Explaining Artifact Evolution

David Kirsh
Dept of Cognitive Science, UCSD
kirsh@ucsd.edu

Introduction: Interpreting an Artifact's Cognitive Life

Much of a culture's history – its knowledge, capacity, style, and mode of material engagement – is encoded and transmitted in its artifacts. Artifacts crystallize practice; they are a type of meme reservoir that people interpret though interaction. So, in a sense, artifacts transmit cognition; they help to transmit practice across generations, shaping the ways people engage and encounter their world. So runs one argument.

Is it plausible? Imagine a wooden flute beached on a desert island. Might it yield a few of its secrets to a dedicated inhabitant with time on his hands? Like a stone axe crafted to afford certain grasps more than others, a flute encourages certain modes of use, certain practices, making others difficult, unnatural. A flute has been designed for specific sorts of interaction. When effective it helps to shape the type of interactive cognition that is characteristic of flute playing. To master a musical instrument completely, of course, requires appropriating the repertoire and culture of that instrument – something not possible in isolation. But the basics of flute engagement – how to make enduring, melodious sounds – are driven by the flute's physical design rather than its specific culture of use. The design of the flute helps to constitute the practices that it supports, generation after generation. This imbues it with a certain sort of cognitive life, something that lives on after its primary users die and which can shape future cognition.

The same might be said of books. A seminal book of mathematics retains the potential to shape the thinking of new readers, once they can interpret it. In the case of Euclid and the Greek philosophers, books can reshape entire generations of thinkers, as they did in Spain and Italy during the Renaissance. They encode knowledge, method, and style. Is it wrong to say they have a 'cognitive' structure that shapes and interacts with human cognitive processes – a form of cognitive life of their own?

Archaeologists interested in this dialectic between artifact and community of users naturally want to understand how artifacts preserve and bias cognition across generations – those aspects of practice that artifacts freeze in their form and affordance. But to infer the function of an artifact and the reasons a culture has sustained its use is an interpretive task of immense difficulty, given how complex the cognitive, social and physical ecology of a culture is.

The problem is the familiar one of underdetermination. Because task and activity are co-constituted by artifact and practice, it is not possible to determine the task an artifact was designed to facilitate without also determining the practices that fashioned the way it was used. Sometimes this is unproblematic. It is easy to infer the gross function of a blade-like implement, since every culture requires cutting, and it is not possible for humans to

Kirsh, D. (2006). Explaining Artifact Evolution. To appear in, Cognitive Life of Things: Recasting the Boundaries of the Mind. Malafouris, L. (Ed.). McDonald Institute for Archaeological Research, Cambridge, April 7-9, 2009. Cambridge: Cambridge University press.

cut without a cutting implement. But what is the function of a dowel, an amulet, a carved shell? What tasks were they designed to support? What practices developed to make effective use of them? What larger activities were these tasks part of?

My concern in this paper is to understand why a particular artifact has the design it has. Insight into this problem ought to allow more evidence to be extracted from the archaeological record. What are the forces that shape how artifacts evolve? The better our understanding of these dynamics the better should be our inference of an artifact's role in the tasks, activities, and practices of its users, all aspects of human cognitive life.

The argument I shall explore as to why more evidence can be extracted from the archaeological record starts by assuming that users, artifacts, practices and tasks have co-evolved. Artifacts typically arise to address a need connected with a task. Task needs prompt creative humans to modify their technology. Once new artifacts are built, however, people are able to address their tasks in different ways. Practices evolve as people appropriate new artifacts and discover their latent possibilities. Users develop new skills, which in turn drive the demand for yet better tools. The result is that tasks, practices, users and artifacts form a tightly coupled co-evolving system.

I expect broad agreement on the interdependence of tasks, practices, users and artifacts, though not perhaps on the details I offer. The real question I explore is whether the evolution of artifacts can be explained on analogy with natural selection. Is it plausible to suppose that the artifacts that dominate the archaeological record are better (in some independently specifiable way) than other culturally available, but less prevalent, artifacts? If so then the reason certain artifacts spread more quickly and broadly in a culture is precisely because they are more effective, and culturally superior. This view would support a functionalist view of the archaeological record: the more prevalent an artifact's design, the more adaptive, fit and effective it must be. Taken to its extreme this adaptationist assumption suggests that artifacts that are widely used will be optimally or near optimally designed. The equivalent claim made about the design of biological organisms is a popular, though frequently challenged, assumption in biological evolution. If it is true for artifacts it has major consequences: archaeologists will be able to use prevalence in the archaeological record to constrain their interpretative problem.

Part A: Clarifying the Argument for Optimal Design

Bounding the adaptationist project

The assumption of artifact optimality is a powerful hypothesis. Even if true, however, this approach to understanding artifactual change cannot be applied to all artifacts and in all circumstances. Specific artifacts are often valued for reasons other than their task effectiveness. Functionality is a powerful force for change; but it is not the only force. And it carries more weight for tools and utilitarian artifacts than for non-utilitarian ones.

For instance, a noble person's crown is not valued for its functional powers, literally construed. It has symbolic value. It is a unique item with history and meaning; often it is irreplaceable. Accordingly, it has a social significance and meaning that goes beyond the pragmatic. This general idea applies to other artifacts. A father's sword may be valued for its sentimental value, its history in the family, the wear marks that show the violence of battle or the marks of success or failure. Gifts and bequeaths tend to have special meaning. And, as we all know, an artifact may be valued for its beauty, the costliness of its material, its grace, or the delight we take in its use. All these aspects of value are non-utilitarian. They contribute to a person's experience with an artifact and a source of its worth. Such artifacts may dominate the record because care was taken with their preservation. But they need not be optimally designed, nor widely duplicated in a population despite their representation in the record.

In our own era, artifacts tend to be mass-produced, more like a commodity than a crafted work fashioned by a skilled artisan. They are non unique. This bodes well for a functional approach to artifacts. The more perfect a market economy the more powerful are the forces driving product diversity and selection for design quality. Better designed artifacts should predominate. But not always. In post-industrialist societies self-expression is a significant force driving preference. The urge to use artifacts to project a social identity, clearly figures when someone chooses jewelry, clothing, or household knickknacks. It can also be a factor in the choice of tools, which are the most functional of artifacts. By owning or using a particular kind of tool a person represents him or herself to be a certain kind of person – a person who uses the best, or the most valued sort of tools, the same sort as used by the royal carpenters, or acknowledged masters. People often feel that a tool's brand reputation confers prestige on its owner, as if the owner inherits attributes from renowned users of the tool. The modern conception of a brand rests on this fact.

The more powerful the non-pragmatic aspects of an artifact are, the more difficult it is to interpret its evolution in the language of optimization. If brands or artifact types are valued because they are reputed to be the best or the most cost effective then the language of optimality may well be appropriate. But when the value derived from an artifact is personal or idiosyncratic, or when an artifact is unique or highly customized, there are few generalizations about its evolution to be found. There are too few instances to test hypotheses as individuals may vary so widely in the way they use those idiosyncratic artifacts. Because science relies on generalizations, looking at single artifacts will only be informative if those objects can be assumed to be instances of a more widely distributed kind. Only when this assumption is justified, is it reasonable to take as a working hypothesis that those artifacts are well designed for the tasks they were meant to support. Testing this working hypothesis is the challenge I accept here. Can we explain how key functional artifacts evolve by assuming they are optimally near optimally designed?

Interpreting the archaeological record by assuming optimality

The argument for optimality in artifactual design is analogous to the argument for optimality in biological design, especially as developed in the modern theory of evolutionary ecology. The initial argument I present has three steps. First, I define what I mean by an artifactual ecology. Because artifacts and organisms have different environments it is important to ensure that our terms are defined and that the analogy is clear. I then explain how an artifact's fitness – its goodness – can be measured independently of its prevalence. This is necessary because fitness or 'good design' must be measured separately from prevalence otherwise prevalence entails good design and the hypothesis of design optimality, or near optimality, would be unfalsifiable. I then restate in more precise manner a version of the adaptationist position that is thought to warrant treating functional artifacts as optimally designed.

1. Defining an **artifact ecology**. Artifacts and their users form a complex ecology where tasks, practices, artifact collections and users co-adapt and co-evolve. The elements of this ecology are:
    a. **Artifact** species. Such things as axes, chisels, trousers, pots and baskets or, in recent times, pens, ovens, TV's and so on. These artifacts compete with and complement other artifacts. It is the evolution of these species or perhaps artifact sets (knife, fork, spoon) that we seek to explain.
    b. **Artifact systems** or collections**.** Only occasionally are artifacts expected to stand on their own. More often a given artifact is part of a constellation of artifacts that mutually support, constrain, and complement each other. A frying pan without a spatula, a screw without a screwdriver, a salad bowl without serving spoons, are artifacts missing their mates. These dependencies can range over dozens of artifacts, as in factories, jigs, and so on.
    c. **User** groups. These groups are sub-populations in a culture that perform tasks and activities, and choose artifacts to use: craftsmen, warriors, blacksmiths, cooks, and householders. They select and keep artifacts. It is their material that is prevalent in the archaeological record. [1]
    d. **Practices** are the techniques and methods user groups have evolved to work with artifacts and accomplish tasks. How do cooks use their pots and stirring instruments? Without practices artifacts are inert, lifeless. Practices explain how users put artifacts to work to get things done.
    e. **Task** environments are the underlying constraints determining what is helpful and unhelpful in accomplishing a task. When a practitioner uses an artifact or collection of artifacts to work toward completing a task it is the structure of the task – the task environment – that determines how successful each artifact-using step is in bringing the practitioner closer to the goal. At an intuitive level, because any artifact may be used to perform several different tasks, the niche of a given artifact consists of a collection of task environments along with all competing and complementary artifacts the practitioner uses in those tasks.

    These five elements – artifact species, artifact systems, user groups, practices and task environments – mutually constrain each other, constituting an artifactual ecosystem, a system of interdependent artifactual niches. Jointly they determine how artifacts are used, when they are used, for what, by whom and how effectively.

[1] Traders also figure in this group when they preselect the artifacts they sell based on their knowledge of their clients' likes. If there are enough traders, however, it may be better to see them as market makers - agents who match buyers and sellers.

2. To determine how **well designed** an artifact is two measures must be distinguished: a theoretical measure based on task performance, and an empirical measure based on prevalence.
    a. **Theoretical Goodness**: when an artifact is used correctly, and it is well designed, it ought to lead to more successful performance, other things being equal, than less well designed artifacts. Measuring effectiveness requires a theory of tasks and metrics for evaluating how well a user performs a task. From a theoretical point of view the 'best' artifact need not be the most popular. It must be the most effective among a set of possible artifacts that a society could have created given its technology and capacity for production. Hence **optimal design** is a theoretical notion concerning the connection between users, tasks, practices, production possibilities and artifacts.
    b. **Empirical goodness**: when an artifact is prevalent in a population we assume that at some time it was popular in the marketplace. Users vote by acquiring artifacts; they reveal their preference by their possessions. The greater the **prevalence** of an artifact **(**or market share if the era supports a market economy) the better users think it is (in some sense), at least when compared with other comparable artifacts in the marketplace. Prevalence should be reflected in the archaeological record, and the history of artifact trends.
3. **Adaptationist assumption:** once theoretical and empirical goodness are independently defined we can state the adaptationist or optimality assumption: prevalence implies optimality, empirical goodness correlates with theoretical goodness. This is a precise statement of the thesis explored here. The more prevalent an artifact is the more adaptive, fit, and effective it ought to be at helping users perform their tasks. If true, archaeologists can use cultural prevalence to help solve their simultaneous problem of inferring the tasks, activities and practices of user groups. Popularity will be an indicator of optimality, thus making it possible to break into the hermeneutic (ecological) circle involved in interpretation.

