Bayesian updating for self-assessment explains social dominance and the winner effect

Abstract
Dominance interactions determine access to scarce resources and are thus major contributors to fitness in many social species. Although in some systems, dominance is a direct reflection of intrinsic resource holding potential, often there can be mismatches between an animal's expected dominance (based on body condition, fighting ability, social connections, etc.) and the actual results of dominance contests and their position in a hierarchy. This can be driven in part by the winner/loser effect, in which prior results bias the outcome of future contests. In many systems, the winner/loser effect is driven not by a measurable change in relative resource holding potential, but by the aggressiveness of experienced individuals, suggesting a shift in self-assessment based on newly gained information. This updating of a prior estimate is ideally described by Bayesian updating, which predicts animal behavior in other contexts, such as foraging and mate selection. Despite the obvious relevance, the winner effect has never been modeled as a Bayesian process. Here we show that Bayesian updating provides a mechanism to explain why winner effects form and makes clear predictions for the behavior of individuals and social groups. We implement an agent-based model to compare Bayesian updating to other possible strategies. We first show that Bayesian updating reproduces empirical results of the behavior of individuals and groups in dominance interactions. We then provide a series of testable predictions that can be used in empirical work to distinguish Bayesian updating from other potential mechanisms. Our work demonstrates the utility of Bayesian updating as a mechanism to explain how changes in self-assessment can determine the outcome of repeated social interactions.

Introduction

Nearly all social groups have measurable dominance hierarchies (Ellis 1995; Shizuka and McDonald 2015; Strauss et al. 2022). These hierarchies can be defined by consistent imbalances in pair-wise aggression and subordination (Rowell 1974; Drews 1993). In the simplest cases, dominance relationships are a measure of contest results, with the stronger or otherwise 'better' individuals consistently winning over weaker ones. If all outcomes were determined by the intrinsic ability to win a fight—often termed Resource Holding Potential (Parker 1974) — then dominance networks would be perfectly linear and predictable. And while dominance hierarchies often are stable and linear, their outcomes, that is, the positioning of different individuals within the hierarchy, are often unpredictable (Landau 1951; Chase et al. 2002, 2022). This is because, in reality, contest outcomes are often stochastic [reviewed in (Strauss et al. 2022)], and are influenced by a combination of factors, including motivation (Arnott and Elwood 2008; O'Connor et al. 2015), strategy (Parker 1974), and social feedback (Chase et al. 2002). Dominance contests also tend to be ritualized to minimize the cost of competition (Searcy and Nowicki 2010; Green and Patek 2018; Tibbetts et al. 2022), such that in almost all cases, contests end not with death, but when one individual yields (Rowell 1974). This dichotomy—often unpredictable but unusually stable outcomes—highlights the gap between our current understanding of the mechanisms of dominance and the reality of natural systems. Despite the importance of social dominance for individual and group outcomes (Dewsbury 1988; Snyder-Mackler et al. 2020; Simons et al. 2022), we still lack models that can predict the formation of social networks over time by incorporating the additional social mechanisms that reinforce or even supersede differences in intrinsic resource holding potential (Chase et al. 2002; Tibbetts et al. 2022).

Winner and loser effects have been hypothesized to be a key mechanism in reinforcing the linearity and stability of dominance networks (Dugatkin 1997; Bonabeau et al. 1999; Hemelrijk 2002). 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 failure leads to more failure [reviewed in (Hsu et al. 2006)]. Although these two phenomena are similar, prior work has shown the loser effect tends to be stronger than the winner effect [reviewed in (Rutte et al. 2006)], and each has been shown to occur independent of the other [(McDonald et al. 1968; Goubault and Decuignière 2012) reviewed in (Hsu et al. 2006)].

Winner/loser effects are typically modulated by a change in an individual's aggressive behavior [reviewed in Hsu, 2006], likely indicating a shift in self-assessment (Rutte et al. 2006). However, how these shifts in assessment occur is unknown. This is an important knowledge gap because without an understanding of the cognitive mechanism, we cannot predict how long such winner effects may last. Indeed, there is contrasting evidence in the literature for the duration of winner effects. Winner/loser effects are most often thought of as short-term phenomena, lasting a limited time (i.e., minutes to days) (Chase et al. 1994; Johnsson 1997) or until some external disruption occurs (Stevenson et al. 2000). This is especially evident in the observations of recency biases, where the most recent win or loss can overrules previous outcomes (Hsu and Wolf 1999). Contrasting these findings, however, some studies have found persistent

winner/loser effects influencing behavior months after the initial conflict (Lan and Hsu 2011; Laskowski et al. 2016). If changes in self-assessment are what underpins winner/loser effects, then a mechanistic understanding how those changes occur can help reconcile these seemingly disparate results about the durability of winner effects.

There was an early boom of relevant theoretical work exploring the formation of dominance hierarchies. Much of this work focused on whether the existence of a winner/loser effects could explain the observed linearity of networks [reviewed in (Lindquist and Chase 2009)]. Current models generally use a linear shift in estimate (and/or ability), (e.g. (Dugatkin 1997; Bonabeau et al. 1999; Hickey and Davidsen 2019)), though other models have developed more complex methods of contest assessment ((Hemelrijk 2002; Van Doorn et al. 2003; Hock and Huber 2006)). All of these models combine intrinsic ability (i.e. resource holding potential) with all other factors (motivation, self-assessment, etc.), into a single variable, which is then modified following a win or a loss. While this increases the ease of modeling, this also makes it impossible to separate the impact of intrinsic ability from informed self-assessment. In other words, these models were useful at describing the impacts of winner/loser effects but have largely treated the mechanisms underlying those winner/loser effects as a black box. This is important because as Lindquist & Chase show, current models often fail to capture the dynamics of social networks, and are not designed to incorporate the range of social-cognition necessary to account for group behavior. A more explicit modeling of the mechanisms underlying winner effects may be the way forward to better explain the empirical complexity of social networks.

If winner/loser effects are mediated through a shift in self-assessment following a contest, then Bayesian updating provides an optimal modeling framework to explain this process. Bayesian updating is a powerful framework for modeling decision-making in animals (McNamara et al. 2006), characterized by an informed prior estimate which is updated by the likelihood of the observed effect. Mary Whitehouse noted that a Bayesian model might be able explain her observations of winner/loser effects in social spiders (Whitehouse 1997), however, to our knowledge, there is no model that treats winner/loser effects as a process of Bayesian learning. In general, animals are expected to at least approximate a Bayesian approach (Okasha 2013), since Bayesian updating is the mathematically optimal strategy to refine estimates (i.e., calculate conditional probability). Indeed, Bayesian models have been shown to match empirical observations of behavior within foraging (Olsson 2006) and mate choice (Luttbeg 1996). In the context of social dominance, after each fight, individuals would be expected to shift their estimated self-assessment of their own intrinsic ability based on the likelihood of having won against an opponent (calculated across the distribution of possible own-RHP values). By modeling the winner effect as a process designed to estimate intrinsic ability, Bayesian updating answers both the how and the why of winner/loser effects: winner/loser effects are an expected side-effect of the central process of animals assessing and establishing their social status. It is important then to identify the predictions of a Bayesian model and understand how we might distinguish Bayesian updating from alternative, possible mechanisms.

