

Agent based models of heuristic decision making related to energy intake and expenditure

Journal Title
XX(X):1-15
©The Author(s) 2021
Reprints and permission:
sagepub.co.uk/journalsPermissions.nav
DOI: 10.1177/ToBeAssigned
www.sagepub.com/

SAGE

Abstract

Obesity is a result of a long term energy imbalance that is a consequence of those decision-making processes associated with energy intake and expenditure. The way in which living organisms make these decisions fits the definition of heuristics: cognitive decision-making processes that have a rapid and effortless implementation which can be very effective in dealing with scenarios that threaten an organism's viability. We study the implementation and evaluation of heuristics, and their associated actions, using agent-based simulations in environments where the distribution and degree of richness of energetic resources is varied in space and time. Artificial agents utilize foraging strategies, combining movement, active perception and consumption, while also actively modifying their capacity to store energy, a "thrifty gene" effect, based on three heuristics from the CONSUMAT model. We show that the selective advantage associated with higher energy storage depends on both the foraging strategy and heuristic used, as well as being sensitive to the spatio-temporal distribution of energy resources, with the existence and duration of periods of food abundance and scarcity being crucial to the existence of this selective advantage. We conclude that a "thrifty" gene is only beneficial in the presence of a "sloth gene" and a "glutton gene"

Keywords

Agent Based Models, Heuristics, Food Decision Making, Energy Intake, Energy Expenditure, CONSUMAT model, Thrifty Gene

Introduction

The current obesity epidemic and its consequences represent one of the world's most challenging public health problems. According to the (World Health Organization, 2020), in 2016 more than 1.9 billion adults were considered overweight and, of these, more than 650 million were classified as obese. Obesity is associated with an increase in both general and specific-disease related mortality (Flegal, Graubard, Williamson & Gail, 2005; Flegal, Graubard, Williamson & Gail, 2007), with the most important obesity related comorbidities including: type 2 diabetes mellitus, coronary heart disease, high blood cholesterol level, high blood pressure and osteoarthritis (Must, Spadano, Coakley, Field, Colditz & Dietz, 1999). Overweightness and obesity are associated with a body composition linked to excess fat and are determined using Body Mass Index (BMI), which is calculated as weight divided by height squared, expressed in metric units and rounded to the nearest tenth. Overweight is defined as BMI between 25.0 and 29.9 and obesity as BMI equal or greater than 30.0. This relatively simple classification of fat levels is widely used and is reliable enough for diagnosis (Bray, 2007).

Fat accumulation associated with obesity is due to a long-term positive energy imbalance, where energy intake from food consumption is greater than energy expenditure and is attributed to the complex interaction between an individual and an obesogenic environment. Although the wide availability of high calorie content food and drink and the prevalence of sedentary lifestyles are widely thought to be the principal causes of the obesity epidemic, modeling of this phenomenon is exceedingly difficult

due to the myriad factors that may influence energy intake and expenditure. For example, (Spiegelman & Flier, 2001) detail expenditure as due to physical activity, base metabolism, and adaptive thermogenesis. However, each one of these is itself constituted and influenced by numerous behavioral, physiological and genetic factors. Moreover, each element relates to others at multiple levels, thereby creating complex feedback loops. Although traditional, reductionist approaches have been the dominant framework for analyzing the problem, recently, more systems-based, or complexity-based, papers have appeared. For instance, (Huang, Drewnowski, Kumanyika & Glass, 2009) propose a systemic, multilevel approach to address energy imbalance, providing a framework that makes possible the individual analysis of processes and social interactions. In this broader scope it is possible to study obesity and the concurrent obesity epidemic as a Complex Adaptive System (Hammond, 2009). Complexity-based approaches have also been used to study other aspects of the phenomena, such as the role of nuclear receptors in lipid metabolism and energy expenditure (Evans, Barish & Wang, 2004), or the role of social relationship networks in the emergence of obesity (Christakis & Fowler, 2007). Certainly, given the importance of the problem, and the fact that, globally, it is worsening, new approaches are to be welcomed.

A recurring question in obesity research is: what has led to an explosion in obesity incidence in the last generation or so? Have certain environmental factors changed radically? If so, which? Have our decision making processes changed? Both? One interesting line of research has been that associated with the notion that we are genetically mal-adapted to the

current obesogenic environment relative to the environment in which our species originated. Associated with this idea is the “thrifty” genotype hypothesis (Neel, 1962; Johnson & Andrews, 2010), which is based on the idea that our evolution, in an environment where food resources were scarce, favoured those genes that allowed for better fat storage. However, within this genomic approach, although there is ample evidence that certain genetic polymorphisms can lead to an altered physiology that favours fat storage, these individual genetic alterations are relatively rare and cannot be used to explain a phenomenon as universal as the current obesity epidemic. At heart, obesity, seen as a consequence of energy imbalance, is not really a physiological problem due to a genetic predisposition to fat deposition but, rather, is a result of human behavior - principally overeating and sedentariness - that are associated with sets of individual decisions - eat/don't eat, seek/don't seek food etc. Furthermore, these decisions are affected by a multitude of factors beyond those that might be associated with a microscopic “ome”, such as the genome, proteome or transcriptome. Indeed, it has been argued that what is required is a “Conductome” (Stephens, 2021), thought of as the universe of factors that influence a particular behaviour, such as overeating. In this sense, if there is to be a genetic component, then, as well as explaining certain physiological predispositions, it should also explain why our behavior is obesogenic.

Related to this notion is the question of whether a Conductome that leads to a positive energy imbalance through overconsumption and lack of physical activity is a purely human phenomenon or not. Recent evidence has identified obesity epidemics in various animal populations (German, 2006; Chandler, Cunningham, Lund, Khanna, Naramore, Patel & Day, 2017). In this case, we may ask: what is the origin of these obesity epidemics? Do they have common roots with the human version? For instance, we may ask: Do other species have the “thrifty gene” behaviour?

Accepting that obesity arises from obesogenic behaviors in a potentially obesogenic environment, and given the tremendous multi-factoriality of the problem, with heterogeneous, dynamic and adaptive risk factors, that range from the genetic to the social, an agent-based modeling approach offers several advantages relative to purely empirical approaches, or more standard mathematical modelling, not least of which is the possibility to compare and contrast behaviours in different environments. In this paper an agent-based model (ABM) is presented in order to simulate decision making in different food environments and show how the heuristic-environment interaction can lead to persistent energy imbalances.

Heuristics and food decision-making

Decision-making is central to understanding energy imbalance. A person decides what, when and how much to eat and how to expend their energy. However, individual choices are often constrained by factors beyond one's immediate control. In a classical decision-making framework, environmental boundaries are tied to the exhaustive and exclusive sets of uncertain events and available actions (Lindley, 1991). This could include, for instance, resource scarcity scenarios or strong competition between species. Besides these factors,

we need to account for the fact that physical and emotional internal states can also modify decisions. A framework for understanding decision making from a complexity perspective needs to consider how every factor shapes every other. A generalized deviation from a consumption-expenditure strategy that keeps fat at healthy levels occurs in a wide variety of societies and cultures, as evidenced by the world obesity epidemic. What is behind this apparent universality? In this context it is vital to better understand the decision making process ((Haselton, Nettle & Murray, 2005) (p. 725)).

Decision-making theories based on utility function maximization are frequently inconsistent with human behaviour. The term “*cognitive bias*” was used in (Tversky & Kahneman, 1974) to refer to these discrepancies, which were first observed in the context of probabilistic judgments, difference estimations and risk evaluation. A common characteristic of all these situations is attribute substitution, *i.e.*, the exchange of a computationally complex problem for a simpler one. One of the main criticisms of this description is that it just considers *cognitive bias* to be an inferior, secondary alternative due to a lack of sufficient computing resources (Gigerenzer & Gaissmaier, 2011; Gigerenzer, Fiedler & Olsson, 2012; Mousavi & Gigerenzer, 2014). A number of studies have found, however, that in particular contexts this kind of decision-making process generates better results than purely utility-based decision making (Cosmides & Tooby, 1994; Reyna, 2004; Gigerenzer & Gaissmaier, 2011). The concept, which we shall refer to as *heuristic*, to avoid any negative connotation, provides an explanation of efficient and immediate responses to risk scenarios.

Eating tends to involve more heuristic judgment than most behaviours, due to its intimate relationship with survival (Cohen & Babey, 2012). Some important heuristics in food decisions are based on only one or two important food attributes, such as portion size (Geier, Rozin & Doros, 2006), which is an important element in the obesity epidemic, and disregard many others (Scheibehenne, Miesler & Todd, 2007). Attentional heuristics, related to eating stimuli (Shafraan, Lee, Cooper, Palmer & Fairburn, 2007), are also important. Much work on marketing and product placement has been carried out guided by these principles, in order to encourage food purchases by changes in the environment (Cohen & Babey, 2012). Heuristics are particularly important in situations where survival is at risk and there is insufficient time to perform a detailed deliberation, or in states of cognitive depletion. These scenarios are not exclusive to human decision-making but are common to many living organisms. As pointed out in (Marsh, 2002), many animal behaviours that fulfil simple decision rules can be viewed as heuristics. For example, when animals try to flee from predators, when mating or when foraging, heuristics can confer advantages if they represent appropriate responses to the corresponding environmental conditions. Many questions arise: Is it reasonable to call decision-making mechanisms observed in animals heuristics? What about even less cognitively developed organisms? Is the mere presence of the risk of death sufficient to generate a heuristic in an organism? How exactly does the environment modify the interpretation of precariousness of an organism?