It is not hard to see how an adaptationist approach would work in practice. Every time archaeologists unearth an artifact they ask 'what tasks would this be well designed for?' 'What practices might have emerged to put this artifact to good use?' These questions cannot be asked in isolation of the record because artifacts evolve from previous artifacts, in the context of other artifacts, specific users and cultural practices, including production techniques for making artifacts. But presumably the evolutionary problem for artifacts is no different than it is for organisms that also must reproduce in a pre-existing ecology. The niche for an organism is the biological ecology it is part of, and evolution has driven a sequence of incremental improvements, speciation and hybridization, generation after generation. The niche for an artifact is the socio-technical ecology of tasks, users, practices and other artifacts it interacts with.

Figure 1. The speciation of chisels both within and across cultures in some ways resembles the speciation of beetles. Both come in a variety of sizes, with differences in appearance and comparative advantage. Both honor the principle that there be 'enough' distance between species. All share a family resemblance.

My inquiry is structured as follows. In Part B I elaborate the nature of optimality arguments prevalent in modern ethology, using optimal foraging theory as an example. I then adapt that model to the artifactual world and show how we might explain the evolution of an artifact. I use as my case study a modern commodity, televisions, but the theoretical account is representative of the general adaptationist view applied to artifacts more generally.

Part C – the bulk of the paper and the part most directly involved in understanding the way artifacts saturate culture – is devoted to discussing the many difficulties and complexities involved in adapting an evolutionary model to human artifacts. I begin with problems related to an artifact's niche. If artifacts, tasks, practices and users interact and co-evolve, why assume they are ever in *stable* equilibrium? The tasks a community wants to perform with their artifacts are constantly changed by those artifacts, as are their practices and skills. Nothing is ever stationary long enough to become optimal. Other problems arise on both the generation and selection side of artifact evolution. For instance, artifacts do not change through random processes or semi-random sexual selection; they are generated as a result of human designers making intelligent choices. This limits and biases design variation, possibly limiting the diversity of designs so greatly that optimality cannot be defended. Meanwhile on the selection side there are deep questions about cultural differences in design sensibility. To prevent the optimality hypothesis from being unfalsifiable we assume that it is possible to give an independent measure for the goodness of an artifact's design. If there is no culturally neutral standard for evaluating optimal design why assume there is an objective measure for determining whether one artifact is better than another? Another challenge to optimality stems from the holism of artifact systems – the way the effectiveness of one artifact depends on many others. How is it possible to determine whether an artifact is well designed itself or is free-riding on the effectiveness of the system of artifacts it is part of. Isn't there an unsolvable problem of credit assignment?

These are huge problems but not, in general, fundamentally different in artifact evolution than in organism evolution. Archaeology and ecology both face the same complications when trying to show that current design is optimal design. The outcome is by no means certain in ecology and if it cannot be defended for biological ecologies it will certainly not fly for artifactual ecologies. But these concerns aside there remains a further and distinctive challenge for archaeology. The artifacts that dominate the archaeological record, especially the record of pre-market societies, are artisanal, few in number and not subject to the forces of large markets that reward diversity and choice. In the world of artifacts, markets run proxy for natural selection. Without markets to ensure diversity and reward artifact quality, however, there is no guarantee that there will be sufficient variation of design and sufficiently many users to drive optimization. I think this failure is insurmountable and ultimately makes the assumption of optimality of limited value for ancient objects. But it remains a live question whether it explains post-artisanal artifacts. And for a discussion of the cognitive life of things it highlights just how closely artifacts are tied to human practices, skills and tasks, the non discursive side of cognition.

Part B: The Case for Optimality in Artifact Design

The Optimality Assumption as used in Optimal Foraging Theory

The viability of this approach depends on making sense of the optimality assumption. To say that design A is better for certain tasks than design B implies there is some metric for comparing design efficiency or effectiveness. Often this is a performance metric, such as the speed or accuracy shown by a representative user when completing a task, or the amount of energy it takes to do a task, or the cognitive effort involved in staying in control, or the utilization rate of some other resource.

In biology, a nice example is found in optimal foraging theory where the metric is expected energy surplus. It is assumed that an organism has optimal, or near optimal behavioral strategies for its niche, given predation probabilities and the metabolic costs and benefits of foraging, searching and sleeping. Optimality here means that strategies yield the greatest chance of the organism maintaining a surplus of metabolic energy. See table 1. If it is known how much energy a creature, such as a chameleon, needs for metabolism when it eats, wanders from its nest, runs away from predators, versus the energy it needs when it rests or sleeps, and the energy reward it gets from eating bugs, the optimal foraging assumption allows an ethologist to come up with an initial guess at how much time that chameleon should spend foraging, searching, sleeping etc. and how far from its nest it should search. To make that estimate, the ethologist must observe the kind of bugs that are available, how much energy they contain, how plentiful they are, how quickly they replenish, how hot and humid it is, how prevalent the different sorts of predators are, and their behavior. Given this characterization of a chameleon's niche, strategies can be ranked by their cost benefit profile. Because selection should reward chameleons with better behavioral strategies, the best ones should proliferate and those with less optimal strategies should die out. The result is that if conditions have been stable for a while the majority of chameleons should behave optimally, or near optimally, relative to their niche.

Optimal Foraging Theory applied to Chameleon			
i	**Behavior/task**	**Benefit**	**Cost**
1	Foraging and Eating	Food (metabolic value)	Predation (medium), Metabolic (high)
2	Searching	Change in food probability	Predation (medium), Metabolic (high)
3	Sleeping	Regeneration	Predation (low), Metabolic (low)
4	Run from predators	Lower probability of predation	Predation (very high), Metabolic (v high)
5	Thermo-regulation	Lower body temp	More sluggish

$$\text{Determine weights } (w_1, \quad w_2, \quad \ldots \quad w_n) \text{ to maximize } \sum_{i=1}^{n} w_i b_i = B_i - C_i$$

Table 1 shows the costs and benefits of different behaviors $\mathbf{b_I}$ for the Costa Rican chameleon. By meticulously observing the environment the chameleon inhabits ecologists build up a chart of the probability of predation during different behaviors, the average metabolic costs of different behaviors, the nutritional value of different bugs and so on. It is then possible to model different behavioral strategies consisting of a weighting of behaviors $\boldsymbol{w_I b_I}$ to determine an optimal weight set where the difference between expected benefits **B** and expected costs **C** is maximal.

In optimal foraging theory, the assumption of optimality serves two functions. It allows predicting the behavioral strategies that a well-adapted creature should have. And it helps indicate the things an ecologist should look for in behavior and niche. This follows because the assumption of optimality implies that well-adapted creatures are in equilibrium with their niche. This is a strong assumption. If it *seems* that a creature is not in *equilibrium* with its niche, either niche or behavior must be incorrectly specified. For instance, if there seems to be a large difference between behavior as observed and behavior deemed optimal, given the current account of the costs and benefits of different behaviors, then it is likely that either observation of behavior or observation of niche – which forms the basis of cost benefit assessments – is incorrect. This assumption may be incorrect, at times. For instance, the current environment of a creature may have suddenly changed from parched to dry, due to climatic shifts. Its behavior, then, would no longer be in equilibrium with current conditions. The creature was well adapted to how things were – when it was extremely dry – before the climatic change. Such cases are assumed to be rare. But when they are true the assumption of equilibrium is incorrect and inferences based on optimality are incorrect.

Despite the possibility of being in a non-equilibrium state, ecologists and arguably anthropologists, normally operate under the assumption of equilibrium. Thus, returning to the chameleon, observations of a chameleon's current behavioral strategies can be used as a system of constraints for biasing the ecologist's interpretation of its niche. Assumptions about niche serve to bias the interpretation of why a creature acts as it does, and assumptions about behavioral strategy bias the interpretation of how a niche helps to shape behavior. This apparent circularity is not vicious as long as the assumption of optimality serves as a methodological directive – an assumption that is not falsified by observations singly but can be rejected if data accumulates and the assumption becomes counterproductive over the long run. Accordingly, until shown to be misguided in a particular case, the optimality assumption can be used to extract more information from data about both behavior and environment. It adds a constraint to the norms of when an interpretation of behavior or niche is adequate.

Clearly, optimality is an *interpretive strategy*, much like the principle of charity, rationality, or benefit of the doubt made by philosophers and economists to justify inferences based on sparse data. Optimization models are rampant in most fields of science. They are used in the study of motor control (optimal sensori-motor control theory, Todorov[2006]), in evolutionary theory (the tradition of D'Arcy Thompson [1961]), and the current adaptationist paradigm in biology, (see claims by Gould & Lewontin [1979] on their prevalence), in linguistics (harmony theory, Smolensky & Legendre [2006]), in most economic modeling , and almost universally in the natural sciences. In the sciences of the artificial, optimization, or rather constrained optimization – satisficing – is the normative goal of design, Simon [1969]. By making the assumption of optimality theorists are able to draw strong conclusions from weak data.

Can we assume that the trajectory of designs an artifact describes as it changes over time is somehow optimal or near optimal given the uses to which it is put and the existing constraints on production that bound the possible designs culturally available?

A Modern Example

To get an idea of why someone interested in design ought to entertain the optimality principle consider the design context in the early days of television design, before remote controls were invented. Televisions are obviously different than axes and simple tools in basic ways. Their technology is more diverse, production methods are post-industrial, and perhaps most significantly, they are mass-produced and distributed through a highly evolved market to millions of users. Still, the forces of adaptation may apply to axes and televisions in similar ways. How might we explain why early televisions evolved to become systems with remote control and picture in picture?[2]

The optimality principle suggests that users of televisions – *consumers* to use economic terminology – try to optimize a set of value dimensions, such as cost and size, picture quality, and ease of use. Further dimensions like appearance, snob or brand appeal, and marketing factors like the number and appearance of knobs, are additional dimensions that go beyond the functional, though arguably many elements of appearance have functional import – look and feel matter to performance. See Table 2. Because consumers can be assumed to select televisions that deliver the highest value, along the dimensions that matter most to them, the population of televisions found in a given community or culture, at any time, should always be near optimal, given what consumers value. As new TV's come on the market, old and less good ones, are gradually thrown out. So, at any moment, the extant TV's will reflect a distribution of valued features, more or less close to the optimal distribution, given community values.

The optimality principle is based on the idealization that with enough time, information and computational power, every consumer makes the best choice, given their ranking of

[2] Readers uninterested in the details of how to evaluate the optimality of an artifact may skip this section.

the dimensions they care about. This broad idea is enshrined in the economic theory of consumer choice. [3]

Major Value Dimensions for TV

i	Value Dimension	Weighting	Value
1	Price		
2	Size		
3	Picture Quality		
4	Appearance/look/feel		
5	Ease of Use		
6	Range of Features		
7	Brand appeal		
8	Number of Buttons		

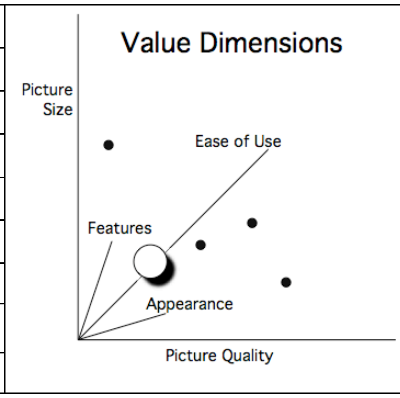

$$\sum_{i=1}^{n} w_i v_i = overall\_value$$

Table 2. In order to make a rational purchase consumers are assumed to evaluate the goodness of each possible purchase by assigning a value to each of their value dimensions, weighting the importance of each, and then computing the sum. For a given price they should always choose the TV that has the highest weighted value – that is, the longest Euclidean distance from the origin. According to classical economic theory their decision concerning how much to spend on an arbitrary good depends on what else they could do with their money. Rational economic agents will spend their next dollar on whatever yields the highest utility for that dollar. This is reflected in the overall value of a TV at a given prince.

