Agent-based model experiments cast doubt on Dunnell's adaptive waste explanation for cultural elaboration

W. Christopher Carleton, Brea McCauley, André Costopoulos & Mark Collard

Published online: 02 Aug 2019.

Submit your article to this journal ⬚

View Crossmark data ⬚

Routledge
Taylor & Francis Group

Agent-based model experiments cast doubt on Dunnell's adaptive waste explanation for cultural elaboration

W. Christopher Carleton [a], Brea McCauley[a], André Costopoulos[b] and Mark Collard [a]

[a]Department of Archaeology, Simon Fraser University, Burnaby, Canada; [b]Department of Anthropology, University of Alberta, Edmonton, Canada

ABSTRACT

Ancient monuments are puzzling from an evolutionary perspective. It is obvious that their construction would have been costly in terms of energy, but it is not clear how they would have enhanced reproductive success. In the late 1980s, Robert Dunnell proposed a solution to this conundrum. He argued that wasting energy on monuments and other forms of what he called "cultural elaboration" was adaptive in highly variable environments. Here, we report a study in which we used an agent-based model to test Dunnell's hypothesis. We found that the propensity to waste was subject to strong negative selection regardless of the level of environmental variability. At the start of the simulation runs, agents wasted ca. 50% of the time but selection rapidly drove that rate down, ultimately settling at ca. 5–7%. This casts doubt on the ability of Dunnell's hypothesis to explain instances of cultural elaboration in the archaeological record.

Introduction

Evolutionary theory predicts that individuals will usually seek to maximize their reproductive success. One implication of this is that energetically expensive behaviours that do not contribute to reproductive success will be selected against and therefore should be uncommon. Yet, the archaeological record of the last 12,000 years is full of items that would have been energetically expensive to make and are not easy to explain in terms of reproductive success. Obvious examples include the henge monuments of Britain, the pyramids of the Maya region, and the beehive tombs of Greece. In the late 1980s, the well-known American archaeologist Robert Dunnell proposed a counter-intuitive solution to this apparent paradox (Dunnell 1989; see also Dunnell 1999 and Dunnell and Greenlee 1999), a solution that has since become known as the "waste hypothesis." According to Dunnell (1989), wasting behaviour produces an evolutionary benefit in highly variable environments. In the present paper we report a study in which we tested the key prediction of Dunnell's (1989) hypothesis with agent-based model (ABM) experiments. These experiments involved manipulating environmental variability and monitoring changes in the prevalence of wasting behaviour in successive generations of agents.

Background

Dunnell introduced the waste hypothesis in a book chapter titled "Aspects of the Application of Evolutionary Theory in Archaeology" (Dunnell 1989). In this paper, Dunnell referred to the construction of monuments, temples, and tombs as "cultural elaboration" (pg. 47). He argued that cultural elaboration has two effects that are adaptive. First, it decreases a population's size directly by diverting energy away from reproduction and this effectively creates a buffer between the population and the carrying capacity of its environment. In a highly variable environment, Dunnell (1989) proposed, the lower short-term birth rates of wasteful individuals result in lower near-term resource requirements, which gives wasters an advantage during environmental downturns because a reduction in carrying capacity is less likely to affect them. Over the long-term, the buffer leads to higher population growth rates among wasters relative to non-wasters. According to Dunnell (1989), cultural elaboration also produces an evolutionary benefit by acting as a sink for excess time. When environmental conditions worsen, time that would previously have been wasted can instead be used for resource acquisition. Given these two benefits, Dunnell (1989) posited, wasting can be expected to have a long-term adaptive advantage over not wasting.

Ten years after Dunnell introduced the waste hypothesis, Madsen, Lipo, and Cannon (1999) outlined an alternate version. They reframed the hypothesis in terms of "bet-hedging," which is an evolutionary strategy long studied by biologists (Gillespie 1973, 1977, 1974; Seger and Brockmann 1987; Roff

1992; Charlesworth 1980). Specifically, they argued that energetically wasteful behaviour is a kind of variance-reducing bet-hedging. They suggested that wasting energy rather than directing it into reproduction improves the chances of parental and offspring survival in highly variable environments because it reduces the number of offspring a parent needs to feed during downturns. While this leads to lower reproductive output in the short-term, it can lead to a greater number of wasters over the long term by lowering through-time variance in the fitness of such individuals.

The notion of adaptive waste has been used to explain a number of instances of cultural elaboration in the past. Hamilton (1999), for example, noted the coincident timing of the building of mounds in Mississippi, Louisiana, and Florida and changes in the El Niño Southern Oscillation (ENSO). Around 5000 BP, ENSO periodicity and storm severity increased creating more variable environmental conditions. According to Hamilton (1999), the increased variability coincided with mound building at Archaic period sites in the three states. Mound building, she suggested, would have been an energy sink that kept the population below carrying capacity – i.e. Dunnell's (1989) buffer, protecting the population from climate downturns – leading to increasing numbers of mound builders over time until the behaviour became common. Similar arguments have been made in relation to mounds elsewhere in the US (Dunnell and Greenlee 1999; Madsen 2001; Peacock and Rafferty 2013).

In another example, Kornbacher (1999, 2002) used the waste hypothesis to explain cultural elaboration in ancient Peru. She pointed out that Peru's environment is highly variable because it contains river valley oases within deserts, and experiences strongly seasonal rainfall, tectonic activity, avalanches, and volcanic eruptions. She also pointed out that the two most pronounced periods of cultural elaboration, Moche III/IV and the Middle Sicán, coincided with or occurred shortly after periods when environmental variability became more pronounced. Kornbacher (1999, 2002) argued that the cultural elaboration was wasteful behaviour and that it mitigated the negative effects of the highly variable environment.

Among the other cases of cultural elaboration that have been explained with the waste hypothesis are passage graves in Ireland (Aranyosi 1999), the Egyptian pyramids (Sterling 1999), and Classic Maya monumental architecture (Sack 2012). Monumentality, tattooing, and carving traditions in Polynesia have also been explained with the waste hypothesis (Graves and Sweeney 1993; Graves and Ladefoged 1995; Hunt and Lipo 2001; Genz and Hunt 2003; Allen 2010).

While the waste hypothesis has been used to explain a range of instances of cultural elaboration in the archaeological record, critical evaluations of the hypothesis have been few and far between. To the best of our knowledge, there has only been one attempt to assess the hypothesis in detail and that analysis only examined the bet-hedging version (Madsen, Lipo, and Cannon 1999). With this in mind, we decided to evaluate Dunnell's (1989) version of the waste hypothesis with an ABM, which is a computer programme that allows selection experiments to be conducted on virtual individuals attempting to live and reproduce in a virtual environment, the characteristics of which can be manipulated (Cegielski and Rogers 2016).

We opted to use an ABM in light of some of the challenges of evolutionary analysis. One of these is that the patterns produced by evolutionary systems can be difficult to predict a priori due to complex interactions between individuals and their environment. Another challenge is that evolution often involves many years, especially for long-lived organisms, making it hard to gather the real-world data required to evaluate evolutionary hypotheses. A third challenge of evolutionary analysis is that it can be difficult to distinguish selection-driven evolution from drift when looking at real-world data (Millstein 2008; Walsh and Lynch 2018). In part, this is because making such a distinction requires that we have a neutral trait to serve as a baseline for comparison, but identifying a neutral trait can be as difficult as identifying one that is under selection (Millstein 2008). ABMs can overcome all three of these challenges. They can incorporate any dynamics of interest so that we can observe the outcome of complex processes directly, and the timescales of evolutionary processes in ABMs are controllable so we can gather data relevant to both the short and long term. Moreover, ABMs can include a truly neutral trait to serve as a benchmark for identifying selective pressure on other traits. Together, these features of ABMs allow us to evaluate evolutionary hypotheses that would otherwise be difficult or even impossible to test with real-world data.

Dunnell (1989, 47) defined waste as energy that is expended in such a way that it "cannot be recovered at a later date" but which nonetheless has an adaptive role during environmental perturbation because it "lowers the birth rate" and creates "a reservoir of time that an organism can devote to subsistence and/or reproduction in difficult conditions." Note the stress on the individual in the latter quotation. This is important for attempts to test the hypothesis because it indicates that Dunnell (1989) viewed waste as adaptive for the individual not the group, i.e. it indicates he believed that the frequency of wasting behaviour in a population is the result of individual-level selection not group-level selection. Dunnell and Greenlee (1999) made this even clearer in a response to Neiman (1998). Neiman framed the waste hypothesis in terms of group-level selection, but Dunnell and Greenlee (1999) rejected this. Individual dietary differences among Late Woodland people, they argued, "serve to illustrate Neiman's

(1998, 286) error in characterizing the waste argument as inherently 'group selection'" (Dunnell and Greenlee 1999, 384). Thus, for Dunnell, waste is energy that is expended rather than stored or invested in a trade-off, and it will be favoured by selection at the individual level as environmental variability increases. Given this understanding of adaptive waste, the central prediction of Dunnell's (1989) waste hypothesis is that when environmental conditions are highly variable individuals who engage in wasteful behaviour should have greater long-term reproductive success than individuals who do not engage in such behaviour.

We tested this prediction by running simulation experiments with an ABM and analysing the resulting data with statistical methods from the field of quantitative genetics that are designed to identify selection. In the experiments, agents inherited a waste trait that affected their reproductive output, and we investigated how different levels of environmental variability affected the frequency of the waste trait in the population of agents over a large number of generations. We also tracked a neutral trait in the experiments, i.e. a trait that was only affected by drift. This trait was included to provide a null distribution for assessing whether significant changes in the frequency of the waste trait had occurred. Importantly, we used the lifetime reproductive success (LRS) of the agents – measured in terms of the lifetime number of adult offspring – to evaluate the fitness of the traits.

Methods

The model

We developed our ABM in NetLogo (Wilensky 1999). The overview, design concepts, and details of the ABM are given in Supplemental Material 1 in the form of an ODD document (Grimm et al. 2010). The code can also be downloaded at https://github.com/wccarleton/abm_waste.

Our model included a set of agents and a dynamic environment. Each agent had several traits (Table 1) and they followed a set of routines that determined how they interacted with their environment and other agents. Their environment was abstract, aspatial, and fluctuated through time according to a simple of model of environmental variability. We explored the dynamic relationship between the agents and their environment by running simulations. Each simulation involved multiple time-steps, which are called *ticks* in NetLogo. Each tick, the agents performed the following operations in the order presented:

(1) *Age.* Each agent incremented their age variable by 1. The founding population began each simulation at maturity, a setting determined at the outset and one that we varied among the values 1, 5, and 10 in the course of the experiments.

(2) *Track offspring.* The agents updated their internal list of offspring, adjusting their counts of juvenile and mature offspring as necessary. The agents kept track of their offspring so they could provision any juveniles they had, and so we could track their lifetime reproductive success.

