Economic Explanation

Antony Eagle

Choices, Models and Morals » Lecture 2

Contents

Scientific Explanation and the Social Sciences

Explanation and the Aims of Science

- A scientific **prediction** involves deducing *what*, *when*, and *where*. We apply our scientific theories, given some background information about the circumstances in which we find ourselves, and try to figure out what will happen.
- The *why* (and *how*) don't feature. We can have an excellent predictive theory without the slightest clue about why the events we are predicting happen.
  - Think of various bits of herbal lore; the use of various plants and extracts to treat medical conditions predates recorded history, and there is a wealth of knowledge of what works accrued through trial and error – e.g., willow bark extracts were used to treat fevers in 400 BCE, but it wasn't until Vane (1971) that the mechanism of action was uncovered.
- Indeed, early 20th century 'logical positivist' philosophy of science was very suspicious of explanation. Asking *why* meets a human need to make nature comprehensible to us; the logical positivists queried whether that is part of science – and even whether asking *why* is too close to asking *for what purpose*, giving an opening to **teleological** (Hempel and Oppenheim 1948: 142–46) or even theological notions that have no place in science.

Theories of Explanation

- Philosophical suspicions notwithstanding, philosophers soon enough turned out accounts of how scientific information is used in giving explanations (Salmon 1989).
- In all cases, this is explanation of a **particular event**, a 'token phenomenonon' (Reiss 2013: 19).

Deductive-Nomological An explanation consists in the **derivation** of the event to be explained ('explanandum') from the **laws of nature** and **background conditions** (Hempel and Oppenheim 1948; Reiss 2013: 23).

Statistical-Relevance An explanation consists in the provision of statistical information showing the event in question to be **probabilistically dependent** on some other event (Salmon 1971).

Causal 'to explain an event is to provide some information about its causal history' (Lewis 1986a: 217).

The D-N Theory in Social Science

- While there are many economic laws, as Reiss (2013: 22) notes, it is nevertheless difficult to fit economic explanation into the D-N framework.
- For one thing, many of the 'laws' of economics are **regularities** that themselves seem somewhat accidental: they are correlations observed in the data, such as Okun's law (Reiss 2013: 24–25), or the 'Pizza Principle' that the price of subway fare in NYC is the same as the price for a slice of pizza.
- If those correlations continue to hold, they suffice for a derivation – they might even have predictive accuracy. But laws simply 'read off' the data need an explanation themselves: *why are these factors correlated?*
- Even theoretically well-supported regularities in social sciences may not be **exceptionless**, and this can block the existence of a successful derivation.
- Often enough, however, there are widespread robust regularities in human behaviour – these 'almost-laws' may be explanatory – perhaps they make the explanandum **highly probable** given the background conditions, which may be enough for explanation (Reiss 2013: 24).

Hume on Psychological Regularities

> It is universally acknowledged, that there is a great uniformity among the actions of men, in all nations and ages, and that human nature remains still the same, in its principles and operations. The same motives always produce the same actions: The same events follow from the same causes. Ambition, avarice, self-love, vanity, friendship, generosity, public spirit; these passions, mixed in various degrees, and distributed through society, have been, from the beginning of the world, and still are, the source of all the actions and enterprizes, which have ever been observed among mankind. Would you know the sentiments, inclinations, and course of life of the Greeks and Romans? Study well the temper and actions of the French and English: You cannot be much mistaken in transferring to the former most of the observations, which you have made with regard to the latter. Mankind are so much the same, in all times and places, that history informs us of nothing new or strange in this particular. Its chief use is only to discover the constant and universal principles of human nature, by shewing men in all varieties of circumstances and situations, and furnishing us with materials, from which we may form our observations, and become acquainted with the regular springs of human action and behaviour. (Hume 1777: §8.7)

General Problems with D-N

- So perhaps, as Hempel argued (1948: 140–41), there is scope for using regularities of human behaveior to support prediction, and hence (if the D-N model is correct) explanation.
- In general, however, the D-N theory is not a good model of scientific explanation (van Fraassen 1980, pp 104–106).
  1. Many laws allow a two-way association between events, enabling the derivation of either from the other; but explanation is **asymmetric**. A flagpole casts a shadow 100m long; the shadow is explained by the height of the flagpole and the angle of the sun, given the laws of optics; but the legnth of the shadow doesn't explain the angle of the sun.
  2. Derivability doesn't ensure **relevance**: suppose it's a law that no one who takes the pill gets pregnant; then we can derive 'John didn't get pregnant' from 'John scrupulously takes the pill'. But this is not explanatory.