Design enters the equation because the better an artifact is designed the more highly it should place on one or more of the dimensions that consumers value. Design determines how easy an artifact is to use, the sort of features that it is given, how it looks, and so on. If buyers desire these features and design values, as we typically assume they do, it is because designers have been successful in identifying what consumers want and in building artifacts that meet those wants. Over time artifacts with better designs should proliferate and push out those with less successful designs.

The optimality principle is a strong assumption that will justify inferences about both the design of artifacts and user values found in a population. But as it stands it is so strong that it implies that the artifacts that are most wanted – the ones that sell best – are *necessarily* the best. Popular artifacts cannot be badly designed. Because most of us believe this to be false, optimality is too strong an assumption. There must be market dominant artifacts that occasionally are not the best of available designs.

The way to deal with this sad truth of civilized life is to accept that optimality in design is better thought of as near optimality, or even better, that optimality reflects a long-term

---

[3] A more recent version of this theory, developed by Lancaster, and sometimes called the new theory of consumer choice allows these dimensions to be made explicit and be the main reason consumers' demand curves are not always smooth and continuous.

equilibrium that can be misleading if applied in the short run. There can be periods – transients – where the prevalent design is a sub-optimal design. Owing to lags and imperfections in the market, products that are well designed might never dominate. At least in the short run. But eventually the market rights itself and good design wins out. This assumption is necessary to save the adaptationist and economic methodology. Reject this assumption and you reject the governing principle of modern economics and evolutionary biology.

Returning to the activity space of television use, the adaptationist principle should explain the long-term dominance of remote control and to a lesser extent the dominance of picture in picture, by examining how both of these features are somehow advantageous to users given the obvious cognitive and physical tasks involved in watching TV.

How might such an analysis of TV's proceed? The major determinants of consumer choice, along with price, picture quality and appearance, are interactive features such as turning the TV on and off, changing channels and changing volume. See table 3. These factors affect the experience of watching television, which we assume is the ultimate source of value. Other costs and determinants of experience, such as those associated with discovering the time and channel of programs, or with controlling audio-video elements such as muting, or color intensity, can be treated as being of secondary importance. Improvements that reduce the costs of these activities, or which increase the benefits, should, accordingly, translate into modest preferences for one model over another. This is an idealization. Marketing and advertising can often manipulate consumer perceptions of ease of use, and desirability, leading consumers to choose TV's that are in fact less easy to use or ones which deliver a less desirable experience on 'objective' grounds than others present in the marketplace. But we shall bracket these concerns right now as further cases of short-term perturbations, and in any case, marketing would not be a factor driving preferences for artifacts in ancient civilizations.

The Major TV Tasks

Task	**Cost pre-remote**	**With remote**	**With PIP**
On/Off	Distance to TV	Click	Click
Change Channels	Distance to TV	Click	Click
Change Volume	Distance to TV	Click	Click
Surfing - Bi-channel - Multi-channel	○ Distance to TV ○ Cost of occluding TV ○ Discomfort of not sitting	○ Click to other channel, ○ Recall of original channel, ○ Constant checking for end of advertisement	○ Two channels for free ○ More channels require clicking but no need to recall
Multi-watch	Cost of sitting in front of TV (see above)	Cost of clicking and recalling	○ Two channels for free

Table 3. The major tasks involved in watching TV and their cognitive and pragmatic costs. The goal of good design is to lower costs. The more frequently a task is performed, the more it should be weighted when determining the overall value of a design. Here we see how adding features such as picture in picture, and remote and control, has an impact on cost structure. As with cost benefit models in optimal foraging, an expert or archaeologist may determine the product that is best for users, and so should be most popular.

To evaluate the ecology of televisions and television use in a quantitative way, an expert must put a price on the experiences a user has when performing the tasks listed in table 3 (as well as other drivers of experience such as picture quality). To determine these values in a scientific manner requires going to the homes of TV users and measuring environmental conditions, user goals, and the frequency of different behaviors.

The methodology for studying a television's adaptiveness is the same as that used to study a chameleon's. How far is the TV from where viewers sit? How fast can the average viewer move? How many channels are there? How often do viewers change channels? How often do they mute the sound, change volume, and turn the TV on and off? The greater the physical distance, or slower the viewer, or the more channels the viewer watches, the greater the average cost of channel changing.

Unless things are terribly skewed the objective of design should be to lower the cost of one or more of these component activities in watching TV. This motivation for invention and improvement has been well studied (see for instance, Petrovski [1992]). The more frequently a component activity or task is performed, the more important it is to lower the cost of that activity. Competing designs can then be evaluated by how much they lower the sum of weighted costs.

For example, a designer might consider moving the couch closer to the TV to lower the time and effort of reaching the TV, or add random access to the channel dial in order to save having to twist the dial through dozens of channels. Since a design change must be effective in different settings, the solution must be one that works given the variety of watching distances, viewer speeds, viewer lethargy, and so on. On this basis, moving a couch is not a good solution. It is too specific, and a feature of the environment that is not reliably under the control of the designer. A more universal and clearly better solution is to provide a hand sized controller that allows a viewer to manage the TV from anywhere in the room. This clearly improves the cost benefit profile across a variety of environments.

The focus of this analysis is on tasks, or as activity theorists would say 'operations' – the key elements of activity. This means that to determine the value of different devices in accordance with the value dimensions shown in table 2 we need a further diagram mapping devices onto task cost. For each task we need a cost metric that shows how good or bad a given design is at helping the user to perform that task. The cost structure of the activity associated with a particular design is found by looking at the region the device occupies on each task dimension. That is, each competing device can be represented as a region in a multi-dimensional task space. The 'best' TV is the one with the lowest overall cost, given the weighting of each task. See figure 2.

Thus, a TV without remote control lies far out on the dimensions associated with changing channels, changing volumes, and other tasks that may have to be done frequently. Setup tasks that are done once and for all are not much more costly to do without remotes. Alternative design candidates that might compete with the type of remotes we currently use are cell phones, Ipods, or computers. And alternative input devices might be ones based on using voice control, gesture, or a laser pointer at a menu, rather than pressing buttons, as we do today. The end result is a plotting of TV designs on a weighted task space that tells us the best TV design.

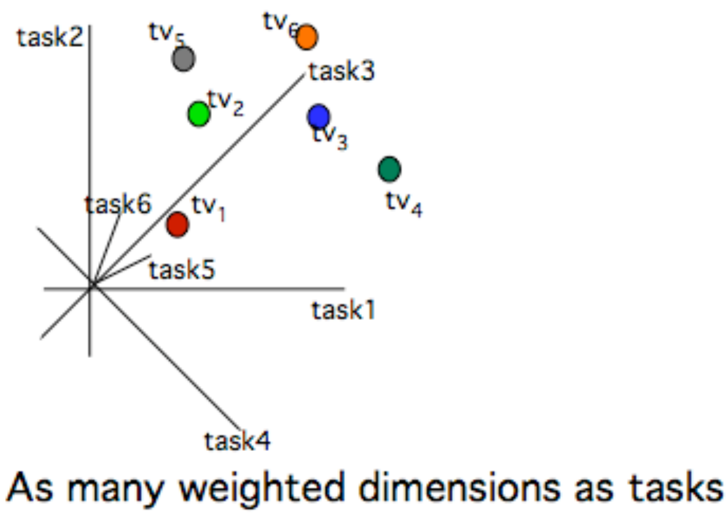

Figure 2. We can represent the cost structure of an activity space in a few ways. The summary view, shown here, displays each task that a user might want to perform on or with a TV as a separate dimension. In principle there could be n dimensions or tasks. Different TV's can then be plotted in this n-dimensional space according to the cost of performing each task with that TV; that is, each TV is represented by an n-dimensional cost vector. To determine the cost of performing a task with a specific TV, however, it is necessary to derive the cost from a collection of dimensions, not explicitly represented here, such as the speed accuracy for the task, the error-cost vs. frequency table, learnability and so on. See Figure 3. Each of these task evaluation dimensions must be assigned a relative weight and then projected onto a single number, which represents a measure of the overall cost of that TV for that one task. This value must then be further weighted to reflect how important that task is to a given user. The best overall TV is the one with the best cost given its weighted costs on all tasks.

As promising as this method is, it is not the whole story. The evaluation of a device is more complex than plotting its position in a single multi-dimensional task space. The space we looked at, so far, is based on speed accuracy observations: how fast can a viewer perform a task, such as changing channels without making unacceptably many errors. But there are other metrics we may care about.

For instance, a given device tends to give rise to a distribution of errors of varying severity. We can therefore ask about the variance of errors that a typical user will exhibit – the reliability of the device. We can plot this variance on a graph that shows how often different types of errors occur, and how bad they are – the risk profile of the device. We can ask how easy it is for a user to recover from errors, and how long it usually takes – the robustness of the device. And we can ask how easy it is to master a device – the learnability of the device – since the cost structure of a device under almost all measures changes with expertise. There are others, such as what is the most complex task a user can readily perform. See figure 3. All these are measures that can determine the overall cost-benefit structure of an artifact.

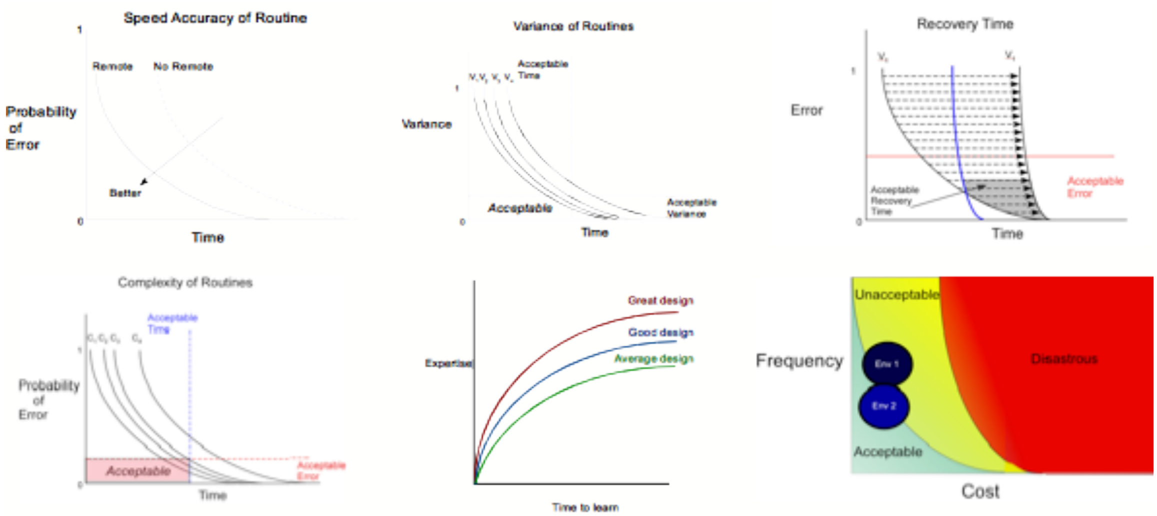

Figure 3. Several scales are useful in determining the cost structure of an activity. Clockwise from the lower left they are: 1) Generic speed accuracy curves 2) Variance in output is the hallmark of quality control and shows us how reliable an artifact is. Viewed as a bull's eye target the goal of an artifact is to reduce the size of errors to the point of being negligible. When this happens users may still err but the consequences are trifling. 3) How long does it take a user to recover once they have made an error? This is a key measure that is rarely considered in determining the costs of use. 4) Speed accuracy vs. routine complexity: how fast routines of different complexity can be performed (these are speed accuracy curves). They show us in some sense how effective an artifact is at facilitating a task. Only certain artifacts enable users to perform complex tasks fast enough to be usable for them. 5). Learning curves tell us how easy it is to master a device. The most effective tools in the world may be unusable for most people because they are too hard to learn. 6) Error frequency vs. cost is complementary to variance curves but now an explicit estimate of cost is included. Variance might be plotted in terms of seriousness of errors but it does not tell us how bad it is if we commit those errors. An otherwise perfect car that blows up every once in a while is not acceptable.