(3) *Gather energy.* Each tick, agents gathered energy, which was an abstract resource the agents needed for survival and reproduction. The amount of energy obtained by each agent was a function of the number of other active agents and the base level of agents supported by the environment (see ODD). This base level was held constant for all of the experiments. As the number of agents increased, the amount of resources available to each agent declined until the population reached the carrying capacity of the environment. Each tick, however, the carrying capacity was perturbed by a simple autoregressive function (see ODD) so that we could explore the impact of through-time environmental variability on waste. When an agent gathered energy, they overwrote their existing energy level with the newly acquired amount. This was done because storage was not permitted in the model. We excluded storage as an option because, as we explained earlier, Dunnell (1989, 47) explicitly stated that waste differs from storage in that it cannot be recovered at a later date.

Table 1. Parameters of the agent-based model, their potential values, whether they were fixed for all experiments, and a brief description of each one.

Variable	Value(s)	Fixed?	Description
Number of ticks	2000	Y	The length of the simulation
Initial number of agents	1000	Y	Number agents at the beginning of a simulation
Base agents	2000	Y	Base agents permitted – determines base carrying capacity before accounting for environmental variation and maturation period
Mature	1,5,10	N	The number of ticks for which agents are immature and require provisioning from parents
Provision	Random	Y	Inheritable provisioning strategy, randomly assigned to the founding population
Consume cost	1	Y	Amount agents have to consume every tick to stay alive
Reproductive cost	1	Y	Amount of energy required for agents to reproduce
Probability of waste	0–1	N	The probability of waste trait. Each agent can have a value for this trait in the range 0–1 and the starting agents are assigned values randomly with a uniform probability.
Waste drift	0.01	Y	Standard deviation of the random normal distribution characterizing mutation in the probability of waste
Climate AR	0.3	Y	Autoregression parameter for the autoregressive climate process
Climate standard deviation	0.3,0.5	N	The standard deviation of the simulated autoregressive climate process.

(4) **Consume and provision.** Agents consumed energy for survival at a rate determined by a variable that we held constant at 1-unit for all experiments. This amount was chosen so that the base level of agents would be in units of energy available to the population and equivalent to the number of agents that could be supported in the environment – e.g. assuming a base level of 2000 and a current census of 1000 agents, each agent would have two units of energy available for consumption and reproduction, before accounting for environmental change. If an agent had less than one unit of energy available when they performed this step, they died. If an agent died while supporting juvenile offspring, those offspring also died. If a given agent collected more energy than they needed for their own survival, they provided energy to their juvenile, dependent offspring following one of four heritable strategies:

  (a) **Even.** The remaining energy was divided evenly among all of the agent's juvenile offspring.

  (b) **Random.** The agent randomly sorted their list of juvenile offspring and provisioned first to the top of the list and then sequentially to the other juveniles in the list. The agent would always provision an amount up to that required for survival for a given juvenile before moving on to the next juvenile in the list.

  (c) **First-born-biased.** The given agent first provided energy to their oldest juvenile offspring. They gave an amount up to that required for survival (1-unit) before provisioning other offspring in order of descending age.

  (d) **Last-born-biased.** The agent first provisioned their youngest juvenile offspring. They provided an amount up to that required for survival (1-unit) before provisioning their other offspring in order of ascending age.

(5) **Waste.** The instances of cultural elaboration that have been explained with the waste hypothesis range from tattooing to pyramid building (Allen 2010; Sack 2012). Because the costs of these clearly vary along a continuum, we decided that the trait that affects wastefulness should be a continuous one. To reflect this, we used a probability to determine wastefulness. The agents performed a Bernoulli trial – i.e. a virtual coin toss – with a probability equal to their heritable waste trait. The possible outcomes of the trial were *waste* and *not waste*, which determined whether the agents wasted or not.

(6) **Reproduce asexually.** Assuming an agent did not waste, they could devote their remaining energy to reproduction. If their remaining energy was equal to or greater than the reproduction cost of 1-unit of energy (also a constant throughout the experiments) they produced a single offspring. When the agents reproduced, they passed on their waste and neutral traits to their offspring with a small mutation, and they passed on their provisioning strategy with no mutation (see ODD). The starting agents (i.e. the founding population) began the simulation with randomly chosen values for their waste and neutral traits drawn from uniform distributions that spanned 0–1, corresponding to the minimum and maximum values for the traits. They were also randomly assigned one of the four provisioning strategies. Differential reproductive success among the agents meant that all heritable traits were potentially subject to selection.

The simulation experiments

In the experiments we explored the impact of two parameters while keeping the others constant. One of the parameters we manipulated was the variance of the carrying capacity. Climate-driven environmental variability was approximated with a time-series from a simple autoregressive process. Although real world climate processes are more complicated than this, we were primarily interested in variability rather than trends and so chose to limit the sources of uncertainty in the model by using a simple abstract climate process. The process produced a random time-series with values drawn from a normal probability distribution where the value at a given time correlates with the value at previous times. Importantly, the mean of the process only fluctuated over short intervals and was itself mean-reverting so that the long-run mean of the process was stable. We set this long-run mean at 1 for all experiments. We then controlled the variance in the time-series by changing the standard deviation of the autoregressive process, which we varied between two values: 0.3 and 0.5. These values represented climatic regimes with lower and higher variability, respectively. We selected these values to assess the effect of increasing environmental variability without allowing it to wipe out the population (initial exploration showed that if the variance was too high, climate shocks killed off all the agents). The absolute value of the climate time-series at a given tick was then multiplied by the base level of agents to produce the carrying capacity for that tick. The nature of the climate process meant there were runs of ticks during which the carrying capacity was high and other runs during which it was low (see ODD document). Hence, agents were exposed to both periods of energetic abundance and leaner periods that would limit their ability to survive and reproduce.

The other parameter we manipulated was the agents' maturation period, which was the length of

time over which they needed to be provisioned by their parent. Dunnell did not discuss provisioning explicitly in any of the papers in which he describes the waste hypothesis (Dunnell 1989, 1999; Dunnell and Greenlee 1999). When describing the adaptive waste mechanism, he referred only to "investment in reproduction" (Dunnell 1999, 246), and to a "sink of 'excess' time and resources that could be devoted to subsistence/reproduction" (Dunnell 1999, 245). So, it is not clear that he thought provisioning was necessary for adaptive waste. Nonetheless, provisioning is a critical form of reproductive investment for humans and it seemed possible that the length of the maturation period could affect selection on waste by increasing the odds that an agent experiences an environmental downturn while provisioning its immature offspring. Given this, we decided to include maturation period in our model. During this period, agents were entirely dependent on their parents for survival and could not reproduce. In order to account for the possible effect of maturation on the adaptiveness of waste, we varied agents' maturation period among three values: 1, 5, and 10. The first of these period lengths – 1 tick – represents the case in which no provisioning occurs because agents become mature in the tick immediately following their birth, before the consume-and-provision step of the model. The other two settings, 5 and 10 ticks, represent short and long maturation periods, respectively.

As each simulation progressed, the propensity for waste trait and the neutral trait evolved. The former was affected by both mutation and selection, while the latter was only impacted by mutation. For both traits, mutations were determined by random draws from a zero-mean normal distribution with a standard deviation of 0.01. Thus, a given agent's inherited waste probability was that of their parent's plus or minus a small random amount of usually less than 0.01. Both traits had to remain between zero and one. So, if the mutation caused the inherited value to be less than zero, the "overflow" was reflected back to produce a value greater than zero. A similar reflection occurred if the value was greater than one. Since the propensity for waste trait was inherited with random mutation, over time the environment selected for a range of optimal values in those inherited traits. These dynamics led to an adaptive change in the distribution of trait values within the agent population. Importantly, because the neutral trait did not impact the agents' behaviour, changes in the population level statistics of that trait were due to fluctuations in population size (i.e. drift) rather than to selection.

There were several sources of randomness in the model. These included the random mutation process that affected trait heritability; the climate series, which was a random realization of an autocorrelated process; and the random assignment of starting values for inheritable traits in the founding population. To account for these sources of randomness, we used long simulations (many ticks) and we replicated each experiment multiple times. It should be noted that there is a trade-off here: achieving greater certainty about parameter values requires longer simulation runs and more replications, which greatly increases computation time. To strike a balance, we set individual simulations to run for 2000 ticks and re-ran each experiment hundreds of times. This approach allowed us to find stable estimates of relevant parameters like LRS and optimal waste probability, while also observing the long-run dynamics of the model. It also allowed us to find distributions for quantitative estimates of selection direction and strength. Occasionally, simulations would end before reaching 2000 ticks because the agent population died out due to extreme environmental perturbations; these undershot runs occurred more frequently when environmental variability was high. When undershoots occurred, we ran whatever number of replications were necessary to obtain a minimum of 100 samples for analyzing the ABMs long-term dynamics and a minimum of 200 samples for obtaining quantitative estimates of selection statistics.

The analyses

The primary goal of our analysis was to determine whether environmental variability selected for waste in the manner proposed by Dunnell (1989). We needed a framework for determining that selection had occurred in a given experiment and for characterizing the nature of that selection. For this, we turned to the field of quantitative genetics (Walsh and Lynch 2018). In quantitative genetics, evolution is defined as a statistically significant change in the distribution of a trait in a target population. Such a change can potentially be indicative of selection, especially if the quantitative trait in question correlates with a measure of fitness. Given this framework, we reasoned that if selection favoured waste we should see (1) a positive correlation between waste and the fitness metric we used, LRS; (2) significant changes in the distribution of waste propensity from the founding population of agents; and (3) the mean of the waste trait distribution should be well above the minimum possible level given the dynamics of the model. Importantly, these patterns should be evident in a high-variability environment even if they are not evident in a low-variability one.

We began by plotting the time-series of waste and neutral trait means for each experiment. The series show the common mean of the trait distributions in a given tick. These mean series were plotted individually – one series for each simulation run in a given experiment. We used somewhat transparent lines so that portions of the parameter space most often visited by the simulations in a given experiment would appear

darker than other regions of the same space. Then, we calculated the pointwise median of all mean series at a given tick in each experiment. The series of pointwise medians – the grand median – was plotted in white to distinguish it from background of individual mean time-series. These time-series plots were used to visually assess the short- and long-term dynamics of the neutral and waste trait distributions for each experiment while accounting for the variation between individual simulation runs.

Next, we produced fitness landscape plots for each experiment. Fitness landscapes show the relationship between fitness and a given trait (Walsh and Lynch 2018). In the present case, these plots are simply point scatters showing the LRS and waste probability trait value of the adult agents from a given experiment. To produce the scatters, we used a sample of simulation runs, one from each experiment.

Subsequently, we determined whether there was evidence that selection had occurred in each experiment. This was done with a recently developed metric called the distributional selection differential (DSD) (Henshaw and Zemel 2017; Walsh and Lynch 2018). This metric measures overall changes in quantitative trait distributions. It is sensitive to changes in mean, variance, and shape, which means it can be used to identify both linear and nonlinear selection. The metric can be calculated in several ways (Henshaw and Zemel 2017); we used the method that involves the "covariance definition" of the DSD. According to this, the DSD can be estimated with the covariance between a target quantitative trait and a measure of fitness. As we mentioned, the measure of fitness we used is LRS, quantified by a given agent's lifetime number of adult offspring. This measure was appropriate because our model and the waste hypothesis focus on selection of parental phenotypes – i.e. the waste trait of a parent not the traits of the offspring. It is also only those offspring who survive to adulthood that will contribute to the waste trait distribution of future generations. Hence, the number of offspring successfully raised to maturity is the most appropriate measure of fitness for present purposes.