It is the adaptive match between cognitive and ecological structures that constitutes the basis of heuristic formation (Gigerenzer, Fiedler & Olsson, 2012). Heuristics are a natural consequence of certain, key environmental properties: Uncertainty, redundancy, finite sample size and variability in the relevance of environmental features. Additionally, as decision-making involves the expenditure of energy, final decisions are related to the environmental availability of energetic resources and the constraints on the execution of a specific physical activity. The environment provides a distribution of food resources with a particular energy density in space and time and an energy cost associated with their consumption. The sensorial perception of the characteristics of a food resource depends on the complexity of the consuming organism. For example, humans perceive energy density through visual and olfactory cues (Chamaron, Chisin, Chabanet, Issanchou & Brand, 2015; van Beilen, Bult, Renken, Stieger, Thumfart, Cornelissen & Kooijman, 2011) that link high energy density with a higher degree of palatability, which is reduced during intake by the sensation of food-specific satiety (Drewnowski, 1998). Sensory stimuli from exposure to food can also generate neurally mediated responses, that, arguably, facilitate the assimilation of nutrients (Mattes, 1997). Uncertainty in the availability of food is possibly one of the main environmental drivers affecting food related heuristics. An organism's learned experience in a given food environment allows it to label it as abundant or scarce in food resources for example. So, a food rich environment may provide enough resources such that consumption is only linked to physiological indicators of internal low energy states, whereas food scarcity may induce consumption beyond satiety so as to accumulate reserves. Uncertainty in this sense is a powerful motive for the development of heuristics and other cognitive capabilities, for example, the metacognitive judgments that have been observed in humans and animals (Smith, Shields & Washburn, 2003; Smith & Washburn, 2005; Kornell, 2009; Dunlosky & Bjork, 2013).

Agent-based modeling of heuristics and the CONSUMAT model

The flexibility and heterogeneity of ABMs makes them a suitable testing ground for studies of cognition and decision making. Their *bottom up* approach makes it possible to see how a particular microscopic specification can lead to emergent macroscopic regularities (Epstein, 2006; Macal & North, 2010). In particular, how a specific interaction with the environment can generate a resilient heuristic decision. Various approaches have been proposed to model decision-making processes in ABMs using purely reactive agents with *if-then* rules that are inspired by psychological and neurological architectures (Balke & Gilbert, 2014). (Jager & others, 2000), for instance, proposes a decision architecture for agents that explicitly addresses heuristics: the CONSUMAT model, which attempts to unify psychological theories of learning and satisfaction. CONSUMAT agents make decisions according to a diverse set of needs that produce antagonistic responses. For example, the desire to taste a meal that is perceived as delicious would promote eating, while the intention to reduce

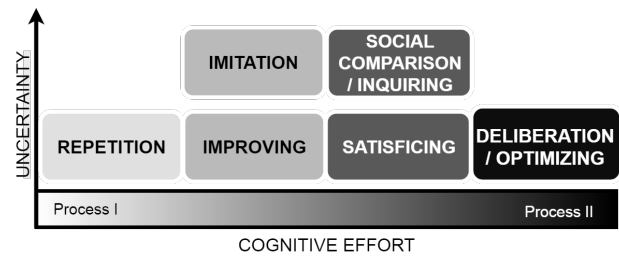


Figure 1. The six decision rules of the CONSUMAT model as a function of a continuous spectrum of dual-process cognition and the uncertainty associated with information gathering (individual or social).

the unnecessary consumption of food in order to achieve a body shape that is considered beautiful would endorse the opposite behaviour. Overall satisfaction is achieved only in the long term, when most needs can be met at a reasonable level without sacrificing one for the other. The agents' dynamical decision system also deals with uncertainty in the environment. As a result, agents create and apply heuristics using basic learning processes.

Several decision-making models (Sloman, 1996; Chaiken & Trope, 1999; Strack & Deutsch, 2004; Evans, 2008; Frankish, 2010) have been designed that divide individual and/or social cognition into two parts: intuition and deliberation. These are portrayed as distinguishable modules and are named: system/process I and II respectively. In this framework, heuristics are developed and implemented by process/system I, which then, in their turn, may or may not be overridden by process/system II (Tversky & Kahneman, 1974; Kahneman, 2003; Frankish, 2010). The CONSUMAT heuristic is characterized by both the amount of cognitive effort required (as in dual approaches) as well as the degree of uncertainty associated with information gathered from the environment. The first version of the model (Jager & others, 2000; Jager, Janssen, De Vries, De Greef & Vlek, 2000) describes four decision rules, which were later extended to six (Jager & Janssen, 2002) (Figure 1). An agent *deliberates* when it considers all the possible action pathways and chooses the one with the greatest utility. This decision rule matches the actions of system/process II in the dual description. The opposite rule, on the cognitive effort dimension, is the blind *repetition* of actions (matching with system/process I). *Satisficing* and *improvement* occur, respectively, when the agent compares the utilities of the distinct actions until one is found that meets their needs, or until an improvement is found (Jager & Janssen, 2002). These latter two strategies are adaptations of the deliberation strategy which involve less cognitive effort. The selection of a cognitive rule reflects two traits of agents - their tolerance to uncertainty and their ambition. For example, deliberative agents have a high aspiration level and low uncertainty tolerance. In terms of uncertainty, imitative or inquiring agents can obtain information from their peers, and by imitating them, or comparing them with others, can reduce the uncertainty they would otherwise have about the potential success of a given strategy.

An update of the CONSUMAT model (Jager & Janssen, 2012) modifies the concept of *social comparison* to that of *inquiring*, which better captures the informative

influence of the social interaction between agents. While normative influence implies group association by similarity, in order to avoid being marginalized from the group, in distinction, informative influence focuses on the interchange of experiences and opinions and the evaluation of the person giving the opinion. The biggest change in CONSUMAT II is the reduction of the agent's needs to a minimal set of three: existential, social and personal needs, which are adapted from the gain, normative and hedonic frames in goal-frame theory (Lindenberg & Steg, 2007). Social needs are balanced between two drives: being similar and being superior. Satisfaction implies a particular degree of accomplishment in the three independent needs and combines both current utility and its future expectations. Uncertainty is calculated only with respect to the variability of information on the existential and social satisfaction dimensions.

The human-like cognitive behavior of CONSUMAT agents is coupled to the environment via their human-like needs. The existential need is their primary bond with the world, while the personal need requires the development of an identity by the agent. A model of primitive living organisms that develop heuristics can still be supported by considering an autopoietic identity (from the closure of a cycle of systems which keep the agent alive), while the incorporation of social needs can be related to the gregariousness that certain groups of animals exhibit.

Social animals find themselves in a society, but without a perfect understanding of its historic development. Ties between the members are likely taken for granted by the agents, which develop a desire to have a role in the group in order to reduce the uncertainties of being alone. For primitive living organisms, the gap between having an autopoietic identity and being able to develop empathy with others seems difficult to explain if we assume that empathy is a reflection of an agent in other agents by a process of mental simulation. This conception has been criticized by (Ratcliffe, 2017), who proposes that the basis of empathy is the openness to interpersonal differences. If social need is considered embodied and enactive, the real distinction between an agent with existential needs only, and another with the full set of needs, is the degree of maturity of their ties with the environment and with their peers.

Agent-based models

The present ABM aims to test how different degrees of uncertainty in the environmental availability of energy sources can promote and induce a positive selection for: i) the emergence of certain preferred foraging and food consumption strategies; ii) the implementation of particular decision-making heuristics; and iii) a preference for higher energy storage capacity. In this model, agents monitor their energy, $E_\alpha(t)$, which determines their existence and ties them to the environment. $E_\alpha(t)$ is degraded according to two factors: the intentional activities performed by the agent and the amount of energy required to maintain basal metabolic functions. This basal energy expenditure per time step is modeled as proportional to the internal energy of the agents which would be a proxy of agents body weight, in analogy with formulae the most relevant factor of the Mifflin-St. Jeor equation (Mifflin, St Jeor, Hill, Scott, Daugherty &

Koh, 1990). However, to avoid non-eating resting agents to live indefinitely the basal energy expenditure is constant for agents with energy below an energy threshold E_T .

Then, the basal expenditure is $E_\alpha(t)M_B$ when $E_\alpha(t)$ is bigger than E_T and E_TM_B otherwise. M_B represents a constant rate of energy used for basic functions and E_TM_B the very minimum energy amount needed to keep an organism alive. Above the energy threshold E_T , the linear metabolic expenditure creates an upper bound on the energy that agents may accumulate for a given value of E_s , where, at that point, the agent's metabolic expenditure is equal to the amount of energy that can be consumed in any given time step:

$$\Delta E_\alpha(t) = E_s - M_b E_\alpha^{(max)} = 0 \Rightarrow E_\alpha^{(max)} = \frac{E_s}{M_b} \quad (1)$$

An agent has a limited set of available intentional actions: i) They can try to consume all the food on a given position at any specific moment; ii) Can increase their capacity for active perception to, in addition to observing the place where they are located, perceive their surroundings to identify food: and/or iii) can forage by moving to one of the closest contiguous spatial positions. A combination of these elements, as shown in Table 1, corresponds to an agent *strategy*. Each strategy has an associated energy cost, depending on whether perception and/or movement are included. In conjunction with the base metabolism of the agent, this leads to a net total energy expenditure at a given time step. If the agent consumes and this energy expenditure is less than E_s then we may say that the agent has "overeaten" at that time step, in that there is a positive energy imbalance. In the case that the agent does not consume or that E_s is less than the total energy expenditure, then there is a negative energy imbalance. This concept of energy imbalance per time step can be aggregated and the consequences of long term energy imbalances explored.