Clearly, it may be beyond the practical abilities of any student of technology to convert all these metrics into a single measure for each task. But the structure and grounds of the method should be clear: use the assumption of optimization to measure as many aspects of behavior, behavioral preference and environment as possible. The closer users resemble ideally rational agents, the more their preferences are based on design grounds and so the more the most prevalent artifacts can teach us about practice, task, and good design.

This, then, is the main outline of an evolutionary or adaptationist approach to artifacts. It is a special type of optimization approach. Will it work? Can we explain the history of artifact change by showing how each new artifact somehow lets its users perform their tasks better? That is, can we show why one artifact is better than another?

Part C

Problems with an Optimization Approach

Although it is difficult, if not impossible, to determine the empirical values of competing artifacts on all the dimensions that matter, the model, as described, has a formal elegance that makes it attractive to some. To others the thought of collapsing preference to metrics measuring things like efficiency, effectiveness and learnability, might seem like a misguided continuation of 20th century positivism. I am sympathetic to both views. Attitudes aside, though, there are good reasons to question how broadly the approach can be applied even when artifacts are regarded to be essentially functional entities. In the end an optimization approach may be valid in only a few studies of technology and culture. This would have implications for a theory of extended mind too as it is valuable to know if our artifacts, practices and tasks are optimally or near optimally matched. The principle of optimality remains a driving ideal in cognitive and neuroscience.

What follows are five deep problems facing an adaptationist approach. Each explores something fundamental about the relation that humans develop to their artifacts. After articulating each I offer what I think is the best response.

Objection One: the artifact-task cycle is never in equilibrium, so artifacts never become optimal.

First, and perhaps of greatest theoretical interest, is that the artifact-task cycle may never reach an equilibrium state: change the artifact and the tasks and their costs soon change. As soon as the cost structure of activity changes, users adjust their behavior, which in turn alters their preferences and drives the demand for further artifactual change. Artifact ecologies are in continuous motion.

Consider TV's again. What happened when remote control units spread throughout the user community and became standard? As users adapted to the new cost structure, they became proficient at using their remotes to change channels, mute, increase volume, and turn the set on and off. At that point, their new behavioral equilibrium – the new distribution of how much time they spent doing this or that – gave rise to emergent behavior: specifically, channel surfing and 'multi-watching', both tasks that before were non-existent, or of little import before remotes.

This constant interaction between task and artifact, between what users *can do* and what they *want to do*, is what can be called the artifact-task cycle. Before remotes existed no one wanted to channel surf. After remotes proliferated a new TV might be judged by how well it handles channel surfing. As soon as a practice stabilizes around an artifact, users become sensitive to improvements in the cost structure of their task. This prompts the invention of new designs, which, in turn, leads to new practices. The cycle never ends. This makes it hard to determine the cost-benefit environment an artifact is supposed to be near optimally designed for: the previous cost structure that prompted invention, the current one that is in flux while users adapt, or the equilibrium one that will emerge once users have fully adapted their practices to the new features?

One side effect of the artifact-task cycle is that tasks may become doable in completely new physical environments. Taller giraffes reach higher branches and so move to areas where higher trees grow. The thicker fur of mastodons allowed them to withstand colder temperatures and live in colder regions. In a parallel manner, televisions with remote controls no longer need be reachable. They can be placed on ceilings or high up on walls, locations no one had considered before. As televisions find their way into bars, kitchens, restaurants, dental clinics and waiting rooms, a whole new set of user communities, with new objectives, needs and preferences, interact with the artifacts and begin to drive innovation. In a dental clinic, for instance, the only television viewer is the patient sitting in a chair, staring at the ceiling. Sound might be distracting to the dentist. So audio now needs to be piped directly to the patient, either via headphones plugged into the dental chair or wirelessly. The need for this feature is multiplied in sports bars where many televisions will be tuned to different channels. Patrons need personal audio equipment (or speakers at their booth) to allow them to choose the TV they want to listen to.

The net effect is that a change in the design of an artifact, may not only change practices and tasks, it may lead to a change in the environments where it is being used and a change in the sub-populations who now make use of it. As more types of users surface, the more differentiated are the needs the artifact must meet. This regularly causes speciation of artifact and segmentation of user community. It allows small changes in initial design to have large impacts on later design. This can lead to rapid or even ‘catastrophic’ changes in design.

This speciation is evident in tools. Generic hammers are useful for nailing, breaking, cracking, smashing, and tapping. But different users, working in different workplaces, have distinct functional needs that drive design variation. Masonry hammers are used on bricks, drywall hammers on drywall, slater hammers on slate roofs. See figure 5. Hammer design bifurcated to meet the specialized needs of dry-wallers, masons and slaterers. Hammer designs pulled so far apart it is best to see them as having speciated.

Figure 5. Although some hammers are general-purpose tools, most hammers are designed for specific jobs such as driving and removing nails or working with drywall or masonry. There were over 500 different hammers made in Birmingham in the 1860’s.

This speciation means that most structurally different hammers were designed to meet the needs of different users. They are not direct competitors vying for supremacy in a single niche.

The same occurs in medical contexts. In surgery the scalpel takes many forms, depending on where it is to be used, how, when and by whom. See figure 6.

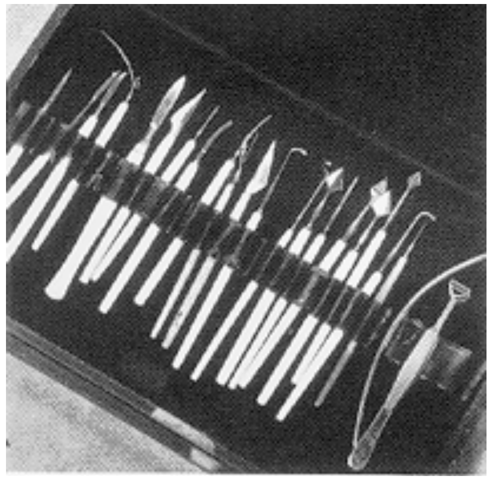

Figure 6. Scalpels all serve as precision devices for cutting tissue. Here we see a collection of ophthalmic scalpels for operating on human eyes. There are many types of scalpels because of differences in the tissue and local conditions in which they are used, and the many different surgical tasks they were designed to aid.

Once design change is seen to be a continual and turbulent process in which practices, tasks, users and even venue may change, the very idea of an artifact in equilibrium with its niche becomes suspect. Small design changes can lead to explosive design changes as new classes of users come online, or as emergent behaviors create unanticipated user needs. This calls into question the assumption of even momentary task-practice-artifact-user equilibria. Since the entire optimization/adaptationist program is based on the premise of long-term equilibrium this is a serious challenge.

Just how serious? Once the assumption of static equilibria is questioned the populations of artifacts found in a community at any given moment may *not* be strongly representative of the *current* needs and practices of that community. The artifacts are in transition. This year's models may be well designed to last year's needs, and next year's models may be very different. Even the population an artifact is well designed for might vary from year to year, making it hard to determine which subset of a population an artifact is supposed to be optimized for.

In the study of animal behavior rapid environmental change is seen as a strong reason to question theories such as optimal foraging. If climate is in transition, or an environment is undergoing rapid change, inferences based on optimization are in trouble. When things are changing quickly organisms play catch up – they spend too much time in the transients and not enough in attractor basins to suppose that their behavior is optimally adapted to current conditions.

In such cases prediction becomes unreliable. Imagine trying to predict the stable resting place of a ball in a bowl immediately after the bowl has been shaken. Once dislodged, the ball can move arbitrarily far from its minimal energy position, and it is hard to predict how long and where it will wobble before settling. Because the ball can move both closer and farther from its long-term equilibrium position, it is unreliable to guess what that equilibrium position will be based solely on the trajectory of the ball. To determine the long term resting position requires an independent measure of the energy landscape. This may be possible in the case of a physical bowl, where the physics of the system is simple enough for us to determine the stable resting place. But it is not plausible in the biological case where the niche-organism interaction is so complex, and where ethologists use the behavior of the organism itself as a probe, helping to identify the aspects of its surroundings that constitute its niche. Without a metric there can be no grounded prediction as to how behavior will equilibrate. Returning to the bowl analogy, it is obvious that the problem is worsened if the bowl is moved a second and third time, as the ball will never settle. It will spend all its time in transients.

Returning to optimal foraging theory, suppose we come across a chameleon population living in a dry climate. Their current condition is dry because, unknown to us, they are living through an unusual period of drought. An ethologist who studies the chameleon in current conditions might be confused by the apparent irrationality of behavior because unknown to him or her, the chameleon's behavior was optimized for a wet climate. This means that without a theory of historical conditions the theory of optimal foraging or near optimal design would be incorrectly applied. The chameleon is, in fact, not optimally designed for current conditions.

An archaeologist, similarly, might be confused by the function of a long handled axe associated with a peaceful community. In fact, the axe may have been optimized to times of war and has endured during peace times because it is prevalent. Without a lengthy analysis of the history of that society how is s(he) to know?

Reply to the challenge of changing cultural conditions and the non-equilibrium nature of the task-artifact cycle

Non-equilibrium conditions make inference based on optimization shaky. This highlights the need to know how prevalent are non-equilibria cases. To determine the probability that a culture's artifacts lag seriously behind the tasks, activities and practices its members perform, it is necessary to look at the persistence of that culture's designs, and the cultural mechanisms driving technological change.

The first step is to determine the historical facts about each artifact in question. This is one reason archaeologists do longitudinal studies. Since long-term stability implies optimization[4], the prevalence of an artifact over generations is partial proof that it was near optimal for the tasks and practices its users had. If the record shows rapid change in

[4] It remains an open question just how long conditions have to remain unchanged to warrant assuming adaptation to have stabilized and have reached a near optimal solution.

design, the challenge is to show that the rate of design change is in step with the new demands of users, their tasks, changed environmental conditions and changing practices.