We calculated the DSD for both the waste and neutral trait in each of the simulation runs for each experiment. This resulted in samples of 200 DSDs per trait per experiment. Then, we used histograms to compare the distributions of DSD samples between the waste and neutral traits. Changes in the neutral trait, measured with DSD distributions, reflected the effect of pure drift on the evolution of traits in the population. Drift occurred when mature agents died of starvation. This only happened when a random environmental perturbation suddenly decreased the carrying capacity, at which point a portion of mature agents were unable to gather sufficient resources for survival, reproduction, and provisioning. The specific agents were chosen randomly – determined by their order in the "gather energy" step of the ABM. Consequently, changes in the population distribution of the neutral trait caused by this random culling are purely a product of drift. Thus, the neutral trait allowed us to determine whether any observed changes in the waste trait – as measured by DSDs – could be differentiated from changes that would be expected if the waste trait were not under any selective pressure. Importantly, we calculated the DSD using the founding population in each experiment. This was because the founding population is the only cohort that represents pre-selection conditions.

While the DSD is a useful general measure for identifying selection, it does not indicate the direction of selection. To determine the latter, we used selection gradients (Walsh and Lynch 2018). A selection gradient is an estimate of the rate of change in fitness given a unit change in the mean of a quantitative trait. This measure naturally includes a sign and is often estimated using a least-squares approach that fits a linear model to fitness and trait data – essentially a linear regression wherein a given quantitative trait is used to predict fitness. In the case of the waste ABM, the gradient estimates the relationship between agent LRS and waste probability. A negative gradient (i.e. a negative slope of the regression line) would indicate a decrease in the trait mean, while a positive gradient would indicate an increase. We calculated the selection gradients for the waste trait in each experiment using the same software we used for the DSD and with the same data from the founding populations. We also plotted the LRS of agents against their waste and neutral traits for each experiment.

Thereafter, we used selection differentials to quantify the strength of selection. A selection differential is the change in a quantitative trait mean after a selection event (Walsh and Lynch 2018). They are measured in units of the trait – probability, in the present case – and so can also be used to estimate the strength of selection. We did this by calculating the ratio of the magnitude of the differential to the maximum possible monotonic change in the mean of the waste trait. We began each simulation with a uniform distribution of waste in the agent population that had a mean of 0.5. The largest possible monotonic change in this mean was the difference between its starting value and the minimum/maximum value possible given the dynamics and constraints of the model. The minimum mean for the waste trait in the experiments we conducted was ∼0.007 and the maximum was ∼0.993 (see the ODD). Thus, the largest magnitude of monotonic change possible for the mean of the waste trait distribution was ∼0.493. This is the upper limit on the response to selective pressure that could have been experienced by the founding population. A negative change of that magnitude, for example, could only

have occurred if every agent with a waste trait greater than zero failed to reproduce, which would indicate extremely strong negative selection against waste. Ratios of the initial mean to this upper limit that are close to one indicate that nearly the strongest conceivable selective pressure was present, while ratios near zero indicate very weak selective pressure was present.

Lastly, we compared the waste trait distributions to the theoretical lower-limiting distribution for that trait. Given the model's dynamics for inheritance and the boundaries of the waste trait, the lower-limiting distribution for the waste trait is half-normal with a scale parameter equal to $(1/\sigma)$, where σ is the standard deviation of the normal distribution used to mutate the heritable traits in the ABM (see the ODD). In the six experiments, the mean of this lower-limiting distribution was ∼0.007. This is what we would expect to see if only agents with a waste trait of zero were permitted to reproduce. The lower limit is a distribution with non-zero mean instead of simply zero because the mutation process is random and continuously injects non-zero values for the waste trait into the system. It is important to point out, though, that this lower-limiting distribution does not account for the random environmental fluctuations in the model. Functionally, the lower limit may be higher than this estimate because of environmental upswings that reduce selective pressure. Still, the theoretical lower limit served as a point of reference against which we could compare the observed level of waste in a given experiment.

All the analyses were conducted in R (R Core Team 2019). We calculated the DSD with the aid of software developed by Henshaw and Zemel (2017) and made available on Github (https://github.com/jjono/DSD).

Results

As we explained earlier, at the core of Dunnell's (1989) waste hypothesis is the idea that, in a highly variable environment, individuals who suppress their reproduction by wasting will have lower offspring mortality than individuals who waste less and have a lot of offspring. To ensure that the dynamics of our model captured this, we selected one of the simulation runs and created scatter plots that compare propensity to engage in wasting behaviour to two key reproductive variables – the lifetime number of immature and mature offspring, or total reproductive output (TRO), and the fraction of offspring who survived to adulthood, or lifetime reproductive success (LRS) (Figure 1). The plots show that low-waste agents have many more offspring over a lifetime than high-waste agents, but also that low-waste agents suffer a much higher immature offspring mortality rate. We examined comparable plots for a large number of simulations involving different simulation parameters and they all showed the same pattern. Additional plots we created demonstrated that the effect is accentuated as environmental variability is increased (Supplementary Material 2). Thus, the model's dynamics captured the waste hypothesis' central idea.

Figure 2 illustrates what happened to the average frequencies of the waste and neutral traits in the course of the experiments. In each panel, the black lines track the frequency of the traits in a single simulation, while
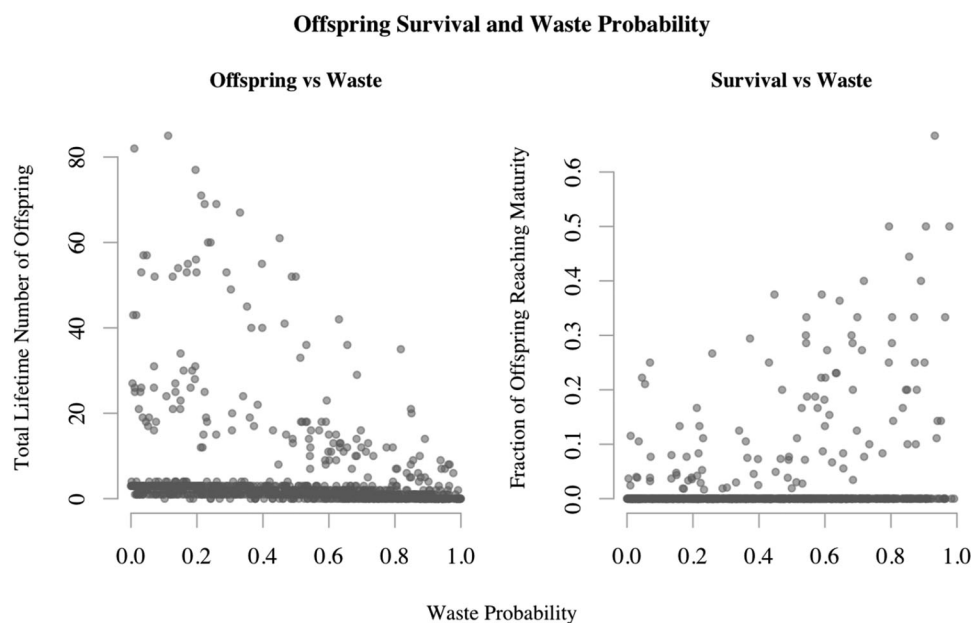

Figure 1. Scatter plots comparing waste propensity to two key reproductive variables. One variable (left plot) compares waste to the lifetime number of immature and mature offspring, or total reproductive output (TRO); and the other (right plot) compares waste to the fraction of offspring who survived to adulthood, or lifetime reproductive success (LRS).
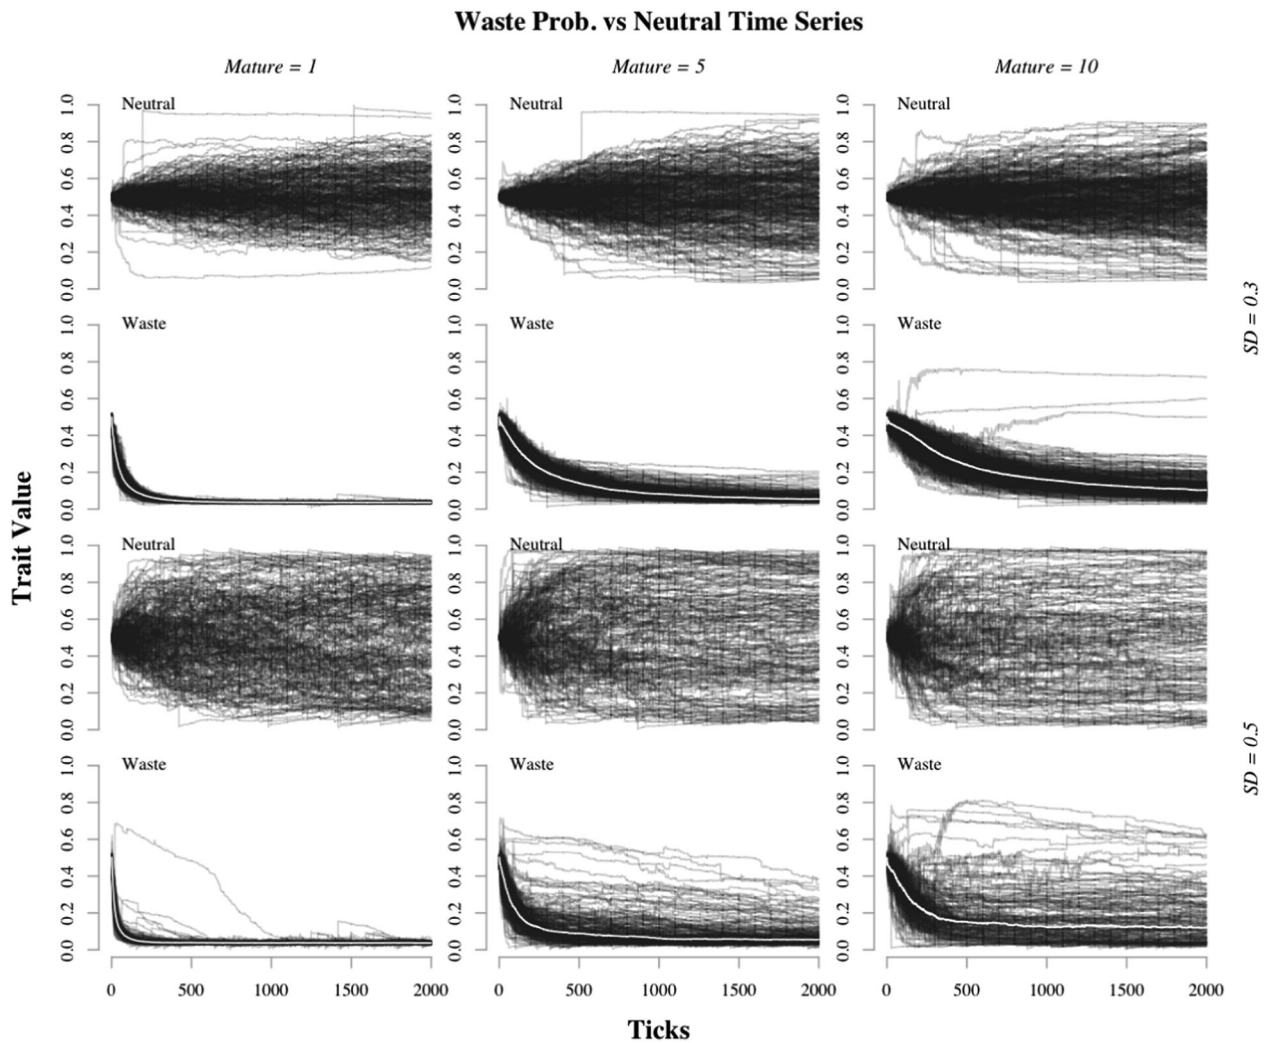
Waste Prob. vs Neutral Time Series

Figure 2. Time-series plots of the mean neutral and mean waste traits at a given tick in a given experiment. The dark lines represent the results of individual simulation runs while the white lines represent the median of the sample of means.

the white line is the grand median of the means of the traits across all the simulations in a single experiment. The grand median approximates the optimal, long-run level for the trait after accounting for model stochasticity and random drift.