The environment in which the strategies are enacted is modelled as a square lattice of spatial cells with periodic boundary conditions, wherein agents occupy only a single cell at a time, but where a cell can accommodate more than one agent. Units of energy (food resources), E_s , equivalent to "portion size", are situated at each cell of the lattice and can be consumed by only one agent at a time. Consumed resources are then regenerated in a given cell with probability p_g per unit time. E_s and p_g reflect the food/energy resource availability of the environment and its corresponding uncertainty.

Within this environment, an agent may decide to consume or not at a given spatial position at a given time. An agent may also perceive, or not, the existence of food resources in the 8 cells adjacent to the agent's current position, the energy cost of perception being constant, C_p , per unit time. Finally, an agent may move to an adjacent cell. This may be done in conjunction with perception, whereupon the agent chooses at random one of the cells that has been perceived to contain food, if such exists, or, in the absence of perception, the agent chooses a cell at random. The cost of movement is the agent's internal energy multiplied by a linear factor C_m to reflect the fact that moving more mass requires more energy.

Actions follow a causal order: i) decide to eat/not eat; ii) decide to perceive/not perceive, and then; iii) decide to move

| | Eat | Perceive | Move | γ_0 | γ_1 |
|------|-----|----------|------|------------------------|------------|
| S1. | 0 | 0 | 0 | $M_b * E_\alpha$ | 0 |
| S2.* | 1 | 0 | 0 | $M_b * E_\alpha$ | 0 |
| S3. | 0 | 1 | 0 | $M_b * E_\alpha + C_p$ | 0 |
| S4.* | 1 | 1 | 0 | $M_b * E_\alpha + C_p$ | 0 |
| S5. | 0 | 0 | 1 | $M_b * E_\alpha$ | C_m |
| S6.* | 1 | 0 | 1 | $M_b * E_\alpha$ | C_m |
| S7. | 0 | 1 | 1 | $M_b * E_\alpha + C_p$ | C_m |
| S8.* | 1 | 1 | 1 | $M_b * E_\alpha + C_p$ | C_m |
| S9.* | 1 | 1 | 0 | $M_b * E_\alpha + C_p$ | 0 |
| | 1 | 1 | 1 | $M_b * E_\alpha + C_p$ | C_m |

Table 1. Strategies for agents and their corresponding energy costs. 1 and 0 means that the corresponding action is executed or not, respectively. * indicates that the final energy state of the strategy is not known a priori.

or not. This generates the eight possible action strategies (S1 to S8) shown in Table 1, to which is added an extra one, (S9), in which moving is decided upon only after processing the information from the perception step. When agents decide not to eat, changes in energy are negative and known a priori. However, when an agent implements a strategy that can result in an increase in their energy, the corresponding result is uncertain, in that it depends on whether they encounter food resources or not.

In this model, in order to make contact with thrifty genotype arguments, we also explored the role of different energy storage capacities. Specifically, we consider agents with different accumulation limits, $L_\alpha(t)$, associated with the maximum amount of energy they may store and consider this as a key parameter associated with the agents' strategies and which may change in time according to the heuristic used by the agent. Although we interpret this in the current model as being associated with a cognitive, decision making process, we believe that our results would be similar if the model implementation was such that it was interpreted as a genetically controlled parameter. What would change would be the relative timescales involved, and the mechanisms by which such a strategy is inherited or learned. In the case where the agent has not reached its energy accumulation limit, the energy change of an agent from one time step to another is given by

$$\Delta E_\alpha(t) = -\gamma_0 - \gamma_1 E_\alpha(t) + E_s f_{eat}(t, x, \alpha), \quad (2)$$

where γ_0 and γ_1 depend on the agent's particular strategy and their previous energy levels (Table 1). $f_{eat}(t, x, \alpha) = 1$ when agent α finds food resources at cell x at time t and consumes it, and is zero otherwise.

While a lack of energy can lead to systemic failure and death, energy surplus affects agent survival indirectly and gradually. Thus, if an agent consumes all their energy they die. However, an excessive accumulation of energy can also damage an agent by increasing their basal metabolism, thereby increasing their risk of future starvation, and by also affecting their ability to move. In the present model we represent this latter effect by reducing the probability that they can effectively move when their internal energy exceeds a limit E_T . Movement in the model then follows a simple rule: Once an agent decides to move, the probability of doing

so is the following:

$$p_m = \begin{cases} \exp(E_\alpha(t) - E_s) & \text{if } E_\alpha(t) < E_s \\ \exp(E_T - E_\alpha(t)) & \text{if } E_\alpha(t) > E_T \\ 1 & \text{otherwise.} \end{cases}$$

Agents with an internal energy $E_\alpha(t)$, between E_s and E_T , will move every time they decide to do so, while those that have an energy bigger than E_T , or smaller than E_s , **Is this E_s Chucho? Por que?** would move less. The aim of this feature, as mentioned, is to model the impact of overweight or underweight on movement.

Model needs have different timescales regarding goal realization (Lindenberg & Steg, 2007). Personal (hedonic) needs are associated with pleasure and with an immediate reward in terms of feelings, while existential needs have a longer but still limited temporal horizon tied to resource management. The normative or social dimension represents what the agent *ought to do* and has the strongest dependence on external support. Contrary to other needs, compliance with social norms never implies that a state of satisfaction is arrived at.

Eating is triggered by hunger, a sensation that has its origin in the ventromedial hypothalamus, which controls food consumption in order to maintain fat stores at a base level. This makes obese individuals more hungry than lower-weight ones as they need to maintain a higher lever of reserves (Nisbett, 1972). This also happens when individuals experience food deprivation. However, eating can also be instigated in the absence of perceived hunger, such as in cases of depression and anxiety (Plutchik, 1976), or with emotional lability and negative feelings towards oneself (Hudson & Williams, 1981). In this case, the decision to overeat is decided in the context of the personality frame. Overeating, then, can be multi-faceted, sometimes being episodic, secretive, linked to different conditions for different individuals and characterized by ingestion of high-calorie or high-carbohydrate food (Ganley, 1989). Moreover, the over consumption of energy can generate a pleasant feeling of fullness and satiety mediated by leptin, which then impairs further consumption (Morton, Blevins, Williams, Niswender, Gelling, Rhodes, Baskin & Schwartz, 2005).

The corresponding model for needs and satisfaction for simpler living organisms uses modified satisfaction and uncertainty functions, keeping only those terms that are directly related to their existential need. This model is implemented in order to contrast the decision dynamic of such an organism with that of a being that has a developed sense of self, as well as a social interaction, and in order to show how this affects their corresponding heuristic.

Decision rules are implemented in line with the CONSUMAT schema, with the algorithms described in Table 2. Although the heuristics can be applied to the choice of strategy, S1-S9, in the current paper we restrict the heuristics to be applied to the energy accumulation limits only, returning to their application to the different foraging strategies in a future paper. The four heuristics we consider are: Repetition, optimization, imitation and inquiring (Jager, Janssen, De Vries, De Greef & Vlek, 2000). Repetition is the decision process that requires the least cognitive effort, and is the simplest to implement, as it only requires that

| Heuristic | Algorithm |
|--------------|---|
| Repetition | $L(\alpha, t + 1) = L(\alpha, t)$ |
| Imitation | $L(\alpha, t + 1) = L$ where $L \leftarrow^R \{L_i L_i \in \{L(\alpha', t)\}_{\alpha' \in C_\alpha}\}$ |
| Inquiring | $L(\alpha, t + 1) =$ $\text{argmax}_{L(S,t)} \{L_i L_i \in \{L(\alpha', t)\}_{\alpha' \in C_\alpha}\}$ |
| Optimization | $L(\alpha, t + 1) = \text{argmax}_{L(\alpha,t)} \{L1, L2\}$ |

Table 2. Agent algorithms used to select an energy accumulation limit (L) at time t according to the heuristics of the CONSUMAT model. Agent α considers limits from other agents $\alpha' \in C_\alpha$, where C_α is the set of agents in cells that agent α can perceive.

an agent keep doing the same thing. When an agent α chooses a strategy by imitation or inquiring at a given time step, they create a pool of possible alternatives (strategies or accumulation limits), by observing those strategies $\alpha' \in C_\alpha$ associated with agents that they can “perceive”. In the case that the agent has no explicit perception element in their strategy, then the set α' is restricted to those strategies or accumulation limits present in agents located at the same position as the agent α . On the contrary, when perception is an element in the agent’s strategy, then the set α' is extended to those strategies or accumulation limits that are also located in adjacent cells. For imitation, the probability to choose a particular energy accumulation trait depends on the number of agents with that trait. **como es esto explicitamente Chucho?**, while, for enquiring, they select that accumulation limit associated with one of the agents with maximum internal energy in the group they perceive. Finally, in the case of optimization the agent always selects the highest limit. **otra vez, como es esto explicitamente Chucho?** In this model we assume that primitive agents are also able to develop these heuristics despite their level of complexity. Such a potentially unrealistic assumption arises from the fact that the heuristics constitute a limited set of rules that could be learned by reinforcement.

Uncertainty in the decision-making process, as described in our model, arises from several sources. Obviously, uncertainty in the distribution of food resources is an important source. There is also uncertainty as to whether an agent will be able to move at all if, for example, their internal energy state corresponds to “overweight” or “underweight”. In terms of the social component, there is uncertainty in calculating the utility of the strategies associated with the community of a given agent relative to the set of all strategies. In other words, an agent does not a priori know if the best strategy in their community is the best possible strategy overall.