Explaining Behaviour

Economic Phenomena

- What should replace the D-N theory? How do we explain economic phenomena?
- First: what is to be explained? Economics covers a great many subjects, from the momentary choices of a single person to society-scale facts about employment and productivity.
- One natural idea however is that many economic facts are ultimately grounded in the decisions of **individuals** and how those decisions aggregate together.
  - 'Individual' doesn't have to mean 'person': it could be an individual business or institution: it can be any entity that is amenable to economic **interpretation**.
  - Nor do we have to accept that *all* economic facts bottom out in individualist explanations (so-called 'microfoundations') – perhaps there is some **autonomy** to high-level economic regularities (Reiss 2013: 111).
- That gives us a tighter focus: what is an economic explanation of individual decision?

Actions and Reasons

- Human beings are physical, and our behaviour is bodily. Yet predicting and explaining our behaviour as a branch of pure physics is a task of unimaginable complexity – nor would it look remotely like economics (because the economic individual is not a distinctive class of entity, physically speaking).
- The economic domain begins with **choice behaviour**, intentional decisions to perform a certain action.
  - Not every bodily movement is behaviour (being knocked off your bike by a car). Not every behaviour is economically relevant (twitches, reflexes). Not every consequence of an intentional action is intended (alerting the prowler to your presence, Davidson (1963), p. 686).
- To explain decisions, once we've narrowed them down from behaviour more generally, is to **rationalize** them: to show that the action is the product of the agent's **reasons**:

> Whenever someone does something for a reason, therefore, he can be characterized as (a) having some sort of pro attitude toward actions of a certain kind, and (b) believing (or knowing, perceiving, noticing, remembering) that his action is of that kind. (Davidson 1963: 685)

Folk Psychology: Actions, Reasons, and Causes

- A rationalizing explanation for an intentional action cites a reason. This is explanatory only if we assume a certain psychological theory, namely, that reasons **cause** actions.

> a person can have a reason for an action, and perform the action, and yet this reason not be the reason why he did it. Central to the relation between a reason and an action it explains is the idea that the agent performed the action *because* he had the reason. (Davidson 1963: 691)

  - E.g., a priest may rationalize his prurient interest in the sex lives of his congregants by a genuine concern for their immortal souls and a belief in the efficacy of confession; nevertheless, the 'real' reason may be to satisfy by proxy his own unresolved sexual desires.
- **Folk psychology** is 'the everyday theory of human rationality' (Hausman, McPherson, and Satz 2017: 55), which explains actions by (i) citing reason-constituting internal states of the individual – **beliefs** and **desires** – and (ii) noting that one such reason caused the action to be explained (Reiss 2013: 30–31). This is thus a species of causal explanation.

Rationalization and Interpretation

- A prior sceptical worry arises: are there any actions to be explained?
  - If Freud or the evolutionary psychologists are right, then lots of our behaviour isn't intentional, and hence isn't the product of reasons.
  - Indeed, people are physical, so there will be a purely physical explanation of all bodily movements, including putative actions.
- With human beings there is a natural presumption that they are like us; but what about other sorts of entities we may want to include in economic explanations?
- **Interpretativism**: if an entity can be globally rationalized by an accepted theory, they are an **agent**:

> We suppose that people tend to behave in a way that serves their desires according to their beliefs. We should take this principle of instrumental rationality to be neither descriptive nor normative but *constitutive* of belief. It enters into the implicit definition of what it is for someone to have a certain system of belief [and desire]. (Lewis 1986b: 36)

- One fundamental constraint: we should ascribe beliefs and desires so as to make a behaviour a rational action by the standards of folk psychology (Lewis 1974: 337–38).

Utility and Choice in Conditions of Certainty

Folk Psychology Elaborated

- Talk of belief and desire is all very well, but it is explantorily limited.
- Suppose you have the true belief that your car contains your daily recommended iron intake, and you have the admirable desire to have sufficient iron. That is a reason for your to eat your car (Schroeder 2005: 8).
- Nevertheless, making the choice to eat your car would not be rational, even though it would be rationalisable.
- Why not? Because it would not be the **best** option for you.
- A rational choice is not merely based on reasons, but on the reasons that the agent thinks best. So, despite the fact that you have a reason to eat your car, you don't.
- Other things being equal, your choices reveal not only your reasons, but your **preferences**: you choose to promote the outcome that you most prefer from among the available options.
- Someone who suffers from **weakness of will** might be caused to act by a reason, but not one that they wish to be motivated by. If you eat a whole tub of icecream, and then reproach yourself thereafter, you might have acted on a reason but prefer that you had acted on another.

Revealed and Intrinsic Preference

- We understand preferences as **psychologically real**, as 'a stable mental ranking of alternatives' (Reiss 2013: 35).
- This is slightly unorthodox; it is common instead to take preference to be **revealed** in choice.
  - This is partly the product of positivist scepticism about theoretical entities: a "properly scientific" economics should eschew reference to mental entities in favour of empirically observable phenomena.
- But a simple reduction of preference to choice cannot succeed.
  - We have many preferences that never get a chance to manifest in choice (I certainly prefer $10 million dollars for free to $5 million dollars for free, but doubt I'll get the chance to show it.)
  - We must admit the reality of weakness of will, i.e., an action that isn't the agent's own most preferred choice – impossible if choice **is** preference.
  - To reveal preference sometimes involves forced choice; but why think that reveals, as opposed to **creates**, preferences?
- There is something to revealed preference. Interpretativism to some extent assigns preferences on the basis of the interpretation of behaviour. But the interpretivist but allows belief and desire to be assigned **holistically** on the basis of the best overall explanation of behaviour.

States, Acts, Outcomes

- If we have an set of possibilities Ω, a **partition** of Ω is a set of possibly true propositions $\mathbf{X} = \{X_1, \ldots\}$ such that (i) at least one member of $\mathbf{X}$ is true, and (ii) only one member of $\mathbf{X}$ is true.
- In a decision problem, an agent faces **indecision** about which act to perform, and may also face **uncertainty** about the consequences of that act. For each problem, then there will be two kinds of propositions of particular importance:
  1. An **act** is a possible result of the agent's decision. Any set of possible acts in some decision problem forms a partition: in each possibility, one and only one act is performed.
  2. A **state of the world** is a condition of the objective world, not under the agent's direct control, which influences the consequences of the agent's decision. Any set of states will be a partition too: each possibility is in just one one actual world state.
- An **outcome** is **a conjunction of an act and a state**. The outcomes of a decision problem will form a partition of the set of possibilities.
- This framework assumes a certain representation of what is possible, and we trade off completeness for feasibility here (Reiss 2013: 36–37).

Preferences over Outcomes

- Suppose we idealise, and let the agent be **certain** of the state of the world.
  - Then an agent's preferences over outcomes are wholly determined by their preferences over their own acts: e.g., they know that the act *choosing icecream* will yield the outcome *choose icecream and have icecream*, etc.
- This give rise to a **preference relation**. It is convenient to express this in terms of **weak preference**: A weakly prefers x to y, symbolised $x \succeq y$, if (roughly) they regard outcome x as no worse than, and possibly better than, outcome y. (Strict preference: $x \succ y$ iff $x \succeq y$ and $y \not\succeq x$.)
- What preferences over outcomes are rational? **Any that satisfy the following** (Reiss 2013: 37; Peterson 2017: 99):

  Completeness For any actions x and y either $x \succeq y$ or $y \succeq x$ (or both): no pair of actions are unranked.

  Negative Transitivity Rational agents who do not strictly prefer outcome x to outcome y, and do not strictly prefer outcome y to outcome z, must also fail to strictly prefer outcome x to outcome z: if $x \not\succ y$ and $y \not\succ z$ then $x \not\succ z$. (If $\succ$ is negative transitive, $\succeq$ is transitive: $z \succeq y$ and $y \succeq x$ entail $z \succeq x$.)

Challenges to Necessity: Negative Transitivity

- Are these conditions **necessary** – are they needed for rational preference? **Sufficient** – do we need some further principles too?
- (Negative) Transitivity is quite plausible, as intransitive preferences can be **money-pumped** (Reiss 2013: 38–39).
  - An agent with negatively intransitive but complete preferences who has begun by opting for x would be happy to opt for the no less preferred y, and having opted for y would be happy to opt for the no less preferred z.
  - But since they have intransitive preferences, they strictly prefer x to z. They regard x together with a **mild penalty** as no less preferred than z. But then the agent has swapped x – by a sequence of choices between outcomes they are **indifferent between** – for x-with-a-mild-penalty. And it is **irrational** to be indifferent between an outcome and a **soured** version of that same outcome!
- One problem with the money pump argument: what if the agent **changes their mind** between the swaps? Then it may not be irrational.
- That is a red herring, since the money pump simply dramatises the fact that someone with negatively intransitive preferences at a single time would regard an outcome and a souring of it as equally valuable, which is incoherent.