The grand medians for the neutral trait are more or less constant. There is an increase in the variance of the neutral trait through time, but the grand medians are always around 0.5. The pattern for the waste trait is markedly different. In all six experiments, the grand median for the waste trait decreases. This suggests that wasting behaviour was subjected to negative rather than positive selection, which runs counter to the predictions of Dunnell's (1989) hypothesis.

The rate of decline in average waste propensity varied depending on the agents' maturation period. It was fastest when the maturation period was 1, intermediate when the maturation period was 5, and slowest when maturation period was 10. This was not a consequence of increasing maturation period moderating the selection pressure against waste. Rather, the patterns of change in the waste trait (Figure 2) and the age-at-

death distributions (Figure 3) show that it was due to slower population turnover when maturation periods were longer. Agents only died of starvation in the ABM; they did not die from old age. Consequently, when agents were able to obtain enough resources to live but not enough to provision offspring, average adult agent mortality decreased, lifespans increased, and the population turned over more slowly. This occurred more often in experiments with longer maturation periods because they provided more opportunities for offspring to die of starvation.

Comparing the results of the low and high variability experiments reveals another pattern that goes against the predictions of Dunnell's (1989) waste hypothesis. Figure 2 demonstrates that regardless of the maturation period, the impact of negative selection on wasting behaviour was larger in the high environmental variability condition than in the low environmental variability one. Given that the waste hypothesis specifically argues that wasting behaviour is an adaptive response to high environmental variability, this is the opposite of what we should see if the hypothesis is correct.
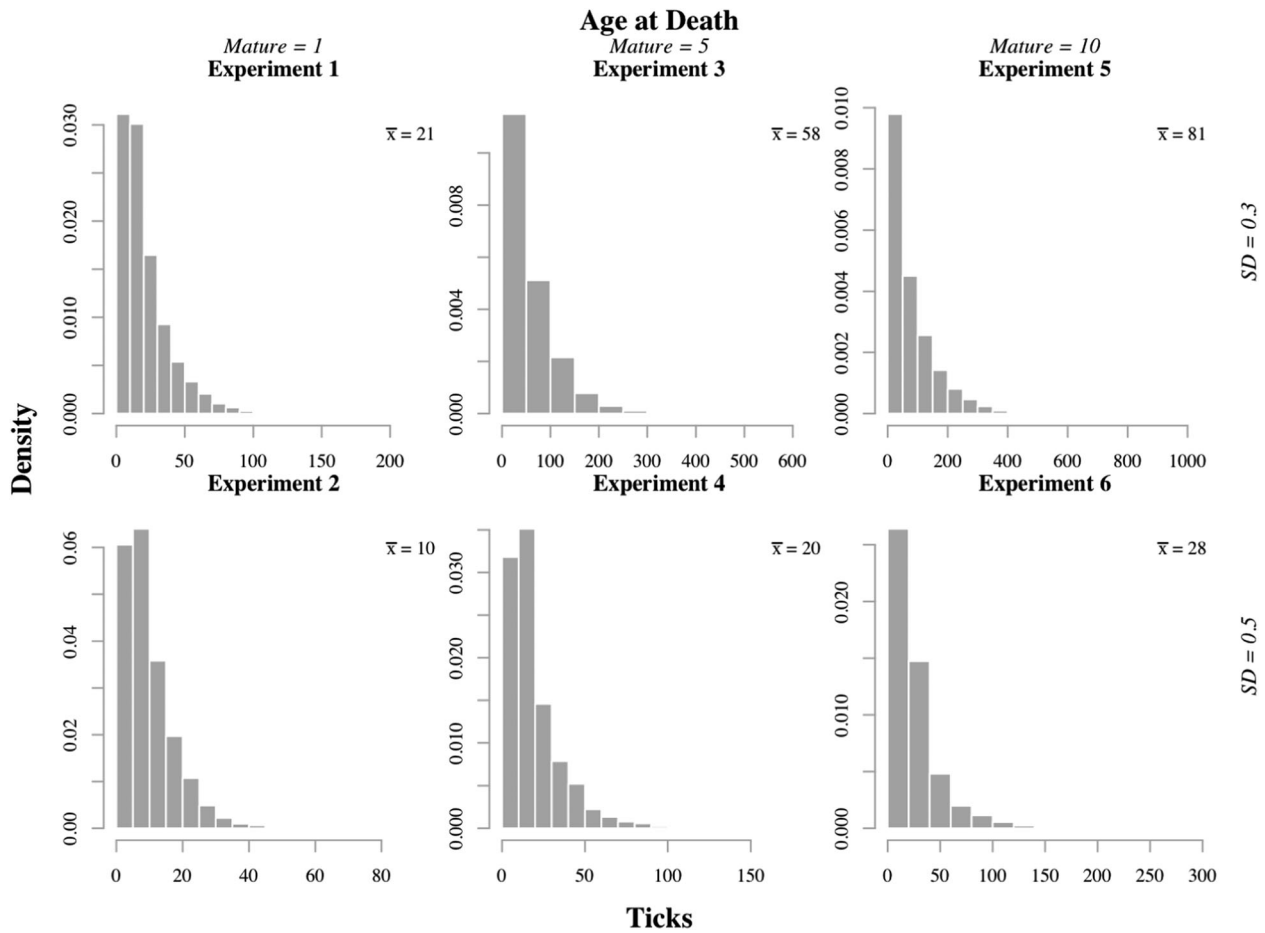
Figure 3. Age at death distributions based on all agents from all simulation runs in a given experiment.

Figures 4 and 5 are fitness landscapes for the neutral and waste traits derived from the data yielded by the six experiments. Figure 4 shows the LRS and neutral trait values for all adult agents that ever "lived" in the experiments; Figure 5 presents the same for the waste trait. In Figure 4 it is obvious that there is no correlation between the neutral trait and LRS, which is as it should be, given that the trait was only affected by drift. In contrast, Figure 5 shows that there is a strong, negative relationship between LRS and the waste trait. While there is substantial variability in LRS in each simulation, and many agents never manage to successfully rear any offspring to adulthood, a greater propensity to waste is associated with lower LRS in all the experiments. Thus, this set of results also contradicts the predictions of Dunnell's (1989) waste hypothesis.

Figure 6 shows the DSD distributions of the neutral and waste traits. To reiterate, when compared to a neutral baseline, DSD values can reveal the occurrence of selection by showing that a change in a trait is sufficiently large to rule out drift as the cause. In Figure 6, all the waste DSD histograms (dark grey) are right-shifted relative to the neutral DSD histograms (light grey). The difference between the two distributions decreases as maturation period increases but it never disappears. As such, the distributions indicate that there is a difference in DSD estimates between the

waste and neutral traits. This in turn indicates that the changes in the waste trait distributions are likely due to selection rather than drift, which is contrary to the predictions of Dunnell's (1989) hypothesis.

Figure 7 presents the distributions of the waste trait selection differentials. As we explained earlier, selection differentials indicate the mean change in fitness per unit change in a quantitative trait. Figure 7 shows that almost all the selection differentials were negative with few exceptions. This means that selection drove down the frequency of the waste trait in nearly all the simulation runs. The ratio of the magnitude of the differentials to the maximum possible change in mean waste probability indicates that the selective pressure was strong (Figure 8). The averages of the ratios ranged from 0.10 to 0.34, which translates into negative changes of anywhere from 10% of the maximum to 34% of the maximum in a single generation. Selection on propensity for waste, therefore, was both negative and strong in the experiments. This is another finding that is strongly inconsistent with the predictions of Dunnell's (1989) hypothesis.

Figure 9 compares the minimum theoretical distribution of waste in the ABM with the observed distributions for each experiment. The light grey curve is the minimum theoretical distribution for the waste trait given the inheritance and mutation dynamics of
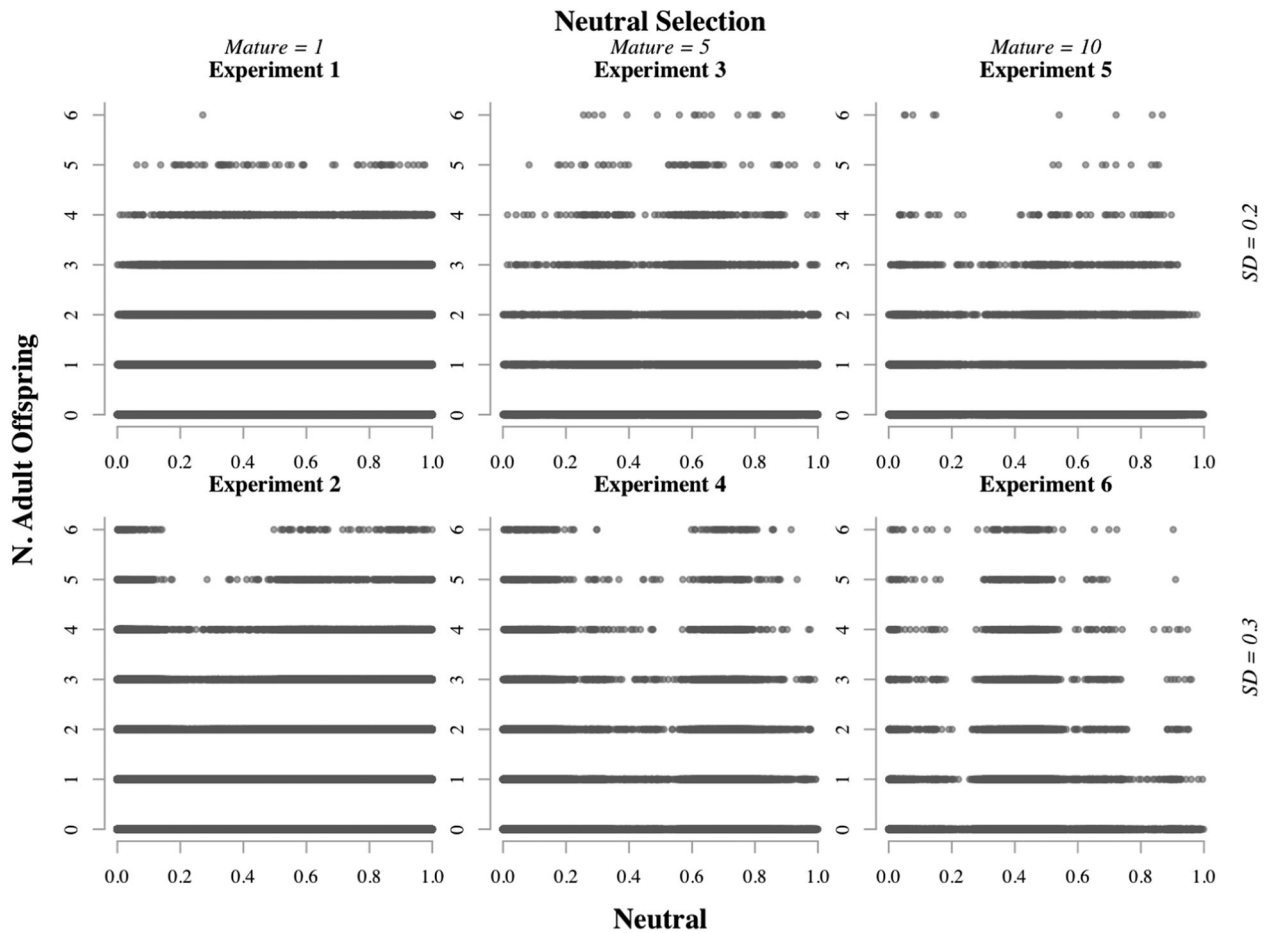
Figure 4. Plots of lifetime reproductive success versus the neutral trait for a randomly selected simulation run from a given experiment.

