Animal Behaviour 224 (2025) 123191

Contents lists available at ScienceDirect

Animal Behaviour

journal homepage: www.elsevier.com/locate/anbehav

Bayesian updating for self-assessment explains social dominance and winner–loser effects

Ammon Perkes *, Kate L. Laskowski

Department of Evolution and Ecology, University of California Davis, CA, U.S.A.

ARTICLE INFO

Article history:
Received 24 May 2024
Initial acceptance 11 July 2024
Final acceptance 9 January 2025
Available online xxx
MS. number: A24-00315R2

Keywords:
Bayes
dominance
hierarchy
loser effect
model
mutual assessment
social
winner effect

In animal contests, winners of previous contests often keep winning and losers keep losing. This coupling of previous experiences to future success, referred to as the winner–loser effect, plays a key role in stabilizing the resulting dominance hierarchies. Despite their importance, the cognitive mechanisms through which these effects occur are unknown. Identifying the mechanisms behind winner–loser effects requires identifying plausible models and generating predictions that can be used to test these alternative hypotheses. Winner–loser effects are often accompanied by a change in the aggressiveness of experienced individuals, which suggests individuals may be adjusting their self-assessment of their abilities after each contest. This updating of a prior estimate can be effectively described by Bayesian updating, and here we implement an agent-based model with continuous Bayesian updating to explore whether this is a plausible explanation of winner–loser effects. We first show that Bayesian updating reproduces known empirical results of typical dominance interactions. We then provide a series of testable predictions that can be used in future empirical work to distinguish Bayesian updating from simpler mechanisms. Our work demonstrates the utility of Bayesian updating as a mechanism to explain and ultimately predict changes in behaviour after salient social experiences.

The winner effect is a widely observed phenomenon where success in one contest leads to increased probability of success in subsequent contests, while the loser effect describes how defeat often leads to more defeats (Chase et al., 1994; reviewed in Hsu et al., 2006). Because of the impact of prior experiences on future success, winner–loser effects can have an important role in the formation of dominance hierarchies (reviewed in Tibbetts et al., 2022). Although it is established that an individual's position within their dominance hierarchy will have important consequences for their success (Dewsbury, 1988; Simons et al., 2022; Snyder-Mackler et al., 2020), predicting any given individual's position in social networks remains a challenge (Chase et al., 2002, 2022; Landau, 1951). This difficulty persists in part because we lack a full understanding of the behavioural and cognitive mechanisms used by individuals to navigate repeated social interactions (Chase et al., 2002; Tibbetts et al., 2022). Identifying the mechanism behind winner–loser effects and how they function in the formation of social hierarchies could provide powerful insight into what determines individual dominance status and social network structure. If we hope to one day demonstrate the underlying mechanisms of winner–loser effects, we must first identify the predictions of specific potential mechanisms so that they can be tested empirically (Supplementary Material 2, Table S1).

The existence of a winner or loser effect implies a mechanistic link between past experiences and future contest outcomes. Assuming winning a contest is the result of some combination of intrinsic ability and individual behaviour, there are three mechanisms that could explain winner–loser effects: (1) winning/losing could modify an individual's intrinsic ability; (2) contest outcomes could change individual behaviour directly by modifying the individual rules governing that behaviour; or (3) contest outcomes could modify behaviour indirectly via some upstream internal state variable (e.g. self-assessment). This state variable could then change behaviour via static rules governing behaviour. While a change in internal state (specifically self-assessment) is commonly cited as the explanation for winner–loser effects (reviewed in Rutte et al., 2006), most existing theory on winner–loser effects has modelled self-assessment only implicitly, modelling the change either as mechanism (1) or (2) described above.

The first wave of models exploring winner–loser effects (e.g. Bonabeau et al., 1996; Dugatkin, 1997; Hemelrijk, 2000), termed

* Corresponding author.
E-mail address: aperkes@ucdavis.edu (A. Perkes).

https://doi.org/10.1016/j.anbehav.2025.123191

type I models by Mesterton-Gibbons (2016), focused on whether changes in the probability of winning could explain the nature of the resultant dominance hierarchies. Social dominance is complex, but empirical evidence shows that hierarchies are generally stable and contain more linear dominance relationships than would be predicted by chance (Jackson & Winnegrad, 1988; Tibbetts et al., 2022). In type I models, the existence of winner—loser effects is a stated assumption, and the focus is to explore the consequences of winner—loser effects for social hierarchies, specifically their role in promoting linearity. In these models, the underlying mechanisms of winner—loser effects often remain something of a black box. Type I models tend to model winner—loser effects directly as a change in intrinsic ability, either by scaling the resource-holding potential (RHP) of the winner (and loser) by some fixed amount (Bonabeau et al., 1996; Dugatkin, 1997; Hickey & Davidsen, 2019; Kura et al., 2016), or changing RHP as a function of the difference between individual and opponent RHP (Hemelrijk, 2000; Hock & Huber, 2006). While these models explore the predictions of winner—loser effects on dominance hierarchies, they remain structurally agnostic as to whether the change in RHP reflects a change in behaviour, intrinsic ability, or both.

A second class of models termed type II models (Mesterton-Gibbons et al., 2016), test how winner—loser effects can evolve, i.e. whether there are evolutionarily stable strategies that give rise to winner—loser effects. Because these models are interested in the evolution of behaviour, they generally distinguish behaviour from intrinsic ability. In some cases, agents' estimates of that ability is only implicit (Leimar, 2021; Van Doorn et al., 2003), but two models (Fawcett & Johnstone, 2010; Mesterton-Gibbons, 1999) treat winner—loser effects as an explicit change in self-assessment. To our knowledge, these are the only two models in which winner—loser effects result from an explicit change in self-assessment, but in both cases, self-assessment is incidental to their respective questions, so they limit agents to two or three possible states (big/small or naïve/post-win/post-loss). Furthermore, for simplicity, both models limit the contest behaviour to a binary all-or-nothing decision, compete or yield (i.e. a hawk—dove contest), rather than allowing individuals to decide how much and how long to invest in a given contest. Additional models are needed to identify the predictions of continuous changes in self-assessment and contest effort, as well as explore how these predictions vary based on the underlying mechanisms of adjusting self-assessment.

If we assume winner—loser effects are driven by modifying uncertain estimates of individual ability, what mechanism best describes this change in self-assessment? While there are many potential mechanisms, Bayesian updating is a clearly relevant approach. In brief, Bayesian updating entails using Bayes' theorem to modify an existing estimate of the state of the world using newly acquired information (McNamara et al., 2006), thereby calculating the precise probability of some given state (e.g. being an individual of size x). Bayesian updating itself thus operates entirely on individual perception, making it well suited to describe assessment-based winner—loser effects. Conveniently, because the features that factor into Bayesian updating are generally explicit models of biological features (e.g. size distribution, probability of winning a contest), researchers can input relevant knowledge of these distributions during model specification. Because Bayesian updating calculates the precise conditional probability of some event, it provides a theoretical best-case scenario, and existing theory and empirical research indicate that many animals at least approximate Bayesian processes when making decisions in other contexts, for example during foraging or mate choice (Luttbeg, 1996; McNamara et al., 2006; Okasha, 2013; Olsson, 2006; Valone, 2006). Despite being so well suited to model postcontest changes in self-assessment (Whitehouse, 1997), to our knowledge, only one model (Fawcett & Johnstone, 2010) has incorporated Bayesian updating when modelling winner—loser effects, and only in the context of binary ability (big/small) and effort (hawk/dove) as described above. Natural contests are generally decided by the relative size, ability and effort of competitors, all continuous traits, but no model has used Bayesian updating for continuous abilities and outcomes. Bayesian modelling can be challenging to implement (see McNamara & Leimar, 2020), particularly as the complexity of the social information increases, but it is likely that the predictions of a continuous model would differ from a binary scenario, and understanding these specific predictions of changing self-assessments via Bayesian updating could provide important insight into the actual mechanisms underlying the behaviour of various systems. Given the potential power of Bayesian updating as a hypothesis to predict and explain winner—loser effects, it is important to identify the specific predictions of a Bayesian model and understand how we might distinguish Bayesian updating from alternative mechanisms.

Here, we present an agent-based model of winner—loser effects, and dominance formation generally, as a process of Bayesian updating for continuous self-assessment, in which individual agents with self-assessed contest ability compete to win social contests, using Bayesian updating to modify their self-assessments. The purposes of this model are (1) to explore whether Bayesian updating is a plausible explanation of winner—loser effects by comparing the features and limitations of Bayesian updating to empirical observations and (2) to identify testable predictions that could assess empirically whether animal behaviour is consistent with Bayesian updating, and ideally distinguish it from alternative mechanisms of modifying self-assessments. Because our model is motivated by and for empirical concerns, it is designed primarily for demonstrating whether Bayesian updating is a plausible mechanism behind winner—loser effects and for generating testable predictions, rather than (for example) being a general description of social behaviour. To this end, we show that Bayesian updating can reconcile disparate empirical observations of the attributes of winner—loser effects, while laying out clear, testable predictions that can be used in future experiments, to better understand the mechanisms that drive winner—loser effects specifically and individual dominance status and social structure generally.

GENERAL APPROACH

We were interested in a scenario where animals are motivated to win contests without overinvesting (Maynard Smith & Parker, 1976; Parker, 1974) and can observe opponent size (Arnott & Elwood, 2009) but cannot accurately observe their own size. (Size is used here as a proxy for any intrinsic trait that drives contest outcomes.) For illustration purposes, in Fig. 1 we depict the agents in our model as a tank of fish of varying sizes, but there is nothing specific about this model to fish, and readers should imagine the individual agents as representing any species where conspecific competition occurs and there is some uncertainty about relative contest ability (e.g. rats, parakeets, baboons, humans). Under these conditions, we would expect individuals to update their self-assessment following a contest to better determine future optimal contest effort. Specifically, we construct a scenario where animals compete for dominance in paired contests and individuals win by outlasting their opponents. This is similar to the war of attrition described by Maynard Smith (1974), which is well suited to describe animal contests where effort and/or persistence determines the outcome, as are common in nature (Enquist et al., 1990; Koops & Grant, 1993; Leimar et al., 1991; Marden & Waage, 1990; Mesterton-Gibbons et al., 1996). Note, however, that this is

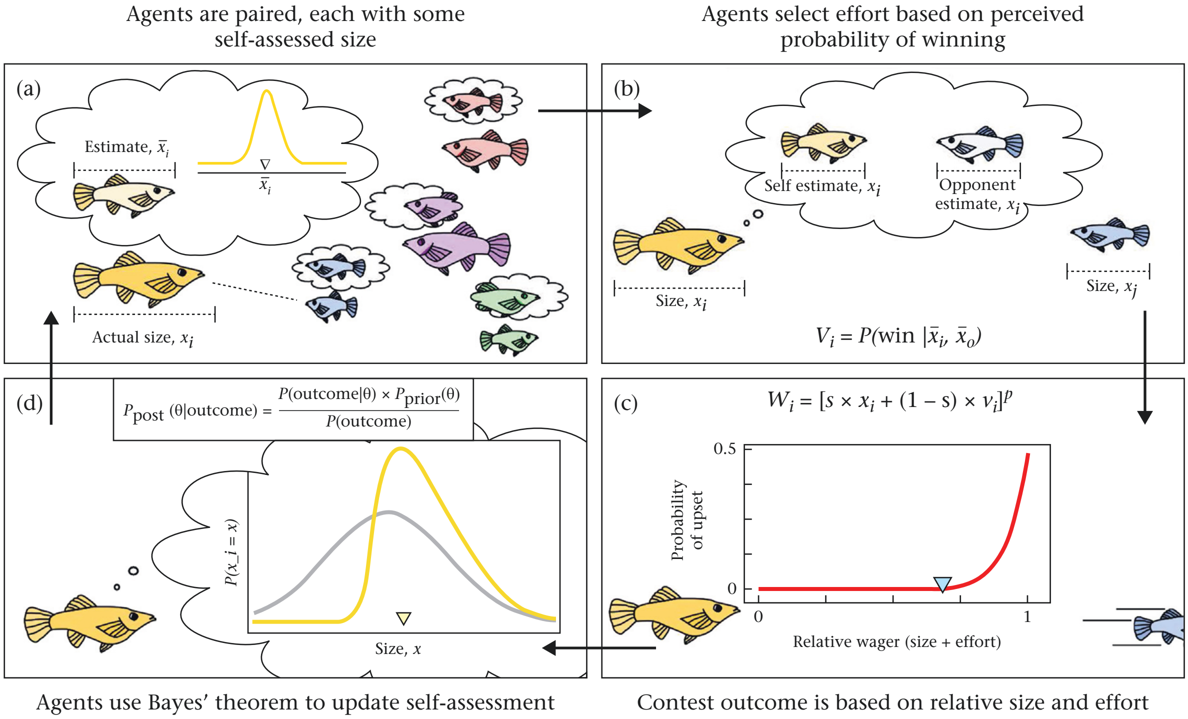

Figure 1. Overview of the model. (a) Individuals with fixed size, representing intrinsic ability (which is unknown to self), and an estimated self-assessment are paired. (b) Each individual estimates their own size and their opponent's size to determine their effort, based on their estimated probability of winning. (c) The size and effort of each agent determines their relative wagers. The relative wager of the smaller individual determines the probability of winning. (d) Based on this probabilistic outcome, agents update their self-assessment of their size, multiplying their prior assessment by the calculated likelihood of the observed outcome.

not quite a game in the mathematical sense, as there are no explicit costs or benefits to winning/losing, and our model is not evolutionary: we specify investment strategies and other aspects of the model based on known features of empirical systems.

Our model consists of four steps, described graphically in Fig. 1. First, we generate agents with random (normally distributed) intrinsic contest ability and imperfect naïve self-assessments of that ability (Fig. 1a). For illustration purposes, we will refer to the intrinsic ability as 'size', although this could describe any combination of intrinsic traits. Second, these agents are paired with opponents and allowed to interact. In each contest, individuals determine the maximum effort they are willing to invest, which is based on their estimated probability of winning, using their existing self-assessment of their size and their assessment of opponent size (Fig. 1b). Third, contest outcomes are determined probabilistically based on the relative size and effort of each participant (Fig. 1c). Finally, following the contest, each agent updates their self-assessment based on their prior estimate of size and their estimated likelihood of the observed outcome (Fig. 1d). We simulate these contests either in controlled one-on-one contests or in small group simulations where all individuals meet via randomized round-robin pairings. For these latter group simulations, groups are closed, without immigration or emigration, and are composed of $N = 5$ agents, lacking individual recognition or social eavesdropping.

Generation of Agents

To model the modification of self-assessment, we generate focal agents with some intrinsic contest ability, called 'size'. The size remains constant during the simulation and is usually set at $x_i = 50$, except for group simulations, where each agent size is drawn from a truncated normal distribution between 1 and 100 (arbitrary units), such that

$$x_i \sim N(\mu_x, \sigma_x), \text{where } x_i \in [1, 100] \tag{1}$$

Each agent starts with a point estimate, $\bar{x}_i$, of their self-assessment of size, itself drawn from a truncated normal distribution centred around their actual size:

$$\bar{x}_i \sim N(x_i, \sigma_a), \text{where } x_i \in [1, 100] \tag{2a}$$

where σ_a reflects individual 'self-assessment error', i.e. the initial precision with which they can estimate their own size. Note that we imagine this initial self-assessment being based on some direct correlate of individual size, rather than being inferred from observing some surrounding population. Because we are interested in scenarios where naïve self-assessment is difficult, for all main results, we set the naïve self-assessment error quite high, $\sigma_a = 20$ (Table 1), but as with all parameters, we vary this value to observe the impact on our predictions (see Supplementary Material 1 for notes on parameter selection and Supplementary Material 2 for sensitivity analyses). Individuals then calculate a starting prior as the truncated normal distribution centred on their estimate, with standard deviation, $\sigma_\theta = \sigma_a$,

$$P_r(\theta) = N(\bar{x}_i, \sigma_\theta), \text{for } 1 \le x \le 100 \tag{2b}$$

For our model, we will use the term 'estimate' to refer specifically to the point value, $\bar{x}$, which is the maximum likelihood estimate of the probability distribution of size, while 'self-assessment' typically refers to the full distribution.

Investment Strategy for Effort

Although the purpose of our model is to explore updating mechanisms, the impact of individual self-assessment is wholly dependent on the rules individuals use to determine behaviour, so we must first establish how agents determine their effort, v_i, for a

Table 1
Overview of parameters and their default values

Symbol	Default	Range	Explanation
$x : (x, x_i, x_o, \bar{x}_i)$	Random	$1 \le x \le 100$	Size of an individual. x_i represents the size for individual i, while x_o is the size of opponent. $\bar{x}$ is the max-likelihood estimate of size
μ_x	50	$1 \le x \le 100$	Mean of the population from which individuals are drawn for group simulations
σ_x	10	$0 \le \sigma_x < \infty$	Standard deviation of population from which individuals are drawn for group simulations. At $\sigma = \infty$, uniform distribution is used
v: (v, v_i, v_o)	None	$0 \le v \le 1$	Effort, the amount (e.g. of time) agents are willing to invest in a contest
w: (w, w_i, w_o)	None	$0 \le w \le 1$	Wager, or contest performance, a combination of the agent's size and effort
s	0.7	$0 \le s < 1$	Relative importance of size vs effort when calculating the weighted sum
p	6.3	$0 \le p < \infty$	Relative wager modifier, controlling predictability, such as p increases, the probability of an upset decreases
σ_a	20	$0 \le \sigma_a < \infty$	Self-assessment error, the accuracy of agents' starting size estimate
σ_c	3.1	$0 \le \sigma_c < \infty$	Opponent assessment error, the accuracy when guessing opponent size
$P(\theta)$	–	–	Used to denote the prior, the estimated probability distribution of size
n	5	$2 \le n < 10$	The number of individuals in a group, when placed in round-robin style tournaments

contest. In our model, effort represents any degree of variable investment in that contest (e.g. the amount of time/energy an animal will expend before yielding). To determine effort (i.e. maximum investment), each agent assesses the size of their opponent, by drawing their estimate of opponent size, $\bar{x}_o$, from a truncated normal distribution centred on their opponent's true size,

$$\bar{x}_o \sim N(x_o, \sigma_c), \text{where } x_o \in [1, 100] \tag{3a}$$

Here the standard deviation, σ_c, represents opponent assessment error. For all main results, $\sigma_c = 3.1$. Agents then set their effort, v_i, based on their estimated probability of winning, such that

$$v_i = P(\text{win} \mid \bar{x}_i, \bar{x}_o) \tag{3b}$$

where effort, v_i, is equal to the conditional probability of winning.

In plain terms, we assume that individuals are willing to bear a greater cost when their perceived probability of victory is higher. The principle of matching investment to the expected reward is a standard assumption of behavioural theory (Enquist & Leimar, 1983; Maynard Smith & Parker, 1976), and empirical work confirms that individuals are willing to invest more in contests they expect to win (Hsu et al., 2008). This increased investment in turn increases their probability of winning, creating a feedback loop between expected outcomes, investment and actual outcomes that is at the heart of assessment-based winner–loser effects.

In equation (3b) above, calculating the probability of winning requires individuals to estimate not just the opponent's size but also predict their relative effort. For simplicity, in our model, agents assume that both their opponent and they themselves will equally invest 0.5 effort to estimate their probability of winning. Of course, we might expect individuals to better anticipate both their own and their opponent's effort, and there is a rich literature investigating the evolution of optimal strategies (Maynard Smith, 1974; McNamara & Leimar, 2020), including for winner effects (Leimar, 2021; Mesterton-Gibbons et al., 2016), but establishing a full evolutionary model was beyond the scope of this paper. As our model is primarily focused on how self-assessment changes following a contest under Bayesian updating, our specified effort strategy serves as a plausible simplifying assumption for our purposes here.

Having established their assumptions of their own and their opponent's size, agents can calculate their effort, v_i, being equal to their perceived probability of winning, via equations (4a–4b) below. Note that for practical purposes, assessment in our model is instantaneous and conducted prior to the contest, but this also captures scenarios where the assessment occurs over the course of the contest, as is the case in many natural systems (Arnott & Elwood, 2008). At some point, whether before or during the contest, individuals must determine when they will yield. While continuous assessment would likely be important if our goal were to predict changes in behaviour during the fight, our research question is focused on how assessment changes as a result of contest outcomes, so this simplification seemed appropriate. Obviously, our approach will not precisely match every animal system, but it models a broad range of systems where animals are uncertain of opponent ability and behaviour, including systems in which size and/or effort are continuous, which may not have been captured by existing models of winner–loser effects.

Determining Contest Outcome

Once individuals determine their maximum effort, the winner is decided probabilistically, based on the size and effort of each agent, combined to define each agent's wager, w_i:

$$w_i = s \times \frac{x_i}{100} + (1 - s) \times v_i \tag{4a}$$

The wager thus scales with their size (normalized to be between $(0, 1)$) and individual effort, with s allowing us to control the relative importance of size and effort. For example, our default parameter value, $s = 0.7$, means that size contributes 70% of the wager value, with effort contributing 30%, but different parameter values can make contests more or less size dependent, as we explore in Supplementary Material 2.

In most natural systems, the contest ends when one individual yields (becoming the loser). For this reason, we calculate the outcome probability from the perspective of the lower-wagering individual, termed the underdog. We thus calculated the probability of an upset (i.e. a win by the underdog) with

$$P(\text{upset}) = \frac{1}{2} \times \left(\frac{w_{\text{min.}}}{w_{\text{max.}}}\right)^p \tag{4b}$$

where p is a bias parameter, 'predictability', controlling how close the two wagers need to be for there to be a meaningful probability of upset, such that as p increases to infinity, the probability of an upset goes to 0 unless the two wagers are equal. Once $P(\text{upset})$ is established, we use a random number generator to simulate whether an upset occurs.

This approach for settling contests allows us to vary the relative contributions of size, effort and stochasticity while observing the (simulated) behaviour of individual agents and groups. For all main results, we set the values of these parameters at $s = 0.7, p = 6.3$, which corresponds to a system where intrinsic ability is the most important factor in determining contest outcome, and larger/better individuals generally win, but sufficient effort and/or luck drives occasional upsets.

Updating Self-assessment Postcontest

In the results below, we compare Bayesian updating with a simpler strategy (linear updating) and a null hypothesis (static estimates, i.e. no updating). Each is described here, with some additional details in Supplementary Material 2. These three strategies do not capture the entire range of possible updating mechanisms, but linear updating is a common approach in existing models of winner–loser effects (Bonabeau et al., 1996; Dugatkin, 1997; Kura et al., 2016) and serves as an instructive contrast with Bayesian updating, allowing us to distinguish what predictions are specific to Bayesian updating and what is simply a consequence of changes in self-assessment.

Static Estimate (No Updating)

The simplest possible approach to maintain a self-assessment is to never vary from your initial self-assessment of size (i.e. intrinsic ability). We model this scenario, keeping each agent's initial assessment of their size fixed throughout the simulation, regardless of contest outcomes. Under this 'no updating' approach, the accuracy of individual estimates depends on their initial awareness, σ_a.

Linear Updating

A common way to implement simple self-assessment updating, first implemented by Dugatkin (1997), is to either increase or decrease the agent's estimate of their size, using some scalar, k, multiplied by its prior estimate.

$$\overline{x}_{i_{t+1}} = k\overline{x}_{i_t} \tag{5a}$$

To make this compatible with our model construction, we set bounds on the maximum and minimum possible size ($X_{\text{min.}} = 1$, $X_{\text{max.}} = 100$), such that an agent's new estimate, $\overline{x}_{i_{t+1}}$, is a function of their previous estimate, shifted by some factor, where

$$X_{\text{min.}} \leq \overline{x}_{i_{t+1}} \leq X_{\text{max.}} \tag{5b}$$

We also set k to vary dynamically as a function of distance from the max./min. size, where

$$k = 0.1 \times \left(X_{\text{max.}} - \overline{x}_{i_t}\right), \text{if win} \tag{5c}$$

$$k = -0.1 \times \left(\overline{x}_{i_t} - X_{\text{min.}}\right), \text{if loss}$$

This dynamic shift more closely matches Bayesian updating and prevents estimates from escaping their bounds. Since this function only acts on the point estimate (defined by the maximum likelihood estimation of the prior distribution), after updating this maximum likelihood estimate, we generate a new, truncated normal distribution, centred around that estimate, based on the agent's self-assessment error, σ_a, as in equation (2b). This approach provides a simple heuristic that, like Bayesian updating, allows for weighted updating of self-assessment but does not change an estimate's confidence and does not calculate a likelihood function.

Bayesian Updating

In contrast with simpler heuristics, Bayesian updating can potentially calculate the true probability of some value or event, based on known/estimated parameters about the state of the world and the conditional probability of outcomes. Bayesian updating is thus defined by a prior assessment and a likelihood function. In our case, the prior is the agent's current self-assessment of its own size, modelled as a discrete probability distribution. The likelihood function provides the probability of the observed outcome (i.e. winning a given contest), conditioned on being some assumed size, calculated across all possible sizes. Then the posterior distribution, i.e. the new probability distribution of an agent's size estimate, θ, is calculated by Bayes' formula

$$P_{\text{post}}(\theta|\text{outcome}) = \frac{P(\text{outcome} \mid \theta) \times P_{\text{prior}}(\theta)}{P(\text{outcome})} \tag{6a}$$

We discuss here the case of a winning outcome, but the formulation is similar following a loss. For each possible size in our discrete size range, the posterior probability of being a given size is

$$P(x_i = x|\text{win}) = \frac{P(\text{win}|x_i = x) \times P(x_i = x)}{P(\text{win})} \tag{6b}$$

where $P(x_i = x)$ is the estimated prior probability of being a given size, while the likelihood $P(\text{win}|x_i = x)$ is the probability of an individual of size x winning the previous contest, which is calculated according to the probability of upset from equation (4b). In this model, we assume that, although their opponent assessment is error-prone while deciding how much to invest, by the end of the contest, individuals have an accurate assessment of the size and effort of their opponent, allowing them to calculate the probability of the observed outcome (winning/losing) for any given own size, x_i, without needing to iterate over all possible opponent sizes and efforts (see Supplementary Material 2). $P(\text{win})$ is thus calculated as the sum of $P(\text{win}|x_i = x) \times P(x_i = x)$ across all possible sizes, x,

$$P(\text{win}) = \sum_{x=0}^{x_{\text{max.}}} P(\text{win}|x_i = x) \times P(x_i = x) \tag{6c}$$

After calculating the posterior estimate, agents update their point estimate of size, $\overline{x}_i$

$$\overline{x}_i = E(\theta) = \sum_{x=1}^{100} x \times P(x_i = x) \tag{6d}$$

where E is the maximum likelihood estimate of the size distribution. This value is then used to compute agent effort in the next contest.

While the above notation can seem somewhat daunting, the general approach is quite intuitive: individuals can observe that they won (or lost) a contest. They can then infer that (for example) it is very unlikely that they would have won a contest against a large individual if they were small (the likelihood), so they shift their existing estimate (their prior) towards larger values (Fig. 1d). Assuming the information provided is accurate, Bayes' formula makes it possible to calculate the 'true' probability of being a given size, thus providing a best-case scenario for uncertain self-assessment.

Parameter Space

Our possible parameter space comprises the contest outcome parameters (s, p), the first reflecting the influence of size versus effort, and the second describing the degree of predictability and the assessment parameters (σ_a, σ_c), representing self-assessment error and opponent assessment error. For all main figures, we use the following values: $s = 0.7; p = 6.3; \sigma_a = 20; \sigma_c = 3.1$. The default values of s and p were chosen to approximate empirical observations that upsets are rare and contests are roughly predicted by size (Supplementary Material 1, Fig. S1) (Beacham, 1988; Bierbach et al., 2012), while the values of σ_a and σ_c were selected to explore a scenario where opponent assessment is somewhat error-prone and naïve self-assessment is poor. We explore the full parameter space in Supplementary Material 2.

Model Analysis

The goal of our model is to answer the following two questions. (1) Does Bayesian updating for self-assessment produce winner–loser effects and dominance hierarchies that match the behaviour of empirical systems? (2) How can we distinguish, empirically, Bayesian updating from alternative models of dominance establishment and winner–loser effects? To answer these questions, we run simulations of model behavioural experiments while extracting behavioural metrics that would be tractable in and relevant to an empirical context, such as contest intensity. These experiments, and the behavioural metrics we measure, are discussed in their relevant results sections.

RESULTS

Part 1: Matching Empirical Observations

Before addressing how we might test for Bayesian updating, we first compare the predictions of this model to known empirical features of winner–loser effects. Here we focus on three well-documented phenomena: (1) winner–loser effects exist; (2) the relative strength of winner–loser effects varies across species; (3) more recent contests tend to be more impactful. We also explore (4) how Bayesian updating might account for the observed differences in the duration of winner–loser effects across systems and (5) the extent to which Bayesian updating promotes stable linear hierarchies. These observations need not be unique to Bayesian updating, indeed there are several simple models that are consistent with these empirical observations, but they are common observations of empirical contest behaviour, so it is important to test whether a Bayesian model can also recreate these effects, and this provides insight into how Bayesian updating functions in the context of winner–loser effects.

Bayesian updating predicts the existence of winner and loser effects

We first test whether agents using Bayesian updating show winner–loser effects. To do so, we model a simulated experiment (designed to emulate standard empirical approaches) by forcing either a win or a loss in a contest between a focal agent and a 'treatment' opponent of known size (in this case, the same size as the focal individual, $x_i = 50$). We then measure winner–loser effects by pairing these focal winners or losers against 'assay' opponents (a second size-matched agent). We repeat this process for 1000 focal winners and losers. We can then quantify the winner (loser) effect as the proportion of focal agents that won (lost) against their assay opponent. In the absence of winner (loser) effects, we would expect 50% of focal agents to win (lose) their assay contest.

As shown in Fig. 2a, we found that agents using Bayesian updating for self-assessment exhibited obvious winner–loser effects, in that previous winners were more likely to win subsequent contests, while previous losers were more likely to lose. These winner and loser effects were broadly observed across the range of parameter values wherever size and effort combined to determine contest outcome, the only exception being where naïve self-assessment was perfect (Supplementary Material 2, Fig. S2).

Bayesian updating can generate variable biases in the strength of winner–loser effects

Given our model's construction (winning increases an individual's estimate, the individual's estimate determines effort and effort determines the probability of winning), it is perhaps not surprising that Bayesian updating produces winner–loser effects. However, the details of how winner–loser effects function under Bayesian updating are not trivial. We found that, as implemented here, Bayesian updating produced a loser effect that was marginally stronger (on average) than the winner effect, but this depended on the parameters chosen (Supplementary Material 2, Fig. S2). It has long been assumed that loser effects tend to be stronger than winner effects (Hsu et al., 2006), but a recent meta-analysis showed high variation (with some species being winner-biased and others loser-biased) but no general effect (Yan et al., 2024). Our model is broadly consistent with this finding, as in our case, the relative strength of winner versus loser effects was highly sensitive to the investment strategy agents used to determine their effort in each contest. Under the simple 'proportional effort' strategy we use for the main results (equation 3b), we found a very slight loser effect bias (Supplementary Material 2, Fig. S2), driven by the fact that winning a contest was determined by the relative size of the smaller individual, and for any given shift, the relative decrease in the estimated size difference was greater than for an increase in estimate. To illustrate this, consider an individual of size 50 that either increases or decreases its estimate by 25 units, then faces another individual of size 50. Decreasing its estimate results in the

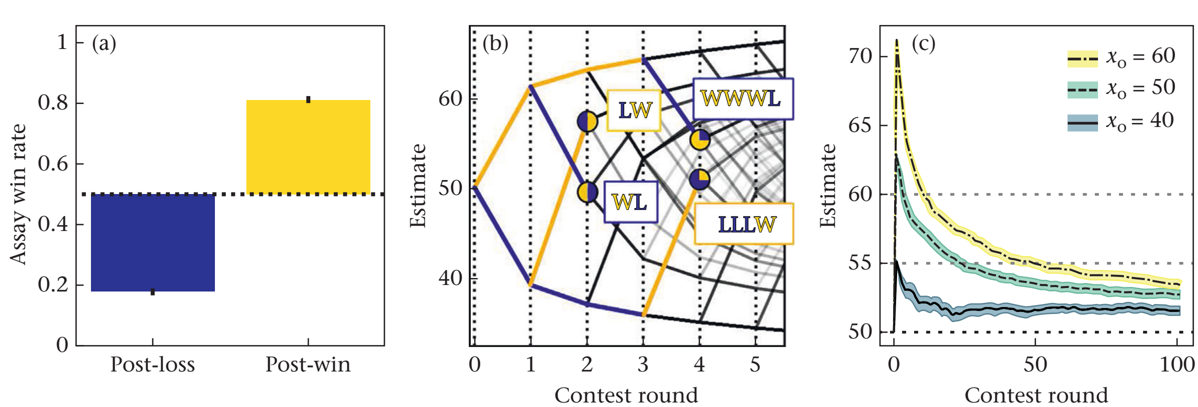

Figure 2. (a) The proportion of focal agents (out of $N = 1000$) winning against a naïve size-matched assay opponent, after a treatment win/loss. (b) Agent estimates over the course of repeated contests against naïve size-matched opponents, with the branching paths showing the estimates following a win (W) or loss (L). The boxes note specific trajectories highlighting the recency effect, which diminishes with increased experience. (c) Focal agents' estimates change following a staged win against agents of varying size. The shift in the focal agent's estimate following a win is greater when their opponent is larger.

individual perceiving itself as (the underdog) being just 50% of its opponent's size (25/50), while increasing its estimate results in the individual perceiving its opponent as (the underdog) being 67% of its own size (50/75). This effect of proportionality, built into the agent's investment strategy, drives the slight loser effect bias in the model described above, but under alternative investment strategies that we explore in Supplementary Material 2 (equation S4), we can produce agents with stronger biases towards winner or loser effects (Fig. S4), even with equal changes in self-assessment. The fact that loser effect bias is mediated by, and highly dependent on, the specific effort strategy used by a given system could account for the high degree of variation in the relative strengths of the winner and loser effects observed across empirical systems, based on their specific evolutionary history.

Bayesian updating produces a behaviour-mediated recency bias in winner—loser effects

Although winner—loser effects vary across species, there is consistent evidence for a recency bias (Benincasa et al., 2023; Hsu & Wolf, 1999), in which the outcomes of more recent contests have a stronger effect on behaviour than do the outcomes of earlier contests. To explore this, we emulate empirical tests of the winner—loser effect, by simulating contests between naïve size-matched agents ($x_i = 50$), then pair these focal winners and losers against new, naïve size-matched opponents, repeating this process over six rounds of contests. We control the outcome of each contest to generate new branching paths with both winners and losers. To limit noise, in this simulated experiment we centre all individual priors on their true size, and agents' opponent assessment is set to match opponent size exactly. Here, we measure winner—loser effects directly as the agents' own size estimates.

In Fig. 2b, we see that our model recreates winner—loser effects with an experience-dependent recency bias. In particular, 'recent losers' (WL), who first won a contest and then lost a second contest, updated their size estimate to be lower than 'recent winners' (LW), who initially lost and then won. This is particularly interesting, since Bayesian updating is known to be 'order invariant' in that the order of identical events should not impact the final estimate. However, in this case the events are not identical: the recent winners (LW) win while investing less in their second contest than the recent losers (i.e. initial winners, WL). Because estimate updating occurs at every step, shifting the amount of effort, these events might be better described as LW^-, WL^+, since later wins and losses occur under different circumstances, which make them more impactful. The strength of this effect depends on the specific parameters used (see Supplementary Material 2, Fig. S5), but under all parameters, recent winners had posterior estimates that were greater than or equal to those of recent losers. Like the other phenomenon described in this section, this recency effect is of course not unique to Bayesian updating, but it confirms that Bayesian updating is consistent with empirical observations.

Under Bayesian updating, experimental methodology drives variation in the persistence of winner—loser effects

In empirical observations of dominance contests, winner—loser effects are persistent in some cases (Lan & Hsu, 2011; Laskowski et al., 2016) while in others they are short-lived (Chase et al., 1994) or even absent (reviewed in Hsu et al., 2006). If we assume Bayesian updating, what mechanism could produce these discrepancies? To assess this, we again simulate a forced-win experimental approach, this time allowing a simulated focal agent of size $x_o = 50$ to win a staged contest against a treatment opponent of variable size: either size-matched to, or smaller or larger than the focal agent ($x_o = [40, 50 \text{ or } 60]$). We then measure the effects of the contest outcome directly by checking the estimate of the agent's own self-assessment after the contest. Following this initial treatment contest, we allowed focal agents to interact with additional naïve individuals so that we could observe subsequent changes in their estimates over time.

We found that the duration of winner—loser effects depended on the size of the initial shift in estimate. The size of this shift was a function of the unexpectedness of the outcome, in this case the opponent's size, with wins against larger opponents generating larger shifts in agent estimates (Fig. 2c). This initial increase in self-assessment attenuated rapidly with additional experience, although it did not fully return to baseline, with the mean estimates of winners (for all opponent sizes) persisting significantly above the true individual size ($x = 50$). The extent to which the change in self-assessment, and the associated winner effects, persist above a given threshold also depended somewhat on the nature of these post-fight experiences (Supplementary Material 2, Fig. S8), but even after 100 contests, the mean self-assessments of these initial winners were higher than their actual size, and those who won against bigger opponents had significantly higher self-assessments than those who won over smaller ones.

Although winner estimates remain higher indefinitely under Bayesian updating, whether such a change in self-assessment could be detected empirically is a function of the power of the experiment used to assay winner—loser effects. If a small shift in estimate (blue line, Fig. 2c) does not lead to visible changes in behaviour, the winner effect may appear to be temporary or even nonexistent. Social experience following the 'treatment' contest and variation in the size and/or effort of the assay opponent would add additional noise that may further mask small differences in focal individual estimates. In a simulated experiment, moderate sample sizes with some intervening social experience were generally insufficient to detect winner—loser effects at the behavioural level (i.e. a deviation from the expected 50% win (loss) rate against a size-matched opponent), even though these effects persisted in individual estimates (Supplementary Material 2, Fig. S7). Indeed it has been shown that in empirical contexts, whether, and for how long, winner—loser effects are observed depends on the experimental methods used (Chase et al., 1994; Huang et al., 2011). Given that, within our model, detecting persistent winner—loser effects requires far greater statistical power than is often feasible, it is not surprising that many studies have found winner—loser effects to be short-lived, even if we assume that the winner—loser effect can be perfectly described by Bayesian updating.

Bayesian updating increases the efficiency of size-based dominance hierarchies

So far, we have only assessed winner—loser effects at the level of individuals, but these effects generally function within group hierarchies. It is thus important to model and test how Bayesian updating is predicted to impact winner effects within social networks since the marginal increases in effort predicted under experimental contexts may not translate to observable changes within dynamic dominance hierarchies. To do this, we established groups of $N = 5$ individuals with sizes drawn from a normal distribution ($\bar{x} = 50, \sigma_x = 10$) that were paired in randomized rounds such that every individual faced every other individual. We then recorded the structure and intensity of dominance contests.

The transition to low-intensity contests is a standard prediction of established social groups (Jackson & Winnegrad, 1988), as individuals should seek to avoid serious injury from escalated social contests, especially when contests are frequent, as is the case in many group-living species. Although individuals in our model were forced to 'interact', agents chose how much effort to invest, between 0 and 1, allowing them to effectively choose whether to engage in a contest. Future implementations could more fully model the costs and benefits of opponent selection, but here we simply assumed individuals were willing to invest proportional to their perceived probability of winning to ask whether increasing the accuracy of individual estimates would decrease the intensity of contests, as well as producing dominance hierarchies that are more linear and stable than we would expect without updating.

We found that our simulated agents using Bayesian updating rapidly reduced the intensity of contests (Fig. 3a). This occurred in our model as the smaller agents began to yield quickly to larger opponents, resulting in lower mean effort across contests. This shift was attributable to the rapid decrease in estimate error, calculated using equation (S5) in Supplementary Material 2, as shown in Fig. 3b. This increase in accuracy, and the associated decrease in intensity, was widely observed across our parameter space (see Supplementary Material 2, Figs S8–S9). For comparison, we repeated these simulated groups using a fixed estimate (Fig. 3) or a linear estimate (Supplementary Material 2, Fig. S9). Compared to these alternative updating methods, Bayesian updating resulted in more accurate estimates which, in turn, resulted in the rapid reduction in high-intensity contests.

This improved accuracy also increased the stability and linearity of networks over time (see Supplementary Material 2, Figs S11–S12), but these networks did not 'self-organize' into stable, linear hierarchies in the absence of intrinsic differences in size. This makes sense, since Bayesian updating results in accurate assessments, so where intrinsic differences are small, individuals accurately assess that they have a high potential to win a contest and invest accordingly. If we modify our model to allow for feedback on agent size (Supplementary Material 2, equation S6), networks do organize into linear hierarchies, and in this context, Bayesian updating causes networks to self-organize more quickly and to a greater extent than with feedback on size alone (Supplementary Material 2, Fig. S13e). In short, the increased accuracy of Bayesian updating has a stabilizing effect on dominance hierarchies, beyond what can occur from inaccurate self-assessment or changes in intrinsic ability alone.

Part 2: Distinguishing Bayesian Updating from Other Mechanisms

We have shown that our Bayesian model is successful at recreating many empirical observations of animal contests. We now explore which testable predictions distinguish Bayesian updating from other potential mechanisms of self-assessment, again comparing three potential updating mechanisms: Bayesian, linear updating or static estimates (no updating). In particular, we focus on the two defining features of Bayesian updating: the likelihood-based updating and the probabilistic prior. We investigate each of these two aspects separately, with simulated experiments using empirically tractable approaches to describe potential experiments in animal systems. Their contrasting predictions are summarized in Supplementary Material 2, Table S1.

The certainty effect is characteristic of Bayesian updating

Because Bayesian updating encodes all previous experiences as a probability distribution, it is able to capture not just the individual estimate but the statistical certainty of that estimate. As such, the statistical certainty gained from increased experience should distinguish Bayesian updating from linear updating, which only maintains the estimate itself. To test for this certainty effect, we simulate focal agents, using either Bayesian or linear updating, who experience between 0 and 50 contests against naïve opponents prior to the treatment win/loss, thereby creating differences in experience. Each contest alternates between a win against a smaller opponent and a loss against a larger opponent, producing individuals with similar point estimates but with different numbers of pretreatment experiences. We then assess the strength of the winner–loser effect as we would empirically, by pairing the focal individual against a size-matched treatment opponent and forcing a win/loss, and then recording whether these focal winners and losers win against a subsequent naïve size-matched assay opponent. This allows us to measure the extent to which certainty gained through experience attenuates the winner–loser effect observed in the final assay contest.

As expected, in agents using Bayesian updating (Fig. 4a, blue line), experienced agents with more prior contests showed weaker winner effects in the final contest than did naïve agents (i.e. agents for whom the number of fights prior to the treatment contest was 0). In comparison, in agents using linear updating (Fig. 4a, green line), winner effects remained strong regardless of previous experience. However, agents that use linear updating can sometimes mimic the same winner–loser effect attenuation seen in experienced Bayesian updaters. This can happen when the intervening

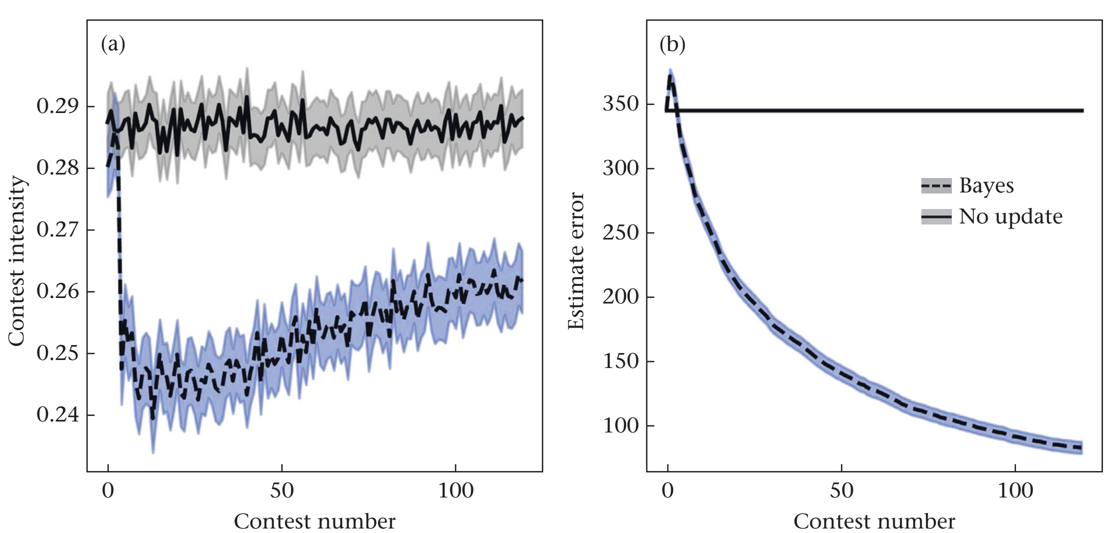

Figure 3. (a) Contest intensity, measured as the effort of the losing individual, decreases with repeated contests as individuals become more confident in their assessments (blue line); however, when agents cannot update, intensity remains high (grey line). (b) Agent estimate error decreases over time with Bayesian updating (blue line), but not when updating is prevented (grey line). The initial spike in error shows the initial winner–loser effects, which cause agents to temporarily over/underestimate their size. Contests over six randomized rounds of paired contests between all individuals in a group ($N = 5$ individuals). Lines and shading show the mean result, ± 1 SEM, averaged across 1000 simulations.

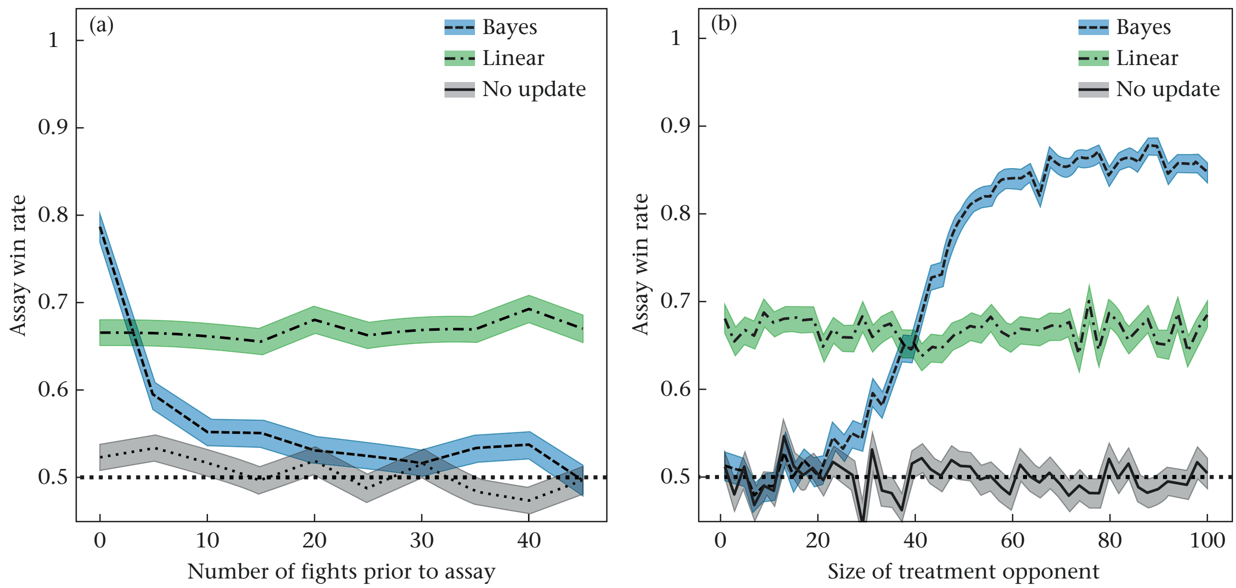

Figure 4. Bayesian updating is characterized by a certainty effect and a discrepancy effect. (a) Under Bayesian updating (but not linear updating), the strength of the winner effect (the proportion of agents winning a size-matched contest after a staged win) depends on the certainty of an agent's self-assessment and is a function of the number of contests prior to the forced win. (b) Similarly, the strength of the winner effect depends on the discrepancy between expectation and outcome (i.e. the treatment opponent size) under Bayesian updating, but not under linear updating.

contests are against randomly sized opponents, which can cause linear updating agents' estimates to strongly diverge over time, making them susceptible to winner—loser effects (Supplementary Material 2, Fig. S14). Thus, depending on the experimental paradigm, this 'certainty effect' may not always distinguish Bayesian updating from simpler mechanisms, but it remains a core feature of Bayesian updating (across a range of parameter values, as shown in Supplementary Material 2, Fig. S15), and failing to find that winner—loser effects attenuate with increased experience would provide strong evidence against Bayesian updating for continuous self-assessment as a mechanism.

The discrepancy effect can distinguish Bayesian updating from simpler models

In addition to the probabilistic prior, Bayesian updating is defined by the likelihood function, which calculates the probability of any given outcome based on various possible conditions. As such, we should expect winner—loser effects to vary based on the extent to which the outcome of the contest is surprising. As was already shown in Fig. 2c, agents using Bayesian updating showed stronger winner effects when their initial opponent was larger, and this effect can be observed across parameters wherever size and effort both affect contest outcome (Supplementary Material 2, Fig. S17). Importantly, this 'discrepancy effect' was not seen for agents using a linear shift (Fig. 4b).

The general observation that more surprising outcomes drive more dramatic shifts in the posterior estimate is a fundamental feature of Bayesian updating (Courville et al., 2006). In our model, it depends on agents conditioning their updates on opponent size specifically, but discrepancy effects could be driven by variation in any observable variable that influences the probability of winning (e.g. duration of the contest, age of the opponent, the quality of some defensive resource), depending on the specifics of the study system. While the simple linear updating mechanism we use for comparison does not incorporate any contest information, and thus does not show a discrepancy effect, linear updating methods could be modified to incorporate opponent size or number of past experiences, which could generate outcomes very similar to Bayesian updating. Thus, the presence of a discrepancy effect does not exclude the other mechanisms, but the lack of any such effect would exclude Bayesian updating as a likely mechanism for modifying self-assessment.

DISCUSSION

Here, we show that Bayesian updating for continuous self-assessment can explain the existence of winner and loser effects, resulting in contest behaviour and dominance hierarchies that are broadly consistent with those observed within social systems in nature. Our model shows that changes in self-assessment alone, whether through linear updating or Bayesian updating, can be sufficient to drive winner—loser effects (measured as changes in contest outcomes), without any change in actual ability or opponent behaviour. This prediction is consistent with previous self-assessment based models (Fawcett & Johnstone, 2010; Mesterton-Gibbons, 1999) and empirical observations of the winner effect occurring independent of changes in individual ability (Hsu & Wolf, 2001). In addition to being broadly consistent with empirical observations, our model generates testable predictions that distinguish Bayesian updating from the simpler, non-Bayesian linear model. These predictions are designed to lead directly to future empirical tests, allowing us and others to assess the mechanisms behind winner—loser effects. Even if these tests do not conclusively establish the underlying mechanism, they could identify features of winner—loser effects (e.g. does the size of the opponent determine the strength of winner—loser effects?) that have not been previously described.

Note that while we have generally discussed Bayesian updating specifically as it relates to how individuals change their assessments following a contest, researchers sometimes use 'Bayesian' to refer to individuals that begin a task with informed assumptions, and this can also be tested empirically. We explore some potential mechanisms that could explain how these initial assumptions may be formed in animals in Supplementary Material 2 (Figs S18—S20) and the predictions of those mechanisms. For example, animals could form these initial assumptions based on some internal cue (e.g. age) or based on observations of their surroundings (e.g. if an individual is in a group of animals of a given size, it is likely that their own size is similar). However, this aspect of Bayesian self-assessment is tangential to the winner effect itself, which is the focus of our model.

Like type I models that explore the consequences of winner effects, our model predicts that Bayesian updating (compared to maintaining naïve estimates) results in increasing the stability, linearity and efficiency of dominance hierarchies over time (Balph, 1979; Senar et al., 1990). But importantly, in our model, stable, linear hierarchies only occur when individuals start with, or can develop over time, intrinsic differences in ability (Supplementary Material 2, Fig. S13f). Our model is intended to explore changes in behaviour in the absence of any changes in intrinsic ability, and we found that Bayesian updating alone did not generate linear networks from identically sized individuals (Supplementary Material 2, Fig. S13d). This makes sense: if individuals can accurately assess their size, and all individuals are the same size, contests should be protracted and unpredictable, as all individuals (correctly) think they can win based on effort. Previous models have shown that when contest outcomes alter intrinsic ability, individuals 'self-organize' into networks that are more linear than expected by chance, even when all individuals start out at the same size/ability (Bonabeau et al., 1996; Dugatkin, 1997; Hemelrijk, 2000). We can extend our model to include direct feedback whereby winning a contest does increase intrinsic ability, and when we do, stratification of initially identical individuals does occur (Supplementary Material 2, Fig. S13b). In this context, Bayesian updating increases the observed linearity of simulated social networks (Supplementary Material 2, Fig. S13e) compared to fixed estimates. Based on these results, we would expect the stratification of dominance hierarchies observed in natural systems (Tibbetts et al., 2022) to require some real intrinsic differences in individuals, which either existed prior to the contest or which come about as a result of contests (e.g. as winners gain better access to resources) (Bonabeau et al., 1999). Accurate self-assessment updating could then act as secondary mechanism that allows individuals to match their effort to their ability, thereby facilitating the stability, linearity and efficiency of these dominance hierarchies.

If we accept that Bayesian updating is a plausible mechanism, how might we test for it? As we have shown, a certainty effect and a discrepancy effect are key features of Bayesian updating. Bayesian updating is expected to produce attenuating winner–loser effects with increasing experience and stronger shifts in behaviour in response to 'surprising' outcomes. These are natural consequences of the probabilistic prior and likelihood functions, respectively. Failing to find either a certainty or discrepancy effect would provide strong evidence against Bayesian updating being behind winner–loser effects, but it is important to note that these effects are not necessarily unique to Bayesian updating, as other updating mechanisms can produce similar results depending on their exact mechanics.

There are two models in particular whose constructions produce similar predictions to ours. The first is the binary Bayesian model from Fawcett and Johnstone (2010); the second is the actor–critic reinforcement learning model of Leimar (2021, extended in Leimar & Bshary, 2022). In both cases, the most obvious difference between these two models and ours is the extent to which effort and individual ability can vary. Our model focuses on continuous changes in self-assessment and effort. In contrast, Fawcett and Johnstone binarize both size (strong/weak) and effort (hawk/dove), whereas Leimar allows for continuous intrinsic ability but keeps contest investment binary (hawk/dove). Fully implementing these models for a direct comparison is beyond the scope of our work here, but we briefly explore some of the relevant predictions of these models and how their explanatory mechanisms differ from our model. A summary of relevant differences can be found in Supplementary Material 2 (Table S1).

In Fawcett and Johnstone (2010), reducing aggression and ability to binary values results in highly stratified estimates, with individuals often becoming 'stuck' playing dove, after which they are unable to gain more information. In our model, because there are many scenarios where individuals win even with very low investment, unlucky individuals tend to recover from initial losses (Supplementary Material 2, Fig. S13c). In natural systems, subordinate individuals often engage with dominants (Guiaşu & Dunham, 1997; Hotta et al., 2021; Rowell, 1974), which could provide an opportunity for recovery, whereby stronger subordinates eventually supplant weaker dominant individuals (Favre et al., 2008; Samuels et al., 1987). In this vein, Fawcett and Johnstone found that their 'juvenile' individuals (comparable to naïve individuals in our model) tended to be hyperaggressive, because of the benefit of correctly identifying themselves as large early on and being able to invest accordingly. Empirically, it does not appear that juveniles are consistently more 'aggressive' (Bernstein et al., 1983; Groves, 1978; but see Baxter & Dukas, 2017; Fortunato & Earley, 2023), but juvenile individuals do famously engage in play fighting (Thompson, 1998; Thor & Holloway, 1984), and as mentioned, subordinate individuals tend to seek out agonistic encounters (Guiaşu & Dunham, 1997; Rowell, 1974), so it does appear that naïve and lower-ranking individuals highly value information and may engage in contests based on this. In our model there is no explicit reward for seeking information (or for winning contests for that matter). If we expanded our model to incentivize information gathering, we might expect individuals to 'overinvest' in potentially informative contests in order to gain valuable information, seeking out information-rich encounters despite the cost of more intense losses.

Whereas Fawcett and Johnstone provided a powerful implementation of Bayesian updating within a highly specific context, Leimar (2021) demonstrated a fully realized simulation of social dominance with a fundamentally different learning mechanism than Bayesian updating, i.e. actor–critic reinforcement learning. Like Fawcett and Johnstone, Leimar found that individuals evolve to be highly aggressive and, interestingly, that this aggression drives the observed loser effect bias. This is consistent with our model, in which shifting the effort function to be more 'hawkish' resulted in far stronger loser effects (Supplementary Material 2, Fig. S5), although the specific mechanisms differ: in Leimar's model, this shift was specified by each individual's policy gradient factor, which was inversely proportional to each individual's level of aggression, while in ours, it was a consequence of how an agent's effort function translated a shift in its estimate to behaviour (Supplementary Material 2, Fig. S21). This comparison is emblematic of these two approaches generally, in which distinct underlying mechanisms (Bayesian updating versus actor–critic reinforcement learning) arrive at similar outcomes, which would require careful modelling and very careful experimentation to distinguish.

Indeed, Bayesian updating and actor–critic learning (and the field of reinforcement learning generally) are closely related: because both algorithms are designed to accurately predict outcome, any effective reinforcement learning paradigm should at least approximate the 'correct' estimate that would be calculated by Bayesian updating. There is a rich field of literature on the application of reinforcement learning to behaviour (Dayan & Daw, 2008; Niv, 2009), how it relates to Bayesian updating (Courville et al., 2006; Kang et al., 2024) and (to a much lesser extent) how reinforcement learning relates to the winner–loser effect specifically (Leimar, 2021). While they are closely related, Bayesian updating functions slightly differently from standard reinforcement learning models (Le Pelley, 2004; Pearce & Hall, 1980; Rescorla & Wagner, 1972). This is because unlike most models of reinforcement learning, Bayesian updating usually establishes a representative model of reality (Courville et al., 2006; Vlassis et al., 2012), for example by estimating the actual size of individuals, based on information about the sizes of opponents and how contests are settled. For complex systems, modelling all relevant features can quickly become intractable (McNamara & Leimar, 2020), but assuming it is possible, Bayesian updating should provide the most accurate possible solution. This is not to say reinforcement learning is less effective, or even significantly less accurate, as prediction error approaches have been shown to closely approximate Bayesian updating (Kolter & Ng, 2009; Poupart et al., 2006). Given their similarity it seems unlikely that there are large, fundamental differences between the two mechanisms, at least in the general sense. There is some theoretical work that has been able to parse the specific predictions of Bayes updating from other forms of reinforcement learning (Courville et al., 2006; Kumaran et al., 2016), but in our view, the more relevant differences are in the practical application of these models. By abstracting the latent features of reality into the behavioural rule and updating mechanisms rather than having to explicitly infer them, reinforcement learning can be much faster to simulate and analytically more tractable. In contrast, Bayesian updating requires making decisions about a range of internal and environmental state variables in order for the prior and likelihood functions to be meaningful. Empiricists often have access to this information and need to identify testable predictions or perform power analyses. In these contexts, the computational run time of Bayesian updating is less of an issue, while the ability to directly input known environmental variables is highly useful.

While model usefulness is a practical question, Bayesian updating is also a hypothesis for the real underlying biological process, and it is useful to briefly discuss how our predictions of Bayesian updating relate to the current understanding of the biological bases of winner–loser effects. First, as mentioned, Bayes' theorem is theoretically the most accurate possible strategy for dealing with uncertain information. It thus represents the target that evolved neurophysiological mechanisms should approximate, whether or not they explicitly infer the underlying variables of interest (Higginson et al., 2018). For example, it is known that the endocrine system is tightly linked with winner–loser effects (Fuxjager, Montgomery, et al., 2011; Fuxjager, Oyegbile, et al., 2011; Zhou et al., 2018); even without cognitive 'learning', we might expect endocrine mechanisms to reflect the predictions of Bayesian updating with larger shifts in the levels of circulating testosterone and other hormones following a dramatic win or loss as well as encoding learning as the establishment of a new, stable state (e.g. high circulating testosterone or changes in receptor densities). Although hormones have been broadly implicated in winner–loser effects, winner–loser effects do not appear to be mediated solely by endocrine state (Fortunato & Earley, 2023; Rutte et al., 2006); for many species they are likely mediated at least in part by processes in the brain, where Bayesian updating could be implemented as the explicit mechanism of learning. The neural correlates of winner–loser effects are even less well defined, but many of the predicted changes in behaviour could be extended to predicted changes in the brain. For example, if we assume winner–loser effects are driven by Bayesian learning, we might expect expression in the brain following a contest to include genes involved in plasticity and learning, with brains adopting a new, stable state reflecting the change in self-assessment. We would also expect priors and likelihood functions to be encoded in the brain (Ashwood et al., 2020; Colombo & Seriès, 2012; Pouget et al., 2013), although identifying where and how these distributions are encoded would be a significant challenge. Regardless of the specific approach, fully identifying the mechanism behind winner–loser effects will require very careful theory paired with careful measures of behaviour and/or the underlying neurophysiological processes involved.

Conclusion

Given the inherent difficulty of inferring internal processes, and the limitations of any single study, what should empirical and theoretical biologists take from this work? First, Bayesian updating provides a plausible and intuitive framework for winner–loser effects, in which winner–loser effects follow naturally from individual attempts to accurately respond to an uncertain reality, and future experiments can test the predictions we have laid out here to identify novel features of winner–loser effects and assess whether they are consistent with Bayesian updating. More broadly, winner–loser effects should be thought of as a combination of multiple pathways, which interact to link prior experience to future outcomes. In this paper, we have explored a few potential mechanisms of individual assessment, updating mechanisms, contest strategies and how contests are settled, but there are likely additional mechanisms like social eavesdropping (Earley & Dugatkin, 2002; Tibbetts et al., 2020), motivation (O'Connor et al., 2015), short- (Zhou et al., 2018) and long-term (Gherardi, 2006) feedback on intrinsic ability and individual variation in behaviour (Laskowski et al., 2022) that influence winner–loser effects and dominance contests generally. All of this can be daunting, but experimentation and theory can isolate individual effects to generate at least some diagnostic behavioural predictions, providing models that can predict individual behaviour as well as powerful insights into the underlying behavioural mechanisms.

Author Contributions

Kate L. Laskowski: Writing – review & editing, Funding acquisition, Conceptualization. **Ammon Perkes:** Writing – review & editing, Writing – original draft, Visualization, Methodology, Funding acquisition, Formal analysis, Conceptualization.

Data Availability

Data are available as Supplementary Material.

Declaration of Interest

The authors declare that they have no conflicts of interest.

Acknowledgments

We thank Jon Aguiñaga, Aviva Blonder and Jeffrey Schank for their valuable feedback on the manuscript, as well as the referees for their careful and insightful suggestions. This work was funded by the U.S. National Science Foundation (NSF PRFB 2209270; IOS 2100625) and the National Institute of General Medical Sciences of the U.S. National Institutes of Health (1R35GM153309-01). The content is solely the responsibility of the authors and does not necessarily represent the official views of the U.S. National Institutes of Health.

Supplementary Material

Supplementary material associated with this article is available at https://doi.org/10.1016/j.anbehav.2025.123191.

References

Arnott, G., & Elwood, R. W. (2008). Information gathering and decision making about resource value in animal contests. *Animal Behaviour, 76*(3), 529–542. https://doi.org/10.1016/j.anbehav.2008.04.019

Arnott, G., & Elwood, R. W. (2009). Assessment of fighting ability in animal contests. *Animal Behaviour, 77*(5), 991–1004. https://doi.org/10.1016/j.anbehav.2009.02.010

Ashwood, Z. C., Roy, N. A., Bak, J. H., & Pillow, J. W. (2020). Inferring learning rules from animal decision-making. *Advances in Neural Information Processing Systems, 33*, 3442–3453.

Balph, M. H. (1979). Flock stability in relation to social dominance and agonistic behavior in wintering dark-eyed juncos. *Auk: Ornithological Advances, 96*(4), 714–722. https://doi.org/10.1093/auk/96.4.714

Baxter, C. M., & Dukas, R. (2017). Life history of aggression: Effects of age and sexual experience on male aggression towards males and females. *Animal Behaviour, 123*, 11–20. https://doi.org/10.1016/j.anbehav.2016.10.022

Beacham, J. L. (1988). The relative importance of body size and aggressive experience as determinants of dominance in pumpkinseed sunfish, Lepomis gibbosus. *Animal Behaviour, 36*(2), 621–623. https://doi.org/10.1016/S0003-3472(88)80042-3

Benincasa, M. D., Earley, R. L., & Hamilton, I. M. (2023). Cumulative experience influences contest investment in a social fish. *Behavioral Ecology, 34*(6), 1076–1086. https://doi.org/10.1093/beheco/arad078

Bernstein, I., Williams, L., & Ramsay, M. (1983). The expression of aggression in Old World monkeys. *International Journal of Primatology, 4*(2), 113–125. https://doi.org/10.1007/BF02743753

Bierbach, D., Klein, M., Sassmannshausen, V., Schlupp, I., Riesch, R., Parzefall, J., & Plath, M. (2012). Divergent evolution of male aggressive behaviour: Another reproductive isolation barrier in extremophile poeciliid fishes? *International Journal of Evolutionary Biology, 2012*, Article 148745. https://doi.org/10.1155/2012/148745

Bonabeau, E., Theraulaz, G., & Deneubourg, J. L. (1996). Mathematical model of self-organizing hierarchies in animal societies. *Bulletin of Mathematical Biology, 58*(4), 661–717. https://doi.org/10.1016/0092-8240(95)00364-9

Bonabeau, E., Theraulaz, G., & Deneubourg, J. L. (1999). Dominance orders in animal societies: The self-organization hypothesis revisited. *Bulletin of Mathematical Biology, 61*(4), 727–757. https://doi.org/10.1006/bulm.1999.0108

Chase, I. D., Bartolomeo, C., & Dugatkin, L. A. (1994). Aggressive interactions and inter-contest interval: How long do winners keep winning? *Animal Behaviour, 48*(2), 393–400. https://doi.org/10.1006/anbe.1994.1253

Chase, I. D., Coelho, D., Lee, W., Mueller, K., & Curley, J. P. (2022). Networks never rest: An investigation of network evolution in three species of animals. *Social Networks, 68*, 356–373. https://doi.org/10.1016/j.socnet.2021.09.002

Chase, I. D., Tovey, C., Spangler-Martin, D., & Manfredonia, M. (2002). Individual differences versus social dynamics in the formation of animal dominance hierarchies. *Proceedings of the National Academy of Sciences of the United States of America, 99*(8), 5744–5749. https://doi.org/10.1073/pnas.082104199

Colombo, M., & Seriès, P. (2012). Bayes in the brain - on Bayesian modelling in neuroscience. *British Journal for the Philosophy of Science, 63*(3), 697–723. https://doi.org/10.1093/bjps/axr043

Courville, A. C., Daw, N. D., & Touretzky, D. S. (2006). Bayesian theories of conditioning in a changing world. *Trends in Cognitive Sciences, 10*(7), 294–300. https://doi.org/10.1016/j.tics.2006.05.004

Dayan, P., & Daw, N. D. (2008). Decision theory, reinforcement learning, and the brain. *Cognitive, Affective, & Behavioral Neuroscience, 8*(4), 429–453. https://doi.org/10.3758/CABN.8.4.429

Dewsbury, D. A. (1988). Copulatory behavior as courtship communication. *Ethology, 79*, 218–234. https://doi.org/10.1111/j.1439-0310.1988.tb00712.x

Dugatkin, L. A. (1997). Winner and loser effects and the structure of dominance hierarchies. *Behavioral Ecology, 8*(6), 583–587. https://doi.org/10.1093/beheco/8.6.583

Earley, R. L., & Dugatkin, L. A. (2002). Eavesdropping on visual cues in green swordtail (*Xiphophorus helleri*) fights: A case for networking. *Proceedings of the Royal Society B: Biological Sciences, 269*(1494), 943–952. https://doi.org/10.1098/rspb.2002.1973

Enquist, M., & Leimar, O. (1983). Evolution of fighting behaviour: Decision rules and assessment of relative strength. *Journal of Theoretical Biology, 102*(3), 387–410. https://doi.org/10.1016/0022-5193(83)90376-4

Enquist, M., Leimar, O., Ljungberg, T., Mallner, Y., & Segerdahl, N. (1990). A test of the sequential assessment game: Fighting in the cichlid fish *Nannacara anomala*. *Animal Behaviour, 40*(1), 1–14. https://doi.org/10.1016/S0003-3472(05)80660-8

Favre, M., Martin, J. G. A., & Festa-Bianchet, M. (2008). Determinants and life-history consequences of social dominance in bighorn ewes. *Animal Behaviour, 76*(4), 1373–1380. https://doi.org/10.1016/j.anbehav.2008.07.003

Fawcett, T. W., & Johnstone, R. A. (2010). Learning your own strength: Winner and loser effects should change with age and experience. *Proceedings of the Royal Society B: Biological Sciences, 277*(1686), 1427–1434. https://doi.org/10.1098/rspb.2009.2088

Fortunato, J. A., & Earley, R. L. (2023). Age-dependent genetic variation in aggression. *Biology Letters, 19*(1), Article 20220456. https://doi.org/10.1098/rsbl.2022.0456

Fuxjager, M. J., Montgomery, J. L., & Marler, C. A. (2011a). Species differences in the winner effect disappear in response to post-victory testosterone manipulations. *Proceedings of the Royal Society B: Biological Sciences, 278*(1724), 3497–3503. https://doi.org/10.1098/rspb.2011.0301

Fuxjager, M. J., Oyegbile, T. O., & Marler, C. A. (2011b). Independent and additive contributions of postvictory testosterone and social experience to the development of the winner effect. *Endocrinology, 152*(9), 3422–3429. https://doi.org/10.1210/en.2011-1099

Gherardi, F. (2006). Fighting behavior in hermit crabs: The combined effect of resource-holding potential and resource value in *Pagurus longicarpus*. *Behavioral Ecology and Sociobiology, 59*(4), 500–510. https://doi.org/10.1007/s00265-005-0074-z

Groves, S. (1978). Age-related differences in ruddy turnstone foraging and aggressive behavior. *Auk: Ornithological Advances, 95*(1), 95–103. https://doi.org/10.2307/4085499

Guiaşu, R. C., & Dunham, D. W. (1997). Initiation and outcome of agonistic contests in male form I *Cambarus robustus* Girard, 1852 crayfish (Decapoda, Cambaridae). *Crustaceana, 70*(4), 480–496. https://doi.org/10.1163/156854097x00069

Hemelrijk, C. K. (2000). Towards the integration of social dominance and spatial structure. *Animal Behaviour, 59*(5), 1035–1048. https://doi.org/10.1006/anbe.2000.1400

Hickey, J., & Davidsen, J. (2019). Self-organization and time-stability of social hierarchies. *PLoS One, 14*(1), Article e0211403. https://doi.org/10.1371/journal.pone.0211403

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: Biological Sciences, 285*(1871). https://doi.org/10.1098/rspb.2017.2411. Article 20172411.

Hock, K., & Huber, R. (2006). Modeling the acquisition of social rank in crayfish: Winner and loser effects and self-structuring properties. *Behaviour, 143*(3), 325–346. https://doi.org/10.1163/156853906775897914

Hotta, T., Awata, S., Jordan, L. A., & Kohda, M. (2021). Subordinate fish mediate aggressiveness using recent contest information. *Frontiers in Ecology and Evolution, 9*, Article 685907. https://doi.org/10.3389/fevo.2021.685907

Hsu, Y., Earley, R. L., & Wolf, L. L. (2006). Modulation of aggressive behaviour by fighting experience: Mechanisms and contest outcomes. *Biological Reviews of the Cambridge Philosophical Society, 81*(1), 33–74. https://doi.org/10.1017/S146479310500686X

Hsu, Y., Lee, S. P., Chen, M. H., Yang, S. Y., & Cheng, K. C. (2008). Switching assessment strategy during a contest: Fighting in killifish *Kryptolebias marmoratus*. *Animal Behaviour, 75*(5), 1641–1649. https://doi.org/10.1016/j.anbehav.2007.10.017

Hsu, Y., & Wolf, L. L. (1999). The winner and loser effect: Integrating multiple experiences. *Animal Behaviour, 57*(4), 903–910. https://doi.org/10.1006/anbe.1998.1049

Hsu, Y., & Wolf, L. L. (2001). The winner and loser effect: What fighting behaviours are influenced? *Animal Behaviour, 61*(4), 777–786. https://doi.org/10.1006/anbe.2000.1650

Huang, S. P., Yang, S. Y., & Hsu, Y. (2011). Persistence of winner and loser effects depends on the behaviour measured. *Ethology, 117*(2), 171–180. https://doi.org/10.1111/j.1439-0310.2010.01856.x

Jackson, W. M., & Winnegrad, R. L. (1988). Linearity in dominance hierarchies: A second look at the individual attributes model. *Animal Behaviour, 36*(4), 1237–1240. https://doi.org/10.1016/S0003-3472(88)80086-1

Kang, P., Tobler, P. N., & Dayan, P. (2024). Bayesian reinforcement learning: A basic overview. *Neurobiology of Learning and Memory, 211*, Article 107924. https://doi.org/10.1016/J.NLM.2024.107924

Kolter, J. Z., & Ng, A. Y. (2009). Near-Bayesian exploration in polynomial time. In A. P. Danyluk, L. Bottou, & M. L. Littman (Eds.), *Proceedings of the 26th International Conference on Machine Learning, ICML, Montreal, Quebec, Canada, 14–18 June 2009* (pp. 513–520). Association for Computing Machinery.

Koops, M. A., & Grant, J. W. A. (1993). Weight asymmetry and sequential assessment in convict cichlid contests. *Canadian Journal of Zoology, 71*(3), 475–479. https://doi.org/10.1139/z93-068

Kumaran, D., Banino, A., Blundell, C., Hassabis, D., & Dayan, P. (2016). Computations underlying social hierarchy learning: Distinct neural mechanisms for updating and representing self-relevant information. *Neuron, 92*(5), 1135–1147. https://doi.org/10.1016/j.neuron.2016.10.052

Kura, K., Broom, M., & Kandler, A. (2016). A game-theoretical winner and loser model of dominance hierarchy formation. *Bulletin of Mathematical Biology, 78*(6), 1259–1290. https://doi.org/10.1007/s11538-016-0186-9

Lan, Y. T., & Hsu, Y. (2011). Prior contest experience exerts a long-term influence on subsequent winner and loser effects. *Frontiers in Zoology, 8*, Article 28. https://doi.org/10.1186/1742-9994-8-28

Landau, H. G. (1951). On dominance relations and the structure of animal societies: I. Effect of inherent characteristics. *Bulletin of Mathematical Biophysics, 13*(1), 1–19. https://doi.org/10.1007/BF02478336

Laskowski, K. L., Chang, C. C., Sheehy, K., & Aguiñaga, J. (2022). Consistent individual behavioral variation: What do we know and where are we going? *Annual Review of Ecology, Evolution, and Systematics, 53*, 161–182. https://doi.org/10.1146/annurev-ecolsys-102220-011451

Laskowski, K. L., Wolf, M., & Bierbach, D. (2016). The making of winners (and losers): How early dominance interactions determine adult social structure in a clonal fish. *Proceedings of the Royal Society B: Biological Sciences, 283*(1830), Article 20160183. https://doi.org/10.1098/rspb.2016.0183

Le Pelley, M. E. (2004). The role of associative history in models of associative learning: A selective review and a hybrid model. *Quarterly Journal of*

Experimental Psychology Section B Comparative and Physiological Psychology, 57(3), 193–243. https://doi.org/10.1080/02724990344000141

Leimar, O. (2021). The evolution of social dominance through reinforcement learning. *American Naturalist, 197*(5), 560–575. https://doi.org/10.1086/713758

Leimar, O., Austad, S., & Enquist, M. (1991). A test of the sequential assessment game: Fighting in the bowl and doily spider *Frontinella pyramitela*. *Evolution, 45*(4), 862–874. https://doi.org/10.1111/j.1558-5646.1991.tb04355.x

Leimar, O., & Bshary, R. (2022). Reproductive skew, fighting costs and winner–loser effects in social dominance evolution. *Journal of Animal Ecology, 91*(5), 1036–1046. https://doi.org/10.1111/1365-2656.13691

Luttbeg, 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

Marden, J. H., & Waage, J. K. (1990). Escalated damselfly territorial contests are energetic wars of attrition. *Animal Behaviour, 39*(5), 954–959. https://doi.org/10.1016/S0003-3472(05)80960-1

Maynard Smith, J. (1974). The theory of games and the evolution of animal conflicts. *Journal of Theoretical Biology, 47*(1), 209–221. https://doi.org/10.1016/0022-5193(74)90110-6

Maynard Smith, J., & Parker, G. A. (1976). The logic of asymmetric contests. *Animal Behaviour, 24*(1), 159–175. https://doi.org/10.1016/S0003-3472(76)80110-8

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

McNamara, J. M., & Leimar, O. (2020). *Game theory in biology: Concepts and frontiers*. Oxford University Press.

Mesterton-Gibbons, M. (1999). On the evolution of pure winner and loser effects: A game-theoretic model. *Bulletin of Mathematical Biology, 61*(6), 1151–1186. https://doi.org/10.1006/bulm.1999.0137

Mesterton-Gibbons, M., Dai, Y., & Goubault, M. (2016). Modeling the evolution of winner and loser effects: A survey and prospectus. *Mathematical Biosciences, 274*, 33–44. https://doi.org/10.1016/j.mbs.2016.02.002

Mesterton-Gibbons, M., Marden, J. H., & Dugatkin, L. A. (1996). On wars of attrition without assessment. *Journal of Theoretical Biology, 181*(1), 65–83. https://doi.org/10.1006/jtbi.1996.0115

Niv, Y. (2009). Reinforcement learning in the brain. *Journal of Mathematical Psychology, 53*(3), 139–154. https://doi.org/10.1016/J.JMP.2008.12.005

O'Connor, C. M., Reddon, A. R., Ligocki, I. Y., Hellmann, J. K., Garvy, K. A., Marsh-Rollo, S. E., Hamilton, I. M., & Balshine, S. (2015). Motivation but not body size influences territorial contest dynamics in a wild cichlid fish. *Animal Behaviour, 107*, 19–29. https://doi.org/10.1016/j.anbehav.2015.06.001

Okasha, S. (2013). The evolution of Bayesian updating. *Philosophy of Science, 80*(5), 745–757. https://doi.org/10.1086/674058

Olsson, O. (2006). Decisions under uncertainty: Bayesian foraging - a meeting held at Lund University in August 2003. *Oikos, 112*(2), 241–242. https://doi.org/10.1111/j.0030-1299.2006.14383.x

Parker, G. A. (1974). Assessment strategy and the evolution of animal conflicts. *Journal of Theoretical Biology, 47*(1), 223–243.

Pearce, J. M., & Hall, G. (1980). A model for Pavlovian learning: Variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological Review, 87*(6), 532–552. https://doi.org/10.1037/0033-295X.87.6.532

Pouget, A., Beck, J. M., Ma, W. J., & Latham, P. E. (2013). Probabilistic brains: Knowns and unknowns. *Nature Neuroscience, 16*(9), Article 11701178. https://doi.org/10.1038/nn.3495

Poupart, P., Vlassis, N., Hoey, J., & Regan, K. (2006). An analytic solution to discrete Bayesian reinforcement learning. In W. W. Cohen, & A. W. Moore (Eds.), *Proceedings of the 23rd International Conference on Machine Learning, 25–29 June 2006, Carnegie Mellon University* (pp. 697–704). Pittsburgh, Pennsylvania: Association for Computing Machinery. https://doi.org/10.1145/1143844.1143932.

Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning : The effectiveness of reinforcement and non-reinforcement. In A. H. Black, & W. F. Prokasy (Eds.), *Classical conditioning II: Current theory and research* (pp. 64–69). Appleton-Century-Crofts.

Rowell, T. E. (1974). The concept of social dominance. *Behavioral Biology, 11*(2), 131–154. https://doi.org/10.1016/S0091-6773(74)90289-2

Rutte, C., Taborsky, M., & Brinkhof, M. W. G. (2006). What sets the odds of winning and losing? *Trends in Ecology & Evolution, 21*(1), 16–21. https://doi.org/10.1016/j.tree.2005.10.014

Samuels, A., Silk, J. B., & Altmann, J. (1987). Continuity and change in dominance relations among female baboons. *Animal Behaviour, 35*(3), 785–793. https://doi.org/10.1016/S0003-3472(87)80115-X

Senar, J. C., Camerino, M., & Metcalfe, N. B. (1990). Familiarity breeds tolerance: The development of social stability in flocking siskins (*Carduelis spinus*). *Ethology, 85*(1), 13–24. https://doi.org/10.1111/j.1439-0310.1990.tb00381.x

Simons, N. D., Michopoulos, V., Wilson, M., Barreiro, L. B., & Tung, J. (2022). Agonism and grooming behaviour explain social status effects on physiology and gene regulation in rhesus macaques. *Philosophical Transactions of the Royal Society B: Biological Sciences, 377*(1845), Article 20210132. https://doi.org/10.1098/rstb.2021.0132

Snyder-Mackler, N., Burger, J. R., Gaydosh, L., Belsky, D. W., Noppert, G. A., Campos, F. A., Bartolomucci, A., Yang, Y. C., Aiello, A. E., O'Rand, A., Harris, K. M., Shively, C. A., Alberts, S. C., & Tung, J. (2020). Social determinants of health and survival in humans and other animals. *Science, 368*(6493), Article eaax9553. https://doi.org/10.1126/science.aax9553

Thompson, K. V. (1998). Self assessment in juvenile play. In M. Bekoff, & J. Byers (Eds.), *Animal play: Evolutionary, comparative and ecological perspectives* (pp. 183–204). Cambridge University Press. https://doi.org/10.1017/cbo9780511608575.010.

Thor, D. H., & Holloway, W. R. (1984). Developmental analyses of social play behavior in juvenile rats. *Bulletin of the Psychonomic Society, 22*(6), 587–590. https://doi.org/10.3758/BF03333916

Tibbetts, E. A., Pardo-Sanchez, J., & Weise, C. (2022). The establishment and maintenance of dominance hierarchies. *Philosophical Transactions of the Royal Society B: Biological Sciences, 377*(1845), Article 20200450. https://doi.org/10.1098/rstb.2020.0450

Tibbetts, E. A., Wong, E., & Bonello, S. (2020). Wasps use social eavesdropping to learn about individual rivals. *Current Biology, 30*(15), 3007–3010. https://doi.org/10.1016/j.cub.2020.05.053

Valone, T. J. (2006). Are animals capable of Bayesian updating? An empirical review. *Oikos, 112*(2), 252–259. https://doi.org/10.1111/j.0030-1299.2006.13465.x

Van Doorn, G. S., Hengeveld, G. M., & Weissing, F. J. (2003). The evolution of social dominance I: Two-player models. *Behaviour, 140*(10), 1305–1332. https://doi.org/10.1163/156853903771980602

Vlassis, N., Ghavamzadeh, M., Mannor, S., & Poupart, P. (2012). Bayesian reinforcement learning. In M. Wiering, & M. van Otterlo (Eds.), *Reinforcement learning. Adaptation, learning, and optimization, 12* (pp. 359–386). Springer. https://doi.org/10.1007/978-3-642-27645-3_11

Whitehouse, M. E. A. (1997). Experience influences male–male contests in the spider *Argyrodes antipodiana* (Theridiidae: Araneae). *Animal Behaviour, 53*(5), 913–923. https://doi.org/10.1006/anbe.1996.0313

Yan, J. L., Smith, N. M. T., Filice, D. C. S., & Dukas, R. (2024). Winner and loser effects: A meta-analysis. *Animal Behaviour, 216*, 15–22. https://doi.org/10.1016/J.ANBEHAV.2024.07.014

Zhou, T., Sandi, C., & Hu, H. (2018). Advances in understanding neural mechanisms of social dominance. *Current Opinion in Neurobiology, 49*, 99–107. https://doi.org/10.1016/j.conb.2018.01.006