

Is impulsive behavior adaptive in harsh and unpredictable environments? A formal model

Jesse Fenneman*, Willem E. Frankenhuis

Behavioural Science Institute, Radboud University, the Netherlands

ARTICLE INFO

Keywords:

Impulsivity
Harshness
Unpredictability
Bayesian inference
Life history theory
Formal model

ABSTRACT

Evolutionary social scientists have argued that impulsive behavior is adaptive in harsh and unpredictable conditions. Is this true? This paper presents a mathematical model that computes the optimal level of impulsivity in environments varying in harshness and unpredictability. We focus on information impulsivity, i.e., choosing to act without gathering or considering information about the consequences of one's actions. We explore two notions of harshness: the mean level of resources (e.g., food) and the mean level of extrinsic events (e.g., being the victim of a random attack). We explore three notions of unpredictability: variation in resources, variation in extrinsic events, and the interruption risk (the chance that a resource becomes unavailable). We also explore interactions between harshness and unpredictability. Our general model suggests four broad conclusions. First, impulsive behavior is not always adaptive in harsh and unpredictable conditions; rather, this depends on the exact definitions of harshness, unpredictability, and impulsivity. Second, impulsive behavior may be adaptive in environments in which the quality of resources is low or high, but is less likely to be adaptive when their quality is moderate. Third, impulsive behavior may be adaptive when resource encounters are likely to be interrupted. Fourth, extrinsic events have only a limited effect on whether impulsive behavior is adaptive. We discuss the implications of these findings for future research, consider limitations, and suggest future directions.

1. Introduction

A common perspective is that self-control helps us make better decisions and achieve valued goals (Baumeister, 2002; Duckworth, 2011; Mischel, Shoda, & Rodriguez, 1989; Moffitt et al., 2011; Verdejo-García, Lawrence, & Clark, 2008). This view regards impulsive behavior as short-sighted, resulting from low levels of self-control. Such a view implies that impulsive behavior is negative, defective, and dysfunctional (Daly & Wilson, 2005; Daruna & Barnes, 1993; Evenden, 1999; Frankenhuis, Panchanathan, & Nettle, 2016; Wilson & Daly, 2006). These judgments are often reflected in definitions of impulsive behavior, for instance, as behavior that is “poorly conceived, prematurely expressed, unduly risky or inappropriate to the situation and that often result in undesirable consequences” (Daruna & Barnes, 1993, p. 23).

The emphasis on dysfunction may be appropriate from a health and wellbeing perspective. Longitudinal studies show that high levels of self-control (low impulsivity) in childhood predict a broad suite of desirable outcomes later in life, including better health, more wealth, and stable social relationships (Duckworth & Seligman, 2005; Duckworth, Tsukayama, & Kirby, 2013; Moffitt et al., 2011). Evolutionary social scientists, however, have offered a different perspective. They note that

there is a large variation in impulsive behavior, both between individuals or within the same individuals over time, and argue that this may reflect a developmental and behavioral adaptation to harsh and unpredictable environments. Of course, this perspective privileges biological fitness (survival and reproductive success), not health and wellbeing. These two notions of adaptation are conceptually independent: behaviors might hurt wellbeing, yet enhance fitness, or vice versa (Frankenhuis & Del Giudice, 2012).

Over the past decade, there has been a surge of papers arguing that impulsive behavior is adaptive in harsh and unpredictable environments (reviewed in section 1.1). However, this “adaptive impulsivity” hypothesis is based on natural language, not on formal modeling. Natural language is often ambiguous, leaving room for different interpretations of concepts and how they relate to one another. As we discuss below, there are multiple interpretations of harshness, unpredictability, and impulsivity, making the truth and scope of this hypothesis unclear. Formal modeling resolves such ambiguity by expressing concepts and their relations mathematically. In addition, formal modeling allows a researcher to analyze under what conditions a hypothesis follows logically from assumptions (Epstein, 2008; Smaldino, 2017). In this paper, we develop a formal model that explores the optimal level of

* Corresponding author at: Behavioural Science Institute, Radboud University, Montessorilaan 3, PO box 9104, 6500 HE, Nijmegen, The Netherlands.
E-mail address: jessefenneman@gmail.com (J. Fenneman).

one type of impulsive behavior (i.e., information impulsivity, see section 1.2.3) in environments varying in both harshness and unpredictability.

1.1. Why impulsive behavior might be adaptive in harsh and unpredictable environments

The hypothesis that impulsive behavior is adaptive in harsh and unpredictable environment has been inspired by life history theory (though not logically derived from it). Life history theory is a framework based on a collection of mathematical models from evolutionary biology and ecology (Stearns, 2000), which is increasingly used in the evolutionary social sciences (Del Giudice, Gangestad, & Kaplan, 2015; Nettle & Frankenhuis, 2019). Central to this framework is the question why species differ in how they allocate finite resource (such as time and energy) across different life history traits, such as the rate of growth, onset of reproduction, quantity and quality of offspring, parenting and mating effort, and delaying senescence. As decisions made early in life influence how life history traits trade-off later in life, scholars have proposed that life history traits cluster on a single dimension. For instance, species that invest in rapid growth at the expense of somatic quality (also called *r*-selected species) may have a reduced lifespan, and delaying reproduction increases the risk of dying before reproduction. As offspring are more likely to die before adulthood, such organisms tend to produce a high number of offspring with lower levels of parental investment per offspring. In contrast, species that invest in somatic quality at the expense of rapid growth (also called *K*-selected species) tend to have longer lifespans, delayed reproduction, and larger investments in fewer offspring.

In the past decade, evolutionary social scientists and biologists have increasingly used life history theory to explain individual differences in behavioral and physiological traits as well. Specifically, they have proposed that underlying patterns of resource allocation may be associated with other behavioral and physiological traits not directly related to life history traits (Réale et al., 2010; Royauté, Berdal, Garrison, & Dochtermann, 2018). According to this pace of life syndrome (POLS) hypothesis, variation in life history traits create suites of correlated behavioral and physiological traits. A widespread (but controversial; see also section 4.2) claim is that these suites fall on a continuum ranging from ‘fast’ to ‘slow’ (Ellis, Figueredo, Brumbach, & Schlomer, 2009; Giudice, Gangestad, & Kaplan, 2005). Fast-paced individuals mature earlier, reproduce earlier, and have shorter lifespans. Due to their shorter expected lifespans, they are thought to focus on immediate (as opposed to long-term) rewards. Impulsive behavior is thought of as a fast-life behavioral trait because it allows individuals to maximize resources in the short term. Impulsivity is a behavioral trait, not a life history trait. Life history traits govern the scheduling and allocation of resources to competing life-history objects (e.g., investing in current or future reproduction). Impulsivity is a behavioral trait that increases available resources as it allows individuals to act swiftly, seize fleeting opportunities, and avoid the costs associated with collecting information.

Two characteristics of the environment are thought to be essential in determining pace of life: harshness and unpredictability (Del Giudice, 2018; Ellis et al., 2009). Harshness refers to morbidity and mortality caused by factors beyond the control of the individual. In a harsh environment the expected lifespan is short, and individuals maximize fitness by reproducing early before they die. Unpredictability refers to stochastic variation in harshness over time or space. In unpredictable environments conditions might worsen at any time. Hence the future is uncertain and difficult to plan for, favoring faster paces of life that are rewarding in the short-term.

As harshness and unpredictability are thought to favor faster paces of life, and impulsivity might be part of a fast pace of life, evolutionary social scientists have proposed that impulsive behavior is adaptive in harsh and unpredictable environments. Empirical evidence for this

hypothesis is mixed. On the one hand, correlational studies show that a faster life-history strategy is associated with higher levels of impulsivity (Lee, DeBruine, & Jones, 2018; Mishra, Templeton, & Meadows, 2017). Similarly, experimental studies show that people who have grown up in harsh environments may be more likely to respond to resource scarcity with higher levels of impulsive behavior (Griskevicius et al., 2013; Kruger, Reischl, & Zimmerman, 2008; but see Pepper et al., 2017 for a replication study that obtains a different result) and residents of countries with low life expectancy show less willingness to wait for a delayed reward (Bulley & Pepper, 2017). On the other hand, fast life history traits such as early maturation do not always correlate with less deliberation, exploration, or future orientation in humans (Copping, Campbell, & Muncer, 2013, 2014) or non-human animals (Royauté et al., 2018).

1.2. Conceptual challenges

The hypothesis that impulsive behavior increases fitness in harsh and unpredictable environments provides an explanation for the variation in impulsivity between individuals, and within individuals over time and across contexts. There remains, however, conceptual ambiguity in this adaptive impulsivity hypothesis: there are different notions of harshness, unpredictability, and impulsivity, making the precise meaning, validity, and scope of this hypothesis unclear. Before presenting our model, we discuss various interpretations of these concepts. Our discussion is not exhaustive; there are other usages, which we do not discuss. Our goal is to familiarize the reader with the common usages in evolutionary social science, and to show there is a need for greater conceptual clarity in theoretical discussions and empirical tests of the adaptive impulsivity hypothesis.

1.2.1. Harshness

Evolutionary social scientists often define harshness as extrinsic mortality-morbidity; that is, the rate at which external factors, which an individual cannot control, cause disability and death. Empirical studies rarely measure extrinsic mortality-morbidity directly. Instead, they operationalize harshness in one of two ways. The first operationalization focuses on exposure to violence and crime (Brumbach, Figueredo, & Ellis, 2009; Griskevicius, Delton, Robertson, & Tybur, 2011; McCullough, Pedersen, Schroder, Tabak, & Carver, 2013; Mell, Safra, Algan, Baumard, & Chevallier, 2018; Pepper et al., 2017) and/or low life expectancy (Aronoff & DeCaro, 2019; Lee et al., 2018; Mell et al., 2018; Pepper & Nettle, 2014; M. Wilson & Daly, 1997). The second operationalization focuses on resource scarcity or low resource quality, often measured as low socioeconomic status (Allen & Nettle, 2019; Griskevicius et al., 2013; Griskevicius, Tybur, Delton, & Robertson, 2011; McCullough et al., 2013; Mell et al., 2018; Pepper et al., 2017; Simpson, Griskevicius, Kuo, Sung, & Collins, 2012). Although extrinsic morbidity-mortality and low resource quality are often empirically correlated (Ellis et al., 2009), each may actually pose different selection pressures on impulsive behavior, as we will show later.

1.2.2. Unpredictability

There are at least three different notions of unpredictability. First, the state of the environment itself can be unpredictable due to random changes in the degree of harshness over time, space, or both. If the environment is only *spatially* unpredictable, the environment does not change over time, but different individuals can experience different environmental states. If the environment is only *temporally* unpredictable, all individuals experience the same environmental state, which varies over time. Empirical studies have measured environmental unpredictability, for instance, as the number of household moves, or the number of different father figures passing through a single-mother household (Belsky, Schlomer, & Ellis, 2012; Mittal & Griskevicius, 2014; Simpson et al., 2012). However, such measurements typically include both spatial and temporal variation. For

instance, people might move house because geographical areas differ in their level of harshness (spatial unpredictability), because the harshness in a geographical area changes over time (temporal unpredictability; e.g., jobs may move away), or both. Second, an environment that does not change over space or time might still be unpredictable if, for a mean level of harshness, there is large variance in possible outcomes (Frankenhuis et al., 2016). That is, even if the mean level of harshness does not change over time or space, the range of experiences in a particular environment may be large or small. For instance, in a harsh but predictable environment, resources might be consistently poor (e.g., money is always tight). In a harsh and unpredictable environment, there is variation in resource availability (e.g., piece workers might have large pay checks in some periods but receive little income in others). Similarly, in unpredictable environments some probabilistic event (e.g., being randomly attacked) might occur (Brumbach et al., 2009; Simpson et al., 2012). Third, resources within stable environment might be predictable but their persistence can be unpredictable. That is, currently available resources (e.g., food or job) might become unavailable due to interruptions (i.e., a decay rate or collection risk; Stephens, 2002). For instance, in an unchanging environment where there is little variation in resource quality (e.g., jobs always provide little income), resources might disappear due to high competition (e.g., other piece workers might ‘steal’ a job site).