the ABM (see ODD). The dark grey histogram shows the sample distribution of waste trait values for the last 500 ticks of a sample simulation from the relevant experiment. While the empirical distributions are not in the lowest possible positions, the plots show that the long-run stable waste distributions were near the lower limit in all cases. This indicates that the long-run optimal level of the waste trait was very low. Once again, this is a finding that is contrary to the predictions of Dunnell's (1989) hypothesis.

Discussion

In the study reported here we used an ABM and statistical methods from the field of quantitative genetics to test the most important prediction of Dunnell's (1989) original version of the waste hypothesis – namely that wasting behaviour should lead to higher lifetime reproductive success for individuals when environmental variability is high. The results we obtained were inconsistent with this prediction. We did not find any evidence that selection favoured waste, even when environmental variability was high. Rather, our analyses indicated that there was always selection against waste, and the selection was strong. We found that average waste levels were reduced by

anywhere from 10% to 34% within a single generation. In addition, the experiments indicated that the optimal level of waste was low regardless of how variable the environment was. In the vast majority of simulation runs, the average waste propensity dropped from the starting value of 50% to between 0% and 5%. Strikingly, the declines in average waste propensity were more rapid when environmental variability was high than when it was low, which is the opposite of what Dunnell (1989) suggested should happen. We contend that, taken together, these results cast considerable doubt on the ability of Dunnell's (1989) version of the waste hypothesis to explain ancient monuments and other forms of cultural elaboration in the past.

Supplementary exploration of the model revealed that it was possible to induce selection for wasting behaviour by manipulating the provisioning strategy variable (Supplemental Material 2). To reiterate, at the beginning of each simulation run in the main experiments, the founding agents were each randomly assigned one of four provisioning strategies: (1) *even*, according to which an agent would divide their energy evenly among their immature offspring; (2) *random*, which involved randomly ordering immature offspring and then provisioning the first in the list up to the amount required for survival before moving onto the
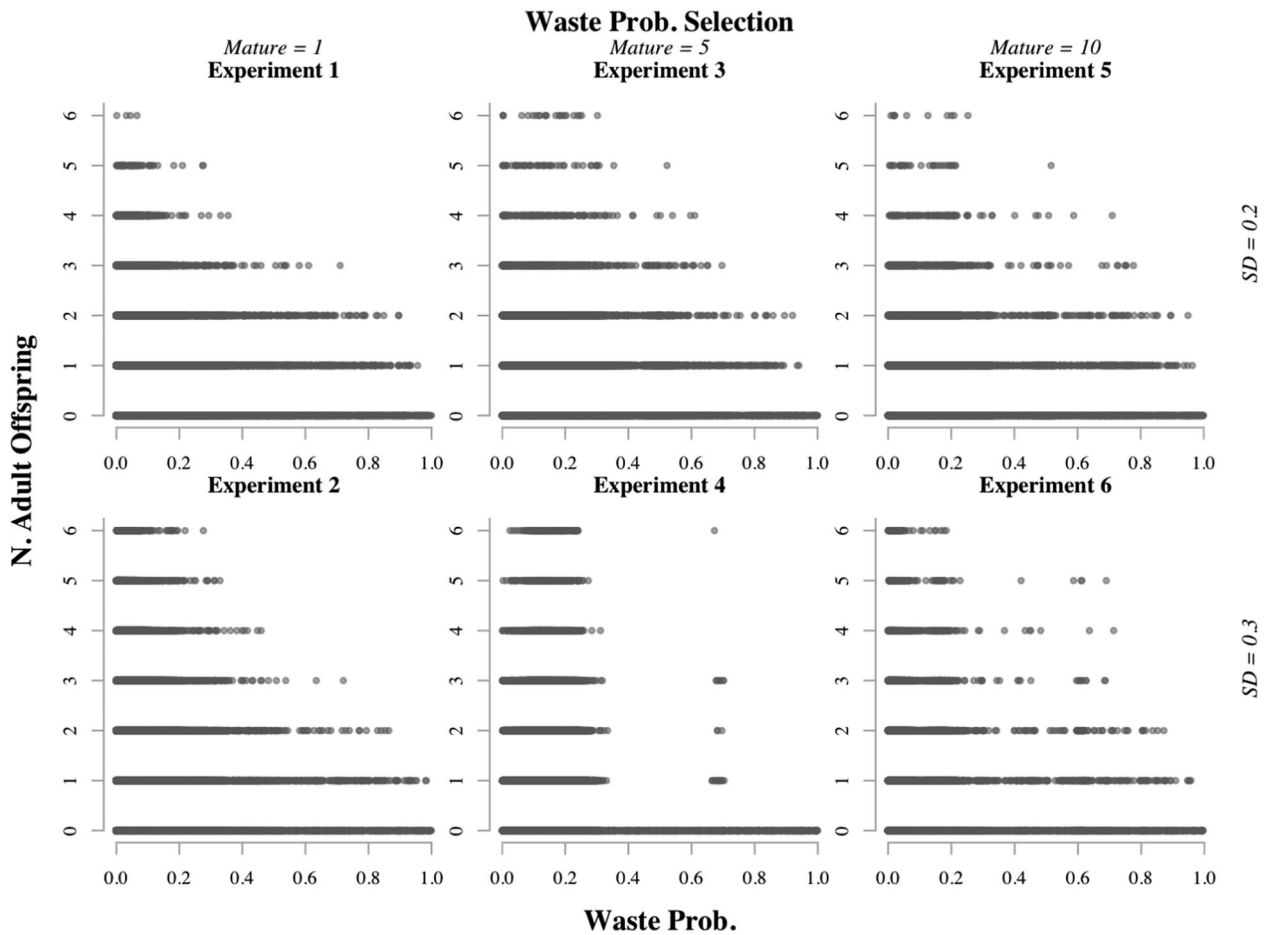
Figure 5. Plots of lifetime reproductive success versus the waste trait for a randomly selected simulation run from a given experiment.

next offspring in the list; (3) *first-born biased*, according to which an agent would first provision their eldest immature offspring up to the amount required for survival before provisioning the next eldest and so on; (4) *last-born-biased*, the opposite of first-born-biased. The provisioning strategy assigned to a founding agent was passed on to its descendants without mutation, so the frequency of the different provisioning strategies in the population was governed by selection and, if population size decreased substantially as a result of environmental downturns, by drift. In all the main experiments, it was the first-born-biased provisioning strategy that was favoured by selection (Figure S1). In the majority of runs, it was the only strategy in use in the population within a few generations. In the supplementary exploration of the model, we fixed the provisioning strategy of the founding population to one of the four possible strategies, thereby making it impossible for selection to choose among them. When we assigned all agents the *first-born-biased* strategy, the simulations returned the same result as the main experiments, i.e. strong selection against wasting behaviour. However, when all agents were assigned one of the other strategies, there was selection for a non-minimum level of waste. This is important because it demonstrates that the model and experiments did not

exclude wasting behaviour as a consequence of our decision-making. That it was possible for us to force the model to favour wasting behaviour means that the main experiments could have in principle supported Dunnell's (1989) hypothesis and that they failed to do so because wasting behaviour was not adaptive.