One important difference between artifactual and biological evolution has to do with the additive nature of artifact functionality, making optimization under conditions of change a better bet for artifacts than creatures. More often than not an artifact keeps most of its previous functionality and just adds more. Remote control may have enabled the emergent activity of surfing, but the television sets they were part of still needed to be turned on and off, channels needed to be changed, volume altered. It is interesting that all these functions remain in manual form, but the point is that the core functionality remained similar despite a change in the technology. This suggests that often a change in an artifact's context creates a need for new features, e.g., better surfing management, but these new needs do not replace the core activities and needs the previous artifact was designed to meet. This means that last generation's design remains a reasonably good fit with most of this year's needs.

There is a second, and more profound reason, however, for supposing the record to reveal near optimal designs: artifactual evolution proceeds much faster than natural evolution because it is driven by invention, and invention is a more powerful mechanism for change than mutation and genetic recombination. Intelligent design is faster, more radical, and more volatile than biological evolution.

There are different sources of this power: the power of analogical thinking, the faster propagation of market signals, the human penchant for fashion.

Take analogical thinking. When an inventor sees a feature in one artifact that might graft well onto another, (s)he has the chance of plucking that 'pre-tested' idea and re-invigorating the existing product. A handle that works well on tennis rackets may be adapted for hammers. Teflon originally designed for space flight spawned a whole new product line of cooking ware. Velcro, originally used as a simple fastener in fabrics, spread across thousands of products, giving rise to new products such as flowerpots that adhere to bricks and other rough surfaces. The common factor two Velcro enabled products share is that both have parts that alternate between being stuck together and (occasionally) being free. Except for this one abstract similarity the two Velcro enabled products may be arbitrarily different. They can have different evolutionary trajectories. For example, bricks, shoes, belts, and archery, all recipients of Velcro late in their life, all have their own quite distinct history of design. Humans are good at appreciating abstract similarities; they are able to see an analogy between using Velcro to stick two ends of a belt together and using Velcro to hold a potted plant on a brick ledge. There is no comparable process of grafting or gene splicing in standard accounts of natural selection, however.[5]

[5] Saltation theories, however, do allow for major jumps in design. According to Goldschmidt "macroevolution would require a different evolutionary mechanism, one that would create the decisive transformational step from species to species, from one higher category into another. It would not take place through a series of atomistic alterations but by way of a far-reaching transformation of intrachromosomal structures.

Why don't we find gene splicing more prevalent in biology? One reason – aside from constraints on reproduction – is that random grafting is rarely good. Inventors must choose their grafts intelligently; and when they find an interesting candidate feature they expect to make a variety of changes in the host design to support its integration. It can require extensive tinkering to make a new feature contribute value to an existing system. Evolution requires that hybrid organisms prove their value early. If it takes generations of adaptations to improve a system, a feature will wither before it takes off.

A similar argument can be given to show why markets are able to select new designs faster than natural selection can: markets propagate signals about the goodness of a design more quickly than reproduction can. In biological evolution negative signals are produced when an organism dies before reproductive age, and positive signals are sent when an organism has reproductive success (the more offspring to reproductive age the bigger the positive signal). Signals travel at the speed of birth and death. With artifacts positive and negative signals are generated by market preference. These preferences may be based on a sampling of groups that are few in number when compared with the population they represent. This is especially the case if some of the market drivers are target groups, focus groups, elites, or design critics. Both these and the population at large can form opinions about products before they buy them. So an item can become 'hot' before it has begun to saturate the market. This means that fashion leaders can set fashion, and design critics can have a disproportionate effect on design.

This difference in the mechanism driving the proliferation of artifacts and organisms bodes well for archaeology. If the popularity of an artifact can be traced among the privileged, the elite, this sub-population may serve as a proxy for the greater population. In modern economies, consumer specialists use opinions of pilot groups as samples of larger group attitudes. In ancient societies, and especially in less consumer-oriented, technologically less advanced cultures, it is doubtful that product testing occurred in a modern sense. But the opinion of experts may have counted powerfully, leading to the widespread adoption of designs that court experts believed were best. Even in early societies the theoretically best, or the near best, might have translated into the popular best faster than any standard model of evolution would assume.

Accordingly, since artifacts can, in principle, change in design much faster than organisms they may adapt sooner to niche forces. In assuming an optimization stance the question is not necessarily whether artifacts evolve the way biological organisms evolve. It is whether the artifacts that are prevalent are well adapted to their environmental niches. In the cultural world, because artifacts are the product of intelligent design and market-like forces, they are quicker to adapt to changes in task, practice and broader

---

This repatterning, or Systemmutation, is attributable to cytologically provable breaks in the chromosomes, which evoke inversions, duplications, and translocations. A single modification of an embryonic character produced in this way would then regulate a whole series of related ontogenetic processes, leading to a completely new developmental type." Schindewolf (1993) p. 352.

culture than simple selection would suggest. Hence there is greater reason to treat extant artifacts as reasonable indicators of optimal artifacts.

Objection two: The co-evolution of systems of artifacts implies technology holism. Holism defeats analysis of optimality of individual artifacts.

The next objection arises because artifacts are not optimal in isolation. You cannot solve one part of the artifact-practice-task-user problem without solving the same problem for the many interdependent artifacts that the relevant user community interacts with.

Take an arbitrary utensil in a kitchen. How can the practices associated with a spatula, for instance, be understood without understanding the stovetop, the pans, the grill or the variety of related tools a spatula might be used with? And how can the functions of a larger device, such as a microwave be understood unless it is known whether the user community is university students living in dormitories or cooking enthusiasts working in near professional kitchens?

Every artifact occupies a place in an ecology of artifacts. The niche of one artifact is defined in part by the others around it. Keys are defined by locks, chairs, in part, by tables, fasteners by materials and tools, and altimeters by the cockpits they are embedded in. See figure 7b. The adequacy of one artifact seems to depend on the adequacy of its partners. This has the consequence that imperfect design in one part of a system may be compensated by excellence in others. You can't know how one artifact was used unless you know how its teammates were used and the way practices developed and co-developed for each. Hence, it is whole systems that must face the optimality condition, rather than each of their parts.

Fig 7a

Fig 7b

Figure 7. There are two senses to artifact ecologies. In 8a and 8b we see the first type: collections of accessories or related artifacts that can be used with a given artifact. In 8a we see accessories for a microwave that allow cooks to brown, crisp, and poach – all techniques not typically done in a microwave. In 8b we see a modern airplane cockpit with its complement of dials, displays and control elements. For any given task a set of artifacts will be relied on. To ask whether a given member of an artifact ecology is, in some sense optimal, we must know its place in its ecology. Add more accessories or dials and this may change.

An illustrative example of co-evolution of practice and artifact is found in the development of the modern oven. Prior to heat controllable ovens, recipes were designed primarily for long heat or a combination of short heat on the 'stove top' followed by long heat. With the advent of stoves that can be set to specific temperatures recipes have changed, the diversity of people who cook according to those recipes has changed, and the whole nature of cooking has, in some sense, changed. It would be difficult to estimate the optimality of cooking artifacts without considering the type of recipes and cooking techniques of their users. The four – oven, recipe, user and practice – have co-evolved. See figure 8.

Fig 8a. Fig 8b. Fig 8c

Figure 8. In 8a we see a turn of the 20th century Welsh range with its many doors and compartments. Cooks could rely on these different compartments to be at different temperatures. Many of our modern recipes rely on setting an oven to a specific temperature. This could be simulated in an expensive range such as this. But most kitchens did not contain such extravagantly expensive ranges. Moreover the temperature in most compartments was much lower than we normally cook at today. Accordingly, recipes called for slow cooking. In fig 8b we see a Kamado oven, similar to the original Japanese design 3000 years ago. Kamados are excellent for slow cooking BBQ style. In fig 8c an Indian tandoor is shown. It too was used to roast meat, though at a much faster rate than Kamado's or ranges. Recipes and ovens co-evolved.

The same pattern of co-evolution shows up in the coordination between food preparation and the utensils for eating. In a world of chopsticks food comes pre-cut in bite-sized pieces. There is no need for a knife if there is nothing to cut. No need for a fork if nothing need be speared or shoveled. Although the Asian cooking process still yields bigger than bite-sized slabs of meat, these are usually cut before serving and only small pieces are laid on the plate. Larger sized elements, such as whole fish, are cooked to the point where they can be picked off the bone at the table, although in more formal contexts a final step of 'plating' takes place to save diners the obligation. Since etiquette, recipes and practices are hard for archaeologists to find solid evidence for, the problem of technology holism is genuine and challenging.

Reply to the problem of technology holism: the credit assignment problem can be solved in AI (with enough iterations) so it can be solved in archaeology.

The problem of holism is a version of what in Artificial Intelligence is called the credit assignment problem: if a metric can be applied only at the periphery of a process – e.g. at the end of a chess game rather than to intermediate moves – it is hard to determine how much credit or blame each component of the process deserves to be given. In a chess game of 35 moves, how much credit should move twelve be given for its contribution to the eventual win? The available data – wins and losses – seem too thin.

The challenge of holism, then, is not that credit assignment is impossible. It is that the obtainable data underdetermine our capacity to determine how to assign credit. Is this true for either biological or artifact ecologies? The simple answer is: not if the artifacts or organisms can be studied ethnographically or experimentally. Biology, ethology, evo-devo, economics and other disciplines have devised techniques for uncovering more data.

For instance, in biology, the prevailing wisdom is that every organism is part of a local ecology – a holistic system – in which the fitness of a single organism cannot be properly understood without understanding how it interacts with the other organisms partly defining its niche. To deal with this increased complexity new techniques have been developed. Some involve simulation, others require new forms of mathematical analysis, and still others rely on manipulating ecological conditions and observing outcomes.[6]

Can the same be said for the study of 'extinct' organisms, cultures and artifacts, which no longer can be studied in their living context?

At present it is difficult, if not impossible, to determine the practices a culture adopted for many of its artifacts because such practices leave no direct trace in the record. A group's diet may be inferred from bone remains; but the way its members prepared their food leaves no trace, except perhaps in the utensils they used to prepare, serve and consume.

Therein lies a hint. The need for a given artifact to fit near-optimally with the rest of the artifacts in its ecological niche makes it possible to look for further constraint. Archaeologists may use their own role-playing or computer simulations of artifactual use to generate 'plausible' scenarios of use. If these scenarios actually took place there may be observable traces that would otherwise be ignored. The way these simulations would work is by running combinations of practices, artifact ecologies and users, under various assumptions and looking for configurations that a) meet the data from the record; b) seem plausible given what is known historically about users and their culture; and c) are optimal in an evolutionary sense – that is they spread in number throughout the population of users. If simulation suggests the prevalence of an artifact's design and archaeologists then find that design in the record, they have prima facie reason to think that it was selected for because it met their needs.

---

[6] See especially work in agent-based and network modeling of biological and cultural systems. Ecological Networks: Linking Structure to Dynamics in Food Webs,
edited by Mercedes Pascual and Jennifer A. Dunne, Oxford University Press.

There is no guarantee that this method will work in archaeology, but without an assumption of optimization, such simulation, experimental and mathematical techniques are unmotivated. It is by now common practice to engage in all of these.

Objection Three: What counts as optimal in design is culturally relative so there is no objective basis for judging how good a design is.

Artifacts are always embedded in a culture, with its historical practices and biases. In modern society it sometimes happens that a design that is successful in one community is not successful in another. If both communities use their artifacts to perform the same tasks why do they prefer different designs? According to the optimality hypothesis the goodness of a design is an objective matter. Prima facie, what is good in one community ought to be good in the other; and what is preferred in one ought to be preferred in the other. The same design should dominate in both cultures. Because this is empirically false, and there are clear cultural differences in design values, more must be at play in artifact choice than the objective notions of efficiency and effectiveness. Perhaps factors that are more subjective, such as style or aesthetics, also play a central role in choice.