1.2.3. Impulsivity

In empirical studies, different measures of impulsivity are only weakly correlated or even uncorrelated (Cyders & Coskunpinar, 2011; MacKillop, Weafer, Gray, & Palmer, 2016; Sharma, Markon, & Clark, 2014), providing limited support (if any) for a single-factor model (Caswell, Bond, Duka, & Morgan, 2015; Stahl et al., 2014), suggesting that impulsivity may not be a unitary construct. There is an active debate on how many and which types of impulsivity exist, with studies typically dividing impulsivity in two to five types (Cyders, 2015; Evenden, 1999; Fineberg et al., 2014; Frederick, Loewenstein, & O’Donoghue, 2002; Hamilton et al., 2015; Hamilton et al., 2015).

A common divide is between choice impulsivity and motor impulsivity. Choice impulsivity refers to preferences that result in impulsive decisions (Hamilton, Mitchell, et al., 2015), and can be further divided into two separate constructs (Caswell, Bond, et al., 2015; Fineberg et al., 2014; MacKillop et al., 2016). The first is temporal impulsivity, or the tendency to prefer immediate rewards over later ones (also known as temporal preference, delay of gratification, or temporal discounting; Caswell et al., 2015; Frederick et al., 2002; Hamilton, Mitchell, et al., 2015). The second type is information impulsivity, or the tendency to act without gathering or considering information about the consequences of one’s actions (also known as non-planning impulsivity, reflection impulsivity, or impulsive processing; Caswell, Morgan, & Duka, 2013; Hamilton, Littlefield, et al., 2015; Kagan, 1966). In contrast, motor impulsivity (also known as impulsive action) occurs after a decision is made, and refers to the inability to control, inhibit, or cancel motor patterns. There are more possible types of impulsivity besides information, temporal, and motor impulsivity. However, these types either result from somatic damage or psychopathology (e.g., ADHD or addiction), or are explicitly defined as negative outcomes (e.g., impulsivity as disadvantageous decision making; Fineberg et al., 2014). In this article we explore under which environmental conditions evolution by natural selection might have favored impulsive behavior. As such, we consider only types of impulsivity that results from evolutionary pressures and that shape preferences. We do not study under what conditions an individual is able to act according to its preference (i.e., motor impulsivity) or how specific (psycho) pathologies influence this ability (i.e., disadvantageous decision making).

Temporal and information impulsivity are both important constructs in life-history research. Temporal impulsivity is associated with a fast pace of life. Individuals with a present-orientation favor

immediate consumption over saving for the future. This behavior may be adaptive in harsh and unpredictable environments, where individuals face increased levels of poverty, violence, disability, and/or experience less control over their environment (Frankenhuis & Nettle, 2020; Pepper & Nettle, 2017). Such hardship reduces life expectancy, increases uncertainty about future rewards, and implies greater opportunity costs to not using resources in the short term (Mell et al., 2019). In such conditions, people might benefit from maximizing short-term rewards.

Information impulsivity – the focus of the present study – is associated with two behavioral traits often considered part of a fast pace of life. The first is poor planning due to a failure to deliberate. Two questionnaires often used in the evolutionary social sciences measure this trait. The mini-K (Figueredo et al., 2006) includes questions such as “I try to understand how I got into a situation to figure out how to handle it” and “I often make plans in advance”. Studies that explore the psychometric properties of the mini-K scale suggest that these items form a single factor that partly measure planning deficits (Richardson, Chen, Dai, Brubaker, & Nedelec, 2017). Other studies in the evolutionary social sciences measure impulsivity using Eysenck’s Impulsivity Scale (Eysenck, Pearson, Easting, & Allsopp, 1985; Mishra & Novakowski, 2016), a scale that explicitly defines impulsivity as “doing and saying things without thinking” (Eysenck, Easting, & Pearson, 1984, p 3.15). The second behavioral trait is the tendency to favor quick and shallow exploration over thorough exploration. When presented with a new situation, a fast-paced individual acts boldly, investing little time in understanding the situation before acting (Baumard & Chevallier, 2015; Del Giudice, 2015; Réale et al., 2010; Sih & Del Giudice, 2012). The Balloon Analogue Risk Task (BART) – used in many fields of social science, including evolutionary psychology (Griskevicius et al., 2013; Humphreys et al., 2015; Lu & Chang, 2019; Mishra, Lalumière, & Williams, 2010) – measures this tendency. In this task, participants earn money by repeatedly inflating balloons. Inflating a balloon makes it more valuable. However, each balloon has a maximum inflation point. Beyond this point, the balloon bursts, ending the trial without reward. Participants thus face a tradeoff between a small-but-certain reward and a larger-but-less-certain reward. Crucially, in most versions of the BART participants are not told what the maximum inflation point is. Rather, they can learn about this point by trial and error, which requires sacrificing a number of balloons. People vary in their willingness to pay for information. Due to task impurity, evolutionary social scientists using these versions of the BART measure information impulsivity, alongside risk preferences.

In the evolutionary social sciences, the term ‘impulsivity’ is used to refer to temporal impulsivity, information impulsivity, or both. Some researchers explicitly use the term to refer to one type of impulsivity: either to information impulsivity (Brumbach et al., 2009; Copping et al., 2013; Del Giudice, 2015; Mishra et al., 2017) or temporal impulsivity (Allen & Nettle, 2019; Daly & Wilson, 2005; Griskevicius et al., 2013; Lee et al., 2018; Mittal & Griskevicius, 2014; Sih & Del Giudice, 2012). Other researchers use separate labels for impulsivity and time-oriented traits such as future orientation or time perspective (Chen & Vazsonyi, 2011; Del Giudice, 2014; Del Giudice & Belsky, 2010; Griskevicius et al., 2013). This distinction suggests they view temporal and information impulsivity as separate constructs. Finally, some researchers regard information impulsivity as a separable component of a broader cluster of impulsivity traits; defining impulsivity, for instance, as “[...] a stable tendency to act without deliberation and without consideration of future consequences, [reflecting] a combination of behavioral disinhibition and future discounting” (Del Giudice, 2018, p. 114). This conceptual manifold is problematic. As MacKillop et al. (2016) noted: “the use of a catch-all term impulsivity to refer to distinct characteristics may foster ambiguity and confusion in the literature” (p. 3362).

1.3. Limitations of existing models

Formal models have explored optimal levels of impulsive behavior as a function of the quality of resources, such as food and mates. For instance, models of mate choice have shown that it is adaptive to collect more information before choosing a mate (i.e., low information impulsivity) when the variance in the quality of mates is high (Collins, McNamara, & Ramsey, 2006; Fawcett & Johnstone, 2003; Luttbegg, 1996; Mazalov, Perrin, & Dombrovsky, 1996). Similarly, foraging models suggest some environmental conditions that promote information impulsivity (Dunlap & Stephens, 2012; Mathot & Dall, 2013). For instance, as food becomes scarcer or the requirements for survival increase, the relative cost of sampling information increases, promoting information impulsivity. Similarly, when the variance in food quality decreases, sampling information becomes less beneficial, promoting information impulsivity. Such models might provide a formal foundation for the adaptive impulsivity hypothesis. However, these models have usually focused on a single type of harshness and a single type of unpredictability, and have not studied the interactive effects of harshness and unpredictability on impulsive behavior.

1.4. Our contribution

We present a mathematical model that computes the optimal level of information impulsivity in environments varying in both harshness and unpredictability. Although the empirical literature on the adaptive impulsivity hypothesis studies both temporal and information impulsivity, we focus on information impulsivity (Copping et al., 2013; Del Giudice, 2015, 2018; Mishra et al., 2017). We explore two notions of harshness: the quality of resources that agents make decisions about (e.g., food) and the quality of extrinsic events over which agents have no control (e.g., being randomly attacked). We explore three notions of unpredictability: variance in the resource quality, variance in the extrinsic event quality, and interruption rate (the probability that a given resource disappears). We leave temporal impulsivity for a future study. Thus, we use the term ‘impulsivity’ here to refer to information impulsivity.

2. Model

In the main text we describe the decision problem. In the supplementary materials we describe how agents solve this decision problem (appendices A to D). Furthermore, we provide the software implementation, programmed in JAVA (version 8.192), with extensive in-code documentation. All appendices and code are accessible at <https://github.com/JesseFenneman/AdaptiveInformationImpulsivity>.

2.1. The decision problem

Fig. 1 provides a graphical overview of the decision problem (see appendix B for an in-depth and formal description). An agent's state has two components: a soma, representing the agent's condition, whether it be physical, social, or material; and belief, representing the agent's knowledge about the resources in its environment. An agent's somatic state can range from 0 to 100 (at 0, the agent dies). This state changes when an agent interacts with its environment.

When encountering a resource (see appendix B.1.1.), an agent makes decisions that influence the outcome of interactions. Although an agent cannot influence the resource quality, it can decide whether to accept or reject a given resource (e.g., eat food or not, accept a job offer or not). If an agent accepts, the encounter is interrupted with a fixed probability before that agent can collect the outcome (e.g., someone else may be hired). If the agent accepts a resource and is not interrupted, the resource is added to its somatic state. This addition can be positive or negative, depending on the resource quality. After this interaction, over which the agent had some control, an extrinsic event

happens, over which the agent has no control (e.g., being randomly attacked, receiving windfall donations). This extrinsic event is added to (if positive) or subtracted from (if negative) the somatic state. A resource encounter is always followed by an extrinsic event, and vice versa. We call the pairing of a resource encounter and the subsequent extrinsic event a ‘cycle’. We want the timescale of a cycle to be short relative to the lifespan of an agent (we do not focus on end-of-life effects), and therefore assume that as long as an agent is alive, there are infinitely many cycles.

2.1.1. Priors, cues, and beliefs

Resources and extrinsic events are drawn randomly from separate and independent normal distributions. An agent does not know the quality of individual resources or extrinsic events until they impact its somatic state. We make two assumptions. First, an agent spends its entire life in a single, unchanging environment. There is no depletion of resources or extrinsic events; the distributions of resources and extrinsic events do not change from one cycle to the next. Second, an agent knows the broader (or meta) parameters of its environment: it knows the mean and variance of the distributions of resources and extrinsic events.

Before deciding to accept or reject a resource, an agent has the option to sample cues that provide imperfect information about the resource quality, and by extension about consequences of its actions (e.g., the food might be healthy or poisonous; a job may be dangerous or safe; see appendix B.1.1.1.). A cue can be positive or negative, and predicts the sign of the resource correctly or incorrectly with a probability known to the agent. This probability is given by the cue reliability, which is conditional on the resource quality: extremely positive (negative) resources are more likely to result in positive (negative) cues than resources that are closer to zero. Positive and negative cues are equally likely if the resource has a quality of 0. The cue reliability depends only on the quality of the encountered resource, not on the mean resource quality in the environment. That is, sampling a cue from a good quality resource in an environment where the mean resource quality is high results in a positive cue with the same likelihood of sampling from a good quality resource in an environment where the mean resource quality is low. However, sampling is costly: it requires an agent to invest some energy in sampling (e.g., gathering information about a potential employer requires effort). Each cue sampled reduces the somatic state by a fixed amount.

After sampling a cue, an agent updates its belief state about the resource quality in a Bayesian manner, the optimal way of information updating (McNamara, Green, & Olsson, 2006). This does not imply that we assume that psychological processes are Bayesian. Rather, we assume that natural selection favored behavior consistent with Bayesian updating. Such behavior might be instantiated in simple if-then rules. Over generations, natural selection might have favored if-then strategies that result in Bayesian-like behavior (Higginson, Fawcett, Houston, & McNamara, 2018; Trimmer et al., 2011). As we study information impulsivity but not temporal impulsivity, we assume that no time passes during an encounter; sampling immediately results in a cue. If sampling takes time, an agent might avoid sampling for two reasons: because somatic costs of sampling outweigh the benefit of information, or because it prefers to have outcomes sooner rather than later. This would make it impossible to separately evaluate the effects of environmental conditions on information impulsivity and temporal impulsivity.

2.1.2. Fitness

The somatic state of an agent at the end of its life determines its fitness (appendix B.1.3.). We explore three different mappings of state to fitness: linear, the marginal fitness increment is constant as state increases; diminishing, the marginal fitness increment decreases as state increases; and increasing, the marginal fitness increment increases as state increases. There are real-world examples for each of these

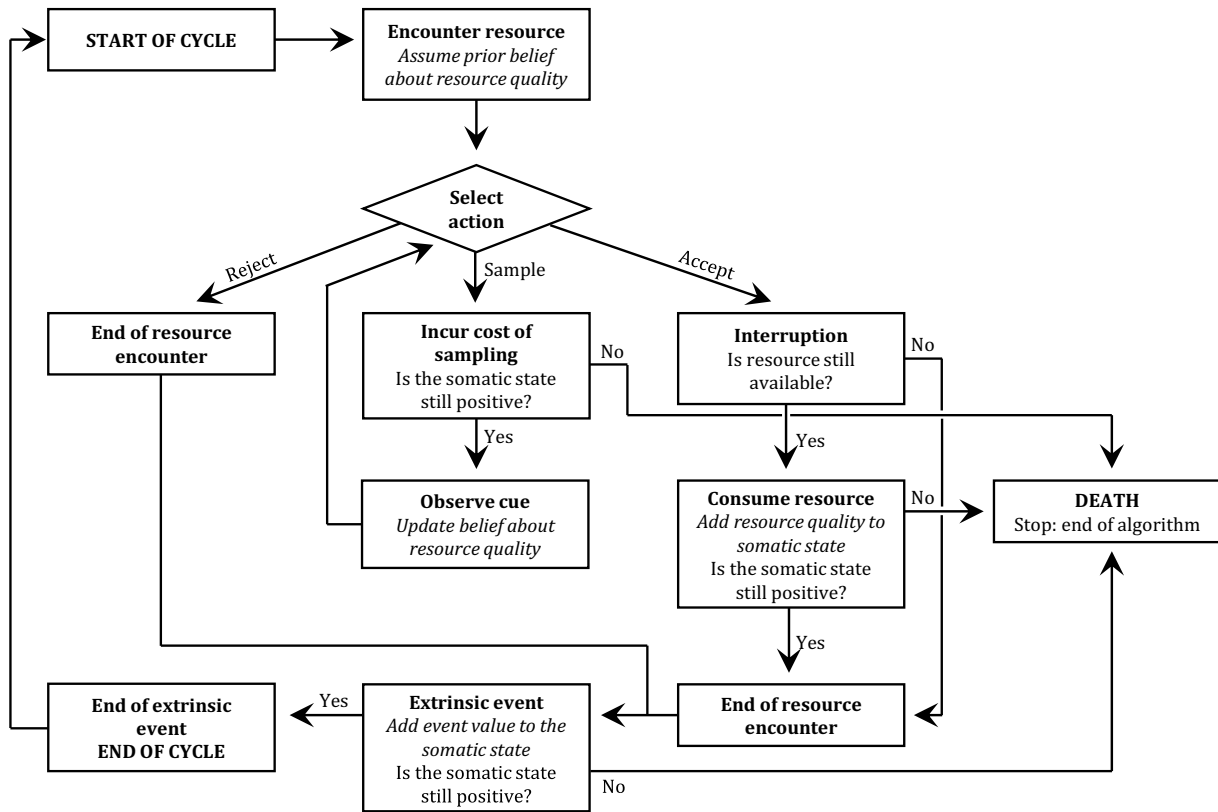


Fig. 1. A schematic overview of the decision problem faced by an agent.

mappings. For instance, in a large and well-mixed population, the reproductive output of a male might increase linearly with each mating. For a well-fed individual each additional morsel of food provides diminishing fitness returns. When there is high reproductive skew, each additional increase in strength might result in larger fitness gains. A marginally diminishing or increasing mapping requires that we set a decay or growth parameter, respectively. We conducted robustness checks to explore different values of these parameters. These analyses showed that the value of these parameters shift quantitative results, but not qualitative results. Therefore, we present results for only one set of parameter values. Readers are welcome to explore other values using the interface we provide with our model (see section 2.4). There are no additional parameters for a linear mapping. We describe results for the linear mapping in the main text, and results for the increasing and decreasing mappings in the supplementary materials.

2.2. Harshness and unpredictability

We vary five dimensions of the environment between agents (appendix B.1.4.):

- The mean resource quality (a kind of harshness);
- The mean extrinsic event quality (a kind of harshness);
- The variance in resource quality (a kind of unpredictability);
- The variance in extrinsic event quality (a kind of unpredictability); and
- The interruption rate (a kind of unpredictability).

2.3. States and optimal policies

A policy specifies the behavior of an agent in all possible states that that agent can be in. The optimal policy instructs an agent to take the fitness maximizing action in all states. We use stochastic dynamic programming and reinforcement learning to find the optimal policy

(Sutton & Barto, 2018). Appendix C provides an in-depth and formal description of how we used these techniques. Because the parameters of its environment do not change, an agent does not change its prior beliefs about resource and extrinsic event quality between cycles. That is, the outcome of a previous cycle does not provide any new information to the agent about these distributions. Therefore, the optimal policy does not depend on cues sampled in previous cycles; it depends only on an agent's somatic state at the start of a cycle.

2.3.1. Somatic damage and the discount rate

For real organisms, a good somatic state does not always guarantee survival. There are other potential causes of mortality (e.g., natural disasters). We capture these causes by assuming that there is a fixed probability of death at each time period. As a result, agents discount future cycles. Specifically, an agent values the outcome of a cycle farther into the future less than cycles closer in time. In our model, discounting is exponential: the outcome of a future cycle is reduced by a consistent factor for each intermediate cycle that agent has to go through. This factor is called the discount rate. It can range from 0 (if future expected outcomes have no value), to 1 (no discounting; all future outcomes are valued equally). Moreover, higher discount rates significantly increase the computational time of our model. We set the discount rate to 0.95. Discounting occurs *between* cycles, not within a cycle: sampling one more cue does not decrease the value of future cycles. This assumption is essential: it ensures that, in our model, information impulsivity is not confounded with temporal impulsivity. Moreover, it ensures that discounting affects the outcomes of all actions equally. An agent weights the long-term consequences of all actions equally.

2.3.2. Quantifying information impulsivity

Information impulsivity is the tendency to act without gathering or considering information about the consequences of one's actions. In our model an agent can only gather information by sampling cues. Each

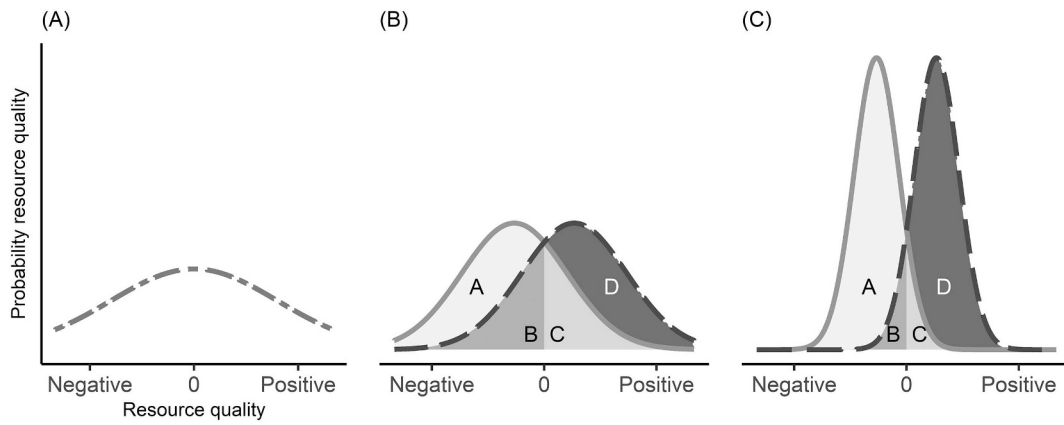


Fig. 2. A simplified overview of how sampling influences beliefs. An agent has a prior belief about the resource value (panel A). After sampling more positive than negative cues (panel B, dashed black line) an agent becomes optimistic about the resource quality. The combination of areas C and D reflect the belief that the resource is positive, whereas area B reflects the probability that the resource is negative. Because the combination of areas C and D is larger than B, accepting has a positive expected outcome. In contrast, sampling more negative than positive cues (panel B, solid gray line) makes an agent pessimistic. Here, area C reflects the belief that the resource is positive, whereas the combination of areas A and B reflect the probability that the resource is negative. In this case accepting is ill-advised: as area C is smaller than the combination of areas A and B, accepting has a negative expected outcome. Sampling more cues further reduces an agent's uncertainty about the resource quality (panel C). Reducing uncertainty increases the expected outcome of a resource encounter as it allows an agent to better differentiate between negative and positive resources. If most sampled cues are positive (dashed black line), the ratio of the combination of areas C and D to area B further increases. Consequently, accepting is more likely to result in a positive outcome. If most sampled cues are negative (solid gray line) the ratio of area C to the combination of areas A and B further decreases. Here, accepting is more likely to result in a negative outcome, and the best action is to reject. However, each additional cue results in a smaller decrease in uncertainty: after sampling a thousand cues, one more cue provides almost no additional information. There is, therefore, a point at which the cost of sampling outweighs the benefits of the reduced uncertainty. Upon reaching this point an agent stops sampling and either rejects or accepts the resource.

time it samples it always receives a single cue. The cue reliability is fixed, and cues are independent. As a result, each cue is equally informative. Because an agent always follows the optimal policy, it always uses this information to update its beliefs about the resource quality. As such, an agent that samples a cue will necessarily act based on more gathered information. Reversely, an agent that does not sample will necessarily base its action on less information. We therefore use the number of cues an agent samples as a measurement of information impulsivity. However, cues are probabilistic. This means that two agents can live in the same environment, follow the same policy, and encounter the same resource, but end up in different belief states and therefore make different decisions. For instance, one agent might sample four positive cues and decide to accept a resource, and another agent might sample two positive and two negative cues and decide to continue sampling before making a decision. We measure information impulsivity as the number of cues an agent *expects* to sample when starting a resource encounter, when it follows the optimal policy. Agents that sample fewer cues are more impulsive.

2.4. Parameters settings

In addition to setting the discount rate parameter, we also have set the cost of sampling and the cue reliability. Some settings would produce obvious or theoretically uninteresting results. For instance, we explored various costs to sampling. If the costs of sampling is very high, no sampling takes place. If, however, the cost of sampling is very low, an agent always samples a maximum number of cues. We therefore assume that there is some cost to sampling, but that this cost is not extremely high. In our view some costs are theoretically plausible. Biologists and psychologists often distinguish between two types of costs: search costs and processing costs. Search costs are paid when an individual actively gathers information (e.g., seeking advice about a potential job). Not all actions have search costs. For instance, a *Daphnia* that passively drifts in a pond may receive cues indicating the presence of a predator. Similarly, humans may receive cues about their environment if they witness many closed storefronts. However, processing this information does incur processing costs, which are the cognitive

costs of understanding ramifications and updating beliefs. Although such costs are unlikely to be high, they are also unlikely to be zero. Similarly, if cues are extremely unreliable, an agent never samples. If cues are extremely reliable, an agent always samples one cue that provides near-perfect information. We assumed that cues carry information, but are not perfect, nor extremely unreliable (c.f. Dall et al., 2005).

Some readers may be interested in different parameter settings. If so, they can use our graphical interface to explore parameter settings other than the ones that we focus on here. Our model can be tailored in at least four different ways. First, the parameters of the decision task can be altered (e.g., vary the cost of sampling). Second, the distribution of resources and extrinsic events within an environment can be changed (e.g., distributions can be made non-normal). Third, a reader might be interested in specific state-to-fitness mappings. Fourth, a reader might be interested in studying the policies in more detail (e.g., what exactly does the optimal policy prescribe?). Our graphical interface allows users to easily adjust the parameters of the decision task, the parameters of the environment of an agent, the state-to-fitness mapping, inspect an agent's policy in detail, and graph results. For more extensive customization, we provide the software implementation of our model online with extensive in-code documentation. Future studies can use this software as a foundation to build models tailored to other decision problems. We discuss possible extensions in section 4.2.

3. Results

We present qualitative results; we do not provide specific numbers and round values to nearest integers. We provide exact results in online appendices E (linear marginal returns), F (diminishing marginal returns), and G (increasing marginal returns). Our results do not qualitatively differ if we assume linear marginal returns, diminishing marginal returns, or increasing marginal returns.

Fig. 2 provides a brief overview of how sampling changes an agent's beliefs. An agent typically accepts or rejects a resource when it is sufficiently confident that the resource is positive or negative. When neither of these conditions is met, an agent samples additional cues. How

much information an agent needs before it is confident enough to make a decision depends on its environment and its current state. To gain insight, we focus on three dimensions that influence an agent's decision: (1) the prior probability that a resource is positive, (2) how many more positive than negative cues an agent needs, on average, before accepting, and (3) how many more negative than positive cues an agent needs, on average, before rejecting. These dimensions explain much of the variation between policies.

The optimal policy for when an agent starts a cycle with an extreme somatic state (i.e., very high or very low) is different from when it starts a cycle with an intermediate somatic state. As we are primarily interested in how the environment (rather than an agent's somatic state) shapes the optimal policy, we first focus on the intermediate starting state (section 3.1). Then we discuss the effect of extreme states (section 3.2). In these two sections, we assume there are no extrinsic events (i.e., the mean and variance of extrinsic events are both 0). In section 3.3, we discuss how extrinsic events shape optimal policies.

3.1. Sampling with intermediate starting somatic states

Fig. 3 shows optimal policies as a function of the mean resource quality, the variance in resource quality, and the interruption rate. We show optimal policies with increasing and decreasing returns in sections 1 and 2 of appendices E, F, and G.

3.1.1. Sampling increases with higher variance in resource quality

As the variance in the resource quality increases, so does an agent's prior uncertainty (Fig. 3, row 1). When this variance is close to 0, there is no uncertainty: all resources have the same quality (e.g., jobs do not differ in quality), and no learning is required (Fig. 3, row 4). When there is moderate variance (e.g., jobs differ slightly in quality), the prior uncertainty is higher, but in most cases an agent knows the sign of the resource before sampling. Specifically, when the mean resource quality is at least moderately positive (or negative), an agent can assume that the resource is likely to be positive (or negative), even though it does not know the exact quality. For example, if jobs are almost always positive (negative), an agent can accept (reject) any encountered job without sampling. When the mean resource quality is close to 0 a resource is equally likely to be positive or negative. In these environments an agent needs to sample a few cues before being confident enough to accept (Fig. 3, row 2) or reject (Fig. 3, row 3). Finally, when the variance in the resource quality is high, both positive and negative resources are possible. Even when the mean resource quality is very low, an agent will sometimes encounter a positive resource. Similarly, when the mean resource quality is very high, an agent will sometimes encounter a negative resource. For example, even in very prosperous (poor) areas some jobs may be unsatisfactory (worthwhile). In high-variance environments the prior uncertainty is high, and an agent needs a large amount of evidence before it is convinced that a resource is positive or negative. Therefore, an agent needs many more positive than negative cues before accepting, and many more negative than positive cues before rejecting. As little to no sampling occurs when variance in the resource quality is low to medium, the rest of this subsection focuses on environments that have a high variance in resource quality.

3.1.2. Sampling is maximal when the resource quality is neither high nor low

Increasing the resource quality has a quadratic (inverted-U) effect on sampling: agents sample few cues when the mean resource quality is extreme (both negative and positive), but sample many when the mean is close to 0. When the mean resource quality is very positive, almost all resources in an environment will be positive. Accordingly, an agent needs little evidence to conclude that a resource is positive and should be accepted. However, an agent needs a lot of evidence before it is convinced that the resource has a rare negative quality and should thus

be rejected. Moreover, because most resources are positive, sampling is more likely to result in a positive cue than in a negative cue. Therefore, an agent typically samples a small number of positive cues, and then quickly accepts. This pattern reverses when the mean resource quality is very negative. In these environments an agent needs only a few more negative than positive cues to reject, but many more positive than negative cues to accept. Because negative cues are more common than positive cues in this environment, an agent is likely to sample a few more negative than positive cues early on. For example, if jobs are on average very rewarding, an agent is quick to accept and rejects only when multiple red flags are raised. Conversely, if most jobs are dangerous, an agent rejects a potential job at the first sight of trouble, but needs strong evidence before accepting. When the mean resource quality is closer to 0, an agent's prior is dispersed: positive and negative resources are about equally likely. Accordingly, an agent needs a medium amount of evidence for accepting or rejecting. In this scenario, an agent needs many more negative than positive cues before rejecting, and many more positive than negative cues before accepting. As positive and negative resources, and hence positive and negative cues, are equally likely, sampling likely provides ambiguous evidence. Therefore, an agent is expected to sample many cues before making a decision.

3.1.3. Interruptions decrease sampling

Increasing the interruption rate increases the cost of sampling, which reduces the number of cues sampled. To see why, consider an extremely unpredictable environment in which half of all encounters are interrupted. For example, it may be that two equally suited candidates apply for the same vacancy. If both accept, one will get the job, whereas the other will experience an interruption. In this situation, accepting a positive (or negative) resource results in a positive (or negative) outcome 50% of the times, and in an interruption (no change in somatic state) otherwise. Because an agent does not know whether an interruption is about to occur, the *expected* outcome of accepting is reduced by 50%. As the expected outcome of a resource encounter depends on the expected outcome of accepting, increasing the frequency of interruptions decreases the profitability of resource encounters. As discussed in Fig. 2, sampling is beneficial because it makes an agent's estimate of the resource value more accurate. However, as interruptions become more frequent and resources become less valuable, the added value of this increased accuracy decreases. As the cost of sampling one additional cue remains the same, the benefit-to-cost ratio of sampling decreases. Hence, an agent needs fewer cues before reaching a point where the benefits of sampling are less than the cost of sampling, reducing the number of cues sampled. As a result, increasing the interruption rate has the same effect as increasing the cost of sampling. Due to this decreased benefit-to-cost ratio, an agent is more tolerant of uncertainty.

3.2. Prospects of starvation and satiation increase sampling

Fig. 4 shows how an agent's somatic state at the start of a cycle shapes the optimal policy (for details, see section 3 of appendices E, F, and G). Consider an agent that starts a cycle with a very low somatic state ('starvation sampling'). If this agent accepts a negative resource, it dies of starvation. To avoid this mortal mistake, this agent is eager to reject, and needs little evidence to conclude that a resource is negative and should be rejected. The agent is also hesitant to accept and demands much evidence to be convinced that a resource is positive. This combination of eager-to-reject and hesitant-to-accept by itself does not result in high sampling. For that to happen, the mean resource quality needs to be positive. If the mean resource quality is negative, sampling often results in a negative cue and in a swift rejection. If the mean resource quality is positive, sampling is more likely to result in a positive cue; hence, the agent is more likely to sample many positive cues. Similarly, an agent is also eager to reject and hesitant to accept when it starts a cycle with a very good somatic state ('satiation sampling'). In

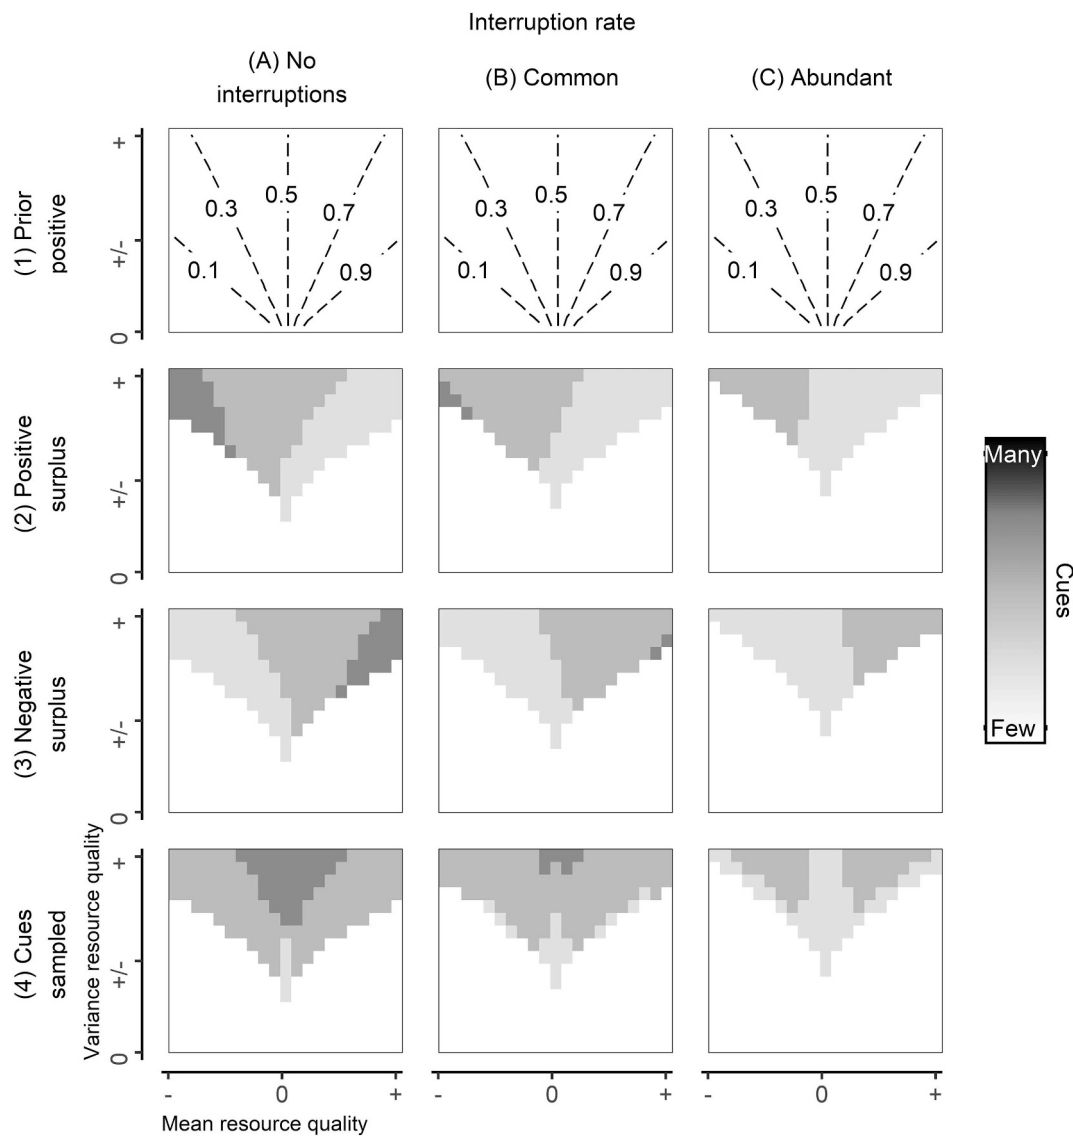


Fig. 3. How the mean resource quality and interruption rates influence the three dimensions that influence an agent's decision and the number of sampled cues. The horizontal axis shows the mean resource quality, ranging from negative (“-”) to positive (“+”). The vertical axis shows the variance in resource quality, ranging from no (“0”) to high (“+”) variance. The columns show different interruption rates (no interruptions, common, or abundant). Row 1, “prior positive”, shows the prior probability that the resource is positive. Row 2, “positive surplus”, shows how many more positive than negative cues an agent needs, on average, before accepting. Row 3, “negative surplus”, shows how many more negative than positive cues an agent needs, on average, before rejecting. Finally, row 4, “cues sampled”, shows the expected number of cues sampled when following the optimal policy. Here extrinsic events are excluded (they always have a value of 0) and the somatic state at the start of the cycle is 50.

this situation accepting a moderately and an extremely positive resources are equally rewarding as both result in the agent reaching the highest possible somatic state, which reduces the potential benefit of accepting. As with starvation sampling, this results in more sampling only when the mean resource is positive and positive cues are more likely than negative cues.

Both starvation and satiation sampling are threshold effects; sharp changes in behavior caused by a steep change in how rewarding outcomes are. Threshold effects are common in the risk sensitive foraging literature, with risk aversion typically increasing when the somatic state approaches starvation levels (Lim, Wittek, & Parkinson, 2015). Likewise, formal models often find increased sampling when an agent's somatic state approaches an upper limit (Mathot & Dall, 2013). In our model, threshold effects exist in part because the probability of death is a deterministic function of the somatic state: it occurs always, but only, when the somatic state is 0; there is no probability of dying when the somatic state is higher than 0. An alternative approach would be to

model the probability of dying as a stochastic function of the somatic state, with the probability of death increasing as the somatic state decreases. Future models can test if our results change qualitatively when death is stochastic, rather than deterministic. However, in the present article we are not interested in how threshold effects shape the optimal policy per se; rather, we are interested in how environmental conditions shape this policy.

3.3. Extrinsic events increase satiation and starvation sampling, but have little influence otherwise

The mean and variance in extrinsic events determine at what somatic state an agent might starve or become satiated (Fig. 5; see also section 1 and 3 of appendices E, F, and G). Consider an environment where extrinsic events are always very negative (e.g., almost all peers are prone to violence and theft). Here the next extrinsic event will immediately and strongly reduce an agent's somatic state. To survive

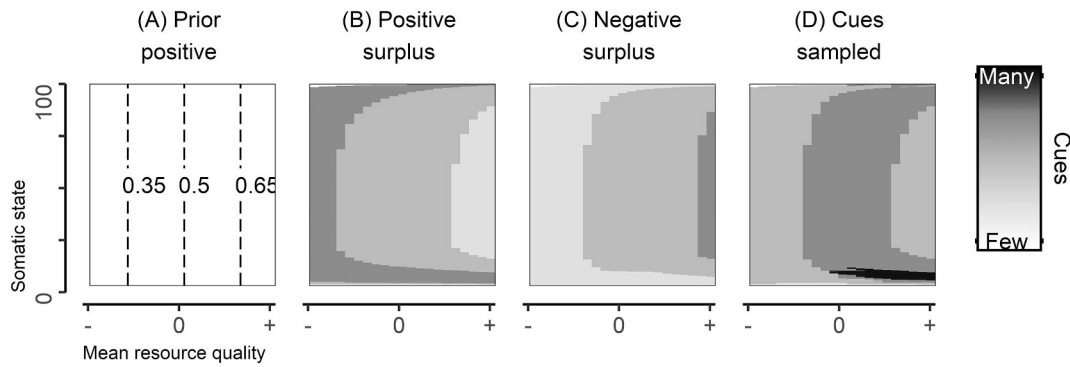


Fig. 4. How the somatic state at the start of the cycle influences the three dimensions that influence an agent's decision and the number of sampled cues. The horizontal axis shows the mean resource quality, ranging from negative (“-”) to positive (“+”). The vertical axis shows the somatic state at the start of a cycle, ranging from 0 to 100. Column A, “prior positive”, shows the probability that the resource is positive. Column B, “positive surplus”, shows how many more positive than negative cues an agent requires, on average, before accepting. Column C, “negative surplus”, shows how many more negative than positive cues an agent requires, on average, before rejecting. Finally, column D, “cues sampled”, shows the expected number of cues sampled when following the optimal policy. In this figure extrinsic events are excluded (they always have a value of 0, interruptions are absent, and the variance in resource quality is high.)

this event, an agent needs to end the current resource encounter with a somatic state high enough to incur this negative extrinsic event. As such, the threshold to avoid starvation is higher when extrinsic events are very negative, prompting an agent to go into starvation mode at higher somatic states. Similarly, if extrinsic events are always very positive (e.g., almost all peers are quick to offer a helping hand), an agent reaches the satiation threshold if it finishes the current resource

encounter with a somatic state close (but not quite at) the upper limit.

Strikingly, extrinsic events have little influence on the optimal policy when an agent has an intermediate somatic state. When an agent's somatic state is neither high nor low, even extreme extrinsic events will not result in starvation or satiation. An extrinsic event can either increase or decrease the somatic state. As long as this change does not put the somatic state under the lower threshold or above the

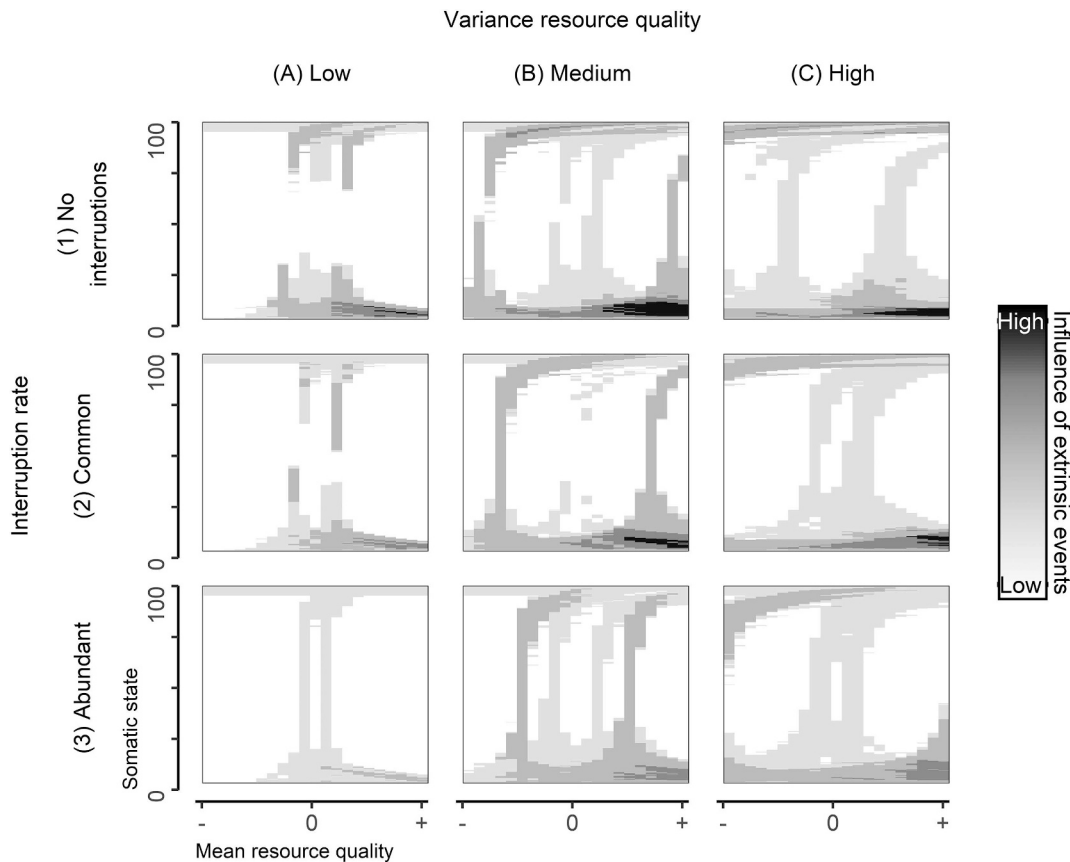


Fig. 5. The influence of extrinsic events. The mean of extrinsic events differed between environments, and was either negative, zero, or positive. Similarly, the variance in extrinsic events was either low, medium, or high. Therefore, for each environment there are eight comparable environments that only differ in the mean and variance of extrinsic events, but have the same mean and variance in resource quality and interruption rates. We compare these nine environments to study the influence of extrinsic events. Specifically, for each somatic state at the start of the cycle (vertical axis), we compute standard deviation in the number of cues sampled in these nine environments. A high standard deviation indicates that extrinsic events have a strong influence on the optimal policy (dark colors); a low standard deviation indicates that extrinsic events have little influence (light colors). We show this influence for different mean resource quality (horizontal axis), different variance in resource quality (columns), and different interruption rates (rows).

upper threshold, it has the same effect as ‘decision noise’. In this situation an extrinsic event affects the expected outcomes of accepting, rejecting, and sampling equally; if it increases the expected outcome of one action, it also increases the expected outcome of all the other actions and by the same amount. Consequentially, as these events do not influence the difference in expected outcome for each action, they also do not influence the optimal policy. The optimal policy therefore does not depend on the mean or variance of extrinsic events.

4. Discussion

Evolutionary social scientists have argued that impulsive behavior is adaptive in harsh and unpredictable environments. We have developed a formal model that explores how commonly used definitions of harshness and unpredictability affect the optimal level of information impulsivity. Our results show that this hypothesis is not universally true, but rather, depends on the exact definition of harshness, unpredictability, and impulsivity; harsh and unpredictable environments can favor high or low levels of impulsivity, or have no effect on impulsive behavior.

Our model suggests five conclusions about how harshness and unpredictability shape the optimal level of impulsive behavior. Two of these are also supported by existing models: individuals should sample more cues when the prior uncertainty of resources is higher (i.e., when the variance in resource quality is high); and individuals that are close to a somatic threshold (starvation or satiation) should sample more information, regardless of the state of their environment. Three other findings may be novel. First, impulsive behavior is adaptive when the resource quality is either low or high, but not when it is moderate. Second, impulsive behavior is almost always adaptive when resources are likely to be interrupted. Models of temporal impulsivity often find that temporal impulsivity increases as interruptions become more common. However, to our knowledge, this is the first model that finds similar effects on information impulsivity. Third, the mean and variance of extrinsic events only affect impulsivity when agents are in a very bad or a very good state. This is surprising because harshness is commonly defined (although not typically measured, see [section 1.2.1](#)) as a high rate in which external factors cause disability and death.

The conclusion that harshness and unpredictability can have multiple influences on impulsivity highlights the need for clear and explicit definitions. Although different interpretations of harshness and unpredictability are typically empirically related (e.g., resource scarcity can increase violence and disease), they are conceptually different. An environment can simultaneously be harsh and unpredictable in some sense, but affluent and predictable in others. Empirical support for the adaptive impulsivity hypothesis is mixed (see [section 1.1](#)). This might be partly due to the jingle fallacy, the erroneous belief that two constructs are the same because they have the same name. However, if empirical results depend on what notion of harshness, unpredictability, or impulsivity is measured, findings from one study might not generalize to other studies or to other populations. This makes it difficult for studies to incrementally build upon each other, stifling academic progress. We therefore strongly recommend that future studies use explicit, ideally formal, definitions of harshness and unpredictability. Such explicit definitions can help improve empirical measurements of harshness and unpredictability. For instance, future measurements of harshness could explicitly differentiate between resource scarcity and high levels of extrinsic morbidity-mortality.

4.1. Formalizing life history theory in the social sciences

Our model also contributes to a larger conversation about how to use life history theory in evolutionary social sciences. A recent bibliometric analysis shows that in the previous decade the life history literature has fragmented into different clusters with dividing lines between the evolutionary psychology, evolutionary anthropology, and non-human animal literatures ([Nettle & Frankenhuis, 2019](#)). Alarming, studies within the evolutionary social science cluster have few ties with formal models

of life history theory. These weak connections are problematic, because references are sometimes used in support of claims that are different, absent, or even contradictory to the source model. One example is the proposed fast-slow continuum. Although its existence is often described as a fundamental prediction of life history theory ([Ellis et al., 2009](#)), formal support for the fast-slow continuum is limited and mixed ([Mathot & Frankenhuis, 2018](#); [Zietsch & Sidari, 2019](#)). Some models show that harsh and unpredictable conditions can favor slow life histories (e.g., [Abrams, 1993](#); [Baldini, 2015](#)). Similarly, our model shows that one kind of impulsivity, which is often viewed as part of a fast life history, is not necessarily favored in harsh and unpredictable environments.

4.2. Empirical predictions, limitations, and future directions

All models are simplifications of reality ([Smaldino, 2017](#)). However, they differ in whether they are general or specific ([Houston & McNamara, 2005](#); [Parker & Smith, 1990](#)). The goal of a general model is to study abstract qualitative patterns. For instance, a prisoner's dilemma model captures the logic of cooperation and defection between two rational players – it does not matter whether the players are people, companies, or rivaling states. The parameters of general models are often difficult to operationalize, predict, and measure. Specific models study the dynamics of a particular real-world system. The parameters of these models are frequently based on empirical data, and these models might provide predictions. We have presented a general model; our goal was to provide a formalization of the adaptive impulsivity hypothesis. As such, we made simplifying assumptions. These assumptions allowed us to explore a decision problem in depth, facilitating theoretic insight about the ways in which key variables interact with each other. Simple models are well suited to producing such insights, but at a cost to realism ([Levins, 1968](#)). We think this is acceptable, because our primary goal is not to make empirical predictions. However, this does not necessarily mean that the conclusions of our model cannot be used as empirical prediction. Rather, this depends on the extent to which the assumptions of our general model capture essential features of real environments. If this match is sufficiently high, the conclusions of our model on how harshness and unpredictability shape impulsive behavior can be used as empirical predictions. Estimating this match is difficult, if not impossible. There are, however, several limitations that reduce realism and limit the scope of our model. These limitations are hierarchical: we can only address some limitations (e.g., our model does not include life history trajectories) after we have addressed more fundamental limitations (e.g., our model does not address development or environmental change). Here we discuss four fundamental limitations. For each limitation we discuss how potential extensions can incorporate more realistic and more complicated assumptions that address these limitations.

First, in order to reduce complexity we assumed that the parameters of the environment are fixed within and between generations. We further assumed that an agent learned the (meta) parameters of its environment through its evolutionary and developmental history. Although extreme outcomes may be unexpected, they do not change an agent's beliefs about its environment. For some organisms a fixed world assumption may be realistic: if the rate of environmental change is slow compared to the lifespan of an organism, the environment might appear to be fixed from that organism's perspective ([Fawcett & Frankenhuis, 2015](#)). However, for species with a longer life span, such as humans, the environment might change both temporally (e.g., due to economic cycles) and spatially (e.g., due to labor or educational migration). In a fixed environment an organism ‘only’ has to infer the value of an encountered resource. In a varying environment, it also has to infer the current state of the environment and forecast what the future might hold. This results in a tradeoff between exploration (sampling information) and exploitation (saving costs by relying on current estimates of the environment). Moreover, in a varying environment, there might be lean years where resources are scarce and/or extrinsic events more extreme. An organism can buffer against such variability by

storing resources. This might increase (to save costs on information gathering) or decrease (to reduce the variance in outcomes) impulsivity. Besides reducing realism, this assumption also reduced the scope of our model: unpredictability is often interpreted as changes in the environmental state (e.g., this kind of unpredictability is the focus of Ellis et al., 2009). Future models could incorporate both temporal and spatial unpredictability.

Second, our model includes no development. We studied organisms that (a) are fully developed at birth, (b) are affected by the environment regardless of their age, and (c) reproduce only at the end of life. These assumptions do not hold for many species, including humans. Rather, individuals typically go through early developmental stages in which they acquire the skills needed to integrate information. If the individual faced early-life adversity, or if this acquisition is costly or time consuming, investing in this skillset might not outweigh the cost. Moreover, both the young and the old might be more affected by resource scarcity and extrinsic events than adults. In hunter-gatherer societies, only adults produce more food than they consume (Kaplan, Hill, Lancaster, & Hurtado, 2000). Consequentially, in lean years the old and young might be more susceptible for starvation than adults. Similarly, negative extrinsic events such as disease and violence might disproportionately affect the young (who are less able to defend themselves) and the old (who might be weakened due to senescence). Finally, we studied an organism that is semelparous, rather than iteroparous. However, in many species fecundity and fertility often peak during middle age. As both reproduction and the subsequent investment in offspring are costly, individuals in middle age might face a higher demand for resources. Future models can build in age structure and reproduction, with survival and fecundity differing at different ages, and explore how such selection regimes shape the optimal level of impulsivity.

Our model can also be extended to include developmental processes in order to explore two empirical patterns. First, impulsivity and risk taking are highest during adolescence, when individuals enter the mating competition market (Figner, Mackinlay, Wilkening, & Weber, 2009; Steinberg, 2007). For risk behavior, a common explanation is that securing a high quality mate requires intense competition for resources and social status (Ellis et al., 2012), which demands high levels of risk taking (e.g., engaging in physical fights). Future models could examine whether this increased need for resources and social status likewise results in more impulsive behavior. Another empirical pattern is the paradoxical (but robust) finding that both behavioral tasks and self-report questionnaires predict real-world impulsivity, yet the two sets of measurement show little to no correlation (Cyders & Coskunpinar, 2011; Reynolds, Ortengren, Richards, & de Wit, 2006; Stahl et al., 2014). A popular explanation is that both sets of measurements tap into separate constructs. Self-reports measure a stable baseline of impulsivity (i.e., trait impulsivity), whereas behavioral tasks measure the capability to flexibly deviate from this baseline in situations that require higher or lower levels. This explanation raises such interesting questions as: Why there is a baseline? Why do we not always adjust our impulsivity to match the current situation? Why do individuals differ in their baseline levels? Is this baseline continuously updated throughout development, or are there sensitive periods in which the baseline is set for the rest of life? Part of the answer to these questions might be that flexibility comes at a cost. For instance, the cognitive machinery needed to make constant adjustments might be expensive to maintain. If we always need the same level of impulsivity – for instance, when our environment is sufficiently stable – the cost of plasticity might outweigh its benefits (Fawcett & Frankenhuis, 2015). Moreover, if the environment is very stable, the best strategy might be to set a fixed baseline early in life (i.e., a sensitive period). Future models can explore these questions by incorporating developmental processes.

Third, we committed to the ‘behavioral gambit’: we studied a single behavioral trait in isolation, and implicitly assumed that the expression of this trait is not hindered by other life history, behavioral, or physiological traits (Fawcett, Hamblin, & Giraldeau, 2013). Furthermore, our model did not address genetic, developmental, physiological, or cognitive limitations that prevent an organism from following the optimal policy. In real life

there are limitations. For example, we assumed that organisms behave as if they perform Bayesian updating. However, Bayesian updating is computationally expensive at the best of times, and computationally intractable in most realistic situations (Trimmer, McNamara, Houston, & Marshall, 2012; van Rooij, Wright, Kwisthout, & Wareham, 2018).

The behavioral gambit is a useful simplification when testing under what environmental conditions impulsivity might be adaptive. However, it limits the scope and realism of our results. Future extensions might explore two different avenues. First, they can incorporate more realistic cognitive processes. For example, future models can study agents that rely on heuristics that human decision-makers are known to use. This extension can study in which environment a specific heuristic performs well, and when it performs poorly. Alternatively, rather than simulating agents that use known heuristics, future models can use the computed optimal policies to explore new heuristics. Specifically, based on modeling results, future research can explore which heuristics would allow animals to approximate optimal decisions. Second, they can increase realism by incorporating other behavioral traits. Such a model can provide novel insights for two different debates. Different notions of impulsivity are only weakly correlated or even uncorrelated (section 1.2.3). A model incorporating multiple types of impulsivity can explore whether environmental conditions moderate the correlation between different conceptualizations. That is, it can explore whether some environments favor high (or low) levels of all types, whereas others favor high levels of one type but low levels of the other. Alternatively, extensions can incorporate other behavioral, physiological, or life history traits that are proposed to cluster on a fast-slow continuum. This extension can test the claim that harsh and unpredictable environments result in faster life-history strategies.

Fourth, we assumed that agents did not interact, nor needed to consider the behavior of other agents (i.e., our model is not game theoretic). This assumption is reasonable for some decisions. For instance, if resources are (practically) infinite, the actions of one agent do not noticeably change the number of available resources (e.g., when job supply is high, accepting a job does not meaningfully decrease the total number of available jobs). In other decisions agents do interact, but only indirectly. In this case, accepting may reduce the resources available for other agents. However, the behavior of other agents does not influence the consequences of an action during a resource encounter. For instance, two predators may share overlapping domains. Although they rarely are in close proximity, resources consumed by one are no longer available for the other. However, whether or not one predator should give chase to prey does not depend on the actions of the other predator. Our model can incorporate some indirect interactions by changing the parameters of an environment. For instance, our graphical interface allows users to increase or decrease the interruption rate (e.g., prey might be more or less easily scared) or to assume that resources are non-normally distributed (e.g., competitors might be more likely to consume positive than negative resources). However, in many real-world decisions an agent does need to consider the actions of other agents. For instance, resources might become scarce if everyone acts impulsively. If so, acting impulsively may be the only way to collect resources. Such policy, where one is impulsive because everybody else is, results in a positive feedback loop that might increase impulsivity. Alternatively, high levels competition may foster selective cooperation, which requires low levels of impulsivity. It can also result in even more complex patterns, where multiple phenotypes coexist, or the population might cycle between multiple phenotypes (Bear & Rand, 2016; Tomlin, Rand, Ludwig, & Cohen, 2015). It will be hard if not impossible to predict outcomes without building the model. Future models might therefore incorporate interactions between agents.

To end, we have presented a formal model of the increasingly common claim impulsive behavior is adaptive in harsh and unpredictable environments. Our results show that this hypothesis is not universally true, but rather, depends on the exact definition of harshness, unpredictability, and impulsivity. We hope our model will contribute to the corpus of formal models of theories that feature centrally in the evolutionary social sciences.

Acknowledgments

We thank Nicole Walasek, Karthik Panchanathan, Ethan Young, Achiel Fenneman, Esther Weijman, and two anonymous reviewers for providing thoughtful comments and insights on earlier versions of this manuscript.

References

- Abrams, P. A. (1993). Does increased mortality favor the evolution of more rapid senescence? *Evolution*, 47(3), 877–887. <https://doi.org/10.2307/2410191>.
- Allen, C., & Nettle, D. (2019). Hunger and socioeconomic background additively predict impulsivity in humans. *Current Psychology*, 1–15. <https://doi.org/10.1007/s12144-019-0141-7>.
- Aronoff, J. E., & DeCaro, J. A. (2019). Life history theory and human behavior: Testing associations between environmental harshness, life history strategies and testosterone. *Personality and Individual Differences*, 139, 110–115. <https://doi.org/10.1016/j.paid.2018.11.015>.
- Baldini, R. (2015). Harsh environments and “fast” human life histories: What does the theory say? *BioRxiv*, 1–10. <https://doi.org/10.1101/014647>.
- Baumard, N., & Chevallier, C. (2015). The nature and dynamics of world religions: A life-history approach. *Proceedings of the Royal Society B*, 282(1818). <https://doi.org/10.1098/rspb.2015.1593>.
- Baumeister, R. F. (2002). Yielding to temptation: Self-control failure, impulsive purchasing, and consumer behavior. *Journal of Consumer Research*, 28(4), 670–676. <https://doi.org/10.1086/338209>.
- Bear, A., & Rand, D. G. (2016). Intuition, deliberation, and the evolution of cooperation. *Proceedings of the National Academy of Sciences*, 113(4), 936–941. <https://doi.org/10.1073/pnas.1517780113>.
- Belsky, J., Schlomer, G. L., & Ellis, B. J. (2012). Beyond cumulative risk: Distinguishing harshness and unpredictability as determinants of parenting and early life history strategy. *Developmental Psychology*, 48(3), 662–673. <https://doi.org/10.1037/a0024454>.
- Brumbach, B. H., Figueredo, A. J., & Ellis, B. J. (2009). Effects of harsh and unpredictable environments in adolescence on development of life history strategies. *Human Nature*, 20(1), 25–51. <https://doi.org/10.1007/s12110-009-9059-3>.
- Bulley, A., & Pepper, G. V. (2017). Cross-country relationships between life expectancy, intertemporal choice and age at first birth. *Evolution and Human Behavior*, 38(5), 652–658. <https://doi.org/10.1016/j.evolhumbehav.2017.05.002>.
- Caswell, A. J., Bond, R., Duka, T., & Morgan, M. J. (2015). Further evidence of the heterogeneous nature of impulsivity. *Personality and Individual Differences*, 76, 68–74. <https://doi.org/10.1016/j.paid.2014.11.059>.
- Caswell, A. J., Morgan, M. J., & Duka, T. (2013). Inhibitory control contributes to “motor”-but not “cognitive”-impulsivity. *Experimental Psychology*, 60(5), 324–334. <https://doi.org/10.1027/1618-3169/a000202>.
- Chen, P., & Vazsonyi, A. T. (2011). Future orientation, impulsivity, and problem behaviors: A longitudinal moderation model. *Developmental Psychology*, 47(6), 1633–1645. <https://doi.org/10.1037/a0025327>.
- Collins, E. J., McNamara, J. M., & Ramsey, D. M. (2006). Learning rules for optimal selection in a varying environment: Mate choice revisited. *Behavioral Ecology*, 17(5), 799–809. <https://doi.org/10.1093/beheco/arl008>.
- Copping, L. T., Campbell, A., & Muncer, S. (2013). Impulsivity, sensation seeking and reproductive behaviour: A life history perspective. *Personality and Individual Differences*, 54(8), 908–912. <https://doi.org/10.1016/j.paid.2013.01.003>.
- Copping, L. T., Campbell, A., & Muncer, S. (2014). Conceptualizing time preference: A life-history analysis. *Evolutionary Psychology*, 12(4), 829–847. <https://doi.org/10.1177/147470491401200411>.
- Cyders, M. A. (2015). The misnomer of impulsivity: Commentary on “choice impulsivity” and “rapid-response impulsivity” articles by Hamilton and colleagues. *Personality Disorders: Theory, Research, and Treatment*, 6(2), 204–205. <https://doi.org/10.1037/per0000123>.
- Cyders, M. A., & Coskunpinar, A. (2011). Measurement of constructs using self-report and behavioral lab tasks: Is there overlap in nomothetic span and construct representation for impulsivity? *Clinical Psychology Review*, 31(6), 965–982. <https://doi.org/10.1016/j.cpr.2011.06.001>.
- Dall, S. R. X., Giraldeau, L. A., Olsson, O., McNamara, J. M., & Stephens, D. W. (2005). Information and its use by animals in evolutionary ecology. *Trends in Ecology and Evolution*, 20(4), 187–193. <https://doi.org/10.1016/j.tree.2005.01.010>.
- Daly, M., & Wilson, M. (2005). Carpe diem: Adaptation and devaluing the future. *The Quarterly Review of Biology*, 80(1), 55–60. <https://doi.org/10.1086/431025>.
- Daruna, J. H., & Barnes, P. A. (1993). A neurodevelopmental view of impulsivity. In W. G. McCown, J. L. Johnson, & M. B. Shure (Eds.). *The impulsive client: Theory, research, and treatment* (pp. 23–37). Washington, D.C.: American Psychological Association. <https://doi.org/10.1037/10500-002>.
- Del Giudice, M. (2014). An evolutionary life history framework for psychopathology. *Psychological Inquiry*, 25(3–4), 261–300. <https://doi.org/10.1080/1047840X.2014.884918>.
- Del Giudice, M. (2015). Self-regulation in an evolutionary perspective. *Handbook of Biobehavioral approaches to self-regulation* (pp. 25–41). New York, NY: Springer New York. https://doi.org/10.1007/978-1-4939-1236-0_3.
- Del Giudice, M. (2018). *Evolutionary psychopathology: A unified approach*. New York, NY: Oxford University Press. <https://doi.org/10.1093/med-psych/9780190246846.001.0001>.
- Del Giudice, M., & Belsky, J. (2010). *The development of life history*. Strategies: Toward a Multi-Stage Theory. The Evolution of Personality and Individual Differences. <https://doi.org/10.1093/acprof:oso/9780195372090.003.0006>.
- Del Giudice, M., Gangestad, S. W., & Kaplan, H. S. (2015). Life history theory and evolutionary psychology. *The Handbook of Evolutionary Psychology* (pp. 88–114). <https://doi.org/10.1002/9781119125563.evpsych102>.
- Duckworth, A. L. (2011). The significance of self-control. *Proceedings of the National Academy of Sciences*, 108(7), 2639–2640. <https://doi.org/10.1073/pnas.1019725108>.
- Duckworth, A. L., & Seligman, M. (2005). Self-discipline outdoes IQ in predicting academic performance of adolescents. *Psychological Science*, 16(12), 939–944. <https://doi.org/10.1111/j.1467-9280.2005.01641.x>.
- Duckworth, A. L., Tsukayama, E., & Kirby, T. A. (2013). Is it really self-control? Examining the predictive power of the delay of gratification task. *Personality and Social Psychology Bulletin*, 39(7), 843–855. <https://doi.org/10.1177/0146167213482589>.
- Dunlap, A. S., & Stephens, D. W. (2012). Tracking a changing environment: Optimal sampling, adaptive memory and overnight effects. In: *Behavioural Processes*, 89(2), 86–94. <https://doi.org/10.1016/j.beproc.2011.10.005>.
- Ellis, B. J., Del Giudice, M., Dishion, T. J., Figueredo, A. J., Gray, P., Griskevicius, V., ... Wilson, D. S. (2012). The evolutionary basis of risky adolescent behavior: Implications for science, policy, and practice. *Developmental Psychology*, 48(3), 598–623. <https://doi.org/10.1037/a0026220>.
- Ellis, B. J., Figueredo, A. J., Brumbach, B. H., & Schlomer, G. L. (2009). Fundamental dimensions of environmental risk: The impact of harsh versus unpredictable environments on the evolution and development of life history strategies. *Human Nature*, 20(2), 204–268. <https://doi.org/10.1007/s12110-009-9063-7>.
- Epstein, J. M. (2008). Why Model? *Journal of Artificial Societies and Social Simulation*, 11(4), 12. <http://jasss.soc.surrey.ac.uk/11/4/12.html>.
- Evenden, J. L. (1999). Varieties of impulsivity. *Psychopharmacology*, 146, 348–361. <https://doi.org/10.1007/PL00005481>.
- Eysenck, S. B. G., Easting, G., & Pearson, P. R. (1984). Age norms for impulsiveness, Venturesomeness and empathy in children. *Personality and Individual Differences*, 5(3), 315–321.
- Eysenck, S. B. G., Pearson, P. R., Easting, G., & Allsopp, J. F. (1985). Age norms for impulsiveness, venturesomeness and empathy in adults. *Personality and Individual Differences*, 6(3), [https://doi.org/10.1016/0191-8869\(85\)90011-x](https://doi.org/10.1016/0191-8869(85)90011-x).
- Fawcett, T. W., & Frankenhuus, W. E. (2015). Adaptive explanations for sensitive windows in development. *Frontiers in Zoology*, 12(1), S3. <https://doi.org/10.1186/1742-9994-12-S1-S3>.
- Fawcett, T. W., Hamblin, S., & Giraldeau, L. A. (2013). Exposing the behavioral gambit: The evolution of learning and decision rules. *Behavioral Ecology*, 24(1), 2–11. <https://doi.org/10.1093/beheco/ars085>.
- Fawcett, T. W., & Johnstone, R. A. (2003). Optimal assessment of multiple cues. *Proceedings of the Royal Society B*, 270(1524), 1637–1643. <https://doi.org/10.1098/rspb.2003.2328>.
- Figner, B., Mackinlay, R. J., Wilkening, F., & Weber, E. U. (2009). Affective and deliberative processes in risky choice: Age differences in risk taking in the Columbia card task. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 35(3), 709–730. <https://doi.org/10.1037/a0014983.supp>.
- Figueredo, A. J., Vásquez, G., Brumbach, B. H., Schneider, S. M. R., Sefcek, J. A., Tal, I. R., ... Jacobs, W. J. (2006). Consilience and life history theory: From genes to brain to reproductive strategy. *Developmental Review*, 26(2), 243–275. <https://doi.org/10.1016/j.dr.2006.02.002>.
- Fineberg, N. A., Chamberlain, S. R., Goudriaan, A. E., Stein, D. J., Vanderschuren, L. J. M. J., Gillan, C. M., ... Potenza, M. N. (2014). New developments in human neurocognition: Clinical, genetic, and brain imaging correlates of impulsivity and compulsivity. *CNS Spectrums*, 19(1), 69–89. <https://doi.org/10.1017/S1092852913000801>.
- Frankenhuus, W. E., & Nettle, D. (2020). The strengths of people in poverty. *Current Directions in Psychological Science*, 29, 16–21. <https://doi.org/10.1177/0963721419881154>.
- Frankenhuus, W. E., Panchanathan, K., & Nettle, D. (2016). Cognition in harsh and unpredictable environments. *Current Opinion in Psychology*, 7, 76–80. <https://doi.org/10.1016/j.copsyc.2015.08.011>.
- Frederick, S., Loewenstein, G., & O’Donoghue, T. (2002). Time discounting and time preference: A critical review. *Journal of Economic Literature*, 40(2), 351–401. <https://doi.org/10.1257/jel.40.2.351>.
- Griskevicius, V., Ackerman, J. M., Cantú, S. M., Delton, A. W., Robertson, T. E., Simpson, J. A., ... Tybur, J. M. (2013). When the economy falters, do people spend or save? Responses to resource scarcity depend on childhood environments. *Psychological Science*, 24(2), 197–205. <https://doi.org/10.1177/0956797612451471>.
- Griskevicius, V., Delton, A. W., Robertson, T. E., & Tybur, J. M. (2011). Environmental contingency in life history strategies: The influence of mortality and socioeconomic status on reproductive timing. *Journal of Personality and Social Psychology*, 100(2), 241–254. <https://doi.org/10.1037/a0021082>.
- Griskevicius, V., Tybur, J. M., Delton, A. W., & Robertson, T. E. (2011). The influence of mortality and socioeconomic status on risk and delayed rewards: A life history theory approach. *Journal of Personality and Social Psychology*, 100(6), 1015–1026. <https://doi.org/10.1037/a0022403>.
- Hamilton, K. R., Littlefield, A. K., Anastasio, N. C., Cunningham, K. A., Fink, L. H. L., Wing, V. C., ... Potenza, M. N. (2015). Rapid-response impulsivity: Definitions, measurement issues, and clinical implications. *Personality Disorders: Theory, Research, and Treatment*, 6(2), 168–181. <https://doi.org/10.1037/per0000100>.
- Hamilton, K. R., Mitchell, M. R., Wing, V. C., Balodis, I. M., Bickel, W. K., Fillmore, M., ... Moeller, F. G. (2015). Choice impulsivity: Definitions, measurement issues, and clinical implications. *Personality Disorders: Theory, Research, and Treatment*, 6(2),

- 182–198. <https://doi.org/10.1037/per0000099>.
- Higginson, A. D., Fawcett, T. W., Houston, A. I., & McNamara, J. M. (2018). Trust your gut: Using physiological states as a source of information is almost as effective as optimal Bayesian learning. *Proceedings of the Royal Society B*, 285(1871), 20172411. <https://doi.org/10.1098/rspb.2017.2411>.
- Houston, A. I., & McNamara, J. M. (2005). John Maynard Smith and the importance of consistency in evolutionary game theory. *Biology and Philosophy*, 20(5), 933–950. <https://doi.org/10.1007/s10539-005-9016-4>.
- Humphreys, K. L., Lee, S. S., Telzer, E. H., Gabard-Durnam, L. J., Goff, B., Flannery, J., & Tottenham, N. (2015). Exploration-exploitation strategy is dependent on early experience. *Developmental Psychobiology*, 57(3), 313–321. <https://doi.org/10.1002/dev.21293>.
- Kagan, J. (1966). Reflection-impulsivity: The generality and dynamics of conceptual tempo. *Journal of Abnormal Psychology*, 71(1), 17–24. <https://doi.org/10.1037/h0022886>.
- Kaplan, H., Hill, K., Lancaster, J., & Hurtado, A. M. (2000). A theory of human life history evolution: Diet, intelligence, and longevity. *Evolutionary Anthropology*, 9(4), 156–185. [https://doi.org/10.1002/1520-6505\(2000\)9:4<156::AID-EVAN5>3.0.CO;2-7](https://doi.org/10.1002/1520-6505(2000)9:4<156::AID-EVAN5>3.0.CO;2-7).
- Kruger, D. J., Reischl, T., & Zimmerman, M. A. (2008). Time perspective as a mechanism for functional developmental adaptation. *Journal of Social, Evolutionary, and Cultural Psychology*, 2(1), 1–22. <https://doi.org/10.1037/h0099336>.
- Lee, A. J., DeBruine, L. M., & Jones, B. C. (2018). Individual-specific mortality is associated with how individuals evaluate future discounting decisions. *Proceedings of the Royal Society B*, 285(1880), <https://doi.org/10.1098/rspb.2018.0304>.
- Levins, R. (1968). *Evolution in changing environments*. Princeton: Princeton University Press.
- Lim, I. S., Wittek, P., & Parkinson, J. (2015). On the origin of risk sensitivity: The energy budget rule revisited. *Animal Behaviour*, 110, 69–77. <https://doi.org/10.1016/j.anbehav.2015.09.007>.
- Lu, H. J., & Chang, L. (2019). Aggression and risk-taking as adaptive implementations of fast life history strategy. *Developmental Science* 12827. <https://doi.org/10.1111/desc.12827>.
- Luttbegg, B. (1996). A comparative Bayes tactic for mate assessment and choice. *Behavioral Ecology*, 7(4), 451–460. <https://doi.org/10.1093/beheco/7.4.451>.
- MacKillop, J., Weafer, J., C. Gray, J., Oshri, A., Palmer, A., & de Wit, H. (2016). The latent structure of impulsivity: Impulsive choice, impulsive action, and impulsive personality traits. *Psychopharmacology*, 233(18), 3361–3370. doi:<https://doi.org/10.1007/s00213-016-4372-0>.
- Mathot, K. J., & Dall, S. R. X. (2013). Metabolic rates can drive individual differences in information and insurance use under the risk of starvation. *The American Naturalist*, 182(5), 611–620. <https://doi.org/10.1086/673300>.
- Mathot, K. J., & Frankenhuis, W. E. (2018). Models of pace-of-life syndromes (POLS): A systematic review. *Behavioral Ecology and Sociobiology*, 72(3), <https://doi.org/10.1007/s00265-018-2459-9>.
- Mazalov, V., Perrin, N., & Dombrovsky, Y. (1996). Adaptive search and information updating in sequential mate choice. *The American Naturalist*, 148(1), 123–137. <https://doi.org/10.1086/285914>.
- McCullough, M. E., Pedersen, E. J., Schroder, J. M., Tabak, B. A., & Carver, C. S. (2013). Harsh childhood environmental characteristics predict exploitation and retaliation in humans. *Proceedings of the Royal Society B*, 280(1750), 20122104. <https://doi.org/10.1098/rspb.2012.2104>.
- McNamara, J. M., Green, R. F., & Olsson, O. (2006). Bayes' theorem and its applications in animal behaviour. *Oikos*, 112(2), 243–251. <https://doi.org/10.1111/j.0030-1299.2006.14228.x>.
- Mell, H., Baumard, N., & André, J. (2019, December 5). Time is money. Waiting costs explain why selection favors steeper time discounting in deprived environments. Doi:10.32942/osf.io/7d56s.
- Mell, H., Safra, L., Algan, Y., Baumard, N., & Chevallier, C. (2018). Childhood environmental harshness predicts coordinated health and reproductive strategies: A cross-sectional study of a nationally representative sample from France. *Evolution and Human Behavior*, 39(1), 1–8. <https://doi.org/10.1016/j.evolhumbehav.2017.08.006>.
- Mischel, W., Shoda, Y., & Rodriguez, M. L. (1989). Delay of gratification in children. *Science*, 244, 933–938. <https://doi.org/10.1126/science.2658056>.
- Mishra, S., Lalumière, M. L., & Williams, R. J. (2010). Gambling as a form of risk-taking: Individual differences in personality, risk-accepting attitudes, and behavioral preferences for risk. *Personality and Individual Differences*, 49(6), 616–621. <https://doi.org/10.1016/j.paid.2010.05.032>.
- Mishra, S., & Novakowski, D. (2016). Personal relative deprivation and risk: An examination of individual differences in personality, attitudes, and behavioral outcomes. *Personality and Individual Differences*, 90, 22–26. <https://doi.org/10.1016/j.paid.2015.10.031>.
- Mishra, S., Templeton, A. J., & Meadows, T. J. S. (2017). Living, fast and slow: Is life history orientation associated with risk-related personality traits, risk attitudes, criminal outcomes, and gambling? *Personality and Individual Differences*, 117, 242–248. <https://doi.org/10.1016/j.paid.2017.06.009>.
- Mittal, C., & Griskevicius, V. (2014). Sense of control under uncertainty depends on people's childhood environment: A life history theory approach. *Journal of Personality and Social Psychology*, 107(4), 621–637. <https://doi.org/10.1037/a0037398>.
- Moffitt, T. E., Arseneault, L., Belsky, D., Dickson, N., Hancox, R. J., Harrington, H., ... Caspi, A. (2011). A gradient of childhood self-control predicts health, wealth, and public safety. *Proceedings of the National Academy of Sciences of the United States of America*, 108(7), 2693–2698. <https://doi.org/10.1073/pnas.1010076108>.
- Nettle, D., & Frankenhuis, W. E. (2019). The evolution of life-history theory: A bibliometric analysis of an interdisciplinary research area. *Proceedings of the Royal Society B*, 286(1899), 20190040. <https://doi.org/10.1098/rspb.2019.0040>.
- Parker, G. A., & Smith, J. M. (1990). Optimality theory in evolutionary biology. *Nature*, 348(6296), 27–33. <https://doi.org/10.1038/348027a0>.
- Pepper, G. V., Corby, D. H., Bamber, R., Smith, H., Wong, N., & Nettle, D. (2017). The influence of mortality and socioeconomic status on risk and delayed rewards: A replication with British participants. *PeerJ*, 5, e3580. <https://doi.org/10.7717/peerj.3580>.
- Pepper, G. V., & Nettle, D. (2014). Perceived extrinsic mortality risk and reported effort in looking after health. *Human Nature*, 25(3), 378–392. <https://doi.org/10.1007/s12110-014-9204-5>.
- Pepper, G. V., & Nettle, D. (2017). The behavioural constellation of deprivation: Causes and consequences. *Behavioral and Brain Sciences*, 40, 1–72. <https://doi.org/10.1017/S0140525X1600234X>.
- Réale, D., Garant, D., Humphries, M. M., Bergeron, P., Careau, V., & Montiglio, P. O. (2010). Personality and the emergence of the pace-of-life syndrome concept at the population level. *Philosophical Transactions of the Royal Society B*, 365(1560), 4051–4063. <https://doi.org/10.1098/rstb.2010.0208>.
- Reynolds, B., Ortengren, A., Richards, J. B., & de Wit, H. (2006). Dimensions of impulsive behavior: Personality and behavioral measures. *Personality and Individual Differences*, 40(2), 305–315. <https://doi.org/10.1016/j.paid.2005.03.024>.
- Richardson, G. B., Chen, C. C., Dai, C. L., Brubaker, M. D., & Nedelec, J. L. (2017). The psychometrics of the mini-K: Evidence from two college samples. *Evolutionary Psychology*, 15(1), <https://doi.org/10.1177/1474704916682034>.
- van Rooij, I., Wright, C. D., Kwisthout, J., & Wareham, T. (2018). Rational analysis, intractability, and the prospects of “as if”-explanations. *Synthese*, 195(2), 491–510. <https://doi.org/10.1007/s11229-014-0532-0>.
- Royauté, R., Berdal, M. A., Garrison, C. R., & Dochtermann, N. A. (2018). Painless life? A meta-analysis of the pace-of-life syndrome hypothesis. *752 Behavioral Ecology and Sociobiology*, 72(3), 64. <https://doi.org/10.1007/s00265-018-2472-z>.
- Sharma, L., Markon, K. E., & Clark, L. A. (2014). Toward a theory of distinct types of “impulsive” behaviors: A meta-analysis of self-report and behavioral measures. *Psychological Bulletin*, 140(2), 374–408. <https://doi.org/10.1037/a0034418>.
- Sih, A., & Del Giudice, M. (2012). Linking behavioural syndromes and cognition: A behavioural ecology perspective. *Philosophical Transactions of the Royal Society B*, 367(1603), 2762–2772. <https://doi.org/10.1098/rstb.2012.0216>.
- Simpson, J. A., Griskevicius, V., Kuo, S. I. C., Sung, S., & Collins, W. A. (2012). Evolution, stress, and sensitive periods: The influence of unpredictability in early versus late childhood on sex and risky behavior. *Developmental Psychology*, 48(3), 674–686. <https://doi.org/10.1037/a0027293>.
- Smaldino, P. E. (2017). Models are stupid, and we need more of them. *Computational Social Psychology*, 311–331. <https://doi.org/10.4324/9781315173726-14>.
- Stahl, C., Voss, A., Schmitz, F., Nuszbaum, M., Tüscher, O., Lieb, K., & Klauer, K. C. (2014). Behavioral components of impulsivity. *Journal of Experimental Psychology: General*, 143(2), 850–886. <https://doi.org/10.1037/a0033981>.
- Stearns, S. C. (2000). Life history evolution: Successes, limitations, and prospects. *Naturwissenschaften*, 87(11), 476–486. <https://doi.org/10.1007/s001140050763>.
- Steinberg, L. (2007). Risk taking in adolescence: New perspectives from brain and behavioral science. <https://doi.org/10.1111/j.1467-8721.2007.00475.x>.
- Stephens, D. W. (2002). Discrimination, discounting and impulsivity: A role for an informational constraint. *Philosophical Transactions of the Royal Society B*, 357(1427), 1527–1537. <https://doi.org/10.1098/rstb.2002.1062>.
- Sutton, R., & Barto, A. (2018). *Reinforcement learning: An introduction*. MIT press (2nd ed.). Cambridge, Massachusetts: The MIT press <https://doi.org/10.1109/tnn.1998.712192>.
- Tomlin, D., Rand, D. G., Ludvig, E. A., & Cohen, J. D. (2015). The evolution and devolution of cognitive control: The costs of deliberation in a competitive world. *Scientific Reports*, 5(1), 11002. <https://doi.org/10.1038/srep11002>.
- Trimmer, P. C., Houston, A. I., Marshall, J. A. R., Mendl, M. T., Paul, E. S., & McNamara, J. M. (2011). Decision-making under uncertainty: Biases and Bayesians. *Animal Cognition*, 14(4), 465–476. <https://doi.org/10.1007/s10071-011-0387-4>.
- Trimmer, P. C., McNamara, J. M., Houston, A. I., & Marshall, J. A. R. (2012). Does natural selection favour the Rescorla–Wagner rule? *Journal of Theoretical Biology*, 302, 39–52. <https://doi.org/10.1016/j.jtbi.2012.02.014>.
- Verdejo-García, A., Lawrence, A. J., & Clark, L. (2008). Impulsivity as a vulnerability marker for substance-use disorders: Review of findings from high-risk research, problem gamblers and genetic association studies. *Neuroscience and Biobehavioral Reviews*, 32(4), 777–810. <https://doi.org/10.1016/j.neubiorev.2007.11.003>.
- Wilson, M., & Daly, M. (1997). Life expectancy, economic inequality, homicide, and reproductive timing in Chicago neighbourhoods. *British Medical Journal*, 314(7089), 1271–1274. <https://doi.org/10.1136/bmj.314.7089.1271>.
- Wilson, M., & Daly, M. (2006). Are juvenile offenders extreme future discounters? *Psychological Science*, 17(11), 989–994. <https://doi.org/10.1111/j.1467-9280.2006.01817.x>.
- Zietsch, B. P., & Sidari, M. J. (2019). A critique of life history approaches to human trait covariation. *Evolution and Human Behavior*, 1–35. <https://doi.org/10.1016/j.evolhumbehav.2019.05.007>.