Given that the supplementary simulations show that the results of the main experiments were heavily dependent on selection having favoured the *first-born-biased* provisioning strategy, it is worth considering whether this scenario is realistic for human populations. There are both theoretical and empirical reasons to think that it is. Theoretically, selection is likely to favour first-born-biased provisioning in species that usually give birth to only one offspring at a time and face intermittent resource shortages, as is the case with humans (Hrdy and Judge 1993; Hertwig, Davis, and Sulloway 2002; Jeon 2008; Bu and Sulloway 2016). Under such conditions earlier-born offspring are more likely to survive to maturity than their later-born siblings. This is because they will experience an initial period when they are the sole recipient of resources, which improves their fitness relative to later-born offspring, making reinvestment the rational choice (Hertwig, Davis, and Sulloway 2002). Additionally, earlier-born offspring will have less time
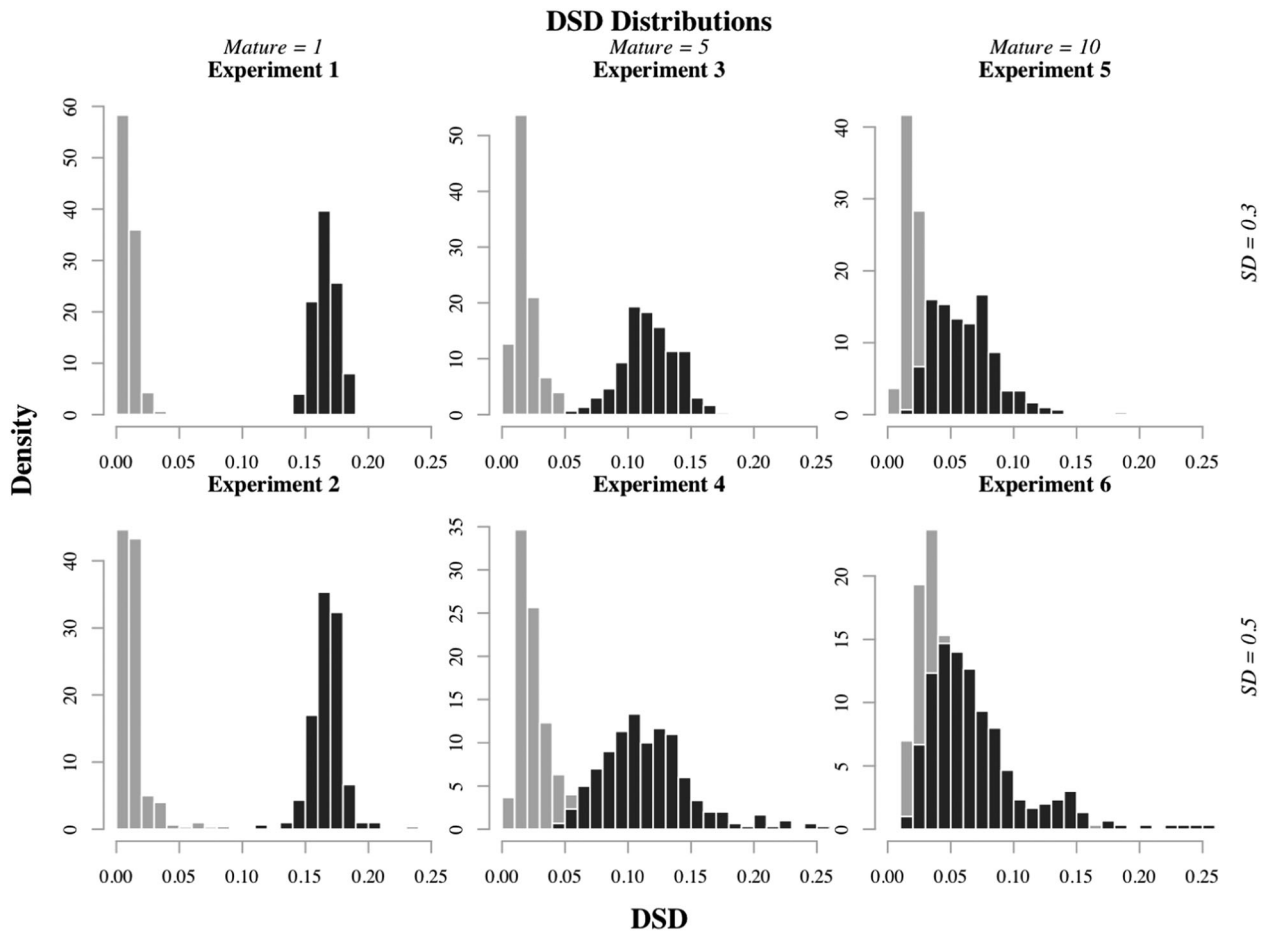
Figure 6. Histograms of DSD values for the neutral trait (light) and waste trait (dark) for a given experiment.

remaining in their maturation period than later-born ones. The time remaining to maturity affects the odds of survival because more time as a dependent makes suffering at least one fatal parental resource shortfall more likely. Thus, even though other provisioning strategies are possible, the one that is theoretically most likely among humans is first-born-biased provisioning. The available empirical evidence is consistent with this theoretical argument. Humans have been found to provide first-borns with more attention (Price 2005), extra nourishment (Horton 1988), better healthcare (Hertwig, Davis, and Sulloway 2002), and/or larger inheritances (Hrdy and Judge 1993; Mechoulan and Wolff 2015). Importantly, recent research has also shown that first-born-biased investment among humans has real fitness consequences. A recent large-sample study of pre-industrial Finns (Faurie, Russell, and Lummaa 2009), for example, determined that first-born males have a higher lifetime reproductive success than later-born offspring males. Thus, selection's preference for the first-born-biased provisioning strategy in our experiments is not only what we would expect based on theory but is also consistent with what we see in many human populations. Needless to say, this adds weight to our contention that our experiments cast doubt on the validity of Dunnell's (1989) hypothesis.

There have now been two formal evaluations of the waste hypothesis – the present study and Madsen, Lipo, and Cannon's (1999) assessment of their bet-hedging version of the waste hypothesis, which also employed an ABM. Intriguingly, the results of these studies point in different directions. While our results imply that Dunnell's (1989) version of the waste hypothesis is not valid, Madsen, Lipo, and Cannon's (1999) results suggest the opposite with regard to their bet-hedging version of the hypothesis. There would seem to be two obvious potential explanations for this difference. One is that the bet-hedging version of the hypothesis captures at least one key factor that Dunnell's (1989) version does not and that factor creates conditions under which waste is adaptive. The other possibility is that Madsen, Lipo, and Cannon's (1999) test of the hypothesis returned a false-positive result. At the moment, we are not in a position to determine which of these possibilities is correct because the software for Madsen, Lipo, and Cannon's (1999) model is no longer available (Carl Lipo, personal communication). However, we do want to draw attention to two features of Madsen, Lipo, and Cannon's (1999) model that could be responsible for the difference in outcomes.

In Madsen, Lipo, and Cannon's (1999) ABM, agents gathered energy from their environment and spent it on survival and reproduction. The agents
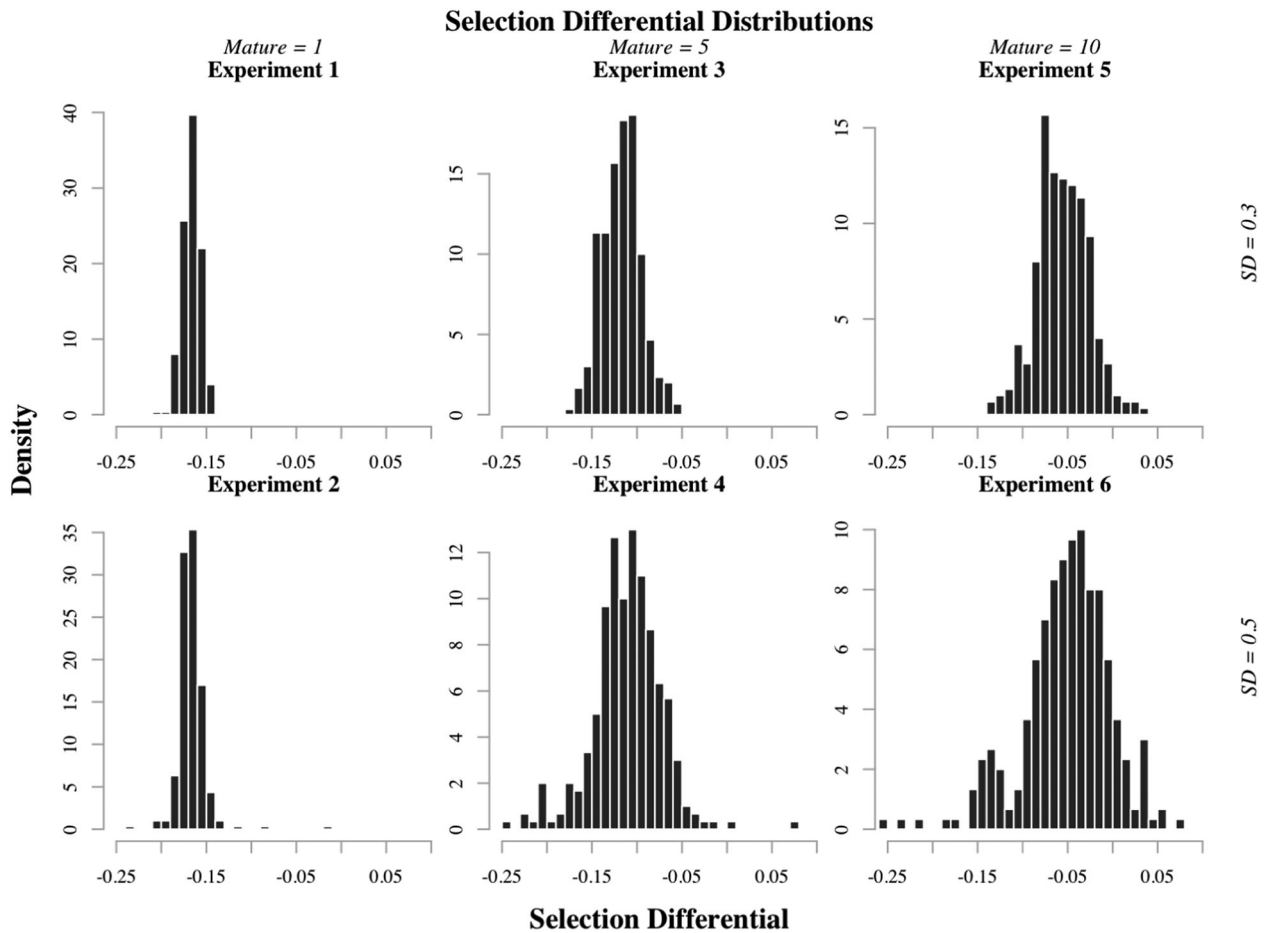
Figure 7. Histograms of selection differentials for the waste trait for a given experiment.

also acquired traits via a dual-inheritance system. One inheritance pathway involved inheritance of a suite of immutable "genetic" traits transmitted from parents to offspring during reproduction. The other inheritance pathway involved a set of potentially mutable "cultural" traits transmitted horizontally between agents during random encounters. Agents initially inherited both their genetic and cultural traits from their progenitors, but they could then change their cultural traits as they interacted with other agents. The constellation of genetic and culturally inherited traits constituted a given agent's phenotype. Importantly for testing the waste hypothesis, one of the cultural traits that could be inherited via transmission was costly to the receiving agent – i.e. the receiving agent lost energy to acquire the trait, representing "wasteful" investment in cultural elaboration. Another culturally inherited trait determined the amount of energy a given agent would be willing to waste during such a transmission.

Madsen, Lipo, and Cannon (1999) ran two main experiments to test the bet-hedging version of the waste hypothesis. One involved a highly variable environment and the other a more stable one. During the experiments, Madsen, Lipo, and Cannon (1999) tracked the lifetime waste and the lifetime number of offspring of each agent. Then, they divided the

agents into high- and low-waste phenotypes based on their lifetime waste indicated by the total amount of energy they spent on the wasteful cultural tokens. Next, they compared the fitness of the two types where fitness was the relative proportion of a given type compared to the other type – i.e. "Fisherian fitness." To compare the fitness of phenotypes, they calculated the arithmetic and geometric mean fitness across all time-steps for each type. They included the geometric mean because it is known to more accurately reflect long-term fitness in variable environments (Orr 2009). In the experiments in the more stable environment, both means were higher for the low-waste phenotype, indicating the low-waste type would have higher long-term reproductive success in low-variability environments. In the more variable environment, in contrast, both types had roughly the same arithmetic mean fitness, but the high-waste type had higher geometric mean fitness. Consequently, in the more variable environment, the high-waste type spread faster over the long-term than the low-waste type. These findings, Madsen, Lipo, and Cannon (1999), argued, demonstrate that the bet-hedging waste model is "theoretically sufficient" to explain wasteful behaviour among humans.