Finally, it is important to emphasize that, given that this is a competitive system, even if an agent chooses the optimal strategy, there is uncertainty as to the outcome, as a competitor may have consumed the desired resource first. To reduce this effect we can begin with a low density of agents, whereupon the effective number of competitors will be small.

The availability of energy in the model is represented by the environmental parameter, p_g , which represents the probability of regeneration of energy resources in any cell.

Other cell parameters include the quantity of food resource per cell, E_s , and the rates at which agents expend energy according to their actions, M_b , C_m and C_p . Although E_s, M_b, C_p and C_m do not play a direct role in the probability that a given agent will eat at a given time, they greatly affect the survival rates of a given strategy at the population level. p_g can be set as a constant parameter for every cell or can be time varying, so as to mimic, for example, the effect of a food glut or a famine, where there are periods of food abundance and periods of food scarcity. We will consider the periods of such “feast and famine” episodes to be chosen from a normal distribution with a particular mean and standard deviation. We choose this particular set of features to model the fact that, as in the real world, uncertainty in the availability of food resources and the regulation of internal energy can generate an ecological pressure that favours certain foraging and energy accumulation strategies, or can cause organisms to restructure their behavior, so as to try different or new heuristics in their decision making.

The lengths of the periods of abundance and famine play a critical role in imposing selection pressure on the strategies and heuristics and, in particular, on the energy accumulation limits. For example, if we consider an agent with the least costly metabolic settings (which implies an S2 strategy) that consumes food every generation in the period of abundance, having started from zero energy, then, given that there are two distinct energy accumulation regimes, $E_\alpha < E_T$ and $E_\alpha > E_T$. In the former, the agent has an energy gain $E_s - M_b E_T$ per time step, while, in the former, the energy gain per time step is $E_s - M_b E_\alpha(t)$. Thus

$$E_\alpha(t) = E_\alpha(t-1) + E_s - M_b E_\alpha(t-1).$$

Considering an initial energy state $E_\alpha(t=0) = E_T$, as $|1 - M_b| < 1$ we have

$$E_\alpha(t) = \frac{E_s}{M_b} (1 - (1 - M_b)^t) + (1 - M_b)^t E_T.$$

The maximum energy that the agent can reach, E_A , is the energy accumulation limit, $L_\alpha(t)$, or the maximum limit E_α^{max} , where metabolic energy expenditure exceeds consumption. With either of these limits we may determine how much time to reach E_A

$$t^{(E_T \rightarrow E_A)} = \frac{\log(E_A - \frac{E_s}{M_b}) - \log(E_T - \frac{E_s}{M_b})}{\log(1 - M_b)}$$

If we take as initial condition $E_\alpha(t=0)$ then we just need to add in the time to get from energy 0 to energy E_T , which is $t^{(0 \rightarrow E_T)} = E_T / (E_s - M_b E_T)$.

creo que hay algo mal con el calculo abajo

identifying the number of generations of abundance that an agent with the least costly metabolic settings (which implies an S2 strategy) needs in order to exceed a certain energy limit, E_A . We assume for simplicity the case of an agent that finds and consumes food every generation in the period of abundance, having started from zero energy. When the energy of the agent is below E_T , the increase in energy is $E_s - \frac{1}{M_b}$ every generation. The time to pass from zero energy to E_T is,

$$t^{(0 \rightarrow E_T)} = \frac{E_T}{E_s - \frac{1}{M_b}}$$

When the agent's energy reaches E_T , the subsequent increase in energy per period is given by $E_s - M_b E_\alpha(t)$. This implies that,

$$E_\alpha(t) = E_\alpha(t-1) + E_s - M_b E_\alpha(t-1).$$

Considering a starting state of $E_\alpha(t=0) = E_T$ we have that

$$E_\alpha(t) = E_A = E_s \sum_{i=0}^{t-1} (1 - M_b)^i + (1 - M_b)^t E_T.$$

As $|1 - M_b| < 1$ we can use the identity $\sum_{i=0}^{n-1} a^i = \frac{1-a^n}{1-a}$, to find

$$E_A = \frac{E_s}{M_b} (1 - (1 - M_b)^t) + (1 - M_b)^t E_T.$$

From this expression we calculate the number of generations needed to pass from E_T to E_A .

$$t^{(E_T \rightarrow E_A)} = \frac{\log(E_A - \frac{E_s}{M_b}) - \log(E_T - \frac{E_s}{M_b})}{\log(1 - M_b)}$$

Finally we have that $t_{ab} = t^{(0 \rightarrow E_T)} + t^{(E_T \rightarrow E_A)}$ which is,

$$t_{ab} = \frac{E_T}{E_s - \frac{1}{M_b}} + \frac{\log(\frac{E_s}{M_b} - E_A) - \log(\frac{E_s}{M_b} - E_T)}{\log(1 - M_b)}. \quad (3)$$

On the other hand, we can also calculate the lowest number of failed attempts to consume that as a consequence produces the death of an agent that has the minimum basal energy expenditure and an energy of E_A . In this case, a period of famine equal to this time or greater will almost surely conduce agents with an accumulation limit of E_A to extinction. This is given by $t_{fam} = t^{(E_A \rightarrow E_T)} + t^{(E_T \rightarrow 0)}$,

$$t_{fam} = \frac{\log(E_T) - \log(E_A)}{\log(1 - M_b)} + \frac{1}{M_b} \quad (4)$$

The above logic will be useful for understanding under what conditions an agent with a given energy accumulation limit may be expected to have a competitive advantage relative to one with a lower limit, both in the capacity to accumulate energy as well as to survive a famine. For example, if the period of abundance is not sufficiently long for an agent with a higher accumulation limit to accumulate more energy than an agent with a lower limit then the higher limit is clearly of no use.

Simulations

Agent systems were simulated with NetLogo 6.0.1 (Wilensky, 1999). The environment consisted of a square grid of 41×41 cells with periodic boundary conditions. The initial agent population was 1680 agents, with an initial energy $E_0 = 2$. The energy limits that determine the probability of movement were set to $E_T = 20$ and $E_{min} = 2$. The agent's parameters, M_b, C_m, C_p , were fixed at values 0.05, 0.02 and 0.01 respectively. Thirty simulations were performed for each combination of parameters.

As with any ABM, there is a potentially large parameter space to be analysed. In the present study, as well as nine foraging strategies and three heuristics, there are several

parameters that can be varied, as discussed in section 2. Additionally, the spatio-temporal distribution of p_g also offers a rich source of variability for determining the relative advantages of one set of agent characteristics versus another. Below, we discuss only a subset of experiments and their results that represent what we believe to be the most important conclusions for understanding the possible origins of the obesity epidemic.

The first sets of experiments were designed to better understand under what environmental circumstances the capacity to accumulate energy (fat) was advantageous, this being related to the "thrifty gene" logic. As the capacity to accumulate energy can be argued to be present in other organisms than humans, we used only the existential need and the heuristic "repetition". In the first set of experiments, we compared different accumulation limits between 5 and 105 units in the context of the different foraging and consumption strategies (S1 - S9) and followed their evolution for 500 generations. The environmental parameter E_s was set to $E_s = 2$. In every simulation, agents were randomly initialized with one of two accumulation limits, corresponding to 840 agents of each limit in the initial population. One of them was chosen to be 5 and the second was chosen to be one of 6, 15, 55 and 105. The environmental parameter p_g was varied from 0.1 (scarcity) to 1.0 (abundance) in intervals of 0.1, and then from 0.01 (extreme scarcity) to 0.15 in intervals of 0.01. In all these cases the environmental parameter p_g was constant. With these parameters, only agents with eating strategies (S2, S4, S6, S8 and S9) can survive.

In a second set of simulations, p_g was allowed to change in order to mimic the effects of "gluts" and "famines", modeled by periods of complete abundance (when p_g is equal to 1.0) and periods of complete scarcity (when p_g is equal to 0.0). The periods of the gluts and famines were chosen from a normal distribution, $\mathcal{N}(\bar{t}, STD)$, with $\bar{t} = 40$ for the famine and 60 for the period of abundance. By varying the standard deviation of the normal distribution we could introduce different degrees of uncertainty into the availability of food resources. We considered $STD = 0$ (constant periods), 1 and 5. System development was followed for 1000 generations, with data collected by population for each action, decision rule and accumulation limit. In these particular simulations agents could have one of two accumulation limits: 55 and 105, and the amount of resources in every cell was increased to $E_s = 3.0$. In these experiments we also considered heuristics other than repetition, the consequence being that an agent could potentially change its energy accumulation strategy over time by imitating the strategy of another agent, or by deducing the existence of a better strategy from concurrent agents (inquiring).

A final set of simulations was associated with presenting a set of ten feast-famine cycles, as above, then continuing with an energy rich environment, $p_g = 1$, for a further 500 generations. The purpose of these simulations was to investigate what happens if a period of resource uncertainty is followed by a period of resource certainty.

Results

The results of the first set of experiments, where agents with different accumulation limits are competing in an environment with a constant p_g , can be seen in (Fig. 2). There, the vertical axis represents the difference in the number of agents present in the population after 500 generations with differing accumulation limits versus agents with a limit of 5. Thus, in the top left figure we see a slight advantage for higher accumulation limits in the case of strategy S2 for $p_g \sim [0.01 - 0.04]$, with there being approximately 20-30 agents more with accumulation limits > 5 than with accumulation limit 5. This relative advantage of a “thrifty gene”-type strategy is present only for those strategies - S2 and S4 - that do not actively forage, or that forage in a random fashion - S6.