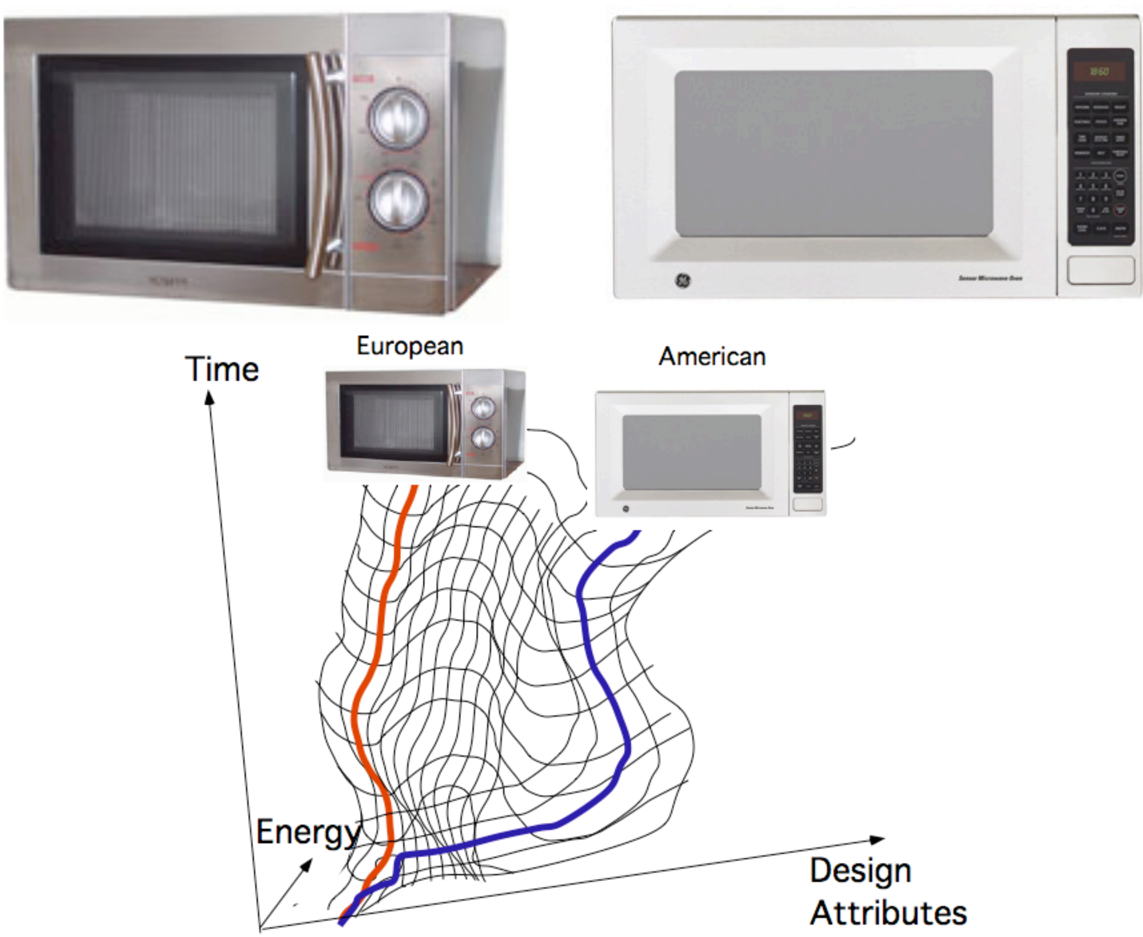

Figure 9. Consumers in France and USA seem to prefer different interfaces for their microwaves. French models have far fewer buttons, relying instead on analog style dials in place of digital number pads. American models strongly bias number pads and add

> as many as ten more buttons for specialized cooking, such as defrosting poultry or fish, heating beverages, or baking a potato. Since all microwave cooking ultimately depends on setting only the power level and duration, the two systems are functionally equivalent (roughly speaking), though they set these parameters in different ways. Why are there these cultural differences in preference? Do the French and American users have different design sensibilities?

In figure 9 the design landscape of a microwave is shown graphically to emphasize how microwave designs in Europe and the US have evolved differently. Although American style microwaves are found in the European market and European style microwaves are sold in the American market the central tendency of design found in the two populations of microwaves is quite different. If this is not the consequence of different practices, needs, or uses then the two communities must either be optimizing along more dimensions than function alone or microwave design is a counterexample to the principle of optimization.

In evolutionary theory there are several mechanisms for explaining why sub-populations might begin to pull away from each other and fork into groups with different attribute distributions. When applied to artifact selection they all explain why there might be cultural differences in design and provide reasons to question the hypothesis of artifact optimality by some objective measure. The three most relevant mechanisms are: genetic drift, the founders effect, and capture by local minima.

Genetic drift explains why, in small, reproductively isolated populations, changes due solely to chance factors can lead to changes in gene frequencies that are unpredictable from simple accounts of recombination, mutation, and natural selection. For instance, in a small population, a few individual organisms may spread disproportionately through a population simply because, by chance, they were the first to thrive, or by chance they mated with others unusually similar to themselves. After a few generations this bias leads to decreased population diversity. The attribute distribution drifts to that found in that randomly successful sub-group.

The cultural equivalent to genetic drift is called *artifact lock-in*. Simply by chance, two societies both using their artifacts for the same tasks, may form divergent preferences because users happen to group into self-reinforcing *opinion cliques.* Each clique has its own preferred artifacts and the members of the society who are not part of a clique soon find themselves choosing between a small set of popular products, whether they like it or not. New designs have a hard time entering the system because the opinion cliques serve to bar entry. They lock out new designs while permitting incremental change in the existing ones. The result is a narrowing of artifact diversity. Returning to microwaves, if the opinion makers of French and American cooking have different views about microwaves, opinion cliques may form and drive the preference of each population away from each other. Even if one microwave is objectively better than the other, the opinion cliques in one of the community will bias the judgment of consumers and lead to the prevalence of a sub-optimal microwave. Thus, French and American consumers may each believe they have the better microwave, but only one may be right.

The founders effect explains how genetic bifurcation can arise when a small sub-population becomes isolated from its parent population. It will not explain why American and French microwaves are different because both populations have access to each other's products. It would explain how the two might have evolved different preferences for subsequent artifacts had they been cut off from each other and if the US had always maintained a restrictive immigration policy. The founders mechanism requires that all the members of a group be descendants of a few founders and that that group never mixes with a different bloodline. In that case, any genetic biases of the founders will be passed along. For example, a small group of bison cut off from the main herd may have genetic idiosyncrasies. If the group is small it will not have all the genes distributed throughout the larger herd. If this small isolated group survives and multiplies its members will pass on to their descendents their genetic biases. This new herd, even if becomes large, may have genetic biases quite unlike the larger herd.

The case for a founders-like effect occurring in artifact populations is a little different. The most obvious analogy is where a few members of a culture isolate themselves and continue using the designs they were accustomed to when they were part of the larger community. Like the Amish people of Pennsylvania the descendents of the original group may perpetuate the past because their culture is profoundly conservative and disapproves of innovation. This Luddite mechanism differs from opinion cliques in that in a reactionary society there are not just a few members who are the fashionistas – the arbiters of good taste – it is the whole community. For cultural reasons the biases of the founders are perpetuated in an authoritarian tradition that is preserved generation after generation. The outcome is similar to drift: a society may evolve with its own traditions and its own sense of what is good. In drift the society slowly moves toward the opinions of a few; in the founders effect, a splinter group with idiosyncratic and authoritarian views prevails leading to artifacts being valued for reasons other than their task optimality.

Drift and founders bias are mechanisms that lose their power as a population increases or if new blood enters the population. This ought to hold for artifacts as well. When there are more people to invent and choose artifacts, people who may buck tradition, there is a greater chance of artifact variation, and greater chance for consumers to choose designs based on their effectiveness. The power of diversity and choice should lead to change and the best winning out. In the long run, a large population ought to generate enough candidates and non-affiliated consumers that opinion cliques or an authoritarian tradition would be eventually overwhelmed and the design that comes out on top should be the best. If there remains a cultural bias in large, enduring populations its explanation must lie elsewhere.

The prime candidate for cultural bias in large mixed populations is the selection mechanism itself, which by its nature is a local maximizing process. In classical evolutionary theory only small mutations or changes tend to take. Although major design revisions might lead to significant performance improvements, these usually require the design to be modified or fine-tuned so that it works effectively with the pre-existing ecology of other artifacts, practices and supporting elements. Until that point organisms based on the status quo design are more likely to be successful despite the long run potential of the new design, assuming it were successfully integrated and fine-tuned. Consequently, there is a bias for small, low risk improvements. Artifacts will evolve along design lineages based on incremental change. This explains why two populations might evolve along different lines if they began with different seed designs. These seeds might have been comparably effective at their outset but have different effectiveness horizons. The lesser design would be incrementally changed until it reached a limit point where only a change in its basic principle could take it higher.

This effect is well known. Status quo designs will often continue to proliferate despite innovations available that would outperform them. Innovative designs languor despite their longer term potential because these benefits would emerge only after appropriate changes had been made in the supporting ecology. Accordingly, anyone who tries producing an innovative design faces an initial period where production costs are not recouped. Without the capacity to take a loss, however, the innovation will die and the population will be forced to accept worse products. If culture is a local optimizer the new artifact will never live long enough to be fine-tuned and fully integrated into surrounding practices and production methods.

Do cultures actually get trapped in local minima in this way? Do designs follow something akin to an evolutionary path that locks them in?

Paul David [unpublished] and Scott Page [2005] in their discussions of path dependence have offered a convincing explanation of why artifacts get trapped in sub-optimal designs. In discussing why QWERTY keyboards continue to dominate the keyboard market despite their being provably sub-optimal from an efficiency standpoint, Page offered the following:

> There were increasing returns to buying a QWERTY typewriter because as more were sold, the costs of marketing QWERTY's fell. The QWERTY was also self-reinforcing. Textbooks for learning how to type were all based on the QWERTY layout. The QWERTY also created positive feedbacks. One QWERTY typist could type on another person's QWERTY. Eventually, once there were enough QWERTYs in use, it became locked in as a technology. We are all part of the QWERTY nation. P5. An Essay on The Existence and Causes of Path Dependence. Scott E Page 2005.

Page and David emphasized the prevalence of QWERTY teaching aids. A system of support technologies and institutions arise that make the QWERTY part of a more comprehensive set of practices and artifacts.

This ties in well with another factor. Consumers suffer from what we may call *learning inertia*. They often prefer the tools they have, which may be second best, to the ones they might have, which are much better, because mastering those new tools requires that they acquire new habits. The cognitive effort involved in studying might be a wise investment given the returns to be enjoyed from better tools, yet people resist acting in their long-term interest because they cannot overcome their resistance to practicing in order to learn new things.

This same resistance to new things may be explained by another generalization: people like what they know. Once people have learned to like the things they have, they can be slow to like new things. A taste for the new, even when it is better, is not always instantaneous. This leads to a distortion in preference functions that is rarely discussed.

In economic theory the preference function of a person or group is usually assumed to be stable and complete. Stability means that tastes do not change in the short run. Completeness means that for any two products P_1 and P_2, either P_1 is preferred to P_2, or P_2 is preferred to P_1, or the two are liked equally. This holds even if one of the products has never been encountered. Both assumptions are psychologically false. People do not know their relative preference for things they have never encountered. And often a measure of their likes or dislikes taken immediately after their first encounter is not a reliable indicator of their longer-term preferences. A new fruit may taste better the second time round (or worse) once its strangeness (or novelty) has worn off. New things take time to *grow* on people. This shows either that people initially do not know their preferences for novel things or that they change those preferences in the short run.

Once the assumptions of completeness and stability are relaxed it is rational for consumers to display resistance to change. Why should they go to the expense of buying a new item if they can't trust their initial reactions on seeing it? Most people are not early adopters anyway. They prefer a wait and see approach. Because this means that new products will be slow starters many good products will never get a foothold and it will be hard to dislodge incumbent products. Once a design became prevalent in a culture it would tend to stick. The outcome is that if two cultures developed artifacts with differences in their *seed* design they would tend to stay apart. Incremental change would likely be better received than major change. And so the designs that proliferated early in a culture would tend to determine the direction of change. Without a change in the tasks people need to perform, cultures that start from different core designs will tend to perpetuate their differences, even if one culture's design is less effective.

Reply to criticism Three: Cultural biases may still lead to artifacts that are optimal or near optimal by some objective measure if we include non-functional attributes in the metric. But this dangerously complicates the analysis and rejects the priority of efficiency.

Of the challenges to an objective, culture independent metric of optimal design only those concerning local optimizing and path dependence are general enough to derail the entire approach. Drift and founders effect are the result of random processes acting in small relatively isolated populations. They do not apply to cases, such as American and French microwaves, where both cultures are huge and have access to the mature designs of each.

The thrust of local optimizing arguments is that major design changes cannot initially compete with status quo designs and so are second best until they are fine-tuned. There are a few reasons why artifacts might escape the second best argument. The first derives from humanity's unique capacity for investment. People will endure short-term costs, such as those associated with fine-tuning, if they expect long term gains to more than compensate them for their short-term losses. Someone who passionately believes that he has a better designed sword may decide to continue making his new design despite his loss of income in the short run, as long as he remains confident that people will eventually discover the potential of his new technology. Natural selection is incapable of this sort of investment. Hence more innovative artifacts may hang around long enough to be tinkered with until consumers recognize their intrinsic effectiveness.

Second, intelligent design is more powerful than the blind watchmaker. As mentioned, designers can tinker with an invention in their workshop until it is better designed than the competition. Fine-tuning can be done offline in the case of artifacts but not in biological creatures. And because inventors can choose features from other designs that have already been proven to work well, invention is not a random process and need not be incremental. The result again is that major innovations can break on the scene in a ready to be accepted manner. The designer paid an offline cost for fine-tuning and refinement.

Both of these arguments suggest that human design can proceed at a faster pace than biological design. Human designers can create interesting variation in design faster and with less immediate reward than biological processes. But consumers are assumed to come around in the end. The arguments from path dependence, learning inertia and incomplete and unstable preference all challenge this last assumption. There are institutional and psychological reasons why members of a population may have entrenched biases. Indeed, the hallmark of culture, it has been said, is that it causes sub-optimal choice [Bednar & Page 2007]. Because there are many ways for a design to be sub-optimal there is little reason to suppose different cultures will converge on the same designs.

These arguments are, to my mind, hard to overcome. Because optimality means there are no technologically accessible artifacts that are *functionally* more effective given a set of tasks, workflows, skills and ecological contexts, then any successful innovations lurking in a culture prove that the status quo artifact is not optimal. At best a culture needs time to reach a new equilibrium where innovative products dominate. At worst the inherently conservative forces of preference and institutional entrenchment will keep a culture from reliably adopting the best innovations.

I think these arguments establish that we cannot assume the artifacts present in a culture are optimally designed for the functions they were meant to serve. All these arguments point to one conclusion: preference is history dependent and to understand the choices of a society at any moment it is necessary to study the trajectory of choices that led up to that moment. Put differently, choices are locally optimal if you know the circumstances of choice well enough. This should be no surprise as it merely restates that evolution is a local optimizer rather than a global one. It does emphasize, though, that the mechanism inhibiting global optimization is different in the cultural and biological case.

The deep question is whether it may be possible to see artifact evolution as a more globally optimizing process if the attributes being optimized include non-functional ones, such as beauty, personal expression, style, and so on. There is no doubt that fashion and trend are often based on attributes unconnected to efficiency and functionality. So why not pots and pans, axes, weapons, carpentry tools and the rest? Can we construct a theory of rational choice if path dependence and imperfections in the preference function mean that users have idiosyncratic judgments of non-functional attributes?

The notion of optimization in economic theories of consumer behavior treats functionality as only one of many grounds for ranking the desirability of artifacts. One consumer may prefer an axe head because of its usability – a key factor in its functional efficiency – another may prefer it for its color, the smoothness of its feel, its appearance or history of use. In economics the diversity of reasons that go into rational preference are not seen as grounds for rejecting the principle of optimizing or satisficing, they are seen as grounds for accepting. People can have very different preference functions and still be rational.

This a dangerous line to follow. First, it does not address the concern that preferences are unstable and incomplete. The foundation of economic theories, accordingly, may be shaky. Even more importantly, the rationale for restricting optimality to functional design was to allow positing a measure of optimality that is independent of revealed preference, and hence independent of the artifacts that are prevalent. Because we used optimal foraging theory as our paradigm of optimality artifacts need to be well designed in the same way that a well adapted creature's behavior is well designed to its foraging tasks and niche. By extending optimality beyond the functional we, more or less have accepted that people's artifacts are often not optimally designed, in a functional sense, for their tasks and practices.

Is it possible to gain any leverage from this weakened approach? In archaeology, as in economics, we are free to allow non-functional considerations such as aesthetics, or cultural preferences for certain types of appearances, to figure among the selective forces that lead to optimal choice. Discovering independent evidence for these preferences, however, is difficult. Only if data can be found for the cultural biases in an ancient society can any part of the optimality principle be saved. Once we recognize that a culture's preference function is partly based on aspects of design unconnected to efficiency and effectiveness we need to complicate our measure of optimality.

One reason to be cautiously sanguine about the archaeologist's plight is that biases have a way of showing up in the frequency of tasks performed. For instance, in France, daily diet and the everyday practice of preparing meals is quite different than in the United States. For one thing, until very recently, French consumers did not have access to the variety and abundance of frozen foods. This meant there was less need to defrost and therefore less need for an interface that simplified the decision problem of choosing the power and time settings for defrosting. From an archaeological perspective, it is reasonable to expect evidence for the prevalence of certain food types and preparation methods. These would support inferences about the reasons certain cooking, cutting and storing artifacts may have prevailed in a culture.

Another way cultural biases may show up is in design similarities across products. Thus if the French show a preference for dials in other artifacts, such as Stereo systems, while Americans show a preference for buttons in other artifacts we have an indicator of bias that is independent of any one product.

The implication is that cultural bias, per se, is not a sufficient reason to overthrow the optimality principle. It remains a helpful methodological directive. But it makes the archaeologist's job so much harder that for many cultures it may offer little advantage over existing methods, which in any event already incorporate many of the same ideas. To be useful the optimality principle should be bounded: conditions where it will likely be helpful should be distinguished from those where it will likely be unhelpful.

Final Objection. In ancient societies the population may have been too small and their artifact choice too thin for selection to deliver near optimal products.

Of all the reasons for questioning the optimality assumption the one I personally find overwhelming has to do with population size and artifact diversity. These are the crucial factors determining when to consider using an optimality principle. For instance, in pre-market societies there is rarely enough variation in design to support optimization – products cover too little of the design space. The user populations, or sub-populations, moreover, are often too small to reliably select optimal designs. As a result, in ancient cultures there may not be enough commissioning of artifacts, or trading between groups to guarantee that evolutionary processes lead to optimality, or near optimality in artifact design, even when history is taken into account. Inventiveness and innovation is be too slow to drive even local optimization.

Failure in artifact diversity and population size can be seen as violations of the following two principles presupposed in evolutionary and economic theory:

1. The principle of sufficient variation in design: The set of artifacts on the market are different enough, and there are enough of them, that they fill up all the dimensions of user value. There must be enough difference in size, appearance, ease of use etc. of the artifacts, in each 'price' range, that every user can, in principle, find his or her personally optimal artifact. If available artifacts only sparsely populate the value or quality space, people too often have to settle for the best available, whether this be second best, third best, or even worse, relative to what they want had they been given more choice. Accordingly, nothing reliable can be inferred about the quality of a design given what is found in the record of a pre-market (non-trading) society, except that the people at the time chose the best of a possibly bad lot. Too little variation means that the design space is thinly covered. The data of user preference will not support inference about preference on any particular value dimension. Users have to make too many compromises.

2. The principle of sufficient population: There are enough consumers demanding a product at any moment that it is possible to infer the design attributes they care most about (assuming there is enough variation in the designs available). If a group is large enough we can get a statistically significant distribution of

> preferences based on their dimensions of interest. If there are not enough consumers relative to quality space then we cannot suppose that a given consumer group will provide us with an unbiased measure of their values. The more dimensions of value there are, the more consumers we need. How much can twenty consumers possibly reveal about the nature of their preference for artifacts having hundreds of design variations?

The challenge these two principles pose is that design optimization may be a statistical anomaly in some cultures. This plays out in two ways.

In non-inventive societies with big populations, tools and artifacts may have more to do with what happened to be invented than with what ought to have been invented, given the needs of users. People end up wanting just the things they're offered. They like what they have because they cannot imagine having more choice. Hence limited choice can lead to limited wants independently of population size.

The story for inventive societies with small populations is not quite as bleak but bad nonetheless. The tools and artifacts that proliferate may have more to do with what catches someone's fancy, or with what hits the market first, than with what has the best design. In small markets, drift and founders effect are real possibilities because the cost to review all the products an inventive group may present can be overwhelming for individual consumers. They are more likely to fatigue and buy something they have already seen than continue searching for what is best. The products at the front of the market often sell most even though they may be inferior. Hence, only some products end up being reviewed and these soon spread throughout the group. Whoever gets to market first has an unfair advantage. A famous example of this outcome occurred in the early 1980's when the VHS format beat out the Betamax format for videotape to become the market dominant tape despite Betamax's technical superiority. But the idea that a product can dominate a market for reasons unconnected to its efficiency, or even its beauty, style and effectiveness, is an old one.

The upshot is that only in societies where both invention is plentiful and the population is large enough to consider all those inventions can we hope for optimal designs to emerge as market dominant.

Reply to Final Objection: This is a compelling argument only partly offset if a culture is both highly inventive and has an expert population.

In societies where inventors are both creative and diligent, and the users selecting artifacts are experts there is a chance that the best candidate designs will proliferate and that these will approach the optimum.

Inventors are intelligent designers. The process of intelligent design no less than the process of natural selection works by generating candidates and selecting the better ones. In problem solving theory this is known as the "generate and test" paradigm. In natural selection the first step, generate design variants, has a random element. In intelligent design, by contrast, it should not. A human designer intentionally modifies an existing design, or creates a new design, not in a random or haphazard process, but reflectively. This generation process is itself the outcome of a selective function, namely the result of reflecting on the likely goodness of different designs. As Popper said, "Our ideas die in our place". So designers may try out a creative idea on paper or in their imagination before committing to physical prototypes. This double process of preselecting a candidate, then realizing it and trying it out in the physical world, means that intelligent design may proceed at a much faster pace and with less random variation than natural selection. Whereas nature would physically generate thousands of variations in order to select a few, human designers can mentally generate thousands of designs, select a few for physical implementation, and let the population of users select the best.

This double selection process is naturally subject to abuse. A designer's intuitions of what will be best may not match the user population's experience of what is the best. A designer may believe (s)he knows what users will like; and in a happy world (s)he will be right; tools will indeed improve rapidly and users will consistently be getting better artifacts. In a less happy world, designers will be wrong about what will work best for users and progress will be slow or non-existent.

When inventors are creative they can meet the requirement of the principle of sufficient variation even with a small number of newly implemented designs providing their own pre-selection is reliable. This is a big if. But even granting their pre-selection is equivalent to having more candidates generated and rejected, inventiveness alone cannot compensate for the problems that arise when the population of users choosing among different artifacts is small.

The problem with a small population is that when designs vary in many dimensions there are not enough choices made to show how the dimensions trade off. What matters most: how well a blade can take an edge or how well it can be thrown? A few people cannot reveal the structure of their preferences unless they can each buy hundreds of tools. This is why the principle of sufficient population matters even for an inventive culture.

When a small population is made up of experts, there is some hope that they can judge the design dimensions of a tool well enough to signal the key attributes of a design to inventors. This might lead to near optimal tools being produced for a small but expert population. For instance, only a few blacksmiths made swords or samurai. But they represent the pinnacle of sword design.

Nonetheless, experts can be wrong or biased. And if they must choose between a set of tools that is too small or the tools may vary along too many dimensions they may still be biased. This highlights the strongest reason large populations are important for design optimization: there is reliability in large numbers. In his famous note on the value of the median estimate Francis Galton [1907] gave a striking demonstration of the remarkable guessing power of large groups. He used as data the recorded estimates of people guessing “the weight of a fat ox after it had been slaughtered and ‘dressed’ ” . The annual West of England Fat Stock and Poultry Exhibition held a contest in which attendees could submit their guess on a ticket they purchased. The winner received the dressed ox. To his astonishment Galton found that of the 787 contestants, most of whom had an average farmer's knowledge of animals, the *median* guess was less than 1% off: "the middlemost estimate is 1207 lb., and the weight of the dressed ox proved to be 1198 lb.; so the vox populi was in this case 9 lb., or 0.8 percent of the whole weight too high". Although there were many guesses that were badly off the mark in one direction or another the median guess was astonishingly accurate. How large a contestant population is needed to ensure a robust estimation is itself hard to estimate. But clearly, the bigger the population the better. The same applies to the robustness of estimating the goodness of designs. The more people there are to look over artifacts, the more reliable will be their judgment of excellence.

The upshot is that cultures with small populations, even those with inventive artisans, may find themselves working with tools and other functional implements that fall well short of optimal designs, even given the constraints of production and cost. I see this as a real limitation on the scope of the optimality assumption. It is all the more significant for archaeology because so many ancient populations are small. There will not be enough carpenters, stone masons, hunters, warriors, blacksmiths and other artisans, to reliably hill climb along the many dimensions of design. In such communities, optimality of design, even local optimality, is more likely to be the exception than the rule. Even when the requirement that artifacts be functionally optimal is relaxed, and beauty and style are added to the attributes that matter to users, and even when optimality is construed in as local optimality and must be understood in a history dependent sense, still there will remain the dimensionality problem that a small number of users cannot span many dimensions of design. The assumption of optimal design is inappropriate for small populations.

Summary and Conclusion

Throughout this inquiry I have been threading together a few themes. The first is ecological: artifacts are part of a dynamic ecology of users, designers, supporting artifacts, tasks and practices, and constrained, moreover, by a culture's momentary technological envelope. Artifacts, tasks, users and practices co-evolve. Each adapts and changes the other. Analyzing this co-evolution is hugely difficult and in all probability impossible without making assumptions about how well designed artifacts are and how well adapted users and their practices are to those artifacts. To break into this holistic complex, I suggested that it is necessary to assume that an artifact's design is a culture's way of adaptively responding to the demands created by its members' tasks and pre-existing practices. By assuming these designs are optimal or near optimal, artifact evolution can be explained using the language and mechanisms of selection, using biological evolution as an analogy.

To explore the validity of this idea I discussed whether it is plausible that the artifacts that are most prevalent in a society are the best designed, given that society's production possibilities. Using optimal foraging theory as an example I explained how artifacts, in ideal conditions, could be see as optimally designed relative to a set of value dimensions which define a cost-benefit space for that artifact. In equilibrium, the artifacts that are mostly widely found in a society are (or ought to be) the ones that have the best cost-benefit profile. Once artifacts are optimal, changes are driven by changes in this cost-benefit space. Thus the more often users engage artifacts and adapt their own practices and skills to the possibilities these artifacts open up, the more they find hassles, imperfections, new uses, and new usage environments for those artifacts. New patterns of behavior emerge leading to new preferences. These changes reshape the cost-benefit structure of activity associated with using each artifact. This in turn creates an opening for new artifact designs. As new artifacts are introduced they find acceptance if they have lower costs, or greater benefits than all others. Hence they are well designed.

This methodology might offer archaeologists a better understanding of the trail of artifacts they find in the record provided they can infer the changing cost-benefit structure of particular artifacts. Of course, a cost-benefit analysis is most plausible for functional artifacts such as tools, where it makes sense to ask how efficient, effective and easy a given artifact makes task performance. But, arguably, it is a start.

I also argued that there are grave challenges to this approach. An evolutionary model of artifacts must also deal with the complexity introduced by rapid innovation and the interdependence of artifacts. When cultural conditions are right humans are capable of rapid innovation. Artifacts then change so fast that a given design is around too briefly for users to reach a reliable equilibrium and so generate a plausible cost structure for that artifact. An ecology that lurches from transient to transient is insolvably complex.

Another challenge comes from the interdependence of artifacts. Rarely do we use an artifact in isolation. Artifacts have their own mini ecology. Pots and pans are used on stoves or fires, and their users count on having spatulas, bowls, plates, knives and cutting surfaces, running water and so on. No artifact is good or bad by itself as the supporting elements of an artifact ecology can compensate for the limitations of any one artifact in isolation. The result is that it is hard, if not impossible, to build a cost structure for a single artifact.

These objections, taken as a whole, threaten the very idea of seeing artifact evolution as an optimizing or near optimizing process. I argued against most of these objections being insurmountable if adaptation is seen as a semi-optimization process. But even then solutions are hard to reach and uncertain. And moreover, in many cultures there are institutional and psychological reasons why the users of artifacts do not quickly gravitate to the best artifacts. Functionality, it turns out, is only one of many attributes driving rational choice.

Even if the hard problems of holism, culture, rapid change, extension of preference to non-functional attributes and other concerns were resolvable there remains limits on the applicability of an optimization methodology. It does not apply to cultures where innovation is sparse. When there are too few artisans creating new artifacts, or when there are too few users clamoring for improvement there is *not* enough diversity or selective force to ensure that artifacts are well designed. The tools and objects that predominate will be worse than a culture is capable of producing. This outcome more or less rules out using an optimization approach on small, technologically quiet cultures – a disappointing outcome. Evidently, there are profound limitations on the usefulness of optimization as a methodological directive in anthropology.

A second theme, implicit throughout, is that things have a certain cognitive life because they mediate how people engage their environment. Without artifacts most of the tasks we perform, and most of the tasks earlier peoples performed, could not exist. Artifacts partially constitute human tasks. They shape how humans encounter the world. Without a flute or a musical instrument there would be no instrumental music or instrumental practices. Musicians and their audiences would be denied a major form of musical experience and activity. Life would be cognitively less rich.

By itself, this does not establish that artifacts have a cognitive life. First, life implies that the living thing tries to maintain itself. Artifacts are not independent from the humans, practices and cultures that use them. Moreover, the ability to support experience, even to be necessary for experience, is not sufficient for cognitive life. If it were then oceans or deserts would have a cognitive life because they support and are necessary for humans to have ocean or desert experiences. The key feature of artifacts, however, is the way they co-evolve with the humans who use them. Artifacts are unlike oceans and deserts in that they evolved to support human experience. They help to preserve the collective practical knowledge of a group. A major function of artifacts, then, is to serve as a repository of knowledge. Because most of the cognitive life of a culture is tied up in the way its members perform tasks and solve everyday problems, the artifacts that partly constitute and frame those practices, and embody the intelligence of generations of designing, serve as partners in task performance. This intelligence in design does not earn them the right to a 'cognitive life' of their own. They are parasitic on human culture. But they are symbiotic and essential to our cognitive life, extending the power, effectiveness and duration of practices beyond the lifespan of individual members of a culture. They are part of humanity's cognitive encounter with the world.

References

Bednar J, Page S, Can Game(s) Theory Explain Culture? Rationality and Society, Vol. 19, No. 1, 65-97 (2007).

David, Paul. Path Dependence – A Foundational Concept For Historical Social Science – unpublished working paper

Galton, F. (1907). Vox Populi. Nature 75: 450–451.

Goldschmidt, R. B. (1960). The material basis of evolution. [Paterson] N.J.: Pageant Books. [©1940]. Reissued with introduction by Stephen Jay Gould, 1982.

Gould, S. J. and R. C. Lewontin (1979). The spandrels of San Marco and the panglossian paradigm : a critique of the adaptat

Gould, S. J. (1977) "Evolution's erratic pace." Natural History 86 (May): 12-16. nist programme. Proc. R. Soc. Lond. 205: 581-598.

Levy, David M & Peart, Sandra, 2002. " Galton's Two Papers on Voting as Robust Estimation," Public Choice, Springer, vol. 113(3-4), pages 357-65

Page, Scott. (2005). An Essay on The Existence and Causes of Path Dependence.

Pascual, Mercedes & Jennifer A. Dunne (2005). (eds). Ecological Networks: Linking Structure to Dynamics in Food Webs, Oxford University Press.

Petrovski, Henry (1992) The Evolution of Useful Things (New York: Vintage Books).

Pierce, G. J.; J. G. Ollason. Eight Reasons Why Optimal Foraging Theory Is a Complete Waste of Time. Oikos, Vol. 49, No. 1. (May, 1987), pp. 111-117.

Prince, A. P. Smolensky (1997). Optimality: From Neural Networks to Universal Grammar. Science 14 Vol. 275. no. 5306, pp. 1604 - 1610

Pyke, G., H. (1984) Optimal Foraging Theory: A Critical Review. Annual Review of Ecology and Systematics Vol. 15: 523-575

Schindewolf, Otto H. (1993), Basic Questions in Paleontology: Geologic Time, Organic Evolution, and Biological Systematics. Edited by Wolf-Ernst Reif. Translated by Judith Schaefer. With a Foreword by Stephen Jay Gould. (German edition 1950).

Simon H. A. (1969) The Sciences of the Artificial. MIT Press, Cambridge, MA.

Smolensky, P. & Legendre, G. (2006). The Harmonic Mind: From Neural Computation To Optimality-Theoretic Grammar Vol. 1: Cognitive Architecture; vol. 2: Linguistic and Philosophical Implications. MIT Press.

Thompson, D'Arcy (1961). On Growth And Form. Cambridge University Press

Todorov E, (2006). Optimal control theory in Doya K (ed), In Bayesian Brain: Probabilistic Approaches to Neural Coding, chap 12, pp 269-298, MIT Press.