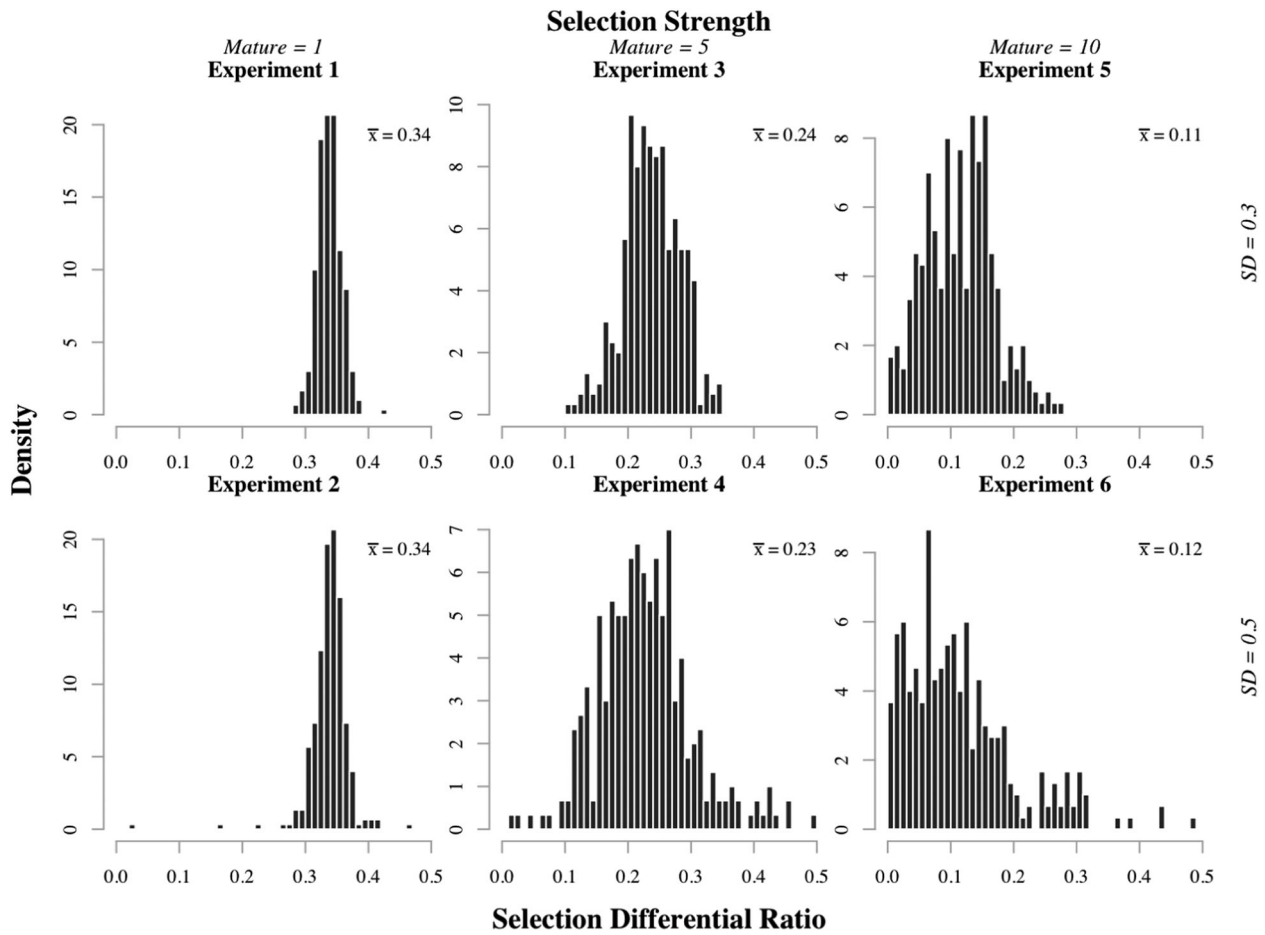One feature of Madsen, Lipo, and Cannon's (1999) ABM that we have concerns about is the

Figure 8. Histograms of the ratio of selection differentials to the maximum possible monotonic shift in the mean for the waste trait for a given experiment.

provisioning strategy employed by their agents. While is it clear that Madsen, Lipo, and Cannon (1999) included provisioning of some sort in their model (pg. 270), they did not specify how parental resources were distributed among dependant offspring. As the results of our supplementary exploration demonstrate, forcing agents to use certain provisioning strategies rather than allowing selection to choose the optimal provisioning strategy can create conditions in which wasting behaviour is adaptive. Thus, it is possible that the results obtained by Madsen, Lipo, and Cannon (1999) differ from those obtained in the present study because they assigned agents a provisioning strategy other than first-born-biased and did not allow the provisioning strategy variable to evolve via selection.

Another feature of Madsen, Lipo, and Cannon's (1999) ABM that we have doubts about concerns storage. In Madsen, Lipo, and Cannon's (1999) ABM, agents inherited traits that determined the surplus energy they would be required to hold before reproducing. This surplus was energy above the level required for survival. By increasing the surplus variable, a given agent would stockpile more energy before reproducing – i.e. they had greater storage. Even if the level of that storage could fluctuate, the agent would be able to

store energy without loss, which seems likely to be a significant adaptive advantage in any environment, especially a highly variable one. However, storage is not a feature of the waste hypothesis. Both Madsen, Lipo, and Cannon's (1999) bet-hedging version of the hypothesis and Dunnell's (1989) version clearly describe waste as energy expended and not stored. Thus, by permitting storage, Madsen, Lipo, and Cannon's (1999) ABM includes a variable that is not only not part of the waste hypothesis but also has the potential to give rise to the results Madsen, Lipo, and Cannon (1999) obtained.

In our view, the uncertainty about provisioning strategies and the inclusion of storage in Madsen, Lipo, and Cannon's (1999) ABM raise questions about the reliability of their results and suggest that further research on their bet-hedging version of the waste hypothesis is required. Specifically, there is a need to recreate and then revise Madsen, Lipo, and Cannon (1999) ABM in such a way that the provisioning strategy variable is subject to selection and agents are not allowed to store energy, and then re-run the analyses they reported. In the meantime, we think it would be wise for archaeologists to also be somewhat skeptical about the ability of the bet-hedging version of the waste hypothesis to explain instances of cultural
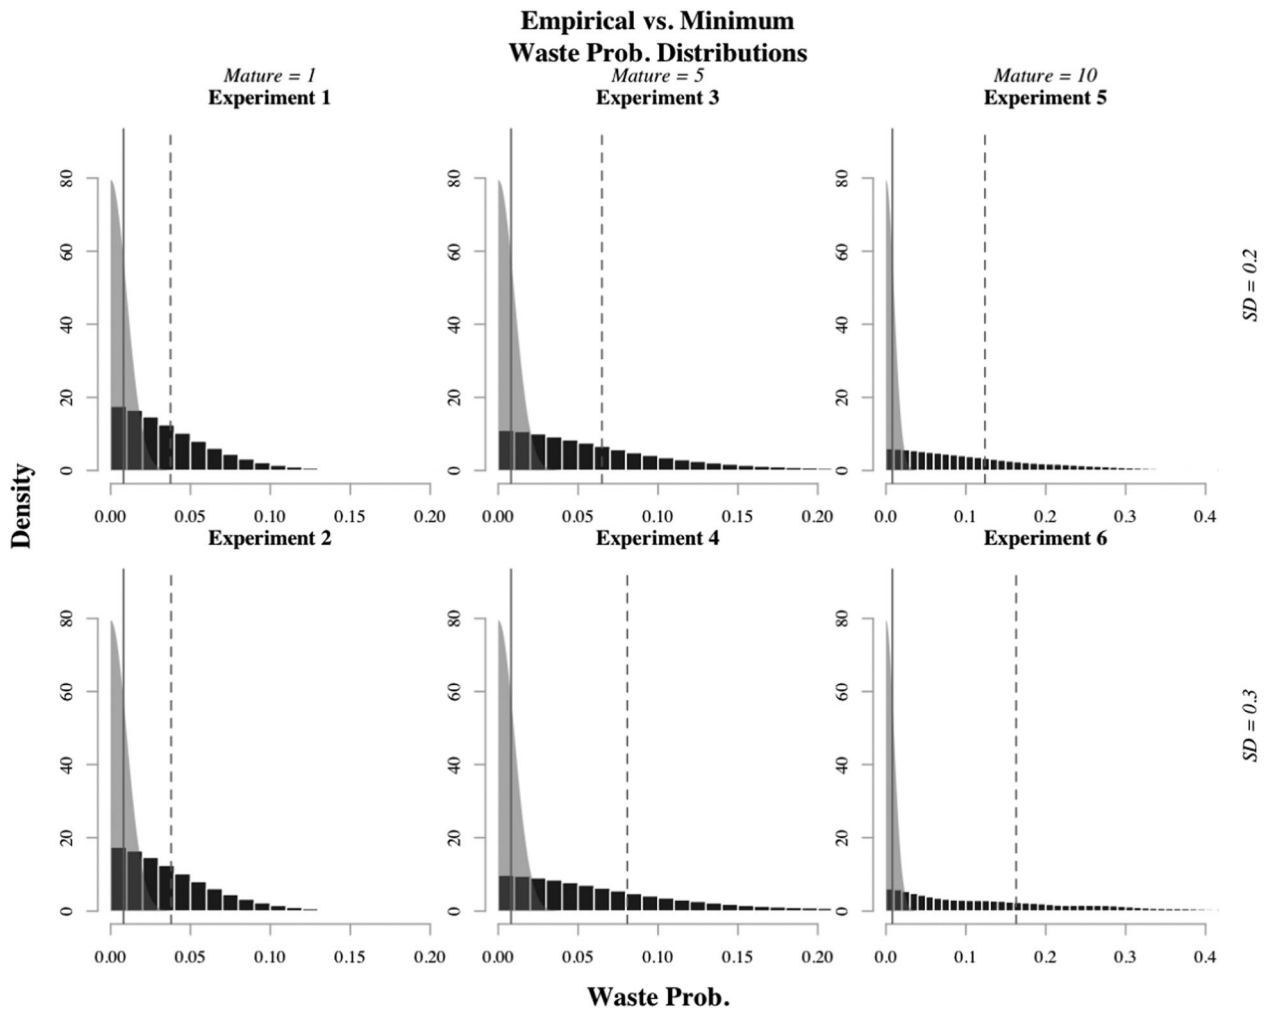
Figure 9. Plots comparing the distribution of the waste trait values in the last 500 ticks of a given experiment (dark histograms) to the theoretical lower-limiting distribution (light grey curve). The samples comprise all waste trait values from the last 500 ticks of all simulation runs of a given experiment. The theoretical lower-limiting distribution does not account for environmental variability, so the functional minimum distribution would have a wider variance and higher mean than this theoretical lower limit. The vertical lines show the sample average of the observed waste trait values (grey dashed) and the mean of the half-normal lower-limiting distribution (grey solid).

elaboration in the past. It could be that, like Dunnell's (1989) version of the waste hypothesis, the bet-hedging version is an unreliable explanation of cultural elaboration in the archaeological record.

Conclusions

In the study reported in this paper we assessed Dunnell's (1989) influential explanation for the occurrence of ancient monuments and other forms of cultural elaboration. Dunnell (1989) hypothesized that wasting energy on cultural elaboration is adaptive when environments are highly variable. To evaluate this proposition, we created an ABM and monitored two variables – the agents' propensity for waste and their long-term reproductive success, which we measured by counting the number of offspring they raised to adulthood in their lifetime. We compared these variables between experiments with different combinations of settings for environmental variability and the agents' maturation period. Toggling between low and high

settings for environmental variability allowed us to determine whether environmental variability selected for agents with higher waste propensities. Including different maturation periods allowed us to account for the potentially confounding effect of differences in provisioning period length. We repeated the experiments hundreds of times to account for randomness in the ABM.