Here, we implement an agent-based model of winner/loser effects—and dominance formation generally—as a process of Bayesian updating. To assess the value of our model, we first confirm that it recreates known empirical features of dominance networks and winner effects. Specifically, we investigate whether our agent-based Bayesian model produces hierarchies that are 1) linear, 2) stable and 3) mediated through low effort interactions. We also test whether this model re-creates 4) recency biases and 5) amplified loser effects as shown in the empirical literature. We then compare the predictions of our Bayesian model with the predictions of alternative models that use different updating mechanisms. Here we quantify the intensity of fights as a potential empirical measure that can help distinguish whether individuals are using Bayesian updating, compared to a linear shift in estimate [as in (Dugatkin 1997) and (Bonabeau et al. 1999), or to a fixed estimate, in which no change in self-assessment occurs. Then, assuming individuals are using Bayesian updating, we quantify how fight intensity is expected to change if individuals are using bayesian (i.e., informed) versus uninformed priors. Our model provides a mechanism to reconcile disparate empirical results to understand why winner/loser effects are sometimes durable and other times, transient. We also generate a series of predictions that can form the backbone of future experimentation to identify the mechanism by which dominance is established.

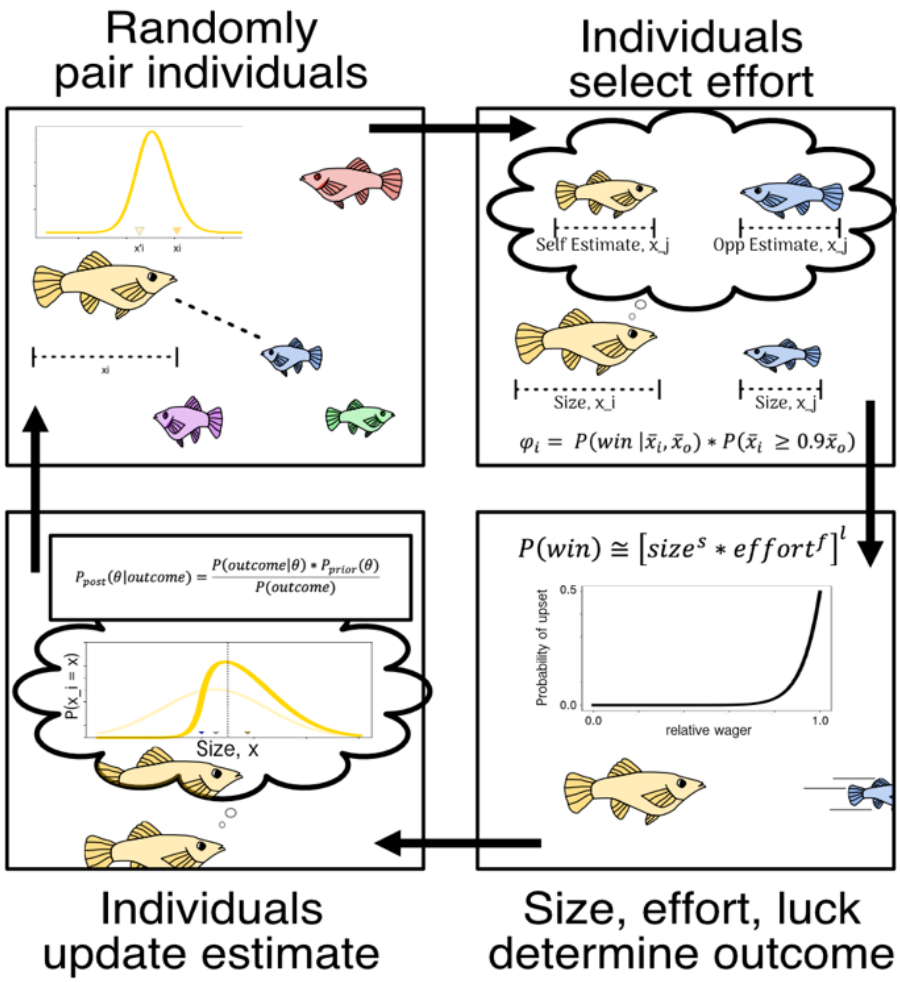

Figure 1: Graphical overview of conceptual framework. Top left. Individuals with fixed size (unknown to self) and an estimated self-assessment are paired randomly. Top right, each individual estimate own and opponent size to determine their effort, based on their estimated probability of winning. Bottom right, The wagers of each agent, based on the relative size and effort of the smaller individual, determine the probability of wining (calculated as the probability of an upset by the underdog in a given contest). Bottom left. Based on the outcome, agents update their self-assessment, multiplying their prior assessment by the likelihood of the outcome calculated across all possible sizes. This process is repeated many times for all possible pairings.

Overview of the model

Our model essentially consists of four steps, described graphically in Figure 1. First, we generate a group of Agents with random (normally distributed) intrinsic ability and self-assessments that ability which we will call their 'estimates' (Fig 1, top-left). For simplicity, we will refer to this intrinsic ability as "size", although in most natural systems other factors also contribute. These agents are then randomly paired and allowed to interact. In each contest, they select their effort, based on their self-assessed probability of winning (based on their initial prior; Fig 1, top-right). Relative size and effort combine to form a "wager" for each agent. Contest outcomes are then determined based the relative wagers of each participant (fig 1, bottom-right). Following the contest, each agent updates their self-assessment based on their Bayesian prior and likelihood (Fig 1, bottom-left). This process is then repeated. Each step of the model is described in detail below.

