


ARTICLE

Intraspecific economics: interpersonal utility comparisons, evolution and culture

Armin W. Schulz 

Department of Philosophy, University of Kansas, Lawrence, KS, USA
Email: awschulz@ku.edu

(Received 07 November 2024; revised 06 March 2025; accepted 06 May 2025)

Abstract

One relatively recent pivot in the discussion concerning the possibility of interpersonal utility comparisons (IUCs) centers on evolutionary biological considerations. In particular, it has been suggested that, since all human beings are part of the same species, we should expect our utility functions to be structured similarly and thus be comparable. However, a closer look at this argument shows that it is not compelling as it stands: since cultural learning plays a crucial role in our cognitive evolution, the conclusion that we evolved to be psychologically similar to each other needs to be revised. This, though, does not mean that evolutionary theory has nothing to say about the possibility of IUCs. In fact, as this paper makes clear, by expanding the evolutionary argument with work in gene–culture–technological coevolutionary theory, it becomes possible to support the contention that IUCs may well be sometimes possible. This has important implications for the analysis and design of social institutions.

Keywords: cultural evolution; evolutionary economics; institution; preference; utility

Introduction

The possibility of interpersonal utility comparisons (IUCs) is one of the foundational issues of economics and is deeply connected to questions of distributive justice (for good overviews of these issues, see, e.g., Jevons, 1911; Harsanyi, 1955; Waldner, 1972; Arrow, 1977; Sen, 1979; Hausman, 1995; List, 2003; Rossi, 2014). Unsurprisingly, therefore, there has been much discussion of this issue since the early days of the subject. One relatively recent pivot in the discussion concerns evolution: in particular, it has been suggested that, from an evolutionary biological perspective, we should expect that IUCs are possible (Goldman, 1995; Rossi, 2014; see also Sorensen, 1998). After all, since all humans are part of the same species, we should expect our utility functions to be built similarly – and thus, we should expect them to be comparable.

However, a closer look at this argument suggests that it is not compelling as it stands. In particular, since cultural learning plays a crucial role in our cognitive evolution, the conclusion that we evolved to be psychologically similar to each other is far from obviously true. This, though, does not mean that evolutionary theory has nothing to say about the possibility of IUCs. In fact, as this paper makes clear, by expanding the evolutionary argument with work in gene–culture–technology coevolutionary theory, it becomes possible to support the contention that IUCs may well be *sometimes* possible. This has important implications for the analysis and design of social institutions.

To show this, this paper begins, in the section ‘The problem of IUCs’, by clarifying the terms of the debate – i.e. how the problem of the possibility of IUCs should be understood. The section ‘The argument from evolution and its problems’ considers the traditional appeal to evolutionary biological considerations in this debate and makes clear why this appeal fails. The section ‘An updated gene–culture–technological coevolutionary argument’ introduces recent work from gene–culture coevolutionary theory to deepen the discussion from the previous section and show its implications for the debate surrounding IUCs. The section ‘Conclusion’ concludes.

The problem of IUCs

The question of whether, and if so, how, we can compare utilities across people can be (and has been) understood in many different ways. To avoid confusion, therefore, it is important to be clear about (i) how the notion of ‘utility’ is to be understood in this debate and (ii) the epistemological nature of this debate.

Starting with point (i), it is clear that ‘utility’ is a famously ambiguous concept (Broome, 1991; List, 2003; Hausman, 2012). In its most general and least substantive understanding, ‘utility’ is just a theoretically grounded measure of a choice ordering. Now, as is widely known, on one of the major economic theories of choice – expected utility theory – IUCs are not possible: we can capture choice behaviors with utilities that are not unique (or only up to certain transformations) and not interpersonally comparable (Hausman, 1995; List, 2003; Rossi, 2014). Nothing in what follows will question this. However, it is important to note that this conclusion is quite limited.

First, expected utility theory is merely one among many different theories of choice. It is true that many other theories – such as prospect theory (Kahneman *et al.*, 1982) and regret theory (Loomes and Sugden, 1982) – do not allow for interpersonal comparisons either. However, some other theories have much less of an issue with this. For example, the theory of simple heuristics (Gigerenzer and Selten, 2001) sees agents as making decisions by following a set of relatively simple rules. There is nothing in this theory of simple heuristics that precludes the evaluations that feature in these rules to be comparable.¹ In this sense, therefore, the fact that expected utility theory does not allow for IUCs merely means that one (albeit major) theory of decision-making does not allow for them – nothing more and nothing less.

¹Though this has not been much explored to date.

Second, it is of course also possible to add something to expected utility theory to allow for expected utility comparisons. For example, we could appeal to ‘extended preferences’ (Harsanyi, 1955, 1977; Arrow, 1977) and scale utilities in reference to an impartial observer. Or we could use the 0-to-1 rule and scale utilities in a linear manner between the most and least preferred option (Hausman, 1995). While there is some controversy over exactly which of these addenda to adopt (if any), the key point to note here is just that they exist – for they imply that the jump from the impossibility of IUCs *in expected utility theory* to the impossibility of IUCs *tout court* is further constrained.

Third and most importantly, when we ask about the possibility of IUCs, we often have something more in mind than just the choice-focused understanding of utility.² In particular, from a public policy perspective, we generally want to know how people’s well-being is affected by a policy. Importantly, to do that, we need to consider how people *evaluate* their world: what it is that they want or value (Dorsey, 2009; Bykvist, 2010; Hausman, 2012).³ For this, though, we need to switch to a more substantive reading of utility than that of a theoretical measure of what is chosen – after all, things are chosen for many different reasons, apart from how highly they are evaluated (such as their placement or what the agent did previously). It is therefore such a more substantive reading of utility that will be at the forefront of the rest of the paper.

It is furthermore the case that substantive readings of utility can be more or less specific. On the most specific end of the spectrum, ‘utility’ is a measure of felt *pleasure* (see, e.g., Narens and Skyrms, 2021).⁴ In this sense, its interpersonal comparability depends entirely on the degree to which different individuals’ felt pleasures are comparable – a particular psychophysical question. Much progress has been made on this latter point, and while there is still much to investigate concerning these issues, there is some reason to think that, in this sense, IUCs are possible to an increasingly specific degree (Narens and Skyrms, 2021).⁵

The problem with focusing on this specific reading of utility is that more than just felt pleasures affect our evaluative outlook. We can evaluate something very highly that involves considerable pain and discomfort, as when we decide to call out others’ racist behavior *even if* it involves painful experiences. Hence, the most common interpretation of ‘utility’ is broader and sees it as a label for a general evaluative outlook (sometimes also called ‘preference’) (Hausman, 2012; Rossi, 2014; Broome, 1991; but see also Angner, 2018). In this sense, utility is a representation of the obtaining of something that is positively evaluated, but the nature of this evaluation can be very varied (Hausman, 2012). In particular, it need not involve much pleasure, but

²There are also other reasons for avoiding a purely choice-focused reading of utility (Hausman, 2008, 2012; Angner, 2018).

³The literature on well-being is large, but for present purposes, a detailed survey is not necessary: on all accounts of well-being – including perfectionist ones – an agent’s evaluative outlook matters at least somewhat for assessments of their well-being.

⁴Often, in these discussions, pleasure and pains are presumed to be on the same spectrum, but this could be questioned. See also Bartoshuk (2014).

⁵This point may be underwritten by pleasures and pains being biologically deeply seated emotions (Ekman, 1989; Ekman and Rosenberg, 1997; Sauter *et al.*, 2010). See also footnote 15.

could be based on our moral commitments or assessments of what is rational. Because of its wide scope and flexibility, this reading of ‘utility’ is the most common in the literature (Hausman, 2012) and will therefore also be the central one in the rest of this paper.⁶

With this in mind, the question at the heart of this paper is whether we evaluate (in this broad sense) the world on *the same scale* (Davidson, 1986; Goldman, 1995). Is it the case that, for you, all states of the world are evaluated nearly similarly and not so differently from a neutral baseline (a modicum of food and shelter, say), whereas I evaluate different states of the world drastically differently from each other, with even minor differences to the neutral baseline leading to major jumps in evaluative strength?⁷ Or is this not the case? If it *is* the case, is there some way of determining the differences in these scales, so that they can be bridged with the appropriate translation function? Or is such a function out of reach for us?

Note that, given the broadness of the notion of utility underlying these questions, they are harder to answer than the more specific psychophysical question of the interpersonal comparability of felt pleasures. Since the sense of utility involved here is broad – and can be based on feelings of discomfort, sadness and pain, as well as no feelings at all – it is not just an issue of whether we feel or display the same amount of pleasure. For this reason, assessing the possibility of IUCs so understood is very difficult: the same behavioral, neural and other physical states seem consistent with many different evaluative scales (List, 2003).

With this in mind, consider point (ii) mentioned earlier: the epistemological nature of the problem. Given the broad understanding of utility just laid out, it is possible to see the problem of the possibility of IUCs as just a specific instance of the problem of other minds.⁸ Is it conceivable, plausible, or deniable that two people have different mental lives with identical behaviors or physical states (List, 2003; Rossi, 2014)? This issue is just a specific version of the general discussion surrounding skepticism or empiricism. Because of this, the latter discussion transfers to the present case of IUCs as well. In particular, on an empiricist reading, if there is no available evidence that could determine whether two agents have different utility functions, positing the existence of such differences is meaningless. This conclusion, though, depends on the acceptance of empiricism. For example, on a Popperian reliabilist reading (Lipton, 1995), it does not go through (see also Goldman, 1995; Rossi, 2014): the availability of evidence (or lack thereof) does not entail anything about what claims we are justified in making. (Of course, strongly skeptical epistemological views would also deny

⁶For present purposes, it is not necessary to further specify the details of this broad sense of utility. As the rest of the paper makes clear, there are ways of answering the question of its interpersonal comparability that leave open its exact nature.

⁷Many other options – including non-linear ones – exist, too, of course.

⁸As suggested earlier, if (a) utility is seen as a measure of felt pleasures, and (b) felt pleasures (unlike other mental states) are physical states which are intersubjectively measurable, the interpersonal comparability of utility is *not* an instance of the problem of other minds. Of course, to the extent that (b) is doubted, the appeal to (a) is not sufficient to remove the question of the interpersonal comparability of utilities from the problem of other minds. The same is true if (a) is denied, as when we move to a broader understanding of utility (as is the case here).

this.)⁹ If so, though, then the problem of the possibility of IUCs may become effectively definitionally unanswerable (or as answerable as general skepticism is).

However, at stake here is a more specific, empirical version of the problem. We may agree with the empiricist that, if IUCs are meaningful at all, they would need to show up somewhere. However, this does not mean that we should presume that the relevant evidence is *easily* obtainable – or indeed ever obtainable *for us*. That is, we may acknowledge that we can ascertain that others have minds, and we may also acknowledge that it may, in principle, be possible to determine whether people's evaluations of the world are fundamentally the same or different – e.g. if we had sufficient neural information (Goldman, 1995). However, we may also note that the relevant (neural) information may not be available for a long time (if ever). What we are looking for are considerations that may help us adjudicate this question *now* – not in some future completed science or social science.¹⁰

On the flip side, the fact that we cannot easily tell *behaviorally or neurologically* whether the scale of our evaluative outlooks is the same does not mean that there is no other (non-controversial) reason to affirm or deny this. There may be other, non-behaviorist or non-neural considerations to which we can appeal to determine whether this is the case. It is precisely here where the consideration of evolutionary biology comes in. (Of course, the appeal to evolutionary biology is not the only way to try to resolve this issue – see also Rossi, 2014. However, it is the one that will be in focus here.)

In the same way, we may also see the problem of the possibility of IUCs as an implication of the general problem of evaluative incommensurability (see, e.g., Chang, 1997). If evaluative attitudes are not generally commensurable *within* individuals, they would not be comparable *across* individuals either. In this reading of the problem, a slightly different version of the previous point applies. In those cases where evaluative attitudes *are* seen as commensurable within individuals – the existence of which is accepted even by those thinking that they are not *always* so commensurable – we would still need to find reasons for thinking that the scales of commensurability within individuals are comparable across individuals as well. This again brings up all the issues just discussed. In short, evolutionary biological considerations may be thought to be a bridging principle to underwrite the empirical possibility of IUCs – assuming this empirical possibility is granted at all.

The argument from evolution and its problems

In particular, several authors have suggested that, from an evolutionary biological perspective, we should expect IUCs to be possible. So, Goldman (1995: 724) (there attributed to Sorensen, 1998; see also Rossi, 2014) writes:

⁹Relatedly, both Goldman (1995) and Rossi (2014) formulate the problem as concerning the justifiability of individual beliefs about utility differences – i.e. as a sort of mindreading problem. While the focus of this paper is on a more social scientific reading of the dispute, many of the arguments of this paper also apply to Goldman's and Rossi's discussions.

¹⁰As just noted, depending on one's epistemology, even then, it may not be ascertainable – though this is then a general question of epistemology and skepticism and does not require a specific discussion here. Note also that the discussion of behavioral signatures of 'sentience' is an attempt to solve the related (though not identical) problem of whether other organisms feel pain (see, e.g., Birch, 2024).

The general hypothesis, then, is that natural selection has molded us into a population of hyper-similar individuals, which helps support the reliability of simulation. It would not be easy to get firm scientific evidence for this hypothesis; but it seems possible, and if accomplished it would strengthen the scientific case for the accuracy of simulation in IU comparisons.¹¹

Similarly, Binmore (2009: 552), when writing about ‘sympathetic preferences’ in the vein of Harsanyi’s impartial observer, states:

It is easy to see why the forces of biological evolution might lead to our behaving as though we were equipped with sympathetic preferences. Mothers commonly care for their children more than they do for themselves—just as predicted by the model that sees us merely as machines that our genes use to reproduce themselves. In such basic matters as these, it seems that we differ little from crocodiles or spiders. However, humans do not sympathize only with their children; it is uncontroversial that they also sympathize to varying degrees, with their husbands and wives, with their extended families, with their friends and neighbors, and with their sect or tribe.

In the background of Binmore’s suggestion here is the claim (similar to that made by Goldman) that it is biologically plausible that we have evolved to be sufficiently similar to each other so that sympathizing with others – putting ourselves in their shoes – is actually accurate. Otherwise, it would just be determining what I would feel if X were true twice over (Nichols and Stich, 2003).

Unfortunately, these appeals to evolutionary biology are not spelled out systematically or in detail, which makes it hard to assess them clearly. The first step in assessing these appeals to evolutionary biology therefore has to be in making the steps of the argument more precise.

To do this, it is best to begin by considering physical traits – such as human livers, hearts, or kidneys – and asking the parallel question to that of IUCs for these traits: would it be reasonable to think that there is intraspecific variation in the ways in which these traits function? That is, is it reasonable to think that in different humans, the processes going on in their livers, hearts, kidneys, etc., are fundamentally different?

Now, in one sense, the answer is straightforward: yes, this is the case. After all, variation is common in biology, and therefore, there will *of course* be intraspecific differences in liver, heart, or kidney function. In this sense, therefore, this question is easy to answer. However, in a different sense, answering this question is not so straightforward. Abstracting away from minor variation around a mean (and leaving aside more extreme cases, such as congenital diseases), it is not unreasonable to see liver, heart, or kidney function as *broadly similar* in humans. Indeed, it is precisely this that makes it reasonable to speak of ‘human liver functioning’ and contrast it with, say, ‘bearded lizard liver function’. Human physiology, as a science, is based on this being the case.

¹¹Recall that Goldman’s (and Rossi’s) target is the justifiability of individual beliefs about utility differences (see also footnote 9), which is why the quote above refers to ‘the reliability of simulation’. Still, the evolutionary argument put forward there is still used to support a premise of psychological sameness and has the same structure as the one presented here (though lacking in some of its details).

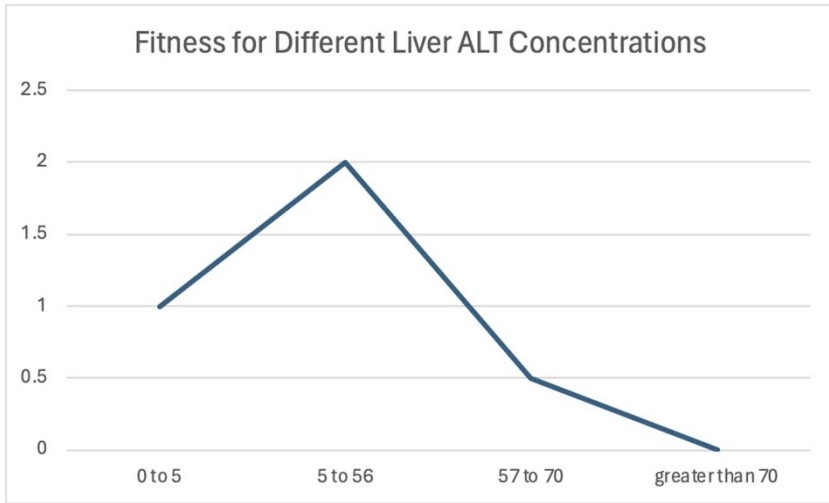


Figure 1. Illustrative fitness values for liver function. ALT, alanine transaminase.