The two highest accumulation limits, $E_A = 55$ and $E_B = 105$, appear to have identical behaviours when using a fixed strategy within an environment with a fixed probability of regeneration. Under these conditions, food is never scarce enough, nor abundant enough, to turn a higher accumulation limit into an ecological advantage. However, it is possible to find temporal patterns of food availability where a higher energy storage capacity can be beneficial and, indeed, where adequate storage becomes fundamental to survival.

In accordance with Eq. (1), for an environment where $E_s = 2.0$, and $M_b = 0.05$, agents cannot exceed $E_\alpha(t) = 40$, thereby explaining the similarity in the populations of agents with limits 55 and 105 (Fig. 2). However, if we increase E_s to 3.0, it is possible to find scenarios, when the period of abundance is large enough, where the agent’s energy exceeds the limit $E_A = 55$. For the given parameters, in Eq. (3) and (4) we have that $t_{ab} = 50.54$ and $t_{fam} = 39.72$. If we then set the abundance and famine periods to 60 and 40 respectively, for two competing populations of agents that use the same foraging and consumption strategy, but have distinct energy storage limits, 55 and 105, respectively, we might expect to see a clear advantage for those agents with a higher accumulation limit. This is the case. However, only strategies S2 and S9 have agent populations that avoid becoming extinct or close to extinct for both accumulation limits in order for this selective advantage to manifest itself.

The results of the second set of experiments can be seen in Fig. 3. In the Top graph, we show box plots of the populations associated with the last generation after completing 10 feast-famine cycles for different decision heuristics. Foraging strategy S9 was used by all agents. We can observe that agents with energy limit 55 end with an average population of around 43 for imitation, but close to zero for the repetition and inquiring heuristics. When the accumulation limit is 105, however, the average population is greater than 100 for every considered heuristic. In other words, survival probability is significantly enhanced for the higher energy accumulation limit agents. On the other hand, the population variance of the lower accumulation limit agents is higher. Furthermore, note that the variance for the imitation heuristic is greater than that of the inquiring heuristic which, in its turn, is greater than that of the repetition heuristic.

In Fig. 3 we see the population average (Middle graph) and the average energy per agent (Bottom graph) throughout the 10 feast-famine cycles for different heuristics and

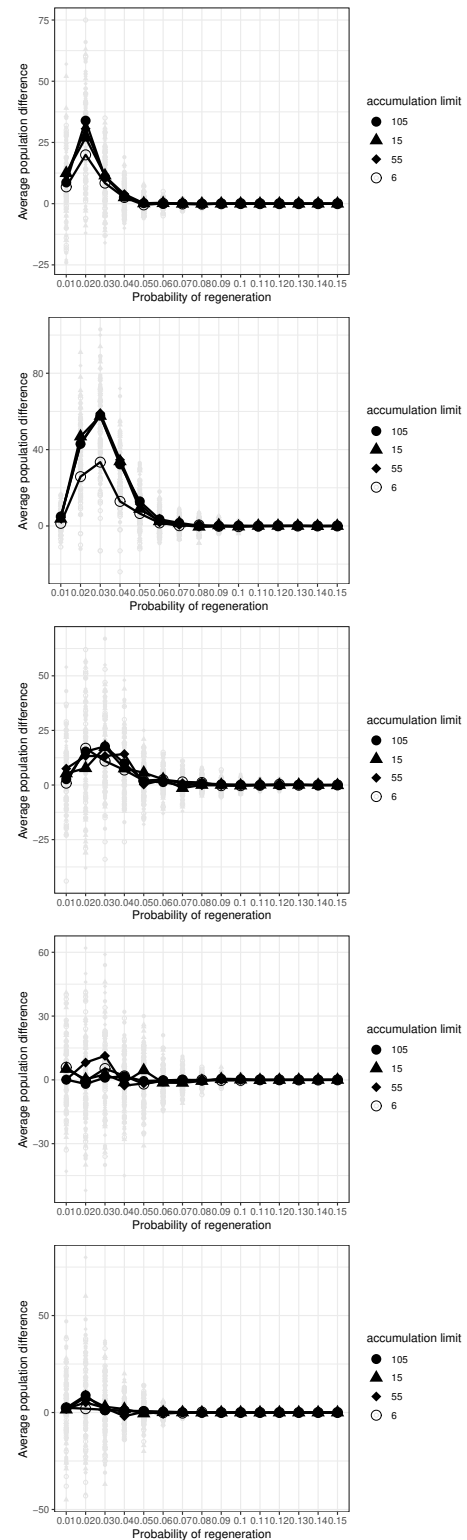


Figure 2. Difference in population size of competing agents with different energy accumulation limits, where agents repeat the same strategy: (Top:) S2, (Center, up:) S4, (Center:) S6, (Center, down:) S8 and (Bottom:) S9, for 500 generations in an environment with a constant value of p_g between 0.01 and 0.15. 50% of the agents start with an energy limit of 5 and 50% start with an energy limit of 6, 15, 55 or 105 respectively. Graphs show the average of 30 repetitions using the same parameters. The gray points in the background represent the result for every experiment.

different initial accumulation limits, where the averages

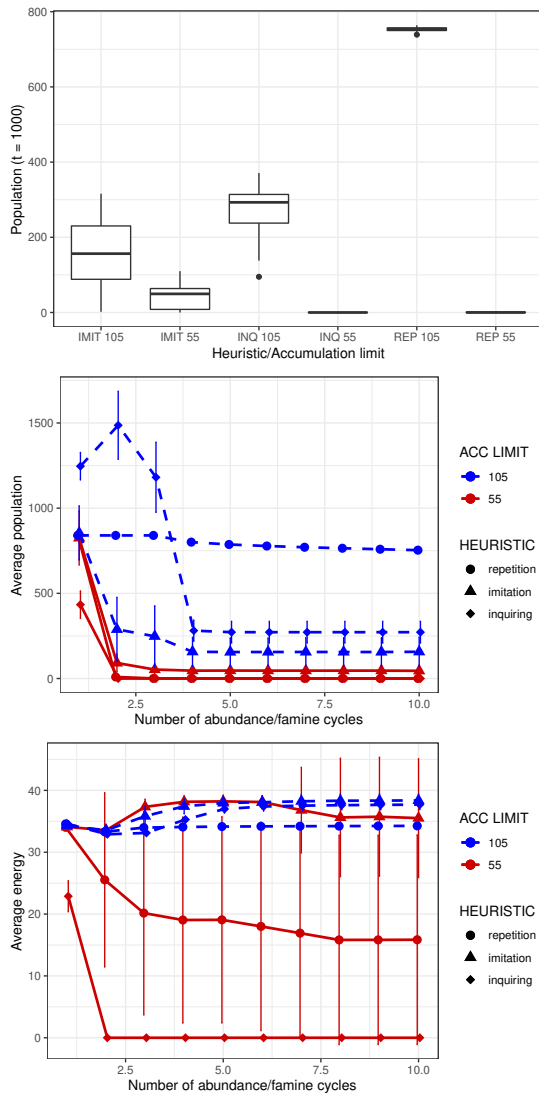


Figure 3. *Top:* Box plots of the population of agents with accumulation limits 55 and 105, at time $t = 1000$ when foraging strategy S9 is performed and the agent's accumulation limit is potentially subject to adaptive change according to one of three heuristics: repetition, imitation and inquiring. Comparison through a set of "feast and famine cycles" of: *Middle:* the average population and *Bottom:* the average energy during a particular cycle for S9 agents with accumulation limits of 55 and 105. In these simulations, agents experience a "glut" period, when food is regenerated immediately after eaten, followed by a "famine" period where food is not regenerated. These periods are fixed to 60 and 40 generations respectively. The average per cycle considers both periods and the 30 repetitions of every set of parameters and error bars on y axis represent the standard deviation on the assemble. Graphs are lightly displaced on the x-axis for visualization.

are computed by considering the abundance phase of each cycle. Again, we show only results for the S9 foraging strategy. We can see that, for each heuristic, the population average is greater for the higher accumulation limit agents throughout the set of feast-famine cycles, although there are significant differences between the results of one heuristic and another. Note that limit 55 agents exhibit steep, monotonic declines after the first feast-famine cycle, with the decrease being particularly notable for the repetition and inquiring heuristics. On the other hand, the behaviour of the

limit 105 agents is quite distinct, with an increase between the first and second cycles for the inquiring heuristic, followed by a very sharp decline between the second and fourth cycles. On the other hand, the imitation heuristic shows a sharp decline from the first to second cycles, followed by a more gradual decline from the second to the fourth. For both the inquiring and imitation heuristics the population average after the fourth cycle is constant. Finally, for the repetition heuristic, the population exhibits a subtle decrease after the third cycle, finishing with an average around 750 after 10 cycles.

Turning now to the average energy (Bottom graph), for the limit 55 agents, for the inquiring heuristic, the average energy goes to zero after the first cycle, while for the repetition heuristic it decreases gradually from 35 to 16 over the full set of cycles. Interestingly, for the imitation heuristic it increases slightly after the second cycle, actually exceeding that of the limit 105 agents for a couple of cycles. For the latter agents, the average energy is relatively constant throughout the set of cycles, $\sim 35 - 40$. However, the ordering of the average energy as a function of heuristic is opposite to the ordering of the population average: repetition, inquiring and imitation for the latter and imitation, inquiring and repetition for the former.

The introduction of uncertainty in the duration of the feast and famine periods produces different results according to the decision rule followed by the agents. We show only accumulation limit 105 agents, as limit 55 agent populations could not survive the added uncertainty. We can clearly see that extra uncertainty has an important detrimental effect, independently of the heuristic used, with average populations decreasing as the uncertainty (standard deviation) increases (Fig. 4). However, this ordering of population sizes as a function of uncertainty is not uniformly present as a function of time but, rather, emerges. Indeed, there is an interesting transient behaviour, such that a higher uncertainty can lead to higher average populations in the initial cycles in the case of the imitation and inquiring heuristics.

In terms of average energy (Fig. 5), the standard deviation 5 results show that the average energy per agent in this case is substantially lower than the corresponding standard deviation 1 or 0 values, independently of the heuristic used. However, the differences between the standard deviation 1 and 0 results are both small and vary between heuristics with the energy being lower for the standard deviation 1 results than the standard deviation 0 ones for the repetition and imitation heuristics, but higher for the inquiring heuristic.

In Fig. 6 we show the results for the average energy (right) and average population (left) for simulations of agents with accumulation limits of 55 and 105, using foraging strategy S9 and the repetition heuristic, in the case where there are 10 feast-famine cycles of 60 and 40 periods respectively, but where the cycles are followed by a period of 500 generations of constant resource availability. The average population follows exactly the pattern set by the 10 feast-famine cycles as seen in Fig. 4 - middle - where there is a gradual degradation in the population of the limit-105 agents and an almost total extinction of the limit-55 agents. However, after the 10th cycle the agent populations are constant. In terms of energy, the surviving agents reach maximum

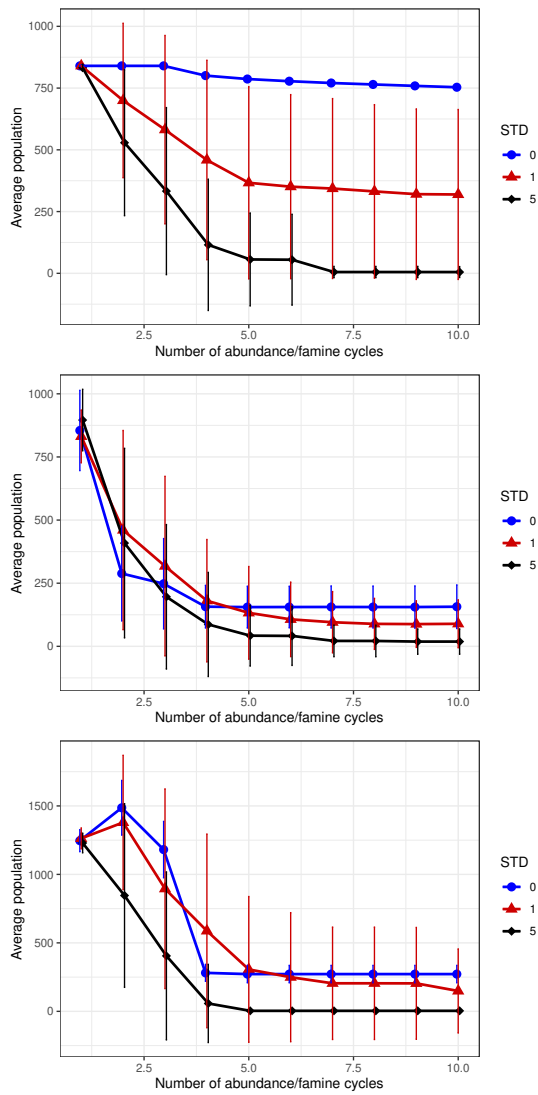


Figure 4. Comparison through a set of 10 “feast and famine cycles” of uncertain duration of the average population at a particular cycle for S9 agents with an accumulation limit of 105 for systems where the employed heuristic is; *Top*: repetition, *Middle*: imitation; or *Bottom*: inquiring. In these simulations, agents experience a “glut” period, when food is regenerated immediately after being eaten, followed by a “famine” period where food is not regenerated. The feast and famine periods are chosen from a normal distribution with means 60 and 40 respectively, and with a standard deviation of 0, 1 or 5. The average per cycle considers both periods and the 30 repetitions of every set of parameters and error bars on the y axis represent the standard deviation of the ensemble. Graphs are lightly displaced on the x-axis for easier visualization.

limits of 62.6 for the limit-105 agents and 55 for the limit-55 agents. Although we are showing here only the S9-repetition combination, analogous results hold for different heuristics or, indeed, other strategies. The combination of strategy and heuristic serves to determine which agents survive the 10 cycles. However, any agent that does survive to the constant period of abundance quickly reaches its corresponding energy limit.




Figure 5. Comparison through a set of 10 “feast and famine cycles” of uncertain duration of the average energy at a particular cycle for S9 agents with an accumulation limit of 105 for systems where the employed heuristic is; *Top*: repetition, *Middle*: imitation; or *Bottom*: inquiring. In these simulations, agents experience a “glut” period, when food is regenerated immediately after being eaten, followed by a “famine” period where food is not regenerated. The feast and famine periods are chosen from a normal distribution with means 60 and 40 respectively, and with a standard deviation of 0, 1 or 5. The average per cycle considers both periods and the 30 repetitions of every set of parameters and error bars on the y axis represent the standard deviation of the ensemble. Graphs are lightly displaced on the x-axis for easier visualization.

Discussion

The results of the experiments seen in Figures 2 and 3 exhibit under what circumstances a particular energy storage strategy has an advantage over others. Figure 2 shows that there is a slight advantage associated with a higher capacity to store energy - a “thrifty” gene - in environments with scarce food resources, but only for strategies S2, S4 and S6. It is stating the obvious to say that the no-eating strategies, S1, S3, S5 and S7, do not survive. However, this is an extreme case of a more general trait - that there is no point having a thrifty gene if there does not exist another

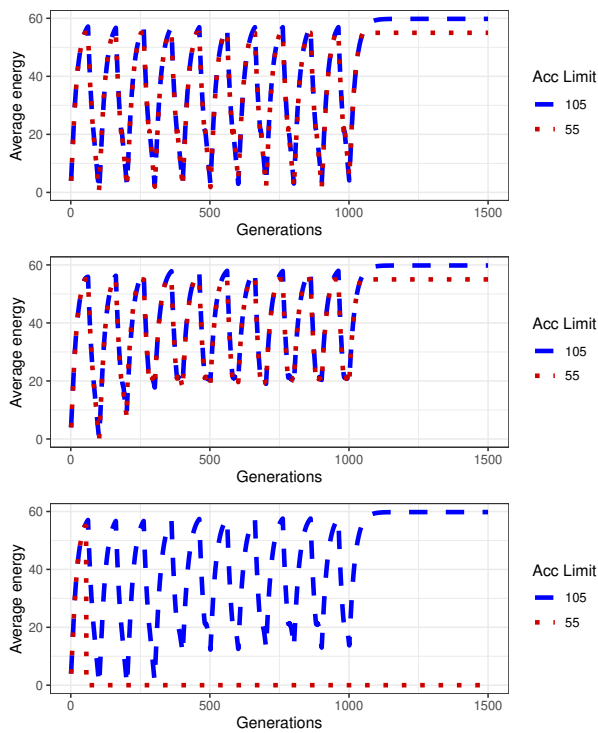


Figure 6. Example of a period of 500 generations of resource abundance after a set of 10 “feast and famine cycles”: Average energy at a particular generation for S9 agents with accumulation limits of 55 and 105, where the employed heuristic is *Top*: repetition, *Middle*: imitation or *Bottom*: inquiring.

fundamental and necessary characteristic: that agents can accumulate energy beyond their immediate metabolic needs. This, in turn, has two requirements: first, there must be enough energy resources present in the environment so as to make it possible to accumulate energy in the first place; and, second, the agents must behave so as to consume more than their immediate metabolic needs. In this sense, they must exhibit a conduct of “overeating”.

So, what is the origin of the small advantage for S2, S4 and S6 and why there is no advantage for S8 and S9? In the case of S2, the relative advantage accrues from the fact that, by chance, some agents manage to consume a higher than average amount of food resources in a given cell, so that the storage of this extra energy allows them to survive in that cell when there is less food than expected. In other words, the uncertainty inherent in the distribution of food resources that can lead to agents dying, by “bad luck”, can be partially offset by an energy storage capacity that is utilized by those agents that, by “good luck”, manage to obtain more food resources than the average. Again, though, this requires sufficient energy resources and the tendency to overeat beyond their immediate metabolic needs. This advantage is exacerbated for strategy S4, as this strategy involves the extra energy cost associated with perception. The advantage is less for S6, as these agents are able to forage, and thus can partially offset the impact of uncertainty at a given location by having access to food resources in other cells. This is even more the case with strategies S8 and S9, where the foraging is done in a more intelligent manner, being directed precisely to those cells where there are food resources.

It is interesting to note that in the case of very scarce resources - $p_g < 0.01$ - the advantage disappears. This is due to the fact that the probability of accumulating sufficient energy above the baseline level of 5 is negligible. In other words, energy storage is of no use if food is so scarce that it is highly improbable to obtain enough energy resources to take advantage of it. On the contrary, above a certain p_g , energy storage is not even necessary, as there are always sufficient food resources available. Finally, such energy storage offers no advantage to those agents that can forage intelligently.

If greater energy storage capacity offers no significant advantage in scarce resource environments, then under what circumstances might it be useful? The results seen in Figures 3,4 and 5 clearly answer this question: Firstly, in those environments where periods of relative scarcity and relative abundance alternate; and, secondly, in those environments that have some degree of regularity in the duration of the feast-famine cycles, with the relative advantage depending on the foraging and consumption strategy used, as well as the precise energy storage capacity relative to the duration of the periods of abundance and famine. Indeed, our choice of famine period is such that an agent with the minimal energy expenditure (without perception or movement) and with an accumulation limit of 55 could not survive without the consumption of energy resources. Agents that forage can potentially survive during the famine, in spite of their extra energy cost, by finding any unconsumed food resources that are left over from the abundance period when the famine began. Of course, such an agent must also survive the competition from its peers in the search for these unconsumed resources.

The results for systems with fixed periods of abundance and famine confirm the survival of S2 (as the system parameters have been chosen to make them do so). We can think of this as the “hibernating bear” scenario, wherein the advantage of energy accumulation accrues only if the agent maintains the lowest energy expenditure possible. Apart from S2, however, only the S9 strategy is able to support large populations of agents that have the higher accumulation limit. The other strategies all have to deal with extra energy costs, and, in a time of famine, this places a great deal of extra selection pressure. S9 agents, however, can regulate their expenditure of energy based on their perception of the environment, remaining static (“sedentary”) during the abundance periods as the resource regeneration guarantees the existence of food in the next time step. In the periods of scarcity, every amount of saved energy is vital. In this case, rather than a “thrifty” gene as being of relevance, we may speak of the possibility of a “sloth” gene, that suppresses physical activity in circumstances where active foraging is not advantageous. Of course, we do not mean to imply that there is a specific gene associated with such inactivity, but that such a behaviour could have a genetic origin.

Thus, higher energy storage alone is not sufficient to survive. Rather, it is the intelligent balance between consumption in place versus movement that permits the higher accumulation limit to exhibit an advantage. More generally, we see that the benefit of energy accumulation limits - the thrifty gene - is dependent on the cognitive strategy used by the agent relative to the temporal availability of food and the competition from other agents. Of course, as

emphasized, a necessary condition for taking advantage of exploiting a higher accumulation limit is also the potential to keep consuming in order to reach that limit above and beyond the short-term energy needs of the agent. This is manifest in the fact that during the periods of abundance, S2 and S9 agents consume more than their metabolic needs for the next generation.

The existence of heuristics introduces an extra level of complexity by permitting agents to change their energy accumulation strategy in the case of imitation and inquiring, with the difference between them being that agents using the imitation heuristic simply copy at random an energy accumulation strategy in their community, whereas an agent using the inquiring strategy will always choose the best one in that community. The effect of this can be seen in Fig. 4. For the inquiring heuristic, many of the 840 original limit-55 agents have imitated the superior 105-limit strategy. Those that don't die. The fact that there is a collapse in the 105-limit population after the second cycle is due to the fact that they are victims of their own success, with the resources available during the famine periods not sufficient to support such a large population. For the imitation heuristic, limit-105 agents may copy limit-55 agents, as well as vice versa, and, in this case, those limit-105 agents that have switched are more at risk. However, there is a bias, in that during the famine the limit-55 agents will have a higher mortality rate and therefore there will be less of them to imitate, thus leading to a relative excess of limit-105 agents. With the repetition strategy, there is no possibility for a limit-55 agent to learn or copy the limit-105 strategy, and therefore the mortality of the limit-55 agents is almost total. The repetition heuristic for the limit-105 agents is the most successful because it avoids both the mistakes associated with the imitation heuristic and the excessive level of initial success of the agents with the inquiring heuristic.

Finally, uncertainty in the resource environment associated with the length of the feast and famine periods in general is detrimental. Certainly it leads to a great deal of variability in both the average population size and the average energy of the agents. For repetition, uncertainty in the availability of energy resources greatly decreases the population size of the superior limit-105 agents that can be maintained, with more than 50% of the agents dying across the cycle for std 1 and almost total extinction for std 5. This harks to the very delicate energy balance in place, where agents can just about make it through a given famine if its duration is 40 but a famine that lasts a bit longer can easily lead to death. If a famine lasts less than the mean however, the positive consequences - the accumulation of a bit more energy - are minimal compared to death. From an evolutionary perspective there is a strong truncation selection in play.

In the case of imitation, we see that the effect on the population average of the extra variability associated with the availability of energy resources is masked by the intrinsic variability inherent in the mistakes made by the limit-105 agents that imitate the limit-55 agents. Interestingly, although the average population sizes are similar over the first four cycles, the average energy decreases monotonically as a function of the standard deviation. This is due to the highly asymmetric nature of the effects of uncertainty in the feast

versus the famine periods. If the feast period lasts longer this is of no significant benefit as the limit-105 agents reach an energy accumulation limit wherein their consumption and their metabolic needs are equal. On the other hand, as emphasized, a longer famine period easily leads to death. This is a "gambler's ruin" type effect.

For inquiring, we see $std = 5$ is definitely inferior in both average population and average energy. Interestingly, the $std = 1$ results are very similar to the $std = 0$ results for average energy and where the average population exhibits a less catastrophic collapse than the $std = 0$ population in that the uncertainty dilutes to some degree the competition between the agents. Additionally, the inquiring agents in the $std = 1$ scenarios have the possibility of storing a bit of extra energy during those periods of abundance that last longer than 60 periods.

Turning finally to the results of Fig. 6, we believe that this result explains the current obesity epidemic if we take the feast-famine cycles to represent our "prehistoric/historic" past, with both predictable and unpredictable components to food resource availability, and the last 500 generations representing our predictable and ample food resource present. In other words, the survivors of the uncertain past that possessed the thrifty genotype were predestined to reach their maximum energy limits. In the case of the limit-105 agents this maximum, 62.6, emerges from an energy balance between constant consumption that is matched by a corresponding metabolic expenditure. This metabolic expenditure involves no energy costs of movement, as the S9 agents just stay in one place, taking advantage of the predictable energy resources that constantly appear. The energy cost is purely associated with their base metabolism, having to maintain a higher body mass. This limit of 62.6 is sensitive to the proportionality factor that determines what proportion of energy is used to maintain the agent metabolically. For any surviving limit-55 agent, however, their energy limit in the period of constant abundance is fixed by the "physiological" limit of 55.

Conclusions

Two important characteristics of the obesity epidemic are its ubiquity and its resilience. With respect to ubiquity: What differs between one country and another is not whether there is a problem with overweightedness and obesity, and its concomitant health problems, but, rather, just how severe the problem is. Effectively, only some sub-Saharan countries have avoided the problem. With respect to resilience, it exists at both the individual level and the group level, where, at the individual level, reversion of the obese state to normal weight is very difficult (Fildes, Charlton, Rudisill, Littlejohns, Prevost & Gulliford, 2015), while at the group level it has been exceedingly difficult to design public health policies that have a significant impact and are widely adopted by the population.

Two potential, complementary explanations for the ubiquity of the obesity epidemic are genetics, such as the thrifty genotype hypothesis, and the recent development of an obesogenic environment. Thrifty genotype explanations blame our genetic heritage, but in a causally indirect way, in that the purported genes are associated with physiological

adaptations that make it easier to get fat. On the other hand, blaming the environment seems to neglect the fact that we ourselves are the creators of that environment. It is undeniable that widespread obesity was not a problem for our prehistoric ancestors but that it is now. But how to you test hypotheses about such changes? ABM are, in spite of their defects, at least one promising avenue for creating and testing such hypotheses, as we have done here. If we take our energy storage parameter as a proxy for a “physiological” thrifty gene, what we have shown here is that, indeed, it can offer a selective advantage in the context of resource environments that mimic feast and famine cycles, with the feast period being essential in order to store up energy in the first place. However, we have also shown that there must be present two important behaviour types - conducts - in order for extra energy storage to be a useful adaptation: overconsumption and sedentariness. By overconsumption we mean that an agent must consume above and beyond their base metabolic level over an extended period of time. By sedentariness, we mean that energy expenditure by unnecessary activity is selected against. This is manifest in the success of the S9 strategy, where movement is initiated only if there is no food resource in the agent’s cell but there is in an adjacent cell.

As our results indicate that there is a selective advantage in overconsumption and sedentariness in energy environments that have periods of abundance and scarcity and, more generally, uncertainty, we must ask how might this selective advantage have manifested itself? Such a strong selective effect must surely have induced a genetic response and, subsequently, left a strong genetic imprint. We believe that the legacy of a genetically imprinted tendency to both overconsume and be sedentary, when combined with a thrifty genotype that gives the opportunity to store the excess energy that accrues from these behaviours, is precisely what has led to the current obesity epidemic. Indeed, our simulation of an energy rich environment after a period of selection through feast-famine cycles clearly shows that agents reach their maximum energy storage through the twin effects of overconsumption and sedentariness.

Finally, although there is much future work to be done, we can see the subtle complexity that enters when adding in heuristics which, here, represent variation, and can be thought of as analogs of mutation or learning. The random imitation of strategies within an agent’s community generates the possibility of an advantageous change for a limit-55 agent changing to a limit-105 strategy and the corresponding disadvantageous change a limit-105 agent changing strategy to limit-55. Compared to the pure repetition of the limit-105 strategy however, such imitation is disadvantageous to the population as a whole. On the contrary, for the inquiring heuristic, the limit-55 agents quickly copy the limit-105 agents which, in turn, leads to a population crash due to excess competition.

Declaration of conflicting interests

The author(s) declared following potential conflicts of interest with respect to the research, authorship, and/or publication of this article: [details omitted for double-anonymised peer review]

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship and/or publication of this article: This work was funded by [details omitted for double-anonymised peer review]

References

- Balke, T. & Gilbert, N. (2014). How do agents make decisions? a survey. *Journal of Artificial Societies and Social Simulation*, 17(4), 13.
- van Beilen, M., Bult, H., Renken, R., Stieger, M., Thumfart, S., Cornelissen, F. & Kooijman, V. (2011). Effects of visual priming on taste-odor interaction. *PLoS One*, 6(9), e23857.
- Bray, G. A. (2007). *The Metabolic Syndrome and Obesity*. Totowa NJ: Humana Press.
- Chaiken, S. & Trope, Y. (1999). *Dual-process theories in social psychology*. Guilford Press.
- Chambaron, S., Chisin, Q., Chabanet, C., Issanchou, S. & Brand, G. (2015). Impact of olfactory and auditory priming on the attraction to foods with high energy density. *Appetite*, 95, 74–80.
- Chandler, M., Cunningham, S., Lund, E., Khanna, C., Naramore, R., Patel, A. & Day, M. (2017). Obesity and associated comorbidities in people and companion animals: a one health perspective. *Journal of comparative pathology*, 156(4), 296–309.
- Christakis, N. A. & Fowler, J. H. (2007). The spread of obesity in a large social network over 32 years. *New England journal of medicine*, 357(4), 370–379.
- Cohen, D. A. & Babey, S. H. (2012). Contextual influences on eating behaviours: heuristic processing and dietary choices. *Obesity Reviews*, 13(9), 766–779.
- Cosmides, L. & Tooby, J. (1994). Better than rational: Evolutionary psychology and the invisible hand. *The American Economic Review*, 84(2), 327–332.
- Drewnowski, A. (1998). Energy density, palatability, and satiety: implications for weight control. *Nutrition reviews*, 56(12), 347–353.
- Dunlosky, J. & Bjork, R. A. (2013). *Handbook of metamemory and memory*. Psychology Press.
- Epstein, J. M. (2006). *Generative social science: Studies in agent-based computational modeling*. Princeton University Press.
- Evans, J. S. B. (2008). Dual-processing accounts of reasoning, judgment, and social cognition. *Annu. Rev. Psychol.*, 59, 255–278.
- Evans, R. M., Barish, G. D. & Wang, Y.-X. (2004). Ppars and the complex journey to obesity. *Nature medicine*, 10(4), 355.
- Fildes, A., Charlton, J., Rudisill, C., Littlejohns, P., Prevost, A. T. & Gulliford, M. C. (2015). Probability of an obese person attaining normal body weight: cohort study using electronic health records. *American journal of public health*, 105(9), e54–e59.
- Flegal, K. M., Graubard, B. I., Williamson, D. F. & Gail, M. H. (2005). Excess deaths associated with underweight, overweight, and obesity. *Jama*, 293(15), 1861–1867.

- Flegal, K. M., Graubard, B. I., Williamson, D. F. & Gail, M. H. (2007). Cause-specific excess deaths associated with underweight, overweight, and obesity. *Jama*, 298(17), 2028–2037.
- Frankish, K. (2010). Dual-process and dual-system theories of reasoning. *Philosophy Compass*, 5(10), 914–926.
- Ganley, R. M. (1989). Emotion and eating in obesity: A review of the literature. *International Journal of eating disorders*, 8(3), 343–361.
- Geier, A. B., Rozin, P. & Doros, G. (2006). Unit bias: A new heuristic that helps explain the effect of portion size on food intake. *Psychological Science*, 17(6), 521–525.
- German, A. J. (2006). The growing problem of obesity in dogs and cats. *The Journal of nutrition*, 136(7), 1940S–1946S.
- Gigerenzer, G., Fiedler, K. & Olsson, H. (2012). Rethinking cognitive biases as environmental consequences. In P. M. Todd & G. Gigerenzer (Eds.), *Ecological rationality: Intelligence in the world*. Oxford University Press.
- Gigerenzer, G. & Gaissmaier, W. (2011). Heuristic decision making. *Annual review of psychology*, 62, 451–482.
- Hammond, R. A. (2009). Peer reviewed: complex systems modeling for obesity research. *Preventing chronic disease*, 6(3).
- Haselton, M. G., Nettle, D. & Murray, D. R. (2005). The evolution of cognitive bias. *The handbook of evolutionary psychology*.
- Huang, T. T., Drewnowski, A., Kumanyika, S. K. & Glass, T. A. (2009). A systems-oriented multilevel framework for addressing obesity in the 21st century. *Preventing chronic disease*, 6(3).
- Hudson, A. & Williams, S. G. (1981). Eating behavior, emotions, and overweight. *Psychological Reports*, 48(2), 669–670.
- Jager, W. et al. (2000). *Modelling consumer behaviour*. Universal Press The Netherlands.
- Jager, W. & Janssen, M. (2002). The need for and development of behaviourally realistic agents. In *International Workshop on Multi-Agent Systems and Agent-Based Simulation* (pp. 36–49).
- Jager, W. & Janssen, M. (2012). An updated conceptual framework for integrated modeling of human decision making: The consumat ii. In *paper for workshop complexity in the Real World@ ECCS* (pp. 1–18).
- Jager, W., Janssen, M., De Vries, H., De Greef, J. & Vlek, C. (2000). Behaviour in commons dilemmas: Homo economicus and homo psychologicus in an ecological-economic model. *Ecological economics*, 35(3), 357–379.
- Johnson, R. J. & Andrews, P. (2010). Fructose, uricase, and the back-to-africa hypothesis. *Evolutionary Anthropology: Issues, News, and Reviews*, 19(6), 250–257.
- Kahneman, D. (2003). A perspective on judgment and choice: mapping bounded rationality. *American psychologist*, 58(9), 697.
- Kornell, N. (2009). Metacognition in humans and animals. *Current Directions in Psychological Science*, 18(1), 11–15.
- Lindenberg, S. & Steg, L. (2007). Normative, gain and hedonic goal frames guiding environmental behavior. *Journal of Social issues*, 63(1), 117–137.
- Lindley, D. V. (1991). *Making decisions*. Wiley.
- Macal, C. M. & North, M. J. (2010). Tutorial on agent-based modelling and simulation. *Journal of simulation*, 4(3), 151–162.
- Marsh, B. (2002). Do animals use heuristics? *Journal of Bioeconomics*, 4(1), 49–56.
- Mattes, R. D. (1997). Physiologic responses to sensory stimulation by food: nutritional implications. *Journal of the American Dietetic Association*, 97(4), 406–413.
- Mifflin, M. D., St Jeor, S. T., Hill, L. A., Scott, B. J., Daugherty, S. A. & Koh, Y. O. (1990). A new predictive equation for resting energy expenditure in healthy individuals. *The American journal of clinical nutrition*, 51(2), 241–247.
- Morton, G. J., Blevins, J. E., Williams, D. L., Niswender, K. D., Gelling, R. W., Rhodes, C. J., Baskin, D. G. & Schwartz, M. W. (2005). Leptin action in the forebrain regulates the hindbrain response to satiety signals. *The Journal of clinical investigation*, 115(3), 703–710.
- Mousavi, S. & Gigerenzer, G. (2014). Risk, uncertainty, and heuristics. *Journal of Business Research*, 67(8), 1671–1678.
- Must, A., Spadano, J., Coakley, E. H., Field, A. E., Colditz, G. & Dietz, W. H. (1999). The disease burden associated with overweight and obesity. *Jama*, 282(16), 1523–1529.
- Neel, J. V. (1962). Diabetes mellitus: a “thrifty” genotype rendered detrimental by “progress”? *American journal of human genetics*, 14(4), 353.
- Nisbett, R. E. (1972). Hunger, obesity, and the ventromedial hypothalamus. *Psychological Review*, 79(6), 433.
- Plutchik, R. (1976). Emotions and attitudes related to being overweight. *Journal of Clinical Psychology*, 32(1), 21–24.
- Ratcliffe, M. (2017). *Empathy Without Simulation* (pp. 199–220). Routledge.
- Reyna, V. F. (2004). How people make decisions that involve risk: A dual-processes approach. *Current directions in psychological science*, 13(2), 60–66.
- Scheibehenne, B., Miesler, L. & Todd, P. M. (2007). Fast and frugal food choices: Uncovering individual decision heuristics. *Appetite*, 49(3), 578–589.
- Shafraan, R., Lee, M., Cooper, Z., Palmer, R. L. & Fairburn, C. G. (2007). Attentional bias in eating disorders. *International Journal of Eating Disorders*, 40(4), 369–380.
- Sloman, S. A. (1996). The empirical case for two systems of reasoning. *Psychological bulletin*, 119(1), 3.
- Smith, J. D., Shields, W. E. & Washburn, D. A. (2003). The comparative psychology of uncertainty monitoring and metacognition. *Behavioral and brain sciences*, 26(3), 317–339.
- Smith, J. D. & Washburn, D. A. (2005). Uncertainty monitoring and metacognition by animals. *Current Directions in Psychological Science*, 14(1), 19–24.
- Spiegelman, B. M. & Flier, J. S. (2001). Obesity and the regulation of energy balance. *cell*, 104(4), 531–543.
- Stephens, C. R. (2021). “Ome” sweet “ome”: From the genome to the conductome. *2019-20 MATRIX Annals* (pp. 287–297).

-
- Strack, F. & Deutsch, R. (2004). Reflective and impulsive determinants of social behavior. *Personality and social psychology review*, 8(3), 220–247.
- Tversky, A. & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *science*, 185(4157), 1124–1131.
- Wilensky, U. (1999). Netlogo. <http://ccl.northwestern.edu/netlogo/> Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL.
- World Health Organization (2020). *Obesity and overweight. Fact sheet*. (accessed November 10, 2020).