon observation you update your beliefs Cr by conditionalization: if you learn that $p \in P$ is true, where $Cr(p) > 0$, then your credence shifts from Cr to a new probability function Cr_p , defined on an arbitrary proposition h as follows:

$$Cr_p(h) = Cr(h|p) =_{\text{def.}} \frac{Cr(hp)}{Cr(p)}$$

(We’ll sometimes write hp for $h \cap p$.) $Cr(h|p)$ is the ‘conditional probability’ of h given p . So much for the background epistemology.

Suppose you observe before choosing and learn p . Given your new credence function Cr_p you now choose $a \in M$ maximizing:

$$EU(a|p) = \sum_{h \in H} Cr(h|p)a(h)$$

If observing means learning which $p \in P$ is true, the *ex ante* expected utility of observation is:

$$(2) \quad \sum_{p \in P} Cr(p)EU(M, Cr_p) = \sum_{p \in P} Cr(p) \max_{a \in M} \sum_{h \in H} Cr(h|p)a(h)$$

Given that a rational person maximizes expected utility, it follows from (1) and (2) that GT holds if this inequality does:

$$\sum_{p \in P} Cr(p) \max_{a \in M} \sum_{h \in H} Cr(h|p)a(h) \geq \max_{a \in M} \sum_{h \in H} Cr(h)a(h)$$

By definition of $Cr(h|p)$, this is equivalent to:

$$(3) \quad \sum_{p \in P} \max_{a \in M} \sum_{h \in H} Cr(hp)a(h) \geq \max_{a \in M} \sum_{h \in H} Cr(h)a(h)$$

Good proves GT by proving (3).

To get a feeling for how, consider this analogy. The US government faces a menu M of nationwide tax regimes that each assigns a fixed rate to each economic class, where individuals occupy the same class if and only if they have the same gross income. The regimes typically have different effects on different classes. The government aims to maximize net income aggregated across all individuals. Equivalently, it aims to maximize the weighted sum of net incomes earned by each class, where the weight of a class is the proportion of individuals within it.

As well as partitioning the population into economic classes, we can partition it into States. Suppose every individual resides in exactly one State. So each tax regime has implications for the income of each State i.e. the aggregate income of its residents. We number the regimes alphabetically by State, so that $a_1 \in M$ is by this measure the optimal regime for Alabama, a_2 for Alaska, ... a_{50} is optimal for Wyoming. (Maybe $a_i = a_j$ for some $i \neq j$.)

Now instead of a single tax regime applied across all States, imagine a scheme a^s that varies the regime across States as follows: it applies a_1 in Alabama, a_2 in Alaska, ... a_{50} in Wyoming. Then (3) is saying that a^s is at least as good, by the government's measure, as any $a \in M$. Good's argument for this is that a^s is at least as good as a in each State. a^s is at least as good as a in Alabama, because a^s implements a_1 in Alabama, which is optimal in M for Alabama, etc. Since the weighted sum of incomes in any State is at least as high under a^s as under any $a \in M$, and since everyone resides in exactly one State, the sum of all incomes, and hence the weighted sum of class earnings, is at least as high under a^s as under any $a \in M$.

Citizens correspond to worlds, income classes to events $h \in H$, States to observations $p \in P$, populations—of classes and States—to credence, and the original menu of uniform tax regimes to your options. Specifically, for a class h , a State p and a tax regime a , the number of individuals in that class in that State corresponds to $Cr(hp)$, the number of individuals in that class across the US to $Cr(h)$, and the net income under a of anyone in that class to $a(h)$.

The intuitive story translates into a literal proof of (3). Since $Cr(h) = \sum_{p \in P} Cr(hp)$, the right-hand side of (3) is:

$$(4) \quad \max_{a \in M} \sum_{p \in P} \sum_{h \in H} Cr(hp)a(h)^2$$

But 'the sum of the maxima exceeds the maximum of the sums'. (4) chooses a single $a \in M$, call it a^* , that maximizes $\sum_{p \in P} \sum_{h \in H} Cr(hp)a(h)$. But the left-hand side of (3) chooses, for each $p \in P$, an $a^p \in M$ that maximizes $\sum_{h \in H} Cr(hp)a(h)$. So if $p \in P$ then $\sum_{h \in H} Cr(hp)a^p(h) \geq \sum_{h \in H} Cr(hp)a^*(h)$; (3) follows.

Let's apply this to *Red Rum*. Before choosing whether to bet you can listen—for free—to a tip. This has two possible outcomes: the tipster forecasts either (p_1) that Red Rum wins or (p_2) that he loses. The forecast needn't be reliable in your opinion; but you have some opinion about its reliability. Suppose you think a win has probability x if forecast and y if not. So $Cr(h_1|p_1) = x$ and $Cr(h_1|p_2) = y$. Since the forecast is free, your payoffs in this new problem—*Red Rum**—are:

² Using $\sum_{x \in X} \sum_{y \in Y} f(x, y) = \sum_{y \in Y} \sum_{x \in X} f(x, y)$.

	h_1p_1 : forecast win, win	h_2p_1 : forecast win, loss	h_1p_2 : forecast loss, win	h_2p_2 : forecast loss, loss
a_1 : bet	1	-1	1	-1
a_2 : no bet	0	0	0	0
a_3 : listen to tip	1 if $x > 0.5$ 0 if $x \leq 0.5$	-1 if $x > 0.5$ 0 if $x \leq 0.5$	1 if $y > 0.5$ 0 if $y \leq 0.5$	-1 if $y > 0.5$ 0 if $y \leq 0.5$

Table 2: *Red Rum**

The expected utilities are:

$$(5) \quad EU(a_1) = 2Cr(h_1) - 1$$

$$(6) \quad EU(a_2) = 0$$

$$(7) \quad EU(a_3) = Cr(p_1)\max\{2x - 1, 0\} + Cr(p_2)\max\{2y - 1, 0\}$$

Now $2Cr(h_1) - 1 = Cr(p_1)(2x - 1) + Cr(p_2)(2y - 1)$. So (7) \geq (5), (6) for any x, y . It's always rational to listen: doing so can't worsen but can improve things, in expectation.³ So much for Good's Theorem.

3 Inexact observation

GT is intuitive and its proof is simple. So counterexamples are interesting. The first, using a case suggested—for different purposes—by (Williamson [2011]), is as follows.

You are facing an unmarked clock with a single digit that can point in any of 60 uniformly separated directions that we (but not the clock) label 0 to 59. Your discrimination of where the digit is pointing is limited; to be reliable in your judgment you must leave a margin of error. If the digit is in fact pointing at 53, you can reliably judge that it is pointing between 52 and 54 inclusive, but you are unreliable about stronger claims. Similarly for every other position.⁴

Plausibly, your evidence is the strongest claim you can reliably get right. This means that the maximal propositions you learn in the various scenarios overlap. Let D be the digit's true direction. If $D = 53$, your evidence is that $D \in \{52, 53, 54\}$. If $D = 52$, it is that $D \in \{51, 52, 53\}$. Each scenario yields evidence compatible with the other. But they yield different evidence.

Call a world w_2 accessible from w_1 if your evidence in w_1 is consistent with w_2 . Accessibility is not transitive. $D = 53$ is consistent with the evidence from looking if $D = 52$. $D = 54$ is consistent with the evidence if $D = 53$. But $D = 54$ is not consistent with the evidence if $D = 52$. More generally the situation is as in Figure 1.




Figure 1: *Accessibility relations on an unmarked clock*

³ GT does not say that you will actually do better by listening. The tip might mislead. GT says that listening is better 'in expectation', from the perspective of your *ex ante* credence.

⁴ In reality, it's presumably vague what margins suffice for your judgments to count as 'reliable'. But the problem arises for any fixed size of the margin above a certain lower bound; hence it arises on both epistemic and supervaluational approaches.

Examples like this have been widely discussed in epistemology.⁵ They are interesting because they show how one might receive evidence without knowing that this is the evidence one received. When (say) $D = 52$, your evidence is that $D \in \{51,52,53\}$. But if $D = 53$, which is consistent with what you learned (namely with $D \in \{51,52,53\}$), your evidence would have been that $D \in \{52,53,54\}$. So, while your evidence is that $D \in \{51,52,53\}$, you do not know this.⁶ If you did know this, you could combine it with your knowledge of the setup to infer that $D = 52$, since that is the only case in which you get this evidence. But it's implausible that you could learn something this exact from looking at the clock—even if you did know the setup.

While widely discussed in epistemology, there has been little discussion of how such cases bear on decision theory, specifically on Good's Theorem.⁷ Yet they do. Let O say that the digit is pointing at an odd number; E that it is pointing at an even number (calling 0 even). And suppose you haven't looked at the clock yet, but know that if you do, you will conditionalize on what you learn.⁸ Then by your own lights *ex ante*, looking is misleading regarding O and E .⁹ Suppose O , say $D = 53$. Then looking rules out all possibilities except that and the 'adjacent' ones $D = 52$ and $D = 54$. Of these, two are in E and one is in O . So your evidence supports E .¹⁰ Similarly, if E is true, looking will support O .

This fact—that, by your own lights *ex ante*, looking is misleading—creates failures of GT. Suppose you are offered two bets. ODD costs 60¢ and pays \$1 if O . EVEN costs 60¢ and pays \$1 if E . You can take either or neither now, or you can look at the clock before deciding. Then you should think it a bad idea to look. Your evidence now supports not betting. Suppose you look before deciding. If D is odd, your evidence will support E to degree 0.67, making it rational to buy only the losing bet EVEN. If D is even, the evidence will make it rational to buy only the losing bet ODD. Either way, you pay 60¢ for a losing bet if you look first. You know all this beforehand. So you know that it's a bad idea to look first.

⁵ (Williamson [1992]) first defends cases with roughly this structure. Much of the responding literature focuses on knowledge, rather than evidence. Exceptions include (Christensen [2010]; Elga [2013]; Horowitz [2014]), which sympathetically discuss the above description. (Cohen and Comesaña [2013]) builds on (Stalnaker [2009], pp. 405-6) to develop an alternative formal treatment, which—when applied to the clock—would make the case consistent with GT. (Hawthorne and Magidor [2009]) and (Williamson [2013], pp. 80-3) criticize this alternative.

⁶ All you know about your evidence is that it is either $D \in \{50,51,52\}$ or $D \in \{51,52,53\}$ or $D \in \{52,53,54\}$.

⁷ Though see (Das [unpublished]) and (Dorst [unpublished]). Other related literature differs in focus. (Geanakoplos [1989]) shows that failures of Good's Theorem arise from imperfections in the processing of partitioned observation: taking a signal at face-value, forgetfulness, wishful thinking, doxastic obstinacy or failures of imagination. By contrast, we argue that observation of the unmarked clock is itself non-partitional, which undermines GT even for agents lacking all those imperfections. (Williamson [2000], pp. 230-7) discusses non-partitional observation and its relevance to decision theory, but focuses on exploitability. Bronfman ([2014]) and Schoenfeld ([forthcoming]) argue that these cases make conditionalization non-accuracy-maximizing. They are opposing the epistemic optimality of processing information in the usual way; we are opposing (and Good is defending) the pragmatic optimality of seeking information in the first place.

The literature also discusses how imprecise probabilities bear on Good's Theorem: Good ([1974]) himself raises this worry, attributing it to Isaac Levi and Teddy Seidenfeld; (Seidenfeld [2004]) and (Bradley and Steele [2016]) discuss it in detail. However, imprecise probabilities and inexact observations raise different issues. Inexact observation, as described above, is consistent with the relevant probability distribution's assigning to each proposition a single real number, hence consistent with standard expected utility theory. By contrast, imprecise probabilities violate standard expected utility theory. Counterexamples to GT that arise from such imprecision are thus more like our second type of counterexample; see note 23 below.

⁸ This assumption, that you know *ex ante* that you are a Bayesian, is essential. But it is one that Good himself both makes and needs: without it we lose motivation for (2), which is crucial to Good's proof. So if our example undermines the assumption, it leaves Good with problems as serious as—though admittedly quite different from—those we want to raise.

⁹ Compare (Horowitz [2014], pp. 735-740), which tentatively defends this feature of the example.

¹⁰ At least if, as we can assume, your initial distribution is uniform i.e. $Cr(D = x) = 1/60$ if $x = 0,1, \dots, 59$.

This can seem puzzling. You know throughout that:

(8) Your evidence on looking is $D \in \{51,52,53\}$ iff $D = 52$ iff ODD is a losing bet.

Why, then, do you pay for ODD upon receiving $D \in \{51,52,53\}$? The answer is just what makes this case epistemologically interesting. Receiving $D \in \{51,52,53\}$ does not tell you that you received $D \in \{51,52,53\}$. So, upon receiving $D \in \{51,52,53\}$, you cannot apply *modus ponens* to (8) to conclude that ODD is a losing bet. Once we understand this crucial, but surprising, feature of the example, the puzzle disappears.

We'll discuss three objections. The first two arise from the fact that, according to our description of the case, it allows you to receive evidence without knowing that you have received it.

First Objection. This shows that we have misdescribed the example. You can always tell what your evidence is; so the evidence can't be as we claimed.

Response. It's natural to think you can always tell what your evidence is. It's also natural to think that your evidence when looking at an unmarked clock is as we said it was. The two conflict. Clearly, we cannot resolve this conflict here.

However, we do not need to. If the debates about inexact perception have taught us anything, it's that it isn't obvious what to say—and hence not obvious that the case does not work as we say. This suffices for our puzzle. GT is supposed to be a platitude. If it depends on a far-from-obvious claim about how inexact perception works, it isn't one.

Second Objection. Grant that our description of your evidence is correct: you receive evidence without knowing that you received it. But then how can it influence you? The argument against Good took for granted that (you know *ex ante* that) if you look at the clock then your choice will maximize *EU* relative to the credence you then acquire. That assumption was needed to argue (n. 8) that you will buy the losing bet if you look. But why think you will maximize *EU* relative to that credence function if looking leaves open what it is?

Response. This objection assumes that to act on a belief you must know what it is.¹¹ The underlying picture of practical rationality is this: the *EU*-maximizer surveys her credences and utilities, compares the *EU* of her options, and chooses one that scores highest.

This picture is wrong. 'Practical rationality' doesn't constrain how an agent chooses in light of her beliefs and desires, but only what choices she makes given that she does in fact have those beliefs and desires. In particular, we want to evaluate the choices of creatures who don't know that they have beliefs, let alone what they are. Chrysippus' dog, faced with three paths, set off down path C after smelling that its quarry was not on paths A or B. We want to say the dog acted rationally; given its belief that the quarry was on path C, it would have been irrational to set off down path A. But the dog does not know it has this belief. So we cannot say this if practical rationality requires doxastic self-knowledge. (Similarly, eliminativists in

¹¹ A less demanding version of the objection says only that you must know at each point what your preferences are. This rules out the particular case (where you won't know your preferences after looking), but not natural variants. Imagine that you'll get to see the hand for slightly longer when it's pointing at an odd number, so that in those cases you can reliably detect its exact position—but the added time is short enough that you can't detect whether you've been given it. Then you know beforehand that looking will give you evidence that the number is odd (raising its probability to 1 if it is, or .67, if not), so that your later preference will be to bet ODD whatever you learned. So you'll also know, after looking, that you then prefer to bet ODD. But your *ex ante* evaluation of looking, hence buying ODD, is still negative: the .5 chance of losing 60¢ outweighs the .5 chance of winning 40¢.

the philosophy of mind deny that anyone has beliefs, so cannot know the contents of their own; this does not protect them from charges of practical irrationality.)

The third objection grants the overall structure of the case, but challenges the details.

Third Objection. Grant that looking eliminates values far from D and leaves open close ones. It doesn't follow that you distribute credence uniformly over the remaining possibilities. More likely you concentrate on central values within the range of still-live possibilities. So a glance when $D = 52$ might shift you to some Cr^* satisfying:

$$(9) \quad Cr^*(D = 51) = 0.1$$

$$(10) \quad Cr^*(D = 52) = 0.8$$

$$(11) \quad Cr^*(D = 53) = 0.1$$

But Cr^* will prompt you to bet EVEN. Generally, if looking makes your credence as sharply peaked as Cr^* about the centre of its support, then looking is a good idea *ex ante*.¹²

Response. Grant that a glance adjusts your credence to some Cr^* symmetric about and modal at the true value of D , and not uniform over its support. Still, if Cr^* is 'close enough' to uniform, then our problem still arises given adjusted charges for the bets.

Let $[j]$ be the proposition that $D = j \bmod 60$; and let $Cr^*([i]) = x$ and $Cr^*([i - 1]) = Cr^*([i + 1]) = y$ if in fact $D = i$. Since we can pick the distance between the positions as we like, there's no harm in assuming that $Cr^*([i - 1] \cup [i] \cup [i + 1]) = 1$, so that $x + 2y = 1$. We can reprice the bets ODD and EVEN at $\$c$. Then, as long as $0.5 < c < 2y$, looking at the clock motivates taking (only) the losing bet on its position. Such a c exists whenever $y > .25$.

Glancing *may* shift your credence to a distribution violating these conditions—like (9)-(11)—and then the example creates no trouble for GT, whatever the charge on each bet. But why *must* it? It is easy to imagine evidence to the contrary. Suppose that you are repeatedly forced to guess D at a glance: you are right one time in three and never out by more than 1. If we like we can build this into the case.

It might seem strange that when $D = i$, $Cr^*([i + 1])$ and $Cr^*([i - 1])$ both exceed 0.25 and $Cr^*([i + 2])$ and $Cr^*([i - 2])$ are both zero, since then the probability that the digit is pointing at j does not decrease linearly to zero with the angular distance between j and i , but rather more quickly as we get further from i . But this is not so implausible. Suppose you have visual receptors corresponding to each second of arc on the clock-face. When you look at the clock and $D = d^*$, visual receptors fire at a rate that is normally distributed about the true position: those corresponding to a position z seconds of arc away from d^* fire at a rate proportional to e^{-z^2/σ^2} for some constant σ . And suppose that looking rules out a position if and only if the receptors corresponding to that position fire at a rate below some threshold $\Delta > 0$. Then $Cr^*(D = x)$ will fall first gently and then sharply – that is, non-linearly – towards zero as $|x - d^*|$ increases, just as in our model.

¹² That you shift to Cr^* is inconsistent with simple Bayesian updating on $D \in \{51, 52, 53\}$ given a uniform prior. A natural model of the shift is Jeffrey conditionalization (Jeffrey [1983], pp. 164ff.). The idea is that perception 'directly' adjusts your confidence in certain propositions ('This cloth is green', ' $D = 52$ ') within the open interval $(0, 1)$. If perception changes your Cr by directly adjusting your credences in X_1, \dots, X_n respectively to $Cr^*(X_i)$, then its indirect impact on your credence is given by $Cr^*(Y) = \sum_{i=1}^n Cr(Y|X_i)Cr^*(X_i)$ for any Y in the algebra. Good assumes ordinary conditionalization; (Graves [1989]) generalizes the argument to Jeffrey conditionalization. Our example shows that, much as Good's argument requires an objectionably narrow definition of 'making an observation', such generalizations require an objectionably narrow view of what makes a Jeffrey shift a 'genuine learning experience'.

Again, we needn't show that our problem case actually arises, only that it would in circumstances that are neither ruled out *a priori* nor excluded from the intuitive principle ('look before you leap'). $y > 0.25$ meets those conditions: nobody ever took GT to hold only if perception doesn't work like that. It is hardly platitudinous that it doesn't. Neither therefore is Good's Theorem itself.

Where does Good's proof go wrong, then? Good identifies 'making an observation' with learning exactly which element of some partition is true: specifically, it assumes you are certain *ex ante* that the maximal proposition learnt by observation is inconsistent with any other maximal proposition you might learn by observation. Without this, there is no reason why (2) should represent the expected utility of observing.

Our example exhibits observations for which no partition P is such that you are certain *ex ante* that making the observation teaches you exactly which $p \in P$ is true. $D \in \{51, 52, 53\}$ and $D \in \{52, 53, 54\}$ are both propositions that might exhaust what you learn by looking; but they are consistent, since both hold if $D = 52$.

So the problem is not that Good's premises fail to entail his conclusion, but that they understand 'observation' too narrowly. Without this narrow understanding, there is no guarantee that a policy that adjusts one's bet to the contents of the observation will do at least as well 'in expectation' as taking whatever bet looks optimal *ex ante*.¹³ Good's proof does not apply to anyone whose observations are inexact. And it's no platitude that there aren't—much less couldn't be—people like that.

4 Independence

The second counterexample is better known: we present it more briefly.¹⁴ Let h_1, h_2 and h_3 partition W and let gambles a_1 - a_4 on these events give you terminal wealth (in millions of dollars) as specified below:

	h_1	h_2	h_3
a_1	0	5	0
a_2	1	1	0
a_3	0	5	1
a_4	1	1	1

Table 3: *Allais*

It follows from *EU*-maximization that a rational wealth-lover prefers a_1 (a_2) from $M_{12} = \{a_1, a_2\}$ if and only if she prefers a_3 (a_4) from $M_{34} = \{a_3, a_4\}$.¹⁵ That is: $a_1 \succ a_2$ iff $a_3 \succ a_4$ and $a_1 \succeq a_2$ iff $a_3 \succeq a_4$.

That is descriptively false. If h_1 says that in a fair lottery over $[1, 100]$ the integer drawn is 1, h_2 that it lies in $[2, 11]$ and h_3 that it lies in $[12, 100]$, most people have $a_1 \succ a_2$ and $a_4 \succ a_3$. This is the Allais paradox.¹⁶

¹³ Partitionality is not necessary for GT: (Geanakoplos [1989]) and (Dorst [unpublished]) identify weaker conditions that suffice. But since such conditions still rule out intuitive examples like the clock, it's no more obvious that observation obeys them than that it is partitional.

¹⁴ For discussion, see (Buchak [2010]).

¹⁵ $EU(a_3) - EU(a_1) = EU(a_4) - EU(a_2)$, so $EU(a_1) > EU(a_2)$ iff $EU(a_3) > EU(a_4)$ and $EU(a_1) \geq EU(a_2)$ iff $EU(a_3) \geq EU(a_4)$.

¹⁶ (Allais [1953]).

That descriptive mismatch *may* be normatively irrelevant. But allow—following some decision theorists—that these widespread ‘Allais preferences’ are rationally permissible. Now suppose you face M_{12} and M_{34} on separate occasions: the draws are independent and you may choose on either occasion without reference to the other. Before betting in each case, you can make an observation P that teaches you one of two things:

$$\begin{aligned} p_1 &= h_1 \cup h_2 \\ p_2 &= h_3 \end{aligned}$$

Think of P as gathering incomplete but accurate information about the draw: conducting P tells you exactly whether the chosen integer exceeds 11.

Should you observe P ? It seems that if *ex ante* you prefer a_1 to a_2 and a_4 to a_3 , and know you will choose rationally after observing,¹⁷ then in at least one situation you should not. Suppose first that you learn p_1 . Then you are indifferent between a_1 and a_3 : those gambles only differ in the event (h_3) that your observation rules out; similarly, you are indifferent between a_2 and a_4 . So on learning p_1 , you have $a_1 \succ a_2$ iff $a_3 \succ a_4$, and $a_1 \succeq a_2$ iff $a_3 \succeq a_4$, given that \succeq is transitive. So, in either M_{12} or M_{34} , you are at least open to choosing what you initially reject.¹⁸ Or, suppose you learn p_2 . Now you are indifferent between a_1 and a_2 , since they assign h_3 to the same payoff; and between a_3 and a_4 for the same reason. So in both cases, learning p_2 makes you indifferent between your options.

So either in M_{12} or in M_{34} it will seem *ex ante* counterproductive to conduct the observation P : whatever you learn will make you open to the option that you now reject, with no compensation. Your position *ex ante* therefore violates GT.

We can put it in terms of our taxation analogy. Suppose the government cares not only about the aggregate but also about the distributional effects of tax. Specifically, it marginally prefers tax schemes that make income more equal. Suppose three economic classes, h_1 , h_2 and h_3 , and four tax regimes a_1 , a_2 , a_3 and a_4 , with net incomes of each class as in Table 3. Then the government may prefer a_1 to a_2 but a_4 to a_3 , because the equality under a_4 more than compensates for its loss of income relative to a_3 ; whereas the equality under a_2 , being imperfect, does not compensate for *its* loss of income relative to a_1 .

Suppose now that the entire US population lives in two States: all those in classes h_1 or h_2 in Alabama, and all those in class h_3 in Wyoming. And suppose the government is choosing between two tax regimes, say a_1 and a_2 . Then Good’s recommendation is effectively to apply in each State the regime that is—by the government’s lights—optimal for that State. This may not yield a result that is at least as good—by those lights—as whichever of a_1 and a_2 is preferred overall. Optimizing on a State-by-State basis ignores the distributional effects across States that made a_4 overall preferable to a_3 .

Returning from the analogy: if the agent marginally disprefers variability in gambles, then the compound gamble, that realizes in each event in a partition the optimal gamble for that event, might look sub-optimal *ex ante* precisely because it ignores her aversion to cross-event variability.

This counterexample differs significantly from the last. The ‘clock’ objection grants (i) that one should learn, before choosing, which element of a partition is true, but denies (ii)

¹⁷ See analogous comments on an analogous assumption at n. 8 above.

¹⁸ Indifference between a_1 and a_3 means $a_1 \succeq a_3$ and $a_3 \succeq a_1$; similarly for a_2 and a_4 . So after you learn p_1 : if $a_1 \succ a_2$, then $a_3 \succeq a_1 \succ a_2 \succeq a_4$; so if $a_4 \succeq a_3$ then $a_2 \succeq a_1$ by the transitivity of \succeq , contradicting $a_1 \succ a_2$. Hence if $a_1 \succ a_2$ then $a_3 \succ a_4$. The converse is provable in the same way: hence after you have learnt p_1 you must have $a_1 \succ a_2$ iff $a_3 \succ a_4$. Similarly if $a_1 \succeq a_2$ then $a_3 \succeq a_1 \succeq a_2 \succeq a_4$; so $a_3 \succeq a_4$ by the transitivity of \succeq . Again the converse is provable in the same way: hence after you have learnt p_1 you must have $a_1 \succeq a_2$ iff $a_3 \succeq a_4$.

that looking is always a way to do that. The Allais-based objection attacks (i). The Allais subject should find it irrational to learn, before choosing from (say) M_{12} , which of $\{h_1 \cup h_2, h_3\}$ is true. Both objections oppose literally looking before leaping; but here that is a side-effect of denying that you should learn from a partition before leaping.

The Allais case highlights Good's reliance on a decision theory that prohibits Allais preferences. What does this in the classical theory is:

Independence Axiom (IA): Let a_1, a_2, a_3, a_4 be acts and let $\{h, h^*\}$ partition W . Suppose:

(i) $a_1(w) = a_3(w)$ and $a_2(w) = a_4(w)$ for each $w \in h$

(ii) $a_1(w) = a_2(w)$ and $a_3(w) = a_4(w)$ for each $w \in h^*$

Then $a_1 \succ a_2$ iff $a_3 \succ a_4$ ¹⁹

Intuitively: a rational person who prefers (say) lottery ticket L_1 over L_2 should prefer any bet with prize L_1 to any with prize L_2 but the same downside. It is easy to see why IA prohibits Allais preferences: putting $h = h_1 \cup h_2$ and $h^* = h_3$, IA implies both $a_1 \succ a_2 \leftrightarrow a_3 \succ a_4$ and $a_1 \succeq a_2 \leftrightarrow a_3 \succeq a_4$, with a_1 - a_4 as in Table 3.

Since the Allais preferences look reasonable, it's not immediately obvious that rationality demands Independence. It's also not obvious on reflection. *EU*-maximization implies that anyone who (i) has concave utility for money and (ii) always declines a 50/50 bet that wins \$110 or loses \$100 should (iii) decline a 50/50 bet that wins \$ x or loses \$1000, for arbitrarily large x .²⁰ This oddity arises from IA. Moreover, plausible alternatives exist: many versions of decision theory, by dispensing with IA, permit Allais preferences²¹, and some are appealing not (only) descriptively but as standards of rationality.²² So it isn't as though abandoning IA leaves no concept of practical rationality at all. And the possibility of abandoning it leaves GT looking shakier than it did.²³

Still, there are well-known pragmatic arguments for IA. They typically infer the rationality of IA from the premise that violating it leaves one open to exploitation.²⁴ Suppose that you start out with $a_1 \succ a_2, a_4 \succ a_3$, but if you learn p_1 then you have $a_3 \succ a_4$. A cunning bookie offers a choice between (a) choosing from a_3 and a_4 after learning which of $\{p_1, p_2\}$ is true and (b) paying \$1 to choose without first learning this. You know that on (a), whatever you learn, your choice will make you worse off (by your own lights *ex ante*) if it makes any difference. So you choose (b) and a_4 . So you are paying \$1 for a_4 when you could have had that gamble for free. Your violating IA is what makes you thus exploitable. Hence, it is argued, rationality demands IA. And we might conclude that it's both harmless and unsurprising that dropping it puts you beyond the scope of Good's Theorem.

¹⁹ Many, such as (Peterson [2009], p. 99), treat **IA** as a single axiom; in Savage's original presentation its content is spread across his definition D1 and postulate P2. (See the endpapers of Savage [1972]; cf. pp. 21-3.)

²⁰ (Rabin [2000]). For other criticisms see (Hansson [1988], pp. 149-51; McClennen [1990], pp. 77-80).

²¹ These include: Prospect Theory (Kahneman and Tversky [1979]), Anticipated Utility Theory (Quiggin [1982]), Generalized *EU* Analysis (Machina [1982]), Disappointment Theory (Loomes and Sugden [1986]), Choquet *EU* Theory (Schmeidler [1989]), Risk-Weighted *EU* Theory (Buchak [2013]), and Counterfactual Desirability Theory (Bradley and Stefánsson [forthcoming]).

²² See especially (McClennen [1990]) and (Buchak [2013]). A complication: the Allais case violates GT only if we're right that rational agents use 'sophisticated choice' when faced with sequential decision problems; McClennen denies, and Buchak entertains denying, this.

²³ These worries about IA depend on the rationality of a form of risk-avoidance not permitted by classical decision theory. Many decision theories designed for imprecise probabilities also invalidate IA and, for that reason, also generate counterexamples to GT. Other decision theories for imprecise probabilities validate IA and may allow a weak version of GT. See (Seidenfeld [2004]) and (Bradley and Steele [2016]) for discussion.

²⁴ For a summary of such arguments see (Machina [1989], pp. 1636-8) and (Buchak [2013], pp. 170-3).

Various responses exist. Some reject the premise. There are, they point out, no cunning bookies who know your preferences and can demand payment to withhold information. And if there were you'd be sure to avoid them. So this 'pragmatic' difficulty with Allais preferences is hardly real.²⁵

Others reject the validity of the argument. Some say that by choosing (b) and then a_4 , you are not really throwing away \$1, since the option of ' a_4 for free' was never genuinely available. You knew in advance that if you were to choose (a) then, whatever you learn about p_1 and p_2 , you would not take the gamble a_4 (unless it made no difference).²⁶ More radically, some say that practical rationality constrains your time-slices, not your temporally extended self.²⁷ Just as the Pareto inefficiency of the Prisoners' Dilemma equilibrium is consistent with both players' rationality, so the diachronic exploitability of the Allais subject is consistent with her time slices' rationality. There is nothing else to count as irrational.

But, setting aside exploitation, someone might argue for IA as follows. It is independently plausible that you should learn from a partition (whether or not you must pay not to). But violation of IA implies that you'd avoid doing that. So we can defend IA on the basis of something like GT itself.

But, having set exploitation aside, is it obvious, independently of an antecedent commitment to IA, that practical rationality demands uptake of free evidence? Clearly, learning will improve your anticipated return from bets on propositions which the evidence settles conclusively. But what about bets on a proposition p which the evidence doesn't settle conclusively? The evidence may move your credence in p closer to the truth (if, say, it is evidence for p , and p is true), but it may also move your credence further from the truth (if, say, it is evidence for p , but p is false). In the first case, it will increase your returns; in the latter, it will decrease them. So a little learning carries risks. Classical decision theory says that the potential benefits always outweigh this risk; but theories which reject IA because of its incompatibility with intuitive risk-avoidance may well reach a different verdict. So the attractiveness of free evidence is not a neutral starting-point from which to launch an argument for IA.²⁸

These responses to the arguments for IA are not decisive. We don't mean to endorse them. Nor, more generally, do we take a stand on whether IA is rationally compulsory. We simply point out that it is not obviously so—not, for instance, in the way that *transitivity* of preference has seemed to be. And if IA is not platitudinous, then neither is GT.

Our two counterexamples show that GT rests on controversial assumptions (i) about observation and (ii) about practical rationality. This is puzzling. We said at the outset that 'look for free before you leap' seems platitudinous. Should we conclude that this is just wrong?—That in fact our confidence in this 'platitude' should not exceed our confidence in the highly unobvious assumptions (i) and (ii)?

The straight answer must be 'yes'. Even when looking is free, looking before leaping is not a platitudinously good idea. It may not be a good idea at all. But we can give a less disappointing, if less straight, answer by showing how to approximate the 'platitude' whilst avoiding the questionable assumptions that underlie it.

²⁵ (Christensen [2004], p. 110). See also (Lewis [1999], p. 133).

²⁶ (Rabinowicz [1995]); (Buchak [2013], pp. 189-90).

²⁷ (Hedden [2013]).

²⁸ See (Buchak [2010], pp. 99-100). Thanks to a referee. Buchak also suggests that we can explain the oddness of avoiding free evidence by saying that it is irrational from an epistemic perspective; (Campbell-Moore and Salow [unpublished]) argue that we cannot say this if we accept rational risk-avoidance.

5 Conditionality

There is a platitude that avoids the counterexamples whilst capturing something of the intuitive content of GT. Section 5.1 states and proves the claim, which we call ‘Conditionality’. 5.2 shows how Conditionality avoids the difficulties facing GT, and that GT is a special case of Conditionality in the kind of environment that Good (mistakenly) took for granted.

5.1 The Principle

First, note that we can represent a question as a partition of W , the elements of the partition being its possible (complete) answers.²⁹ ‘Who came to the party?’ can be represented as the partition containing ‘No one’, ‘Only John’, ‘Only Mary’, ‘John and Mary and no one else’, etc. Moreover, for every partition, there is a corresponding question (‘Which element of this partition is actual?’); so we can move freely between the two ways of talking. In these terms, our platitude is:

Conditionality (C): A rational agent will make her action depend on the (true) answer to any question, whenever she can freely do so.

What does it mean to make your action depend on the answer to a question? Suppose you expect to choose from menu M , consisting, say, of (H) a bet on heads and (T) a bet on tails on the next toss of this (possibly biased) coin. Then to make your decision depend on Q —say, whether it rains tomorrow—is to choose instead from the menu M^Q , consisting of all the ways for the choice from M to depend on the answer to Q . In our example these are: (HH) Bet on heads if rain, bet on heads if no rain; (HT) Bet on heads if rain, bet on tails if no rain; (TH) Bet on tails if rain, bet on heads if no rain; (TT) Bet on tails if rain, bet on tails if no rain. So Conditionality says that you are *ex ante* better off choosing from this second menu.

Note that ‘making her action depend on the answer to a question’ means choosing from all the ways in which the action can depend on the answer. In particular, it includes vacuous dependence: realization of the same option given any answer. (HH) and (TT) illustrate such vacuous dependence. Allowing for this is what makes C platitudinous.

Note also that C does not require that the agent be in a position ever to learn the answer to the question, let alone before choosing. C only tells her to employ some (free) mechanism that makes the option realized dependent on the answer. This needn’t involve checking for herself what that answer is. For instance, suppose you can use a computer to place your bet; you can program the computer to make the bet dependent however you like on some condition to which the computer is sensitive but which you cannot observe and might never learn—say, the material composition of the coin. C advises you to use the computer (whereas GT is silent on this).

This illustrates that C, unlike GT, is not especially about observing or learning before acting. It deals with a more basic activity: making one’s act depend (optimally) on how the world is (in some respect). Intuitively, learning before deciding is a way to achieve such dependence—and as we’ll see, GT follows from C when it is. But as we’ll also see, learning is not a way to achieve such dependence in our counterexamples; that’s why GT can fail there.

Because C isn’t about learning, its restriction to questions (= partitions) isn’t justified by a theory of learning. Instead it rests on the observation that a conditional act is ill-defined

²⁹ (Hamblin [1958]).

unless specifiable relative to a partition—a set of subsets of W that are disjoint and jointly exhaustive. If the conditions are not disjoint, the conditional act could make impossible demands: the conditional act of buying this car if it is Japanese but not if it is a four-wheel drive cannot be realized if the car is a Toyota Land Cruiser. If they are not exhaustive, the conditional act may specify no prize, for instance if the car is a Citroen 2CV. C 's focus on partitions flows from the nature of acts, independently of issues about learning.

So much for what C says; on to its proof. Its mathematical basis is this principle, which is trivial when Y is finite:

Principle of Maximization (PM): If $X \subseteq Y$ and $V:Y \rightarrow \mathbb{R}$ is a function then $\max_{x \in X} V(X) \leq \max_{y \in Y} V(Y)$

We informally sketch the route from PM to C . If you can make the option realized from M dependent as you like on the answer to Q , then you can make it vacuously dependent: so everything in M is available if conditionalization is available. So by PM, the conditional option that maximizes whatever you care about (for instance, EU) does at least as well as any that maximizes it in M . If Alice can choose whichever of $\{HH, HT, TH, TT\}$ maximizes her expected utility, she will do at least as well (*ex ante* in expectation) as when choosing from $\{H, T\}$, because the former set includes the latter.

Now a formal argument. We first prove the formal version of C :

Formal Conditionality (FC): Suppose a finite menu M of acts such that each $a \in M$ takes each h in some partition H of W to an outcome $a(h) = z \in Z$. Let Q be an arbitrary question. Let $F = M^Q$ be the set of all functions from Q to M and for any $f \in F$ and $p \in Q$ let f_p be the act $f(p)$. Define M^* to be the following set:

$$\{a \in A \mid \exists f \in F \forall w \in W \forall p \in Q \forall h \in H: w \in p \cap h \rightarrow a(w) = f_p(h)\}.$$

Then for any function $V:A \rightarrow \mathbb{R}$, $\max_{a \in M} V(a) \leq \max_{a \in M^*} V(a)$.

Proof: we need to show that (i) M^* exists and is a set of acts; (ii) that for any V , $\max_{a \in M} V(a) \leq \max_{a \in M^*} V(a)$. (i) Since H and Q both partition W , so does $H \otimes Q =_{\text{def.}} \{h \cap p \mid h \in H, p \in Q\}$. Hence any $w \in W$ lies in exactly one element of $H \otimes Q$. So M^* is well-defined and its elements are all functions from W to Z i.e. acts. (ii) For any $a \in M$ consider the function $f^a: Q \rightarrow M$ defined as follows: $f^a(p) = a$. Then $f^a \in F$ and if $h \in H$ then $f_p^a(h) = a(h) = a(w)$. So $a \in M^*$. Therefore $M \subseteq M^*$. Since M is finite, so is M^* . So PM applies to yield the result.

To get from FC to C we need this decision-theoretic assumption:

Representability (R): \succ is representable, i.e. \succ is asymmetric and \succeq is transitive.

If R holds and A is denumerable (which we assume), then some value function V from acts to real numbers satisfies the following condition: for any $a, a': a \succ a' \leftrightarrow V(a) > V(a')$.³⁰ By FC, we know that if the agent faces M then for any question Q , maximizing V on M^Q —that is, making the option realized from M optimally dependent on the answer to Q —attains at least

³⁰ For proof see (Kreps [1988], ch. 3).

as high a V -score as maximizing V on M . By R, the agent is doing at least as well in the former case as in the latter, by her own lights *ex ante*.

The only substantial assumption here was Representability. This fragment of classical decision theory is far less demanding than that necessary for GT. We have already seen that the latter involves IA, which many decision theories reject. But most of these non-classical decision theories still endorse R.³¹

Conditionality is therefore platitudinous and relatively weak. Platitudinous because the mathematical idea behind it is—withstanding the tedium of FC itself—the utterly simple idea that a function attains a (weakly) greater maximum over a finite set than over any of its subsets. Weak because the decision-theoretic idea behind it—Representability—is widely accepted and intuitively quite plausible. It is not universally accepted. So Conditionality itself is not a complete triviality. Someone who rejects R might perhaps reject C.³² But C should appeal more widely than any principle that prohibits the Allais preferences.

5.2 Revisiting the counterexamples

Neither of our two examples conflicts with C. It's tempting to think they must because they conflict with GT, and GT seems to follow from C. Roughly, this is because it seems that observing and then deciding is always a way to make one's decision depend optimally on the answer to a question.

More carefully, suppose that a rational agent is choosing from M , and has the opportunity to conduct an observation. It's tempting to accept:

Questions: Any observation can be thought of as providing the complete answer to some (antecedently specifiable) question.

So let Q be the question which the observation would answer. By C, our agent prefers choosing from M^Q instead of M . It is also perhaps natural to think

Equivalence (E): For a rational agent, choosing from M after learning the answer to Q is equivalent to (leads to the same outcome as³³) choosing from M^Q .

Putting these together, it follows that our agent prefers choosing from M after learning the answer to Q over choosing from M . Since choosing from M after learning the answer to Q is just what it is to look first, our agent prefers looking before leaping.

So GT follows from C given initially plausible principles. Nonetheless, neither of our examples conflicts with C. For, as we're about to show, each of the examples conflicts with a different premise in this derivation.

³¹ This includes Anticipated Utility Theory, Choquet Expected Utility, Risk-Weighted Expected Utility, and Counterfactual Desirability Theory, but not Disappointment Theory (see n. 21).

³² For alleged counterexamples to the transitivity of \succ and hence (given asymmetry of \succ) to that of \succeq , see (May [1954], pp. 6-7; Gehrlein and Fishburn [1976], p. 1). Decision theories employing imprecise probabilities may reject the transitivity of \succeq while accepting that of \succ . For they may allow incomparable options; and if a is incomparable to both b and c , yet b is preferable to c , we have $c \succeq a$ and $a \succeq b$ without $c \succeq b$. Many working on imprecise decision theories nonetheless maintain that if it's unacceptable to choose a from M , it must also be unacceptable to choose a from any $M' \supseteq M$ – see e.g. Axiom 1a in (Seidenfeld et al. [2010], p. 164). C follows straightforwardly from this – and hence appeals even to some who reject R.

³³ Or at least the same up to indifference given the answer.

5.2.1 The Unmarked Clock

The unmarked clock resists the slide from C to GT by being a counterexample to Questions: there is no question such that looking at the clock guarantees that you learn exactly the complete answer to it. Some things you might learn are consistent with others. So the things you might learn don't partition W. So they aren't all complete answers to a single question.

Note that looking is, even here, a way to make your decision depend on the answer to some question or other. Let q_i be the proposition that you learn $[i - 1] \cup [i] \cup [i + 1]$. Then $Q = \{q_i | 0 \leq i \leq 59\}$ is a question, and which post-glance option you choose depends on its answer: you bet ODD if the answer is q_i for even i , and EVEN if the answer is q_i for odd i . But this dependence is not optimal *ex ante*: it differs from the pattern 'bet ODD if the answer is q_i for odd i , bet EVEN if the answer is q_i for even i ', which is obviously what you'd choose if allowed to make your bet depend optimally on the answer to Q . But this non-optimality is unsurprising: why should looking at the clock make your bet optimally dependent on the answer to Q , when it doesn't teach you that answer?

The unmarked clock would refute C only if looking at the clock, then doing what you think best, were equivalent to choosing an *ex ante* optimal pattern of dependence of your choice on the answer to some Q . But it can't be. For as we argued, looking at the clock, and then doing what you think best, means betting ODD (EVEN) if D is even (odd), hence a sure loss. Yet one initial option was not to bet, which is strictly better. So for every Q , the (vacuous) pattern of dependence 'take no bet whatever the answer to Q ' is preferable to looking and then doing what you think best. So the latter compound action can't generate an *ex ante* optimal pattern of dependence on Q . C is safe from the unmarked clock.

5.2.2 Allais

C is also safe from the Allais case. Suppose a subject prefers a_1 to a_2 and a_4 to a_3 , and she actually faces the choice (i) between a_1 and a_2 . Intuitively, C says that she is no worse off by her own lights if she now makes the final selection dependent on which of $h_1 \cup h_2$ and h_3 is true in whatever way now seems optimal to her. Inspection of the payoffs shows that she will choose the (vacuously) dependent act from $\{a_1, a_2\}^{\{h_1 \cup h_2, h_3\}}$ that realizes a_1 in either circumstance. Similarly, suppose she faces the choice (ii) between a_3 and a_4 . If choosing from the conditional acts, she will now choose the dependent act that realizes a_4 in either case. The important point is that in (i) and (ii), these dependent acts make her no worse off by her own lights *ex ante*. So C is consistent with the Allais preferences. By contrast and as we saw, GT is not. Given the Allais preferences $a_1 \succ a_2$ and $a_4 \succ a_3$, it must hold either in (i) or in (ii) that choosing after learning which of $h_1 \cup h_2$ and h_3 is true must make her worse off by her own lights *ex ante*.³⁴

Returning one more time to the taxation analogy: suppose again that every US citizen resides in Alabama or Wyoming (not both), with those in h_1 or h_2 resident in Alabama and those in h_3 resident in Wyoming, and that the tax regimes a_1 - a_4 affect these classes as in Table 3. Then C says only that a government choosing from a_1 and a_2 is rational to make the regime faced by a citizen somehow dependent on the citizen's State of residence. This is consistent with the government's having distributional aims that make it prefer a_4 to a_3 . It is consistent with C that the government prefers a_1 -in-Alabama-and- a_1 -in-Wyoming to a_2 but

³⁴ But C does not itself counsel the agent to (pay to) avoid free evidence in Allais-type cases. Allais-type agents will indeed do that, and this may be objectionable (see the discussion towards the end of section 4). But this is a consequence of their preferences, not a consequence of C itself, which is entirely consistent with IA. C, like transitivity, no more mandates the Allais preferences than it prohibits them.

also a_4 -in-Alabama-and- a_4 -in-Wyoming to a_3 , because each preferred system represents some way of making the tax regime conditional on the State of residence, which is all that C recommends.

More formally: C implies that anyone facing M_{12} (M_{34}) is (by her own lights *ex ante*) no worse off when choosing from the menu of dependent acts M_{12}^Q (M_{34}^Q), where Q is the question whether h_3 (as opposed to $h_1 \cup h_2$) is true. That is trivial, since choosing from M_{12} (M_{34}) is the same as choosing from M_{12}^Q (M_{34}^Q). For the options in M_{12} are a_1 and a_2 while the options in M_{12}^Q are (in an obvious notation):

- (i) $h_1 \cup h_2 \rightarrow a_1, h_3 \rightarrow a_1$
- (ii) $h_1 \cup h_2 \rightarrow a_1, h_3 \rightarrow a_2$
- (iii) $h_1 \cup h_2 \rightarrow a_2, h_3 \rightarrow a_1$
- (iv) $h_1 \cup h_2 \rightarrow a_2, h_3 \rightarrow a_2$

Since options are functions from possible worlds to prizes, and a_1 and a_2 have identical prizes throughout h_3 , we may identify (i) with (ii) and (iii) with (iv). Clearly (i) is equivalent to a_1 and (iv) to a_2 . So M_{12}^Q collapses into M_{12} . Similarly M_{34}^Q collapses into M_{34} .

So C holds in the Allais example. Since the observation does exactly and completely answer the question whether h_3 is true, the relevant instance of Questions also holds. So the Allais case must be a counterexample to E, the principle that for an agent who acts rationally, choosing from M after learning the answer to Q yields the same result as choosing from M^Q . And this is exactly what we found. When the agent learns the answer to Q before choosing, she is liable to pick options from $M = M^Q$ that she initially rejects.³⁵

So neither counterexample to GT threatens C, supporting our claim that C is platitudinous. We've also seen how to derive GT from C using the initially plausible, but on reflection non-obvious, principles Questions and Equivalence, thus revealing the sense in which Good's Theorem is the special case of Conditionality that we get whenever those controversial principles do in fact hold.

6 Conclusion

Good's paper began with Ayer's question: 'why, in the theory of logical probability (credibility) we should bother to make new observations'.³⁶ GT answers that observation pays when its cost is negligible. We've shown: first, that it need not pay unless (i) observation is partitional and (ii) rational choice respects Independence—neither of which is obvious. But second, that Conditionality holds irrespective of these assumptions, requiring only that preference be representable. 'Look before you leap' blends assumptions about perceptual epistemology and decision theory, making it only questionably true and clearly not platitudinous. 'Make your actions depend on the world when you can' incurs a lighter, purely decision-theoretic commitment that gives it, in our view, a better claim to being both.

³⁵ It's no coincidence that we needed counterexamples to IA to get a counterexample to E. E follows from IA given R and a weak assumption about diachronic preference change (roughly: if there are acts a which the agent weakly prefers to all acts in M conditional on p – i.e. the agent prefers the act ' $p \rightarrow a, \sim p \rightarrow x$ ' to the act ' $p \rightarrow b, \sim p \rightarrow x$ ' for any $x, b \in M$ – the agent will weakly prefer one of these acts to every act in M unconditionally, after learning p). For space reasons, we omit the proof.

³⁶ (Good [1967], p. 319).

Acknowledgements

AA wishes to thank Julien Dutant, Christian List, and an audience at the London School of Economics to which he delivered an early version of this paper. BS wishes to thank Nilanjan Das, Kevin Dorst, and Ian Wells. We both wish to thank two referees for this *Journal* for insightful comments that led to many improvements.

References

- Allais, M. [1953]: ‘Le Comportement de l’Homme Rationnel devant le Risque: Critique des Postulats et Axiomes de l’École Américaine’, *Econometrica*, **21**, pp. 503-546.
- Bradley, R. and Stefánsson, H. O. [forthcoming]: ‘Counterfactual Desirability’, *British Journal for the Philosophy of Science*.
- Bradley, S. and Steele, K. [2016]: ‘Can Free Evidence be bad? Value of Information for the Imprecise Probabilist’, *Philosophy of Science*, **83**, pp. 1–28.
- Bronfman, A. [2014]: ‘Conditionalization and Knowing that one Knows’, *Erkenntnis*, **79**, pp. 871-92.
- Buchak, L. [2010]: ‘Instrumental Rationality, Epistemic Rationality, and Evidence-gathering’, *Philosophical Perspectives*, **24**, pp. 85-120.
- Buchak, L. [2013]: *Risk and Rationality*. Oxford: OUP.
- Campbell-Moore, C. and Salow, B. [unpublished]: ‘Avoiding Risk and Avoiding Evidence’
- Cohen, S. and Comesaña, J. [2013]: ‘Williamson on Gettier Cases and Epistemic Logic’, *Inquiry*, **56**, pp. 15–29.
- Christensen, D. [2010]: ‘Rational Reflection’ *Philosophical Perspectives*, **24**, pp. 121-140.
- Das, N. [unpublished]: ‘Externalism and the Value of Information’
- Dorst, K. [unpublished]: ‘Evidence: A Guide for the Uncertain’
- Elga, A. [2013]: ‘The Puzzle of the unmarked clock and the new rational reflection principle’, *Philosophical Studies*, **164**, pp. 127-139.
- Geanakoplos, J. [1989]: ‘Game Theory without Partitions, and Applications to Speculation and Consensus’, Cowles Foundation Discussion Paper no. 914.
- Gehrlein, W. V. and Fishburn, P. C. [1976]: ‘Condorcet’s Paradox and Anonymous Preference Profiles’, *Public Choice*, **26**, pp. 1-18.
- Good, I. J. [1967]: ‘On the Principle of Total Evidence’, *British Journal for the Philosophy of Science*, **17**, pp. 319-21.
- Good, I. J. [1974]: ‘A Little Learning can be Dangerous’, *British Journal for the Philosophy of Science*, **25**, pp. 340-42.
- Graves, P. [1989]: ‘The Total Evidence Principle for Probability Kinematics’, *Philosophy of Science*, **56**, pp. 317–324.
- Hamblin, C. L. [1958]: ‘Questions’, *Australasian Journal of Philosophy*, **36**, pp. 159-168.
- Hansson, B. [1988]: ‘Risk Aversion as a Problem of Conjoint Measurement’, in P. Gärdenfors and N.-E. Sahlin (eds), *Decision, Probability, and Utility: Selected Readings*, Cambridge: CUP, pp. 136-158.
- Hawthorne, J. and Magidor, O. [2010]: ‘Assertion and Epistemic Opacity’, *Mind*, **119**, pp. 1087–1105.
- Hedden, B. [2013]: ‘Options and Diachronic Tragedy’, *Philosophy and Phenomenological Research*, **90**, pp. 423-451.

- Horowitz, S. [2014]: ‘Epistemic Akrasia’, *Noûs*, **48**, pp. 718-744.
- Jeffrey, R. C. [1983]: *The Logic of Decision*, 2nd ed, Chicago: Chicago UP.
- Kahneman, D. and Tversky, A. [1979]: ‘Prospect Theory: An Analysis of Decision under Risk’, *Econometrica*, **47**, pp. 263-91.
- Kreps, D. M. [1988]: *Notes on the Theory of Choice*, Boulder: Westview.
- Lewis, D. [1999]: ‘Why Conditionalize?’, in his *Papers in Metaphysics and Epistemology*. Cambridge: CUP, pp. 403-7. Reprinted in A. Eagle (ed.) [2011]: *Philosophy of Probability: Contemporary Readings*, London: Routledge, pp. 132-4.
- Loomes, G. and Sugden, R. [1987]: ‘Disappointment and Dynamic Consistency in Choice under Uncertainty’, *Review of Economic Studies*, **53**, pp. 271-82.
- McClennen, E. F. [1990]: *Rationality and Dynamic Choice*, Cambridge: CUP.
- Machina, M. [1982]: ‘“Expected Utility” analysis without the Independence Axiom’, *Econometrica*, **50**, pp. 277-323.
- Machina, M. [1989]: ‘Dynamic Consistency and Non-Expected Utility Models of choice’, *Journal of Economic Literature*, **27**, pp. 1622-68.
- May, K. O. [1954]: ‘Intransitivity, Utility and the Aggregation of Preference Patterns’, *Econometrica*, **22**, pp. 1-13.
- Peterson, M. J. [2009]: *An Introduction to Decision Theory*, Cambridge: CUP.
- Quiggin, J. [1982]: ‘A Theory of Anticipated Utility’, *Journal of Economic Behavior and Organization*, **3**, pp. 323–343.
- Rabin, M. [2000]: ‘Risk-aversion and Expected Utility Theory: A Calibration Theorem’, *Econometrica*, **68**, pp. 1281-92.
- Rabinowicz, W. [1995]: ‘To Have One’s Cake and Eat it too: Sequential Choice and Expected Utility Violations’, *Journal of Philosophy*, **92**, 586-620.
- Savage, L. J. [1972]: *Foundations of Statistics*, 2nd ed., New York: Dover.
- Schmeidler, D. [1989]: ‘Subjective Probability and Expected Utility without Additivity’, *Econometrica*, **57**, pp. 571–587.
- Schoenfield, M. [forthcoming]: ‘Conditionalization does not Maximize Expected Accuracy’, *Mind*.
- Seidenfeld, T. [2004]: ‘A Contrast between two Decision Rules for use with (Convex) Sets of Probabilities: Gamma-maximin versus E-admissibility’, *Synthese*, **140**, pp. 69–88.
- Seidenfeld, T., Schervish, M., and Kadane, J. [2010]: ‘Coherent Choice Functions under Uncertainty’, *Synthese*, **172**, pp. 157-176.
- Stalnaker, R. C. [2009] ‘On Hawthorne and Magidor on Assertion, Context, and Epistemic Accessibility’ *Mind*, **470**, pp. 399-409.
- Williamson, T. [1992]: ‘Inexact Knowledge’, *Mind*, **402**, pp. 217-242.
- Williamson, T. [2000]: *Knowledge and its Limits*, Oxford: OUP.
- Williamson, T. [2011]: ‘Improbable Knowing’, in T. Dougherty (ed.), *Evidentialism and its Discontents*, Oxford: OUP, pp. 147-64.
- Williamson, T. [2013]: ‘Response to Cohen, Comesaña, Goodman, Nagel, and Weatherson on Gettier Cases in Epistemic Logic’, *Inquiry*, **56**, pp. 77–96.