Indeed, the existence of species-typical physical traits is underwritten by two, mutually reinforcing evolutionary biological reasons: common ancestry and natural selection. Consider these in turn.

First, adaptively important traits (like livers) should be expected to spread through the relevant population. While this need not be the case if the trait only recently came under selection, it is very plausible for older, highly conserved traits – like hearts, livers, or kidneys. That is to say, since liver function matters for *all* humans, there is selective pressure for livers to operate well in *all* humans. On the (reasonable) assumption that the fitness of liver functioning is single-peaked – so that there is one specific way for livers to function that is adaptively optimal – this thus suggests that human livers function in the same way for all humans (modulo minor differences). Concretely, if there is selection for livers to produce alanine transaminase at about 30 units (5–56, say) per liter, then we should expect all human livers to do this – after all, this is what is optimal for all humans. Figure 1 sketches a – purely illustrative – fitness function for this case.¹²

Second, even if a given physical trait has not been under strong selection, the fact that all humans are related to each other can provide a reason to think that the trait works quite similarly across people. Because of interbreeding across humans, humans are all quite alike genetically: genetic diversity in humans is estimated to be about 0.1% (Understanding Human Genetic Variation, 2007). So, just in virtue of the fact that humans are part of the same species, we should expect physical traits that have been part of the human lineage for a long time to be largely similar. The same is true for

¹²Technically, fitness function of livers would need to be relativized to an appropriate life-history stage (and thus be seen as dynamically changing across life-history stages). However, this complication can be left aside here, as it does not affect the substance of the point made in the text.

other organisms: while every bearded dragon is a little different, we still expect bearded dragon spikes and tails to function in roughly the same way across the species (independently of the adaptive histories – or lack thereof – of these traits). While this need not be the case for traits that do not have a long history of intermixing – e.g. because they are only recently culturally inherited – it is plausible for many physical traits (like hearts, livers and kidneys).

For present purposes, noting all of this matters, as it might be taken to suggest that something similar is true for human utility functions, too. In particular, we may think that psychological traits like our evaluative attitudes operate like physical traits so that we can directly transfer the two evolutionary biological points just made from the physical case to the psychological case.

First, how we evaluate the world matters for our decision-making: it determines what we choose to do. In turn and for obvious reasons, the latter greatly impacts our fitness. Hence, it is plausible that utility functions are adaptively very important. By the same reasoning as the one just laid out for physical traits, we thus may expect all humans to have utility functions that are broadly similarly structured. Note that this does not imply that we cannot differ over *what* it is that we want: the content of our wants (plausibly) needs to be adaptively responsive to the specifics of our situation – what resources we hold, which are available, etc.¹³ Still the *nature* of our wanting – the scale of our preferences – could be thought to be the same: the function of wanting (arguably) is to get to us to obtain the object of the want (Millikan, 1984, 2002; Papineau, 1987, 2003); this may thus be seen to be same across people. Hence, we may be led to expect IUCs to be possible.

This argument seems to be further strengthened by the second evolutionary point sketched earlier: the idea of common ancestry. Even if our utility functions have not been under strong selection, the fact that humans form one common gene pool may be taken to suggest that these utility functions are structured in the same way. The intermixing of human traits – through sexual reproduction and the formation of a largely similar human genome – could be taken to suggest that differences in utility functions would be washed out. This thus may be seen to further support the possibility of IUCs.

Overall, therefore, the standard set of evolutionary biological arguments in the discussion of IUCs center on establishing the claim that we should expect all humans to have evolved minds that work in similar ways. It is this pan-human psychological similarity that is then used to support the possibility of IUCs: while the existence of some variation in our evaluative structures may need to be expected, the overall functioning of our evaluative systems should be basically the same. In this regard, this evolutionary argument for the possibility of IUCs is related to the classic nativist evolutionary psychological ‘argument from Grey’s anatomy’ for a human psychic unity (Tooby and Cosmides, 1992; Buss, 2014; Barrett, 2015). Just like it is uncontroversial that we can compile books like a ‘Grey’s Anatomy’ that detail human physiological structures and

¹³Of course, this does not deny that there may *also* be communalities in our wants (Buss and Schmitt, 1993; Shackelford *et al.*, 2005). The point here is just that the evolutionary argument in the text should not be seen to be committed to the claim that *all* of our wants need to be the same.

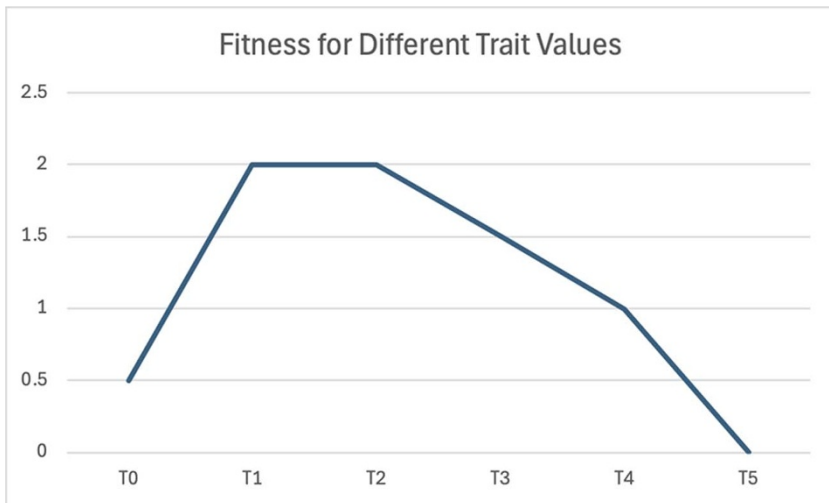


Figure 2. A fitness function with a plateau.

their functions, it should also be uncontroversial that we can lay out a ‘Grey’s Anatomy of the Mind’. After all, minds are just instantiations of biophysical systems, and should thus follow the same principles as all such systems.

However, this argument faces some relatively clear concerns (some of which are also noted by Rossi, 2014). The first concern is that, even accepting the broad outlines of the premises of the argument, questions about the scope of the conclusion of the possibility of IUCs remain. In particular, even if we agree that it is reasonable that psychological evaluations are *largely* intraspecifically similar (just like leaf size in a sugar maple tree is generally between 3 and 6 inches), this does not mean that the remaining ranges are not evaluatively significant. It could be the case that the satisfaction you get from playing the piano is *largely* the same as the one that I get from playing the guitar; however, it could also be that the remaining differences are still quite sizable. Consider again leaf size in sugar maples: while most leaves may be about 4.5 inches across, some may be 3 and some 6 – effectively double in size. The same may be true for utility functions: even if we accept that human evaluative processes are largely similar across people, this does not mean that the remaining variation is not significant. At least, the evolutionary biological considerations appealed to above have provided no reason to rule this out.

The second concern with the above two-pronged evolutionary argument targets the plausibility of the premises on which it rests. On the one hand, as far as the selection-based argument is concerned, this argument presupposes that the fitness function over human psychological evaluative structures is not multi-peaked or containing plateaus. For if either of this is true, then even with strong selection, we cannot rule out the evolution of a polymorphism in the relevant set of traits. As made clearer in Figures 2 and 3, in such cases, the evolution of trait values T1, T2, or T3 will heavily depend on chance and the starting place of the relevant (sub-) populations – there is no unique trait value that these populations need to be expected to converge to.

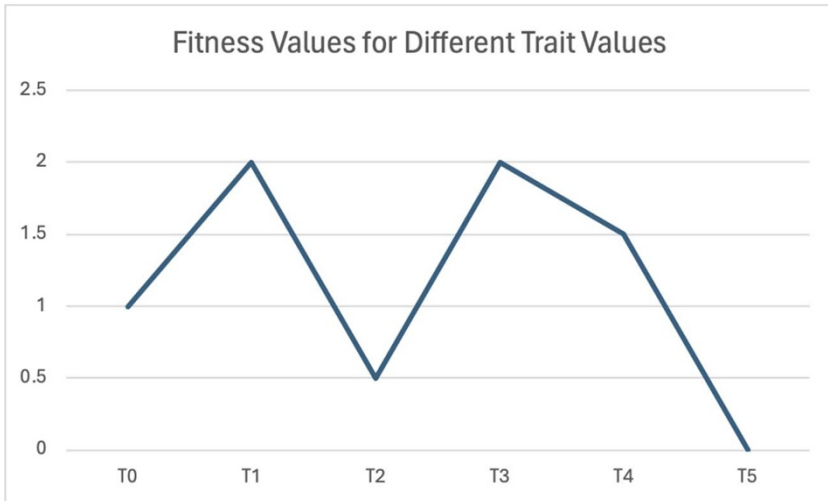


Figure 3. A multi-peaked fitness function.

Moreover, this is not a far-fetched possibility, as plausibly adaptive polymorphisms are common in the biological world. Indeed, they are found in human physiological traits, too: for example, it appears that hemoglobin production in Tibetan highland populations is different from that in other human populations (see, e.g., Huerta-Sanchez *et al.*, 2014). This concern thus cannot be ruled out *a priori* when it comes to human psychological evaluative structures.

On the other hand, as far as the argument from common ancestry is concerned, there is the (related) concern that this argument rests on an overly simplistic picture of human cognitive evolution, as it underplays the role of culture in human evolution. In general, much about human cognition has evolved in response to cultural factors. While there is debate about the strength and nature of cultural influences on human cognition – with some placing most of the emphasis on culture (Tomasello, 1999; Sterelny, 2012; Heyes, 2018), and some seeing cultural influences as merely ‘triggering’ different innate biases (Tooby and Cosmides, 1992; see also Schulz, 2021; Ward, 2022) – there is no debate concerning the fact *that* culture matters for human cognition.¹⁴ Given this, it is entirely reasonable to expect that there are cultural influences on our evaluative systems as well. Put differently, our evaluative system may have evolved to be a culturally mediated cognitive polymorphism – we may not have evolved just one, culturally universal set of preferences, but a system for acquiring many preferences from our culture. Given this, though, the above common ancestry-based argument for the possibility of IUCs does not go through: the utility functions need not have a long history of intermixing, as they are likely to be heavily culturally determined.

¹⁴There are many examples for this point; a particularly salient one – with some plausible close connections to utilities – is a cultural difference in affect and emotion: see, e.g., Scollon *et al.* (2011); Matsumoto (2001).

All in all, therefore, as presented, the above evolutionary argument for the possibility of IUCs cannot be seen as compelling. However, it would be premature to end the discussion here: for while the evolutionary argument as commonly conceived is not compelling, it is possible to use some of the ideas in the background of the argument to build a more compelling, gene–culture coevolutionary case for the *sometime* possibility of IUCs.

An updated gene–culture–technological coevolutionary argument

To understand the impact that broadly evolutionary considerations can have on the debate surrounding the possibility of IUCs, it is best to begin where the criticisms above have left off: namely, by noting that it is very plausible that utility functions – in both content and structure – will, at least partially, culturally evolve.¹⁵

On the one hand, as noted earlier, it is plausible that much of what we want has no particular biological significance outside of the cultural realm – e.g. when it comes to fashion, foods, or forms of music. Perhaps within some very broad limitations, there is no adaptive significance to shirt designs, the eating of cooked shellfish, or the use of microtones in music. Putting this the other way around, our evaluative system has evolved in such a way that we can acquire wants that match our particular culture (Boyd and Richerson, 2005; Sterelny, 2012; Henrich, 2015; Heyes, 2018). In turn, which wants are culturally reinforced will depend on the specifics of that culture – a point to which I will return momentarily.¹⁶

On the other hand and most importantly for present purposes, this is likely to be especially true for the *intensity* of our wanting. That is, different cultures may favor different utility *scales*. So, even though different utility scales may not make a material cultural or biological difference to our behaviors – in line with the common contention that they are very difficult (or even impossible) to detect – they may still be differentially *culturally favored*: it is plausible that we have evolved in such a way that we culturally learn how intensely to want things. To see this, it is best to proceed in two steps.

First, note that we should see human cultural evolution as being both enhanced by and enhancing of technological evolution (Schulz, 2020, *forthcoming*; Tomasello, 2021). In particular, human cultural learning is made powerful in part due to the fact

¹⁵This is consistent with it being the case that *some* aspects of human utility functions have been shaped by biological evolution – for example, it is reasonable and indeed widely accepted that desires for food, sex, shelter and company (among others) have been shaped in these ways and that they are therefore culturally universal (Brown, 1991; Tooby and Cosmides, 1992; Shackelford *et al.*, 2005; Buss, 2014; Barrett, 2015). Indeed, this may well be plausible for the narrower, pleasure/pain-focused reading of utility, for which a purely biological evolutionary explanation may be more compelling (Ekman, 1989; Ekman and Rosenberg, 1997; Sauter *et al.*, 2010) – though this continues to be debated. At any rate, the point in the text is just that, *in general*, our evaluations have plausibly evolved in response to cultural pressures.

¹⁶In the background here is the long dispute over the plausibility of one of the main contentions of work in human behavioral ecology: that humans (just like other organisms) have evolved to make decisions in a largely adaptive manner (Nettle *et al.*, 2013). As noted in the previous section, the reasoning for this is based on the idea that, almost by definition, there is selective pressure on making adaptive decisions. However, while this contention may be plausible in some cases, it is not plausible in *all* cases: the evolution of the psychological apparatus underlying human (or non-human) decision-making is also likely to have been heavily influenced by chancy processes like drift – as well as cultural learning.

that it can be underwritten by artifacts and institutions that aid this learning (and which are themselves the product of cultural learning). So, writing, myths and calculating devices like the Plimpton 322 tablet and quipus can make the storage, acquisition and processing of *much* cultural information possible.

In the present context, these kinds of tools are important, as they can make different utility scales visible in ways that pure behavior cannot. Indeed, this present paper is a good example of this: we can *write or talk about* different utility scales, even if we cannot easily show them in our behavior. For example, in some cultures, people may hear the Goldilocks tale, Stoic lectures, or Buddhist teachings on the importance of keeping one's exuberance in check. In other cultures, people may hear stories like that of Jacob in the Bible, who, in his grief, tears his clothes and refuses to be comforted. These different depictions of emotional intensity can be used to differentially emphasize different utility scales in different cultures: in some cultures, we may be taught less intense preference satisfaction, in others, we may be taught more intensive preference satisfaction.

More specifically, cognitive technology like myths, dances and songs can help define what we should expect as the *baseline* of our evaluative systems. We may learn that 'pain is unavoidable and pleasure is rare', or we may learn that we should 'expect a pain-free life'. Similarly, myths, dances and songs can help define what a 'unit' of evaluation is. We may learn that 'if you feel slightly uncomfortable about something, don't do it', or we may learn to 'ignore anything unless it leads to major damage to body or property'. Importantly, this is exactly what it takes to define utility scales: we need to set out a zero and a unit (Binmore, 2009; Rossi, 2014). By being taught what we should take our baseline evaluative expectation to be, as well as which things we should take into account in our evaluations (and which not), we can thus learn about which utility scales to adopt – even though these scales are otherwise behaviorally inert and thus bio-culturally neutral.¹⁷ Putting this the other way around, we may agree that behavioral data may not be compelling *evidence* for whether utility scale A or B is the right one in the case at hand – but these data can still be used as *pedagogical tools* to instruct others whether to adopt scale A or B. In this way, cultural artifacts can not only scale the utility of an outcome *within* an agent, but also *across* agents.¹⁸

This matters, as the fact that these different evaluative schemes can be taught through cognitive technology like myths means that they can be correlated with other bio-cultural traits – like what to wear when, or how to grow which crops. Myths, stories, songs and rituals rarely if ever only emphasize or teach *one* cultural trait – they mostly or always teach several at once (Boyd and Richerson, 2005). The Goldilocks tale is not just about the value of moderation – it has implications for how to treat others

¹⁷Of course, this does not mean that the content of our preferences is behaviorally irrelevant – as noted earlier, how we evaluate the world matters for what we choose to do. It is just that they are often largely relevant for situating us better *within our culture*: controlling for what is culturally advantageous, listening to Beethoven over the Beatles does not enhance our expected reproductive success very much. On top of this – and this is key here – the *scale* of our evaluative attitudes (our utility function) may not show up in our behavior much or at all: the latter may depend more on the contents of our wants, not their scaling. However, as the rest of the paper tries to show, given the way cultural inheritance works, we may still culturally evolve a quite narrow range of utility functions.

¹⁸This thus also relates to the discussion of the commensurability point made earlier.

(‘the Golden Rule’). Jacob’s tale is not just about the intensity of grief, but also emphasizes the value of forgiveness. Importantly, these other culturally transmitted traits may well be bio-culturally adaptive: how an individual treats others can be a major factor in how well they fare in society and also contribute to the stability of that society (Whiten and Byrne, 1997; Sterelny, 2003; Machery and Mallon, 2010; Mikhail, 2011; Tomasello, 2021; Westra and Andrews, 2022; Westra *et al.*, 2024).

It is important to be clear about the argument here. The key conclusion to be stressed is that it is theoretically and empirically plausible that we can and do culturally learn what evaluative intensities to adopt. Of course, it is also entirely possible that we merely culturally learn the *expression* of our evaluations, not their intensities themselves. However, the latter possibility should not be seen to rule out the former. Indeed, the very fact that it is also entirely theoretically and empirically plausible that we culturally learn the contents of our preferences – and not just particular behavioral choice patterns – should be seen to underwrite this point. It is widely accepted that we can learn the *psychological causes* of behavioral dispositions, and not just the behaviors themselves (Boyd and Richerson, 1985; Fodor, 1989; Barrett *et al.*, 2002). Given that we think this is true for the contents of our evaluations, it is plausible that the same is true for their intensities as well. While this does not amount to the clear *establishment* of the conclusion that we can culturally learn our evaluative intensities (and not just their expression), it is at least *a reason* in favor of the truth of this conclusion. Importantly, this is also useful, as it points to new work that can be done to improve our public policies. I return to this later.

For now, what matters is just to note that a given utility scale can spread through a given culture even though it is otherwise behaviorally inert and not inherently adaptive, biologically or culturally: it may be co-transmitted with other cultural traits that *are* bio-culturally adaptive.¹⁹ That is to say, we may (culturally) teach about utility scales by accident: myths do not spring up to support the transmission of specific such scales – but the latter can be embedded in myths that spring up to support the cultural transmission of other traits. To understand this form of path-dependent cultural evolutionary hitchhiking better, consider the following simple model.²⁰

Assume there is a trait A1 – a particular behavior, like a certain version of slash-and-burn agriculture – that is correlated with A2 – a particular utility scale, like a form of Stoicism that emphasizes a flat evaluative outlook – in culture C. This correlation can stem from the fact that A1 and A2 are both components of the same myth: a myth that teaches both the importance of slash-and-burn agriculture and the importance of the Stoic evaluative outlook. In that case, if A1 evolves in C – because slash-and-burn agriculture is in fact adaptive in C – it can drag A2 along with it.

¹⁹Of course, it is also *in principle* possible that a utility scale just happens to get embedded in a myth that contains no other bio-culturally meaningful information and is then stably transmitted across generations. However, the fact that there is no bio-cultural support for the persistence of this myth makes this a less plausible scenario *in practice*.

²⁰The main point behind the model is quite well understood in the literature, but the model is useful for bringing out the issues of importance here particularly clearly. Indeed, as noted later, the model can really be seen as providing a *framework* for the further investigation of the details of the cultural evolution of utility scales.

In more detail, in the model, there is an optimal agricultural practice out of 100 available practices (the number of available practices and the optimal value can be changed).²¹ However, the society in question does not know what that optimal agricultural practice is. In the initial setup stage of the model, when the first families of the society settled the area, each family (out of the 10 that initially settled the area – an arbitrary number that can be changed) picks an agricultural practice at random out of the 100 available practices, and each family has their own unique, randomly determined utility scale (also out of 100). In the initial stage, the society further creates myths and stories, which (among other things perhaps) contain more or less explicit references to agricultural practices ('do not till fields that cannot be seen by the sun') and utility scales ('Let your sadness run deep when the world is not right, but let your joy run free when it is'). In the model, it is assumed that the agricultural practice of the families that happen to be closest to the optimal agricultural practice gets to determine the content of the myth – both when it comes to the agricultural part *and* the utility scale.²²

Each generation then reproduces according to how close their agricultural practices are to the optimum. That is, each family *i*'s fitness is determined by the following:

$$w_i = (100 - |(\text{Optimal agricultural practice} - \text{Agriculture practice}_i)|)/100. \quad (1)$$

We then allow family *i* to reproduce with probability *p*, where

$$p_i = w_i + w_0, \quad (2)$$

where w_0 is the baseline fitness.

Importantly, subsequent generations initially adopt their parent's agricultural practices and utility scales, but with a normally distributed mutation rate with a standard deviation of 1. However, as they grow up, they then change their agricultural practices and utility scales by moving them toward the founding myth of the society. That is, they set:

$$\begin{aligned} \text{Agriculture}_{\text{updated}} = & [(1 - b) * (\text{best current agriculture} - \text{agriculture}_{\text{initial}})] \\ & + \text{agriculture}_{\text{initial}} \end{aligned} \quad (3)$$

$$\begin{aligned} \text{Utility scale}_{\text{updated}} = & [(d * (1 - b)) * (\text{Founding myth utility scale} - \text{utility scale}_{\text{initial}})] \\ & + \text{utility scale}_{\text{initial}} \end{aligned} \quad (4)$$

where $b \in [0, 1]$ is a bias parameter that models how attached they are to their own utility and agricultural scale, and $d \in [0, 1]$ is a parameter modeling the independence of utility-scale learning and agricultural learning (on which more later).

The final part of the model consists in the cultural evolution of the myth: every generation, the agricultural part of the myth (and only that) moves a little closer to the optimal agriculture practice. This is to model the fact, over time, the people can recognize what agricultural practices work and which do not, which then (perhaps subtly) affects the myths. Then, the older generation dies, the younger generation becomes old, and the model repeats.

²¹The model was coded in NetLogo; the code is available upon request.

²²This last point is not essential to the main conclusions reached here. See also the discussion later.

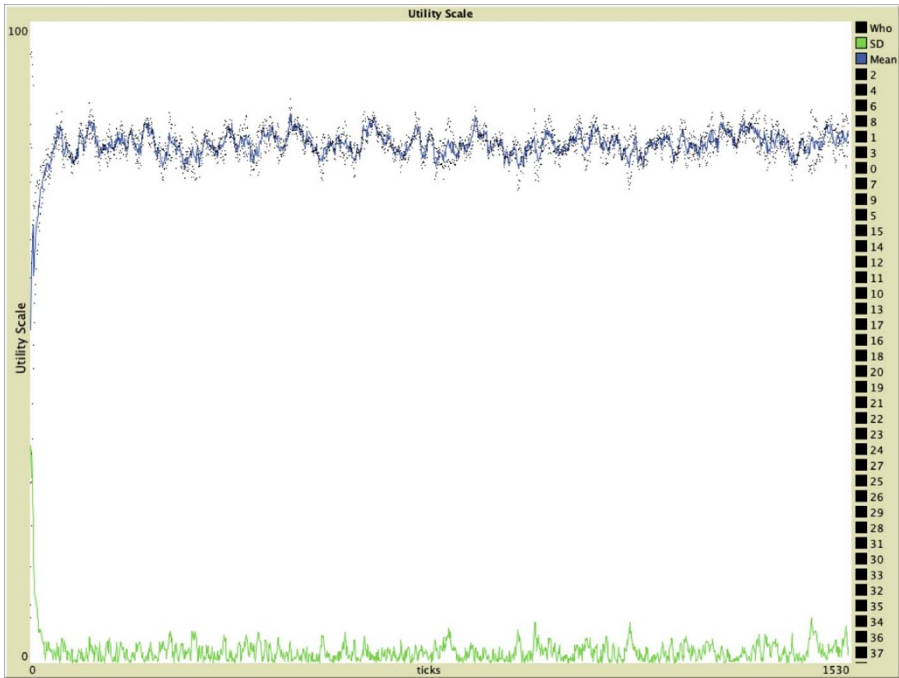


Figure 4. Cultural evolution of a set of utility scales (black) and their variance (green).

For present purposes, what the model shows very clearly is that, even though utility scales are completely irrelevant for bio-cultural success, everyone in the society ends up having a narrow range of utility scales (Figure 4).

As Figure 4 shows, the standard deviation of utility scales plummets over time.

Importantly, also, different societies can – though need not – end up with very different utility scales (even if they inhabit very similar ecologies) (Figure 5).

In the first case, the society ends up occupying a narrow range of utility scales between 70 and 80; in the second case, the range is confined to 25 to 35. These cases are completely distinct: they have no utility scales in common. Of course, in other cases, they may well have overlapping scales. In this model, this is purely driven by chance: what utility scale happens to be embedded in the initial founding myth. This is in line with the fact that utility scales are here assumed to be behaviorally inert – they can be learned, but they do not affect behavior otherwise.

It is important to be clear about the role that the model plays in the argumentative context here. Without a doubt, this is a simple, idealized model, and many of its detailed conclusions depend on the specifics of its assumptions. However, it still can add much of the value to the discussion about the plausibility and possibility of IUCs. Most directly, it can make this discussion more precise; what may seem reasonable informally often turns out to be more complex when put in the more precise setting of a model. More specifically, the model clarifies the argument of the paper along two main dimensions.

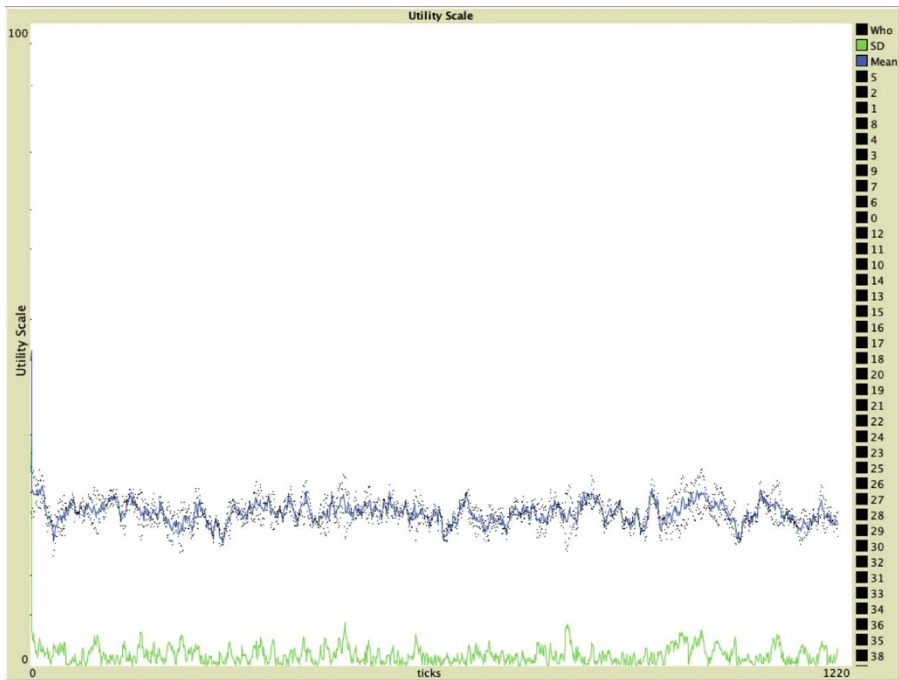


Figure 5. Cultural evolution of a different set of utility scales (black) and their variance (green).

First, the model makes clear that it is reasonable that variation in utility scales will not go to zero. Of course, this is built into the model as a result of the fact that there are errors and biases in the transmission of utility scales – but these are reasonable assumptions that are implausible to give up. Importantly, though, the model also shows that the variation in utility scales *lessens* over time. Put differently, as was the case in the basic evolutionary biological argument, there is still room for individual variation in culturally acquired traits. The point to note here is just that this individual variation is now likely to be significantly *narrower*, since one of the main sources of this variability is controlled for: cultural differences. While this thus does not ensure that IUCs are possible, it makes it more likely. In this way, the model is helpful for making more precise the idea that at least *intra-cultural* IUCs may be possible: within a culture, our evaluative intensities tend to be closer to each other since they are often acquired as by-products of the cultural transmission of other traits.

Second and relatedly, the model shows that the size of the remaining variation in utility scales depends on how tightly utility scales are linked to the – bio-culturally adaptive – agricultural practices. If they are very tightly linked – i.e. if $d = 1$ – the variation in utility scales is very narrow. However, in cases where people pay less attention to the utility scale in the myth than to the agricultural practice – i.e. if $d < 1$ – more variance in utility scales results (Figure 6).

Interestingly, initial simulations suggest that the variation remains low even for small d – though it eventually does get substantial (Figure 7).

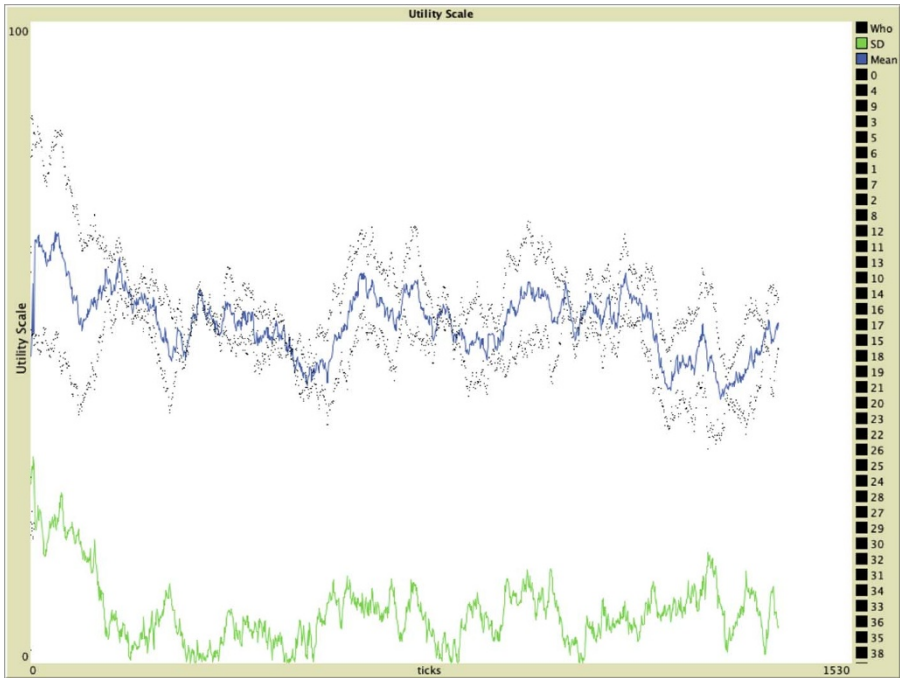


Figure 6. Cultural evolution of a set of utility scales (black) with larger variance (green).

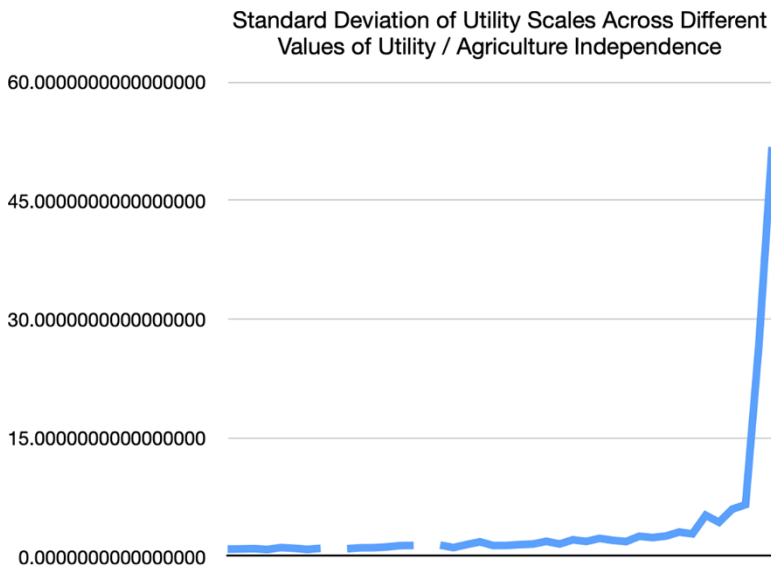


Figure 7. Variation of utility scales across different values of agricultural independence (gaps are cases where the culture went extinct).

This last point is one of the key results here, so it is useful to spend a bit more time on it. Even with all the presuppositions made in the model, significant variation in utility scales may remain, making IUCs difficult even within cultures. In this way, the model again brings out the complexities in this situation: it is not that cultural evolution will definitely lead to easy IUCs within cultures. Whether it does so or not needs to be assessed on a case-by-case basis. Still, the model makes clearer *which* cases are conducive to the possibility of IUCs: namely, those with sufficiently high levels of d . In this way, the model is useful for making the empirical investigation of the possibility of IUCs easier: we now understand better what to look for. Here, it is especially helpful to consider the relationships among its various assumptions. For example, as noted in Figure 5, the variance in utility scales and independence from the agricultural myth are positively correlated – but not linearly so. This also helps us get clearer on where and why IUCs are possible, and thus aids in the design of improved public policies. This a point to which I return momentarily.

In this context, it is also important to note that, while many details of the model can be changed, its overall thrust rests on solid foundations. In particular, the model combines aspects from other, well-established work in gene–cultural coevolutionary theory. So, the model shares some foundations with work on the cultural evolution of maladaptive traits like celibacy (Boyd and Richerson, 1985, 2005; Boyd, 2018) – but unlike in the latter case, utility scales are assumed not to be maladaptive, but just selectively neutral. Similarly, in Boyd and Richerson's (1987) model of the evolution of ethnic markers, ethnic markers become adaptive as a guide to appropriate teachers – though they are not initially adaptive. In the model presented here though, it is not even the case that utility scales function as markers: indeed, since they are assumed to be behaviorally inert, they *cannot* do so. However, as the model makes clear, they can still hitchhike on the evolution of other traits.²³

Putting all of this together, the model can be seen to support and make more precise the key contention of this paper: given the place of culture and technology in human cognitive evolution, it is plausible – though not guaranteed – that utility scales, while behaviorally inert, can hitchhike on the techno-cultural evolution of other, bioculturally adaptive traits. In turn, this provides a reason for thinking that IUCs are sometimes possible: members of a culture converge to a small set of utility scales, simply because these scales techno-culturally evolve as by-products of other techno-culturally evolving traits.²⁴ Put differently, given what we know about the general impact of culture and technology on even fundamental features of human minds (e.g. what we value or how we make decisions), the possibility of IUCs becomes more credible (Tomasello, 1999; Csibra and Gergely, 2011; Sterelny, 2012; Heyes, 2018; Schulz, 2020, 2021; Schulz, *forthcoming*).

²³In this way, it becomes clear that, while the model is drawing on some prior work, it still has many substantially novel aspects: in particular, unlike in previous work, the model shows that it is possible for a behaviorally neutral trait to evolve by hitchhiking on culturally adaptive traits.

²⁴Note again that this is not about mindreading, for which IUCs may not be necessary (Rossi, 2014). In the same way, it is not about 'mindshaping' (Zawidzki, 2013; see also Godfrey-Smith, 2002), where culture shapes minds to make them easier to read. Rather, the issue here concerns the cultural–technological evolution of evaluative intensities as cultural–technological by-products of the cultural evolution of other traits.

This model-cum-argument has several important novel, policy-related implications. First, in its very nature, the present, culturally restricted perspective on IUCs is helpful for moving the debate surrounding the possibility of IUCs away from an ‘either/or’ framing. The question should not be whether or not IUCs are *ever* possible; the question should be *when* IUCs are possible. That is, even if we grant that IUCs are sometimes possible, it is still a further question when they are possible; the former does not entail that they are always possible. This makes the debate more precise and opens up new avenues of investigation. As noted in the introduction, a key motivating reason for why we are interested in the possibility of IUCs centers on moral and political considerations: in making complex policy decisions, we may want to ensure that the losers of a proposed policy do not unduly suffer, compared to the benefits accrued by the winners. What the present argument shows is that it is plausible that it is *sometimes* possible to make these kinds of assessments – but also that this is sometimes *not* possible. We cannot rely on the possibility of IUCs in making all kinds of policy assessments, but nor should we rule out this possibility in all cases.

Second, the present account predicts that IUCs will be harder *across* cultures.²⁵ This matters, as it might help us understand the pervasiveness of some of the difficulties with cross-cultural communication. It may be true that it is not just the case that people in different cultures *express* their evaluations differently (though, as noted earlier, this may also be true); their evaluations might *in fact* differ. Again, as far as we know, this is just a possibility. However, the point here is that this possibility is at least worth investigating. Cross-cultural communication is geopolitically of major importance for many different issues. Anything that can help us better understand its difficulties – and, thereby, point us toward ways of ameliorating these difficulties – is thus important.

Third, the present account suggests that appropriate technology can help improve cross-cultural IUCs. Put differently, appropriate technology can help us determine the *transformation function* between utility scales in different cultures.²⁶ Importantly, this technology is not restricted to better neural imagining (say), but also involves the study of writing, myths and art. If people learn how intensely to evaluate the world, then anything that can help make it clearer how they do so can help us triangulate on the nature of others’ evaluative scales. In turn, this can be used in the creation of improved ways of cross-cultural communication.

Finally, the present account brings into view the question of the origin of the cultural technologies that shape our utility functions. In the above model, that origin was whichever founding family happened to be closest to the most adaptive agricultural practices. However, especially in more complex circumstances, it is possible that the cognitive technology that shapes utility functions could be provided by a multinational corporation or governmental agency. This matters, as such an entity could shape our

²⁵The individuation of cultures is a difficult question; the present account assumes that an appropriate such individuation schema is adopted. In the same way, the difference between cultures and sub-cultures is not always clear.

²⁶Note that some authors may consider the determination of these kinds of transformation functions *the* major challenge behind IUCs. If so, then the present paper shows that (a) within a culture, there is less reason to think that such a transformation function is needed, and (b) across cultures, the determination of this function *may* be possible.

utility function in problematic ways – e.g. so as to make us more or less satisfied with our current situation, which may make us more or less likely to change it (e.g. by buying more or less of certain products or be inclined to vote for or against the incumbent party) (see also Nussbaum, 2000). The model brings this clearly into view, gives us tools to assess the efficacy of this, and thus enables us to change that situation if needed.

Conclusion

It is possible to make progress concerning the question of the possibility of IUCs by considering broadly evolutionary biological insights. However, to do so, three things need to be noted. First, this question needs to be understood as concerning the intraspecific variation in evaluative structures in humans. Second, a simple appeal to evolutionary biology to resolve this question does not fit the fact that there is much variation in biological traits – especially ones that also respond to cultural pressures. Third, it is possible to answer the question by drawing on the facts that many human psychological traits evolve in response to cultural pressures and that it is possible to use technology to teach about evaluative intensities. This analysis makes clear that the possibility of IUCs is not an either/or affair, that differences in evaluative intensity may underlie cross-cultural communicative difficulties and that these difficulties may be ameliorated through the deployment of appropriate cognitive technology.

References

- Angner, E. (2018), 'What preferences really are', *Philosophy of Science*, **85**(4): 660–681. [10.1086/699193](https://doi.org/10.1086/699193)
- Arrow, K. J. (1977), 'Extended sympathy and the possibility of social choice', *The American Economic Review*, **67**: 219–225.
- Barrett, H. C. (2015), *The Shape of Thought: How Mental Adaptations Evolve*, Oxford: Oxford University Press.
- Barrett, L., R. Dunbar and J. Lycett (2002), *Human Evolutionary Psychology*, Princeton, NJ: Princeton University Press.
- Bartoshuk, L. (2014), 'The measurement of pleasure and pain', *Perspectives on Psychological Science*, **9**(1): 91–93. <http://www.jstor.org/stable/44290162>
- Binmore, K. (2009), 'Interpersonal Comparison of Utility', in H. Kincaid (ed.), *The Oxford Handbook of Philosophy of Economics*, Oxford: Oxford University Press, 540–559.
- Birch, J. (2024), *The Edge of Sentience: Risk and Precaution in Humans, Other Animals, and AI*, Oxford: Oxford University Press.
- Boyd, R. (2018), *A Different Kind of Animal: How Culture Transformed Our Species*, Princeton: Princeton University Press.
- Boyd, R. and P. Richerson (1985), *Culture and the Evolutionary Process*, Chicago: University of Chicago Press.
- Boyd, R. and P. Richerson (1987), 'The evolution of ethnic markers', *Cultural Anthropology*, **2**(1): 65–79.
- Boyd, R. and P. Richerson (2005), *The Origin and Evolution of Cultures*, Oxford: Oxford University Press.
- Broome, J. (1991), 'Utility', *Economics and Philosophy*, **7**: 1–12.
- Brown, D. E. (1991), *Human Universals*, New York: McGraw-Hill.
- Buss, D. M. (2014), *Evolutionary Psychology: The New Science of the Mind*, 5th edn, Boston: Allyn & Bacon.
- Buss, D. M. and D. P. Schmitt (1993), 'Sexual strategies theory: an evolutionary perspective on human mating', *Psychological Review*, **100**(2): 204–232.
- Bykvist, K. (2010), 'Can unstable preferences provide a stable standard of well-being', *Economics and Philosophy*, **26**: 1–26.
- Chang, R. (ed.) (1997), *Incommensurability, Incomparability, and Practical Reason*. Cambridge, MA: Harvard University Press.

- Csibra, G. and G. Gergely (2011), 'Natural pedagogy as evolutionary adaptation', *Philosophical Transactions of the Royal Society B: Biological Sciences*, **366**(1567): 1149–1157. [10.1098/rstb.2010.0319](https://doi.org/10.1098/rstb.2010.0319)
- Davidson, D. (1986), 'Judging Interpersonal Interests', in J. Elster, and A. Hylland (eds), *Foundations of Social Choice Theory*, Cambridge: Cambridge University Press, 195–211.
- Dorsey, D. (2009), 'Headaches, lives and value', *Utilitas*, **21**(1): 36. [10.1017/s095382080800335x](https://doi.org/10.1017/s095382080800335x)
- Ekman, P. (1989), 'The Argument and Evidence about Universals in Facial Expressions of Emotion', in H. Wagner, and A. Manstead (eds), *Handbook of Social Psychophysiology*, Hoboken, NJ: Wiley, 143–164.
- Ekman, P., and E. Rosenberg (eds) (1997), *What the Face Reveals: Basic and Applied Studies of Spontaneous Expression Using the Facial Action Coding System (FACS)*. Oxford: Oxford University Press.
- Fodor, J. (1989), *Psychosemantics*, Cambridge, MA: MIT Press.
- Gigerenzer, G., and R. Selten (eds) (2001), *Bounded Rationality: The Adaptive Toolbox*. Cambridge, MA: MIT Press.
- Godfrey-Smith, P. (2002), 'On the evolution of representational and interpretive capacities', *The Monist*, **85**: 50–69.
- Goldman, A. (1995), 'Simulation and interpersonal utility', *Ethics*, **105**: 709–726.
- Harsanyi, J. C. (1955), 'Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility', *The Journal of Political Economy*, **63**(4): 309–321.
- Harsanyi, J. C. (1977), *Rational Behavior and Bargaining Equilibrium in Games and Social Situations*, Cambridge: Cambridge University Press.
- Hausman, D. M. (1995), 'The impossibility of interpersonal utility comparisons', *Mind*, **104**: 473–490.
- Hausman, D. M. (2008), 'Why Look under the Hood?', in *The Philosophy of Economics*, Cambridge: Cambridge University Press, 183–187.
- Hausman, D. M. (2012), *Preference, Value, Choice, and Welfare*, Cambridge: Cambridge University Press.
- Henrich, J. (2015), *The Secret of Our Success: How Culture Is Driving Human Evolution, Domesticating Our Species, and Making Us Smarter*, Princeton, NJ: Princeton University Press.
- Heyes, C. M. (2018), *Cognitive Gadgets: The Cultural Evolution of Thinking*, Cambridge, MA: Harvard University Press.
- Huerta-Sanchez, E. et al. (2014), 'Altitude adaptation in Tibetans caused by introgression of Denisovan-like DNA [Letter]', *Nature*, **512**: 194–197. [10.1038/nature13408](https://doi.org/10.1038/nature13408). <http://www.nature.com/nature/journal/vaop/ncurrent/abs/nature13408.html#supplementary-information>
- Jevons, S. (1911), *Theory of Political Economy*, London: Macmillan.
- Kahneman, D., P. Slovic and A. Tversky (1982), *Judgment under Uncertainty: Heuristics and Biases*, Cambridge: Cambridge University Press.
- Lipton, P. (1995), 'Popper and Reliabilism', in A. O'Hear (ed.), *Karl Popper: Philosophy and Problems*, Cambridge: Cambridge University Press, 31–44.
- List, C. (2003), 'Are interpersonal comparisons of utility indeterminate?', *Erkenntnis*, **58**: 229–260.
- Loomes, G. and R. Sugden (1982), 'Regret theory: an alternative theory of choice under uncertainty', *The Economic Journal*, **92**(368): 805–824.
- Machery, E., and R. Mallon (2010), 'Evolution of Morality', in *The Moral Psychology Handbook*, New York: Oxford University Press, 3–46. [10.1093/acprof:oso/9780199582143.003.0002](https://doi.org/10.1093/acprof:oso/9780199582143.003.0002)
- Matsumoto, D. R. (2001), 'Culture and Emotion', in D. R. Matsumoto (ed.), *The Handbook of Culture & Psychology*, New York: Oxford University Press, 171–194.
- Mikhail, J. (2011), *Elements of Moral Cognition: Rawls' Linguistic Analogy and the Cognitive Science of Moral and Legal Judgment*, Cambridge: Cambridge University Press.
- Millikan, R. (1984), *Language, Thought, and Other Biological Categories*, Cambridge, MA: MIT Press.
- Millikan, R. (2002), *Varieties of Meaning*, Cambridge, MA: MIT Press.
- Narens, L. and B. Skyrms (2021), *The Pursuit of Happiness: Philosophical and Psychological Foundations of Utility*, Oxford: Oxford University Press.
- National Institutes of Health. (2007), *Understanding Human Genetic Variation*, Bethesda, MD, <https://www.ncbi.nlm.nih.gov/books/NBK20363/> (accessed on May 29, 2025).
- Nettle, D., M. A. Gibson, D. W. Lawson and R. Sear (2013), 'Human behavioral ecology: current research and future prospects', *Behavioral Ecology*, **24**(5): 1031–1040. [10.1093/beheco/ars222](https://doi.org/10.1093/beheco/ars222)
- Nichols, S., and S. Stich (2003), *Mindreading: An Integrated Account of Pretence, Self-Awareness, and Understanding Other Minds*, vol. 114, New York: Oxford University Press.
- Nussbaum, M. C. (2000), *Women and Human Development*, Cambridge: Cambridge University Press.

- Papineau, D. (1987), *Reality and Representation*, Oxford: Blackwell.
- Papineau, D. (2003), *The Roots of Reason*, Oxford: Oxford University Press.
- Rossi, M. (2014), 'Simulation theory and interpersonal utility comparisons reconsidered', *Synthese*, **191**(6): 1185–1210. <http://www.jstor.org/stable/24021575>
- Sauter, D. A., F. Eisner, P. Ekman and S. K. Scott (2010), 'Cross-cultural recognition of basic emotions through nonverbal emotional vocalizations', *Proceedings of the National Academy of Sciences*, **107**(6): 2408–2412. [10.1073/pnas.0908239106](https://doi.org/10.1073/pnas.0908239106)
- Schulz, A. (2020), 'Enhancing thoughts: culture, technology, and the evolution of human cognitive uniqueness', *Mind & Language*, **37**: 465–484. [10.1111/mila.12320](https://doi.org/10.1111/mila.12320)
- Schulz, A. (2021), 'Explaining human diversity: the need to balance fit and complexity', *Review of Philosophy and Psychology*, **14**(2): 457–475. [10.1007/s13164-021-00582-1](https://doi.org/10.1007/s13164-021-00582-1)
- Schulz, A. (forthcoming), 'The origins of distinctively human mindreading: a bio-social-technological coevolutionary account', *British Journal for the Philosophy of Science*.
- Scollon, C. N., S. Koh and E. W. M. Au (2011), 'Cultural differences in the subjective experience of emotion: when and why they occur', *Social and Personality Psychology Compass*, **5**(11): 853–864. [10.1111/j.1751-9004.2011.00391.x](https://doi.org/10.1111/j.1751-9004.2011.00391.x)
- Sen, A. (1979), 'Interpersonal comparisons of welfare', in M. J. Boskin (ed.), *Economics and Human Welfare: Essays in Honor of Tibor Scitovsky*, New York: Academic Press, 183–201.
- Shackelford, T. K., D. P. Schmitt and D. M. Buss (2005), 'Universal dimensions of human mate preferences', *Personality and Individual Differences*, **39**(2): 447–458. [10.1016/j.paid.2005.01.023](https://doi.org/10.1016/j.paid.2005.01.023)
- Sorensen, R. (1998), 'Self-strengthening empathy', *Philosophy & Phenomenological Research*, **58**: 75–79.
- Sterelny, K. (2003), *Thought in a Hostile World: The Evolution of Human Cognition*, Oxford: Wiley-Blackwell.
- Sterelny, K. (2012), *The Evolved Apprentice: How Evolution Made Humans Unique*, Cambridge, MA: MIT Press.
- Sterelny, K. (2021), *The Pleistocene Social Contract: Culture and Cooperation in Human Evolution*, Oxford: Oxford University Press.
- Tomasello, M. (1999), *The Cultural Origins of Human Cognition*, Cambridge, MA: Harvard University Press.
- Tomasello, M. (2021), *Becoming Human: A Theory of Ontogeny*, Cambridge, MA: Harvard University Press.
- Tooby, J., and L. Cosmides (1992), 'The Psychological Foundations of Culture', in J. Barkow, L. Cosmides and J. Tooby (eds), *The Adapted Mind*, Oxford: Oxford University Press, 19–136.
- Waldner, I. (1972), 'The empirical meaningfulness of interpersonal utility comparisons', *The Journal of Philosophy*, **4**: 87–103.
- Ward, Z. B. (2022), 'Cognitive variation: the philosophical landscape', *Philosophy Compass*, **17**(10): e12882. [10.1111/phc3.12882](https://doi.org/10.1111/phc3.12882)
- Westra, E. and K. Andrews (2022), 'A pluralistic framework for the psychology of norms', *Biology & Philosophy*, **37**(5): 1–30.
- Westra, E., S. Fitzpatrick, S. F. Brosnan, T. Gruber, C. Hobaiter, L. M. Hopper, D. Kelly, C. Krupenye, L. V. Luncz, J. Theriault and K. Andrews (2024), 'In search of animal normativity: a framework for studying social norms in non-human animals', *Biological Reviews of the Cambridge Philosophical Society*, **99**(3): 1058–1074. [10.1111/brv.13056](https://doi.org/10.1111/brv.13056)
- Whiten, A., and R. W. Byrne (eds) (1997), *Machiavellian Intelligence II: Extensions and Evaluations*, Cambridge: Cambridge University Press.
- Zawidzki, T. W. (2013), *Mindshaping: A New Framework for Understanding Human Social Cognition*, Cambridge, MA: MIT Press.