The results we obtained run counter to the central prediction of Dunnell's (1989) hypothesis. We found that the propensity for waste was strongly selected against irrespective of environmental variability. At the start of each experiment agents wasted 50% of the time on average, but selection rapidly drove down that average rate to around 5–10%, and the final distribution of waste propensities was close to the minimum possible. This suggests that Dunnell's (1989) version of the waste hypothesis is not, in the words of Madsen, Lipo, and Cannon (1999, 252), "theoretically sufficient." Its proposed dynamics do not lead to the proposed outcome. Accordingly, we think

archaeologists probably should stop using it to explain monuments and other forms of cultural elaboration in the past.

Notes on contributors

W. Christopher Carleton has a B.A. in Classical and Near Eastern Archaeology from the University of Saskatchewan, an M.A. in Anthropology from Trent University, and a Ph.D. in Archaeology from Simon Fraser University. Currently, he is a postdoctoral research assistant at Simon Fraser University. Chris' research focuses on the long-term impact of climate change on human society. He is particularly interested in developing statistical tools for testing hypotheses about human–environment interaction using palaeoenvironmental and archaeological data.

Brea McCauley holds both a B.A. and an M.A. in Archaeology from Simon Fraser University. Currently, she is enrolled in the Ph.D. program of the Department of Archaeology at Simon Fraser University. Her primary research interest is body modification practices in historically- and archaeologically-documented societies.

André Costopoulos obtained a B.A. from McGill University and an M.Sc. from Université de Montréal, both in Anthropology. He then completed a Ph.D. in Archaeology at the University of Oulu in Finland. Currently, André is the Vice-Provost and Dean of Students at the University of Alberta. André is an expert in the application of agent-based modelling to problems in anthropology and archaeology. He has also carried out extensive fieldwork-based research on Holocene hunter-gatherer societies in Finland and northeastern Canada.

Mark Collard completed a B.A. in Archaeology and Prehistory at the University of Sheffield and a Ph.D. in Hominin Palaeontology at the University of Liverpool. Currently, he is the Canada Research Chair in Human Evolutionary Studies and a professor of archaeology at Simon Fraser University. Mark has published on a wide range of topics in archaeology and biological anthropology, including the origin of genus *Homo*, the transition to farming in Europe, and the impact of male-biased operational sex ratios on Viking society.

ORCID

W. Christopher Carleton http://orcid.org/0000-0001-7463-8638
Mark Collard http://orcid.org/0000-0002-2725-4989

References

Allen, Melinda S. 2010. "Oscillating Climate and Socio-Political Process: The Case of the Marquesan Chiefdom, Polynesia." *Antiquity* 84 (323): 86–102.

Aranyosi, E. Floyd. 1999. "Wasteful Advertising and Variance Reduction: Darwinian Models for the Significance of Nonutilitarian Architecture." *Journal of Anthropological Archaeology* 18 (3): 356–375.

Bu, Feifei, and Frank J. Sulloway. 2016. "Birth Order and Parental Investment." In *Encyclopedia of Evolutionary Psychological Science*, edited by Todd K. Shackelford and Viviana A. Weekes-Shackelford, 1–6. Cham: Springer International Publishing.

Cegielski, Wendy H., and J. Daniel Rogers. 2016. "Rethinking the Role of Agent-Based Modeling in Archaeology." *Journal of Anthropological Archaeology* 41: 283–298.

Charlesworth, Brian. 1980. *Evolution in Age-Structured Populations*. Cambridge: Cambridge University Press.

Dunnell, Robert C. 1989. "Aspects of the Application of Evolutionary Theory in Archaeology." In *Archaeological Thought in America*, edited by Clifford C. Lamberg-Karlovsky, 35–49. New York: Cambridge University Press.

Dunnell, Robert. 1999. "The Concept of Waste in an Evolutionary Archaeology." *Journal of Anthropological Archaeology* 18 (3): 243–250.

Dunnell, Robert C, and Diana M Greenlee. 1999. "Late Woodland Period 'Waste' Reduction in the Ohio River Valley." *Journal of Anthropological Archaeology* 18 (3): 376–395.

Faurie, Charlotte, Andrew F. Russell, and Virpi Lummaa. 2009. "Middleborns Disadvantaged? Testing Birth-Order Effects on Fitness in pre-Industrial Finns." *PLoS ONE* 4 (5): e5680.

Genz, Joseph, and Terry L. Hunt. 2003. "El Niño/Southern Oscillation and Rapa Nui Prehistory." *Rapa Nui Journal* 17 (1): 7–14.

Gillespie, John. 1973. "Natural Selection with Varying Selection Coefficients—a Haploid Model." *Genetical Research* 21 (2): 115–120.

Gillespie, John. 1974. "Natural Selection for Within-Generation Variance Inoffspring Number." *Genetics* 76 (3): 601–606.

Gillespie, John. 1977. "Natural Selection for Variances in Offspring Numbers: Anew Evolutionary Principle." *American Naturalist* 111 (981): 1010–1014.

Graves, Michael W., and Thegn N. Ladefoged. 1995. "The Evolutionary Significance of Ceremonial Architecture in Polynesia." In *Evolutionary Archaeology: Methodological Issues*, edited by Patrice A. Teltser, 149–172. Tuscon: University of Arizona Press.

Graves, Michael W., and Maria Sweeney. 1993. "Ritual Behaviour and Ceremonial Structures in East Polynesia: Changing Perspectives on Archaeological Variability." In *The Evolution and Organisation of Prehistoric Society in Polynesia*, edited by Michael W. Graves and Roger C. Green, 106–125. Auckland: New Zealand Archaeological Association.

Grimm, Volker, Uta Berger, Donald L. DeAngelis, J. Gary Polhill, Jarl Giske, and Steven F. Railsback. 2010. "The

ODD Protocol: A Review and First Update." *Ecological Modelling* 221 (23): 2760–2768.

Hamilton, Fran E. 1999. "Southeastern Archaic Mounds: Examples of Elaboration in a Temporally Fluctuating Environment?" *Journal of Anthropological Archaeology* 18 (3): 344–355.

Henshaw, Jonathan M., and Yoav Zemel. 2017. "A Unified Measure of Linear and Nonlinear Selection on Quantitative Traits." *Methods in Ecology and Evolution* 8 (5): 604–614.

Hertwig, Ralph, Jennifer Nerissa Davis, and Frank J. Sulloway. 2002. "Parental Investment: How an Equity Motive can Produce Inequality." *Psychological Bulletin* 128 (5): 728–745.

Horton, Susan. 1988. "Birth Order and Child Nutritional Status: Evidence from the Philippines." *Economic Development and Cultural Change* 36 (2): 341–354.

Hrdy, Sarah Blaffer, and Debra S. Judge. 1993. "Darwin and the Puzzle of Primogeniture - An Essay on Biases in Parental Investment After Death." *Human Nature* 4 (1): 1–45.

Hunt, Terry L., and Carl P. Lipo. 2001. "Cultural elaboration and environmental uncertainty in Polynesia." In *Pacific 2000: Proceedings of the 5th International Conference on Easter Island and the Pacific.*, edited by Christopher M. Stevenson, Georgia Lee, and Francis J. Morin, 103–115. Los Osos: Easter Island Foundation.

Jeon, Joonghwan. 2008. "Evolution of Parental Favoritism among Different-Aged Offspring." *Behavioral Ecology* 19 (2): 344–352.

Kornbacher, Kimberly D. 1999. "Cultural Elaboration in Prehistoric Coastal Peru: An Example of Evolution in a Temporally Variable Environment." *Journal of Anthropological Archaeology* 18 (3): 282–318.

Kornbacher, Kimberly D. 2002. "Horsemen of the Apocalypse: The Relationship between Severe Environmental Perturbations and Culture Change on the North Coast of Peru." In *Natural Disasters and Cultural Change*, edited by Robin Torrence and John Grattan, 104–243. New York: Routledge.

Madsen, Mark. 2001. "Evolutionary Bet-Hedging and the Hopewell Cultural Climax." In *Posing Questions for a Scientific Archaeology*, edited by Terry L. Hunt, Carl P. Lipo, and Sarah L. Sterling, 279–304. Westport: Bergin & Garvey.

Madsen, Mark, Carl Lipo, and Michael Cannon. 1999. "Fitness and Reproductive Trade-Offs in Uncertain Environments: Explaining the Evolution of Cultural Elaboration." *Journal of Anthropological Archaeology* 18 (3): 251–281.

Mechoulan, Stéphane, and François Charles Wolff. 2015. "Intra-household Allocation of Family Resources and Birth Order: Evidence from France Using Siblings Data." *Journal of Population Economics* 28 (4): 937–964.

Millstein, Roberta L. 2008. "Distinguishing Drift and Selection Empirically: "The Great Snail Debate" of the 1950s." *Journal of the History of Biology* 41 (2): 339–367.

Neiman, Fraser D. 1998. "Conspicuous Consumption as Wasteful Social Advertising: A Darwinian Perspective on Spatial Patterns in Classic Maya Terminal Monument Dates." In *Rediscovering Darwin: Evolutionary Theory and Archeological Explanation*, edited by C. Michael Barton, Geoffrey A Clark, and Douglas B. Bamforth, 267–290. Arlington: Archaeological Papers of the American Anthropological Association.

Orr, H. Allen. 2009. "Fitness and its Role in Evolutionary Genetics." *Nature Reviews Genetics* 10: 531–539.

Peacock, Evan, and Janet Rafferty. 2013. "The Bet-Hedging Model as an Explanatory Framework for the Evolution of Mound Building in the Southeastern United States." In *Barrows and Beyond: An Environmental Archaeology of Barrows*, edited by David R. Fontijn, 253–279. Leiden: Sidestone.

Price, Joseph. 2005. "Parent-Child Quality Time: Does Birth Order Matter?" *Journal of Human Resources* 43 (1): 240–265.

R Core Team. 2019. "R: A Language and Environment for Statistical Computing." Vienna, Austria.

Roff, Derek A. 1992. *The Evolution of Life Histories: Theory and Analysis*. New York: Chapman & Hall.

Sack, Nancy. 2012. "Explaining Maya Monumental Architecture." Unpublished MA Thesis, University of Hawaii at Manoa.

Seger, Jon A., and H. Jane Brockmann. 1987. "What Is Bet-Hedging?" In *Oxford Surveys in Evolutionary Biology*. Vol. 4., edited by Paul H. Harvey and Linda Partridge, 182–211. Oxford: Oxford University Press.

Sterling, Sarah. 1999. "Mortality Profiles as Indicators of Slowed Reproductive Rates: Evidence from Ancient Egypt." *Journal of Anthropological Archaeology* 18 (3): 319–343.

Walsh, Bruce, and Michael Lynch. 2018. *Evolution and Selection of Quantitative Traits*. Oxford: Oxford University Press.

Wilensky, Uri. 1999. *NetLogo*. Evanston, IL: Northwestern University.