Table of Abbreviations:

Symbol	Default	Range	Explanation
$x: \ (x, x_i, x_o, \bar{x}_i)$	Random	$1 \leq x \leq 100$	Size of an individual. x_i represents actual size for individual i, while x_o is the actual size of opponent. $\bar{x}$ denotes an estimated size
μ_x	50	$1 \leq x \leq 100$	The mean of the population from which individuals are drawn
σ_x	10*	$0 \leq \sigma_x < \infty$	Standard deviation of population from which individuals are drawn. At $\sigma = \infty$, uniform distribution is used.
$\phi: \ (\phi, \phi_i, \phi_o)$	None	$0 \leq \phi \leq 1$	Effort invested in a contest, as with size, subscripts distinguish effort for the focal individual ϕ_i from opponent's, ϕ_o
s	1*	$0 \leq s < \infty$	Relative size dampener, acts as an exponent such that as s increases, so does the importance of relative size
f	1*	$0 \leq f < \infty$	Effort dampener. Acts as an exponent, such that as f increases, so does marginal benefit of increased effort relative to the opponent
l	13*	$0 \leq l < \infty$	Relative wager dampener, essentially the importance of luck, acting as an exponent on relative wager, such that as l increases, the probability of upsets decreases

σ_a	10*	$0 \leq \sigma_a < \infty$	*awareness*, the accuracy of individuals' starting guesses for their self-assessment.
σ_c	10*	$0 \leq \sigma_c < \infty$	*acuity*, the accuracy of individuals when guessing the size of opponents prior to contests
$P(\theta)$	-	-	Used to denote the prior, an estimated, discrete distribution of the probability of being some size.
Φ	-	-	Used to denote the cumulative distribution function used to calculate confidence
p	1.1	$0 \leq x \leq 100$	Proportional shift after a fight, used in linear updating
k	0	$0 \leq x \leq 100$	Constant shift after a fight, used in linear updating
r	10	$1 \leq x < \infty$	Proportional shift after a fight, used in linear updating

* In our model, for ease of use, the actual input parameters are entered in ranges from $(1, -1)$, and then scaled using (COTANGENT EQUATION) to be between $(0, \infty)$ for use in the equations (such that when input variable $L = 0$, equation parameter $l = 1$. We use a similar scaling function for the standard deviations, so they it can be input from $(0, -1)$

Agents

In order to model the development of self-assessment, we first randomly generated agents with some intrinsic ability, hereafter simplified as *size*, drawn from a normal distribution, such that

Equation 1: $$x_i \sim N(\mu_x, \sigma_x)$$

Each agent had an estimate of their own size, $\bar{x}$, formed by drawing from a normal distribution centered around their actual size:

Equation 2a: $$\bar{x}_i \sim N(\mu_x, \sigma_x)$$

Individuals then calculated a starting prior as a normal distribution centered on their estimate

Equation 2b: $$\mathrm{P_r}(\theta) = N(\bar{x}_i, \sigma_a)$$

All agents used the same *effort strategy*—the function for choosing effort in any given fight—and *update strategy*—the function for updating their self-assessment, both described below. Simulated agents were grouped in sets, comprising *n* agents, and paired across r rounds of (n-1) fights with every other agent, for a total of (n-$1 * r$) fights, with the order of fights being shuffled for each round.

Effort

The purpose of our model was to explore updating strategies, but in order to do this, we needed a functional effort strategy. Empirical observations suggest that individuals generally invest proportional to their probability of victory [ref], so we designed the effort strategy to match this. To determine their effort, each agent first estimated the size of their opponent based on their *acuity,* σ_c ,

Equation 3a: $$\bar{x}_o \sim N(x_o, \sigma_c)$$

Agents then set their effort based on their estimated probability of winning—assuming equal effort—multiplied by their confidence of being with 10% of the size of the opponent, in order to account for the degree error around their estimates.

Equation 3b:

$$\varphi_i = P(win\ |\bar{x}_i, \bar{x}_o) * P(\bar{x}_i \ \geq 0.9\bar{x}_o)$$

Here effort is simply equal to the probability of winning—given an agent's estimates about their own and their opponent's sizes—multiplied by the probability that those estimates are approximately correct.

Because the probability of winning is a function of the relative value of the smaller wager, in equation 3b, the probability of winning $P(win\ |\bar{x}_i, \bar{x}_o)$ varies depending on whether an agent estimates itself to be bigger or smaller than its opponent,

Equation 3c:

$$\varphi_i = \left(\frac{\bar{x}_i}{\bar{x}_o}\right)^{sl} * P(\bar{x}_i \ \geq 0.9\bar{x}_o), \qquad if\ \bar{x}_i \leq \ \bar{x}_o$$

$$\varphi_i = \left[1 - \left(\frac{\bar{x}_o}{\bar{x}_i}\right)^{sl}\right] * P(\bar{x}_i \ \geq 0.9\bar{x}_o), if\ \bar{x}_i > \bar{x}_o$$

where, s and l are simulation parameters that account for the importance of size and luck, respectively. This parallels the actual calculation of contest output, but here individuals do not account for effort, or equivalently, assume that they and their opponent will invest the same amount, removing effort as a factor.

The second term in Eq. 3b and 3c, $P(\bar{x}_i \ \geq 0.9\bar{x}_o)$, is a confidence correction representing the probability that the individual is within at least 10% of the opponent size, and is calculated using a cumulative distribution function of the combined normal distributions:

Equation 3c

$$P(x_i \geq 0.9x_o) = \Phi_{\mu_\delta, \sigma_\delta}(-0.1 * x_o)$$

with

Equation 3d

$$\mu_\delta = x_i - x_o$$

Equation 3e

$$\sigma_\delta = \sigma_c + \sigma_\theta$$

Determining contest outcome

Contest outcomes were determined probabilistically (Figure 1, bottom right), based on the 'wagers' that each agent made at the contest start. Wagers are a combination of the relative size and effort of each competitor. In this way the agent with greater relative size and effort (termed the *favorite)* usually won., but upsets, in which the agent with the smaller wager (called the underdog) were possible, particularly if a relatively smaller agent contributed far higher effort. We calculate each agent's wager, w_i:

Equation 4a:

$$w_i = \left(\frac{x_i}{x_{max}}\right)^s * \varphi_i{}^f$$

Here, s and f allow us to control the relative importance of size and effort, respectively. Because these exponents always operate on fractions between 0 and 1, we can conceptualize them as "bias" parameters, where, as they increase, larger inputs (of size or effort) are exponentially more successful than smaller inputs (Supplemental Figure S1a). The probability of upset was then calculated by

Equation 4b:

$$P(upset) = \frac{1}{2} * \left(\frac{w_{min}}{w_{max}}\right)^l$$

Here, an upset is defined as when the smaller *wagering underdog* wins over the larger *wagering favorite* (so a smaller sized individual beating a larger individual investing minimal effort would often not be considered an upset). This equation features an additional "bias" parameter, l, which controls the importance of luck in upsets, or equivalently, by determining how close the two estimates need to be in order for there to be a meaningful probability of upset. For all values of l, when the two wagers are equal, $P(upset)$ is equal to $1/2$.

This fight approach allowed us to vary the relative contributions of size, effort, and luck in order to observe how this changed the behavior of the group. For most purposes, we set the values of these parameters (at s = 5, $e = 1, l = 13$) in order to match empirical observations (Beacham 1988; Bierbach et al. 2012), i.e. that upsets were rare, except when the two agents were similar in size (Supplemental Figure S1a).