Challenges to Necessity: The Dinner

› Do we really have complete preferences between arbitrary outcomes?

> It is dinner-time. Should we go to the Indian restaurant or the Chinese restaurant? We have visited both many times. We know their pluses and minuses. The Indian restaurant is less far to walk. It serves up a sublime mango lassi. The Chinese restaurant is cheaper. Its raucous atmosphere is more child-friendly. All in all it is a wash for me. I have no all-things considered preference between:
>
> (A) Our going to the Indian restaurant.
>
> and
>
> (B) Our going to the Chinese restaurant.
>
> And learning that it is dollar-off day at either restaurant will not give me an all-things-considered preference. When I compare B to:
>
> (C) Our going to the Indian restaurant and saving $1. …
>
> it remains a wash for me. I have no all-things-considered preference between C and B … though I do prefer C to A…. (Hare 2010: 238; cf. Peterson 2017: 184–86)

Completeness and the Dinner

- In the case of the dinner, the agent must have incomplete preferences. If their lack of preference means they are **indifferent** between A and B, then they prefer C to B; but they don't.
- So they neither prefer A to B, nor prefer B to A, nor do they rank them equally. **They simply have no ranking at all** (Reiss 2013: 40).
- We might opt for something weaker:

  (Completability) For any actions x and y, either $x \succeq y$ or $y \succeq x$ (or both) can be coherently added to one's current preferences: no pair of actions are unrankable.
- That leaves it open which preference should be added – maybe both are compatible with what your current preferences are.
- So mere completability does not give rise to a **preference ranking** – completeness and transitivity do (Hausman, McPherson, and Satz 2017: 58).

Ordinal Preference and Utility

- In any case, **if** an agent has transitive and complete preference among alternative outcomes, **then** they **can** have the outcomes **ranked numerically** (Reiss 2013: 38).

  Theorem 1 Given a finite set of alternative outcomes, A's preferences are negatively transitive and complete if and only if there is a function U from outcomes to real numbers such that: $U(x) \geq U(y)$ iff $x \succeq y$ (Peterson 2017: 100–101)

- Such a function U is called an **ordinal utility** function, and it represents a preference ordering on outcomes: higher numbers are more preferred.
- This is not of interest in itself (the conditions on preference are very strong) – rather (I think) it shows how my **values** are reflected in my preferences, if my values have a numerical representation.
- Some terminology: my personal or subjective **utility** in some outcome is the value I personally assign to the state of affairs in which that outcome obtains. It is the numerical representation of **degree of desire** – that is why, under certainty, it licenses preference.

Utility and Rational Choice

› For example, these (partial) rankings of possible actions for A:

Table 1: Some utility functions for A's preferences

Act\Outcome	Utility-1	Utility-2	Utility-3
donate to family planning charity	17	2.7	900
keep money for self	8	2.6	20
donate to 'just say no' drug education	-7	2	1

› The **numbers don't matter**: if there is a utility function representing A's preferences, then there are many, because any other function U' where $U'(x) > U'(y)$ iff $U(x) > U(y)$ would also give the same **order**.

› Our agent acts rationally iff they **can be represented as** choosing the outcome which has the highest utility.

» That doesn't tell us the **priority** of preference vs. value: it doesn't tell us whether (a) they choose it **because** they regard it as of high utility, or (b) its the high utility reflects its **choiceworthiness for them**.

Uncertainty and Risk in Decision

Ignorance and Action

- If A is ignorant of the actual world state, then A can't choose to act based **wholly** on utility.
- A given act will have different outcomes as consequences, **depending on the unknown state of the world**, and those outcomes might have quite different values to the agent. Consider:

> Jill is a physician who has to decide on the correct treatment for her patient, John, who has a minor but not trivial skin complaint. She has three drugs to choose from: drug A, drug B, and drug C. Careful consideration of the literature has led her to the following opinions. Drug A is very likely to relieve the condition but will not completely cure it. One of drugs B and C will completely cure the skin condition; the other though will kill the patient, and there is no way that she can tell which of the two is the perfect cure and which the killer drug. What should Jill do? (Jackson 1991: 462–63)

- The intuitive answer is: *prescribe drug A*. How do we obtain it?

Risk or Uncertainty

- The economic literature has distinguished two kinds of ignorance (Reiss 2013: 42–43):

  Risk A decision is made under risk when the agent does not know the actual world state, but does know the **objective chances** of the possible world states that could be actual.

  Uncertainty A decision is made under uncertainty when the agent does not know the actual world state, and need not know of any chances either.

- These can be given a **unified** treatment in the Bayesian framework:

  Bayesianism Every agent, in every decision situation, has a **subjective degree of belief** (mathematically, a probability) in each possible world state which reflects what their **total evidence** indicates about which is actual (Eagle 2016: §3). Subjective uncertainty is measured by this degree of belief.

- Since an agent's evidence about the chances informs their degrees of belief (e.g., if you think a coin is fair, your degree of belief in heads should equal the 0.5 chance) we can treat **all** decisions as decisions under uncertainty with subjective probabilities over outcomes.

Preference Between Subjective Gambles

- If one is uncertain, then every action is a **prospect** (or gamble, or lottery): which outcome eventuates depends on something you don't know (Reiss 2013: 43).
- But it can be quite **reasonable to prefer some prospects to others**.
  - E.g., suppose you are hungry, and you come across someone else's leftovers in food court. You are presented with a prospect: eat the leftovers, and possibly end up sick (or possibly end up fine); or refrain.
  - Depending on your other options and values – you may have money to buy your own food; you may be uncaring about your own health – you may quite reasonably prefer not to take the prospect of eating the leftovers.
- What is rational preference between prospects?
- Just as in the case of certain decision-making, it is supposed that preference should be negatively transitive and should be complete.
- It is also supposed that rational preference satisfies the **independence of irrelevant alternatives**:

  Independence Principle A rational agent prefers prospect x to prospect y iff they prefer the complex prospect [x with probability p, otherwise z] to the complex prospect [y with probability p, otherwise z]; i.e., if two prospects differ just in one prize (one sweetens the other), then your preference between them reflects your preference between the prizes. (Reiss 2013: 43;

Expected Utility

- If your preferences among prospects meet these conditions, then you can be **represented as** acting to maximise the **expected utility** of your actions (Peterson 2017: 175).
- The **expected value** of a quantity is, more or less, the probability-weighted mean of possible values for the quantity.

(EU) The **expected utility** of action a for an agent A is defined as follows, where $s_1, \ldots, s_n$ are the possible world states, Pr is A's degree of belief function, and U is A's utility function:

$$EU(a) = \Pr(s_1)U(s_1 \wedge a) + \ldots + \Pr(s_n)U(s_n \wedge a)$$
$$= \sum_{i=1}^{n} \Pr(s_i)U(s_i \wedge a).$$

Theorem 2 A's preferences between possible prospects are rational (i.e., complete, transitive, satisfy Independence) iff there is a utility U such that: $EU(x) > EU(y)$ iff $x \succ y$ (Hausman, McPherson, and Satz 2017: 63).

Unique Utility

- The credence is uniquely fixed. The resulting utilities are unique up to linear transformation – any function U which yields these preferences is related to any other U' by an equation of the form $U(x) = aU'(x) + c$, where $a > 0$ (Reiss 2013: 44; Peterson 2017: 107–8).
- This entails that **utility differences** are significant.
  - Suppose U is a utility that fits A's preferences. Let $\delta(x, y) = U(x) - U(y)$.
  - Suppose U' is a linear transform of U, and δ' the corresponding difference between utilities. Then $\delta'(x, y) = U'(x) - U'(y) = (aU(x) + c) - (aU(y) + c) = a(U(x) - U(y)) = a\delta(x, y)$.
  - So if $\delta(x, y) > \delta(v, w)$, then also $\delta'(x, y) > \delta'(v, w)$.
- The numbers thus matter **a bit**: they have no absolute significance in themselves, whether the utility assigned to an act is 10 or 1000 or whatever, and whether the differences between acts are 10 or 1000 or whatever. But if A prefers x to y **more than** they prefer v to w, that will be true on every utility that represents their preferences.

The Drug Case

- Given the drug case, Jill is indifferent between prescribing drug B and drug C, and strongly prefers prescribing drug A to either of them. Let her possible actions be A, B and C; her preferences are $A \succ B$, $A \succ C$, and $B \sim C$ (she is indifferent).
- These acts are prospects: while action A 'very likely' leads to an outcome where the condition is relieved but not cured, action B leads to a complete cure with probability p, or death with probability $1 - p$; *vice versa* for action C.
  - These probabilities are not chances; if Jill knew the chances, she'd know which drug to prescribe. Rather, they are her estimates of the chance that B is the killer drug, or that C is. Those chances are fixed independent of her decision.

Utilities in the Drug Case

- Here are some utilities which license these preferences; we simplify to assume that A leads with certainty to a cure:

Table 2: Drug utilities

Acts\States	Drug B kills/C cures	Drug C kills/B cures	EU
A	90	90	90
B	0	100	$100p \approx 50$
C	100	0	$100 - 100p \approx 50$

- $EU(A) > EU(B)$, mirroring $A \succ B$; $EU(B) \approx EU(C)$, mirroring $B \sim C$; etc.
- So we can represent Jill's preferences as the product of her belief (degrees of belief) and some attribution of utilities to outcomes (conjunctions of acts and states).
- Again, interpretivism says that we **attribute** these 'subjective desirabilities' as real psychological states of Jill on the basis of her total preferences, revealed globally in her behaviour.

Realism about Belief and Desire

- Just as in the case of decisions under certainty, we can ask: is preference between prospects **more basic**, while the utility/credence representation is a convenient model drawn from it? Or are utility and credence to be accepted as psychologically real, and preference **derivative** from them in line with expected utility theory?
- From the point of view of rationalizing explanation, credences and utilities are a sophisticated **refinement** of the categories of folk psychology (Lewis 1974: 337). If folk psychology is credible in the mental states it posits, then these refinements are not obviously less plausible.
- The action is then explained as chosen because it follows from the agent acting on their **strongest** reason, i.e., the act that has highest expected utility, assuming the agent is again caused to act by that reason (and not the victim of some deviant causal chain leading to that act).
- There seems little reason to think preference must be **prior** to these mental states; if we are realist about preference as a mental ranking, why not be realist about the mental states which arguably explain preference?
  - Still, a vestigial positivist suspicion of mental states, and an implausible confidence in the observability of preference, could motivate someone to believe in preference but reject the existence of subjective credence and utility.

References

References

Davidson, Donald (1963) 'Actions, Reasons, and Causes', *The Journal of Philosophy* **60**: 685–700. doi:10.2307/2023177.

Eagle, Antony (2016) 'Probability', in Paul Humphreys, ed., *The Oxford Handbook of Philosophy of Science*: 417–39. Oxford University Press.

Hare, Caspar (2010) 'Take the Sugar', *Analysis* **70**: 237–47. doi:10.1093/analys/anp174.

Hausman, Daniel, Michael McPherson, and Debra Satz (2017) *Economic Analysis, Moral Philosophy, and Public Policy*, 3rd edition. Cambridge University Press.

Hempel, Carl G and Paul Oppenheim (1948) 'Studies in the Logic of Explanation', *Philosophy of Science* **15**: 135–75.

Hume, David (1777/2022) *An Enquiry Concerning Human Understanding*, Amyas Merivale and Peter Millican, eds. https://davidhume.org/texts/e/.

Jackson, Frank (1991) 'Decision-Theoretic Consequentialism and the Nearest and Dearest Objection', *Ethics* **101**: 461–82. doi:10.1086/293312.

Lewis, David (1974) 'Radical Interpretation', *Synthese* **27**: 331–44. doi:10.1007/bf00484599.

Lewis, David (1986a) 'Causal Explanation', in *Philosophical Papers*, vol. **2**: 214–40. Oxford University Press.

References (cont.)

Lewis, David (1986b) *On the Plurality of Worlds*. Blackwell.

Peterson, Martin (2017) *An Introduction to Decision Theory*, 2nd edition. Cambridge University Press.

Reiss, Julian (2013) *Philosophy of Economics*. Routledge.

Salmon, Wesley C (1971) 'Statistical Explanation', in Wesley C Salmon, ed., *Statistical Explanation and Statistical Relevance*: 29–87. University of Pittsburgh Press.

Salmon, Wesley C (1989) 'Four Decades of Scientific Explanation', *Minnesota Studies in the Philosophy of Science* **13**.

Schroeder, Mark (2005) 'Instrumental Mythology', *Journal of Ethics and Social Philosophy* **1**: 1–13. doi:10.26556/jesp.v1i1.10.

van Fraassen, Bas C (1980) *The Scientific Image*. Clarendon Press.

Vane, J R (1971) 'Inhibition of Prostaglandin Synthesis as a Mechanism of Action for Aspirin-Like Drugs', *Nature New Biology* **231**: 232–35. doi:10.1038/newbio231232a0.