Assessment Update

For updating strategies, we compared three updating strategies: Bayesian updating, linear updating, and fixed (no updating), described below. We decided not to model the ELO updating scheme used in (Hemelrijk 2002; Hock and Huber 2006). The ELO framework requires that individuals observe and recall every contest that takes place among all individuals in a group, which seemed implausible in most contexts. More importantly, in an ELO paradigm, winner/loser effects act on intrinsic ability, not on self-assessment, or at least these two mechanisms are treated as a single variable. As such, it could not be easily implemented into our modeling framework, and any comparisons with our model—which is designed to explore fixed-ability, limited-information scenarios—did not seem meaningful.

Bayesian updating

Bayesian updating requires a prior and a likelihood function. In our case the prior, as described above, is the agent's current estimate of size, modeled as a discrete probability distribution across all possible sizes. The likelihood represents the probability of the observed outcome (i.e. winning a contest), given some assumed size, calculated across all possible sizes. Thus, in this case, the posterior distribution, i.e., the new probability distribution of size, is calculated by

Equation 5a:

$$P_{post}(\theta|outcome) = \frac{P(outcome|\theta) * P_{prior}(\theta)}{P(outcome)}$$

We will explore this equation where outcome = win, but the formulation is similar for outcome = loss. This equation states that to update the probability distribution, at each size x, we multiply the existing probability of being that size by the likelihood of winning at that size, forming a distribution. The probability at every possible size is divided by the sum of all these values (i.e., the total probability of winning). Thus, for every possible size in our discrete size range, the posterior probability is given by

Equation 5b:

$$P(x_i = x|win) = \frac{P(win|x_i = x) * P(x_i = x)}{P(win)}$$

Here, $P(x_i = x)$ is given by the prior estimate, while the likelihood $P(win|x_i = x)$ is simply the probability of an individual of size x winning, or more precisely

Equation 5c:

$$P(win|x_i = x) = P(upset), if\ w_i|\ x \leq w_o$$
$$P(win|x_i = x) = 1 - P(upset), if\ w_i|\ x > w_o$$

Note that there is an inflection point based—not on relative size—but on relative wager, such that agents must factor their own and opponent effort into their calculations of the P(win). For simplicity, we assume that after the contest, agents have perfect knowledge of the size and

effort of the opponent, however this calculation could be extended to allow for additional distributions of opponent size and opponent effort (see Equation S1).

The prior and likelihood are then divided by the $P(win)$, which in our case is simply the sum of prior $*$ likelihood for all possible sizes,

Equation 5d:

$$P(win) = \sum_{x=0}^{x_{max}} P(win|x_i = x) * P(x_i = x)$$

After calculating the posterior estimate, agents update their expected size, $\bar{x}_i$, and their estimate confidence, σ_θ, which are used to compute their effort in the next contest.

Equation 5e:

$$\bar{x}_i = E(\theta) = \sum_{x=1}^{x_{max}} x * P(x_i = x)$$

Equation 5f:

$$\sigma_\theta = \sum_{x=1}^{x_{max}} P(x_i = x) * \ (x - \bar{x}_i)^2$$

Linear Updating

A simpler way to model updating, as implemented by (Dugatkin 1997), is to scale each agent's estimate, using some proportion, p of its prior estimate and/or adding or subtracting some constant, k. We modeled this, with limits on the maximum and minimum possible size, such that

Equation 6a:

$$\bar{x}_{i_{t+1}} = \ p\bar{x}_{i_t} + k$$

After updating the mean estimate, we generated a new normal distribution, centered around that estimate, based on their *awareness*, σ_a, as in Eq. 2b.

No updating

We also modeled a scenario where agents performed no updating at all. For this, we kept the initial distribution and/or point estimate fixed after each fight. In the fixed case, an agent's estimate is simply their initial guess, with awareness, σ_a.

Total parameter space

Taken together, our possible parameter space comprises the fight parameters s,f,l, the effort parameters *acuity*, and *awareness*, and, for linear updating, the value of the parameters *k,p*. Even with our simplifying assumptions, this represents a somewhat daunting parameter space. To aid in exploring the parameter space, for those parameters that varied between [0, ∞) we used a scaling function (Supplemental Equation S1) so that we could input parameters as values between either [1,-1] or [0,1], depending on whether there was an inflection point in the real value at 1. By doing this, it was possible to explore the full range of possible parameter values.

Model Analysis

The goal of our model was to answer two questions. 1. Does Bayesian updating for self-assessment produce social networks that match the behavior of empirical systems? 2. How can we distinguish, empirically, Bayesian updating from alternative models of dominance establishment and winner effects?

Our core approach to answering these questions was to run many iterations of social simulations, and measuring the output using metrics that would be tractable in and relevant to an empirical context (see below). Because it was a simulation, we were also able to measure the hidden variables (e.g., self-assessment) that were driving the behavior. We organize the metrics according to the level of complexity at which they are observed.

Agents

For each agent, we recorded true intrinsic ability (referred to simply as size), as well as their current, estimated own-size distribution at every timepoint t, with the calculated mean and variance of the (not necessarily normal) distribution.

Contests

For every contest, we record the outcome, as well as the intensity of the contest, determined by the minimum effort. We also recorded the effort and wager invested by each agent for each contest.

Networks

For the full history of social interactions, we track the linearity, stability, accuracy, and efficiency of networks. *Linearity* is calculated using Appleby's formula (Appleby 1983), which calculates the proportion of triads that are transitive. For *stability*, we first calculate binary dominance-subordinance status for each pairwise relationship, then calculate the proportion of time bins for which the binary dominance network, calculated within that time bin, matches the binary dominance network calculated for the entire simulation. For *accuracy*, we calculate ordered dominance, using an ELO ranking (Albers and De Vries 2001; Neumann et al. 2011), then calculate the spearman correlation between actual size and dominance rank. For our metric of *efficiency*, we simply measure the mean intensity of contests within each time bin. For *efficiency (*intensity), we averaged across each round of all pairwise contests. For *linearity* and *stability*, we used a sliding window of 3 rounds.

Implementation
All code was written in python, using standard libraries. For ELO ranking we used an *elosports* module (pypi.org/project/elosports). All code and results are available on github (github.com/aperkes/BayesBattleBots).

Results

<u>Bayesian updating produces linear, stable, and efficient dominance networks</u>
We found that our simulated agents using Bayesian updating recreated known empirical observations of dominance hierarchies. In particular, simulated networks stabilized over time into hierarchies as individuals became more confident in their relative estimates (Fig 2a). At the same time, the linearity of hierarchies rapidly increased (Fig 2b) while the effort each agent expended per contest decreased (Fig 2c).

In general we found that our simulated agents recreated stable linear hierarchies across the range of parameters tested (Supplemental Figure S2). However, there were certain edge cases where other outcomes occurred, for example, when effort had no impact on outcome (e.g., *s* near infinity, or *f* near 0). Additionally, if the agents' initial awareness was low and the size differences between individuals were small, there were more frequent upsets resulting in less-stable hierarchies. Where size differences were large (Figure 2a,2b), or if the *s* parameter were set such that any size difference could entirely determine outcome (i.e., *s* approaches infinity and there is no luck component), then hierarchies were entirely deterministic based on size and stable from the start, with no network formation period characteristic of many empirical dominance hierarchies ((Tibbetts et al. 2022)). This stabilization effect was sufficiently strong that simulated social networks formed linear dominance hierarchies even in the absence of size differences (Figure 2b, yellow line), particularly where small differences in effort could strongly determine outcomes, and in these contexts we observed the strongest contrast between the initially chaotic network and the eventual stable state.

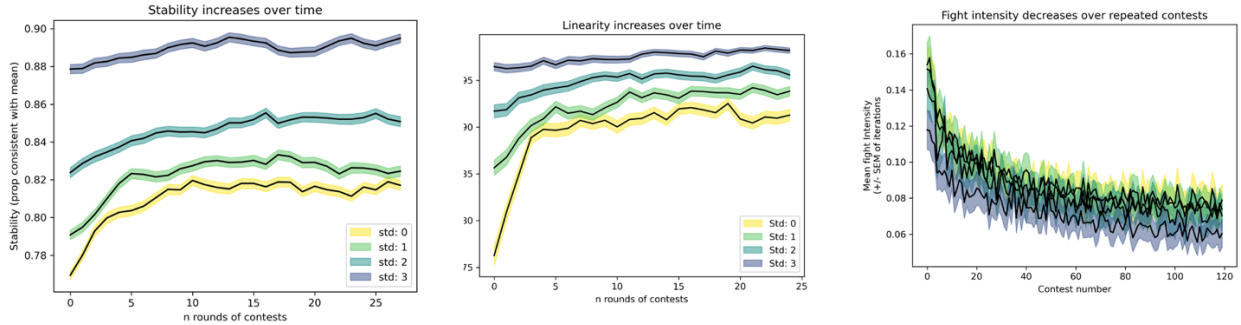

Fig 2: Emergence of stable dominance hierarchy through Bayesian updating. A. Network stability increases over time. B. Linearity of networks increase over time. C. Fight intensity decreases over time.

Bayesian updating produces a recency bias in the winner effect

Although winner/loser effects vary across species, there is strong empirical support that species show a strong recency bias with the outcomes of more recent fights having a stronger effect on current behavior. We found that our model recreated these empirical observations. We first staged contests between naive, size-matched agents and then paired the subsequent winners and losers against new, naïve, size-matched individuals. To simplify this simulation, all individual priors were centered on their true size, and had perfect acuity for opponent assessment. Individuals faced 5 size-matched, naive individuals, where each contest generated a new branch in the individual path. Exploring all possible outcomes, we observed that winner effects have an experience-dependent recency bias (Fig 3a), such that first winning a contest and then losing a contest (WL) results in an agent updating to having a lower estimate their size than if they had lost and then won (LW). This results from the fact that effort is determined by an agent's estimate, so the WL agents experienced invested more heavily in their second contest (after having won their first) and so experienced a steeper likelihood function that LW agents which invested comparatively less in their second contest (after having lost their first). Thus, in a Bayesian updating paradigm, the known recency bias comes not from any particular temporal bias, but from a difference in the recent experience itself. This means that we may expect winner and loser effects to persist until the individual receives a new experience, irrespective of whether that takes hours or days. This also means that there are some conditions (WWWWL vs LLLLW) where the recency bias is not sufficient to overcome the dramatically different prior experience. In general, this recency bias becomes less apparent for individuals with more confident priors, although this is driven by a decrease in winner/loser effects generally as experience increases (see Figure 6b below).

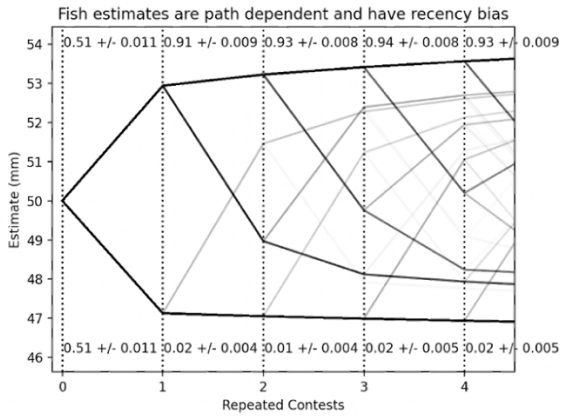

Figure 3: Divergent experience leads to stochastic formation of self-estimates. A. This plot shows the estimates of agents as they compete in repeated contests against naive, size-matched opponents. Each dotted line represents a contest, with the branching paths showing the estimates following a win or lose.

Bayesian updating predicts that the loser effect is stronger than the winner effect.
There is strong empirical support documenting that the effects of losing a contest tend to be stronger than the effects of winning a contest. To test the strength of winner/loser effects, focal agents are given a staged challenge, in which we control the outcome (win/loss) and the effort of the opponent. Following this staged contest, agents face a fair contest against a naïve, size-matched individual. We simulate this process 1000 times for both a winning focal individual, and a losing focal individual, measuring the mean probability of winning. Under these conditions, iur model successfully recreates the empirical observation that the loser effect is stronger (Fig 3a). In our model, this phenomenon is driven by two mechanisms, stemming from our assumption that contests are determined by the relative size of competitors, rather than absolute size difference. *<<This section needs to be a lot more clear, I will honestly be surprised if you can understand why the loser effect is stronger from my current explanation.>>*

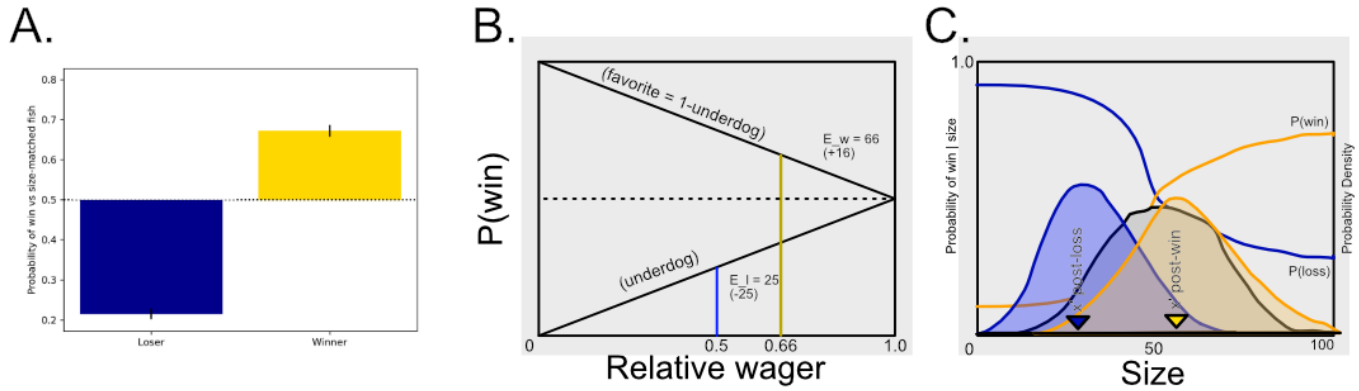

Figure 4. *A. Loser effects are stronger than winner effects. This plot shows the probability of winning, with the dotted line representing the random expectation (0.5).* B. The effects of relative scaling are such that an equal decrease in estimate leads to a disproportionate decrease in effort. C. The flattening of the likelihood function at larger sizes means losses are more impactful.

To illustrate why the loser effect is stronger, we will consider a simplified scenario where, following a size-matched contest, 50cm individuals either increase or decrease their estimate by 25cm. The first mechanism which amplifies the loser effect follows naturally from the fact that during contests, individual effort is proportional to the probability of winning. As Figure 4b illustrates, after a focal win or loss, the perceived probability of winning an additional fight is 25% for the loser and 66% for the winner (The probability for a naïve agent would be 50%). We see that the loser decreases their effort by 0.25, while the winner only increases their effort by 0.16. This effect is sufficient to drive differences in the winner and loser effect, but Bayesian updating further increases the relative strength of the loser effect. As Figure 4c shows, asymmetry in the likelihood functions means that losers actually decrease their estimate by more than the winners increase there's. Likelihood functions are asymmetrical because, once an agent is larger than the opponent, relative size increases more slowly than actual size. In the

extreme case, the size of the minimum individual is 2% of the mean, while the maximum sized agent would only be 50% bigger than average. This flattens the right side of the likelihood function, strengthening the loser effect relative to the winner effect.

Bayesian updating explains the observed variation in the persistence of winner/loser effects

In empirical observations of dominance contests, winner/loser effects are persistent in some cases but short-lived, or even non-existent, in others. Our model can also explain this variation. Within our model, winner/loser effects are based on changes in self-assessment, and so whether winner/loser effects are observed is a function of whether individual experience results in a shift in assessment that is sufficiently large to be observed in an experimental context. Whether winner/loser effects persist is thus a function of individual information—specifically the intensity of the initial experience and the quality and quantity of additional experiences prior to being assayed—and the experimental power with which individual behavior is assessed.

To illustrate the importance of the initial experience, we simulated a set of (size = 50) individuals and again staged a win against an external 'treatment' agent of variable size x_o) thus creating a set of 'staged winners'. Because we could inspect their internal estimates, we could directly quantify the strength of the winner effect. Because Bayesian updating functions through a likelihood function, the strength of the winner effect (i.e., the magnitude of the shift in self-assessment) is determined by the size of the opponent. As the size of the treatment agent increases, so does the strength of the winner effect on the focal agent, due to a larger shift in self-assessment by the focal agent (Figure 5a). It is clear that the size of the "treatment" agent drives the strength of that initial shift. To observe whether this also impacted its persistence, we then placed each 'staged winner' in a group of naïve individuals of set sizes (i.e., 30,40,60,70), All agents went through series of rounds of contests as described previously, in which each round, every agent fought every other agent in random order. We repeated this process until each individual had experienced a total of 20 contests (5 rounds). Here the persistence of the winner effect is driven by the initial shift, which is itself a function of the size of the opponent. The scale and duration of the winner effect is also a function of the confidence of the individual prior, such that as the standard deviation of the starting prior increases, so does the strength of the winner effect. The results described here for winner effects function similarly for loser effects as well.

To test the impact of available information, we again created staged winners of size x=50, this time only against treatment opponents of size 60, and placed with groups of varying numbers of naïve agents (i.e., 0,3,5,10). Here, the sizes of naïve-agents was sampled randomly from a normal distribution (mean = 50, std=10). This allowed all agents to compete against several other agents, thus gaining new information and the ability to update their estimates with each contest. Here, the amount of the diversity of experience (i.e., the unique number of individuals faced over the course of 20 fights), had a strong effect on how quickly the effects of the initial focal fight attenuated, as individual more accurately estimated their true size (Figure 5b).

Regardless of the strength or duration of winner/loser effects, we also show that whether they can be detected depends on the experimental power used to test it. Testing winners and losers against size-matched individuals makes it possible to detect even fairly weak winner effects (figure 5c). In these cases, without using size-matched contests, it is extremely unlikely to observe winner effects in a feasible number of replicates (Figure 5c). Thus observed variation in winner effect duration (and whether it is detected generally) can be ascribed to differences in how the winner effect was generated, how it is assayed, and the experience of individuals in the interim.

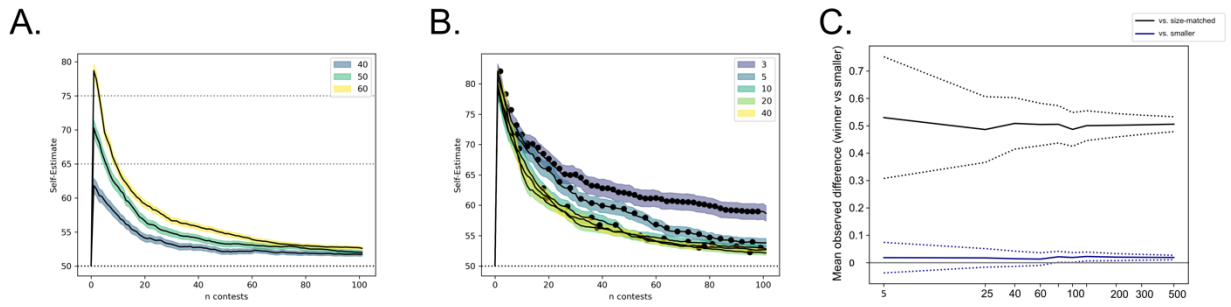

Figure 5: Strength and persistence of the winner effect is a function of information content and experimental power. A. This plot shows focal agents' estimates following a staged win against agents of varying size. The shift in the focal's estimate following a win is greater when their opponent was larger. B. Here, the focal winner is placed in groups of varying size. The x axis shows the number of contests, with circles along the line to show each new round of all agents within that tank. The persistence of this shift depends on the variation in the actual sizes of the surrounding agents. C. This plot shows the mean estimated winner effect in 1000 simulated experiments, calculated by measuring the difference in win rate for focal winners and losers against a size-matched agents (black) or a smaller agent (blue), with the dotted lines showing the standard deviation of estimates.

Distinguishing Bayesian updating from other models

We have shown that our model is successful at recreating many empirical observations of animal contests, consistent with the hypothesis that Bayesian updating may be the mechanism behind dominance hierarchies and winner/loser effects. Bayesian updating provides a potential mechanism to explain and predict how individuals respond to (repeated) contest outcomes, but being consistent with empirical behavioral observations does not confirm that winner/loser effects are driven by Bayesian updating. To provide an approach that can empirically test for Bayesian updating, we explore which testable predictions distinguish Bayesian updating from alternate models, based on the two defining features of Bayesian learning—the likelihood-baed updating and the informed prior. We investigate each one independently, with simulated experiments to model future approaches in animal systems, focusing on measurable contest outcomes within empirically tractable scenarios.

<u>Bayesian updating is characterized by discrepancy and experience effects.</u>

First, to test how different updating mechanisms might be experimentally distinguished, we simulated contests against individuals of different sizes (either size-matched, or smaller) and measured the strength of winner effects using the same, size-matched contest assay used to assess the strength of the winner effect in Figure 4a. In this context, compared to linear and fixed updating, only Bayesian updating showed a significant difference for winner (and loser) effects against size-matched individuals compared to smaller individuals, which is intuitive, since linear updating has no mechanism to incorporate opponent size. This is in line with the discrepancy effect, which is a well-known characteristic of Bayesian updating processes whereby more surprising (and therefore informative) outcomes lead to a more dramatic shift when updating the estimate (Figure 6a).

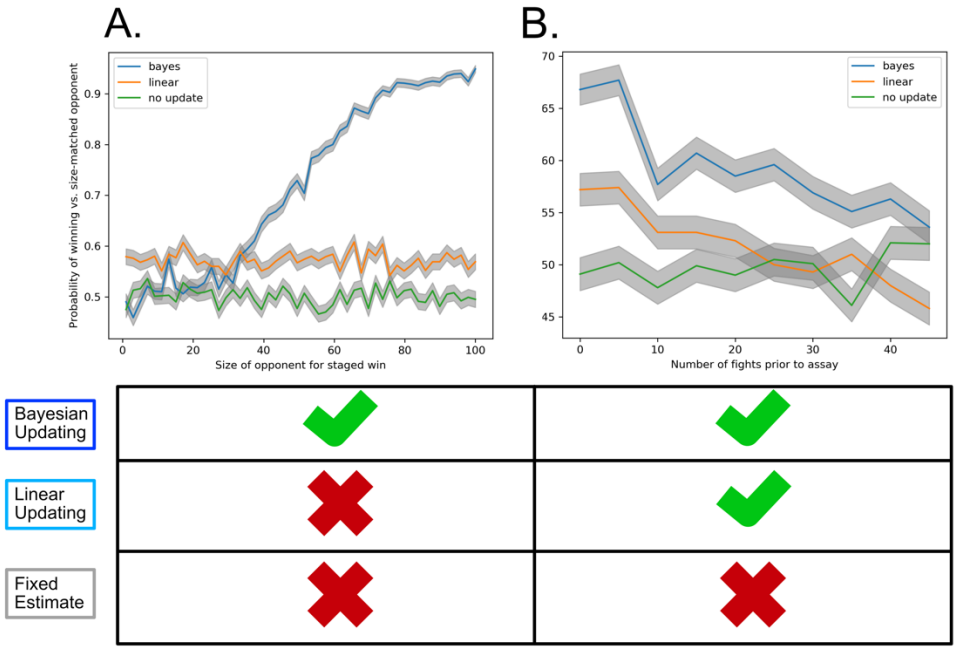

Figure 6: Updating strategy can be distinguished by discrepancy and experiential effects. A. This simulated plot shows the strength of the winner effect (measured as the probability of winning a subsequent size matched fight after a controlled win), plotted as a function of the size of the opponent during the staged win, with each line representing a different updating strategy. B. This plot shows the strength of the winner effect, (measured again as the probability of winning against a size matched agent), based on the number of prior contests. Below: This table shows the predictions for each updating approach, showing that we should expect a significant effect of the size of the opponent on the strength of the winner effect for Bayesian updating, while both Bayesian updating and linear updating predict a significant effect of experience.

We further explored how the outcomes of experience, or repeated contests, should change based on whether individuals are using Bayesian, linear or no updating mechanisms. To do so, we forced wins or losses against size-matched agent, this time for both naïve and experienced individuals, then assayed the winner effect using a second size-matched fight, in which any deviation from 50% defines a winner or loser effect.

We found that the winner effect is weaker in experienced individuals that are using Bayesian updating. This is expected given the well known behavior of Bayesian updating in which more experienced individuals are expected to have more confident priors and thus be less influenced by a single win or loss. However to our surprise, this same experience effect was also evident in individuals that were using linear updating schemes (Figure 6b). (There was, not surprisingly, no experience effect for individuals using no updating mechanism.) On further inspection, this similarity between Bayesian and linear updating is due to the divergent effects of linear updating. With linear updating, any individual outside the median size is either winning or losing the majority of their contests. This means their estimate will diverge over time, with all individuals bigger than the median eventually assuming they are extremely large (and all below-average individuals concluding they are extremely small). Because linear updating does not have a "discrepancy effect" described above, a single outcome will have a minimal effect on the behavior of experienced individuals. A single loss will decrease the estimate of past winners by 1 step, however they are still only 1 step less than "extremely large", so they will continue to behave with extreme aggression. In this scenario, no outcome has a dramatic effect on individual behavior, so on average we will observe no winner or loser effect. In any linear updating scheme, this will occur, with runaway increases and decreases in individual estimates leading to attenuation of winner and loser effects (as well as dramatically inaccurate predictions of size). Thus, while an experience effect is a core prediction of Bayesian updating, and failing to find one would exclude Bayesian updating as a model, finding an experience-attenuated winner effect does not rule out other possibilities. The discrepancy effect is the most straightforward single test for Bayesian updating.

<u>Informed priors result in consistent individual differences that are context dependent.</u>
The second key feature of any Bayesian process is the presence of an informed prior. Here we provide a series of scenarios under which experiments can be used to explicitly test whether individuals are using different types of initial priors. In our model we defined each individual's prior by a Normal distribution centered on the individual's initial estimate of their own size, with a standard deviation equal to the precision (*awareness*) of that estimate. In a natural context, priors could be informed by internal information like the animal's age or viscerosensory information to external information such as observing the distribution of other individuals in the population. We generated a series of predictions that can help distinguish which types of information, internal versus external, individuals are using to inform their priors, based on measuring individual effort during contests. We additionally tested how the intensity of contests are expected to change if instead individuals were not using Bayesian priors at all but instead using uninformed priors such as assuming a uniform distribution of sizes (no information) or a normal distribution, centered around a random point (incorrect information). In order to identify the different predictions of these priors, we conducted simulated experiments of dominance contests. To approximate an empirical approach, here we measured effort against a size-matched opponent over 3 repeated trials. Effort could be assessed in empirical settings by measuring time spent engaged in aggressive behavior (Schlinger et al. 1987; Francis 1990). We then tested for persistent individual differences and context-

dependent variation in aggressive behavior, using mixed linear models to identify the relative contribution of these two sources of variation.

In our simulated, size-matched contests, we find the use of different priors generates key differences in the presence or absence of consistent individual differences in aggressive behavior and whether that behavior is dependent on some relevant context (i.e. the size of the the agent and/or the surrounding individuals, see Figure 7). Specifically, if individuals are using either self-informed or random (likely incorrect) priors, we would expect to see consistent individual differences in effort across repeated contests, whereas if they were using uniform or externally informed priors, we would expect very little variation (since these distributions should be equal for all individuals). When testing for an effect of context on effort, only self-informed priors should produce a significant effect of individual size on effort, while with externally informed priors individual effort should correlate with the size of surrounding individuals.

To simulate externally informed priors, we simulated that some individuals determined their initial prior based on the body sizes of the surrounding individuals. Notably, individuals from the same social environment (large or small surrounding individuals) have extremely low variation, looking essentially identical to individuals with uniform priors. Distinguishing this type of prior estimation is only possible by varying the social context, i.e., by surrounding some individuals with small conspecifics, and others with large conspecifics, we can test for an effect of social context on effort, which can differentiate those with externally informed priors from uniform prior individuals. Drawing prior information from internal and external sources are not mutually exclusive, however, but here we show that by conducting experiments that controls for variation in both individual size and surrounding distribution, it is possible to identify whether individuals use Bayesian priors, and what information they incorporate, in order to inform their self-assessment in the absence of fights.

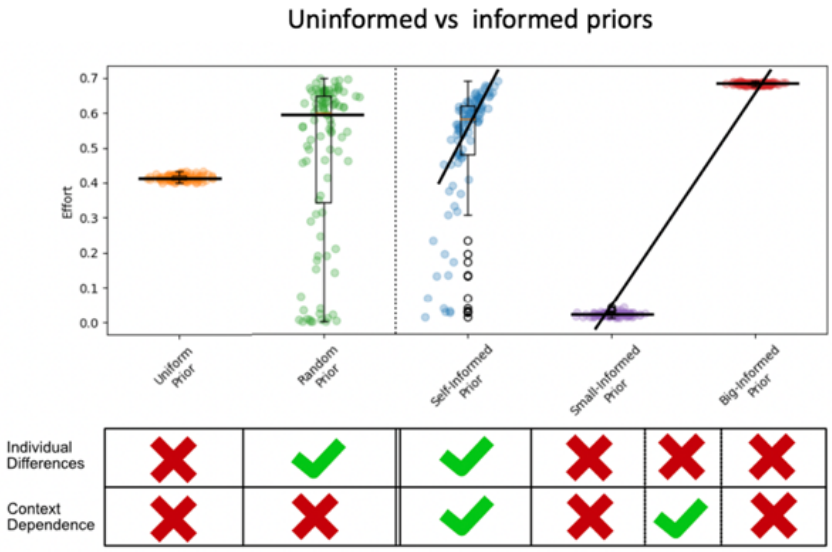

Figure 7: Individual priors are distinguishable through population variation and context-dependence. Top: This plot shows the measured effort from simulated dominance contests, as a function of individual size, in five populations, each using a different prior. The two left most populations use non-Bayesian approaches, either all individuals having a uniform prior, or each individual having their own randomly selected prior. In the right three groups, individuals use information to inform their prior, either internal information (central cluster, blue) or surrounding information (two rightmost clusters). Black lines show the best fit line between effort and size for each population (the dotted line is the best fit line between size and fit when combining the small-informed and large-informed populations into a single population). Below: The prior formation strategy of a given population can be distinguished by testing for the presence of individual differences, and an effect of context of individual effort. (either internal or external). Green checks denote a significant effect, while red x's indicate no effect.

(What follows is a rough first draft of the discussion, don't look too closely at this, but I would love to hear your take on which ideas you find most interesting)

Discussion

Here, we have shown that the winner and loser effects can be described by Bayesian updating for self-assessment, and that agents performing this Bayesian updating generate dominance hierarchies that are consistent with those observed in social systems. More importantly, our model provides testable predictions to evaluate whether species' behavior is consistent with Bayesian updating for self assessment. Our key findings here lay the groundwork for future work, both empirical and theoretical, in understanding the biological drivers of social dominance and the winner and loser effects.

1. Bayesian updating reinforces intrinsic differences to promote stable hierarchies

Intuitively, the winner effect should provide an opportunity to disrupt size-based hierarchies. An individual who wins receives a boost from winning one contest could go on to dominate another larger individual. As we have shown, Bayesian updating *can* generate linear hierarchies in the absence of differences, however Bayesian updating tends towards accuracy, so where size remains a factor and accurate information is provided, all individuals should arrive at their actual size, such that their relative effort in fights matches their relative size, reinforcing intrinsic differences rather than disrupting them.

While Bayesian updating of self-assessment is sufficient to explain the linearity and stability of dominance hierarchies, future work could expand on this to incorporate additional social information. Prior empirical work has shown that allowing individuals to observe contests reinforces linearity, and it is known that most social animals are acutely aware of their social environment. Future models could incorporate additional social information by allowing each agent to keep a Bayesian estimate of the size of others. While adding social eavesdropping to self-assessment increases the complexity of the model, it would make it possible to explore how these effects add and interact.

2. Bayesian updating explains variation in the strength of the winner effect

In our model, the strength of winner and loser effects is a direct output of the information content of the contest. Winning against a larger individual produces a larger change than winning against a smaller individual. Agents with high confidence in their size will change their behavior less following a contest than naive agents. This behavior is consistent with empirical observations, and is a strong prediction of the winner effect being driven by this rational, informed self assessment.

While Bayesian updating does produce individuals that weigh recent fights more strongly, in general, our results show that for most fights (i.e., those predictable outcomes between two experienced individuals), any winner effects should be vanishingly small. In contrast, empirically observed winner effects can be quite dramatic. One explanation for this may our assumption that size (i.e., intrinsic ability) remains fixed for any given individual. In reality, contest ability likely changes over both short and long time scales, and thus while individuals may have high confidence in their average ability, there could be considerable uncertainty in their current ability, causing Bayesian updating to favor large shifts in behavior the short term. This would be especially true when winning opens up access to reproductive or survival resources. By expanding our model to incorporate local uncertainty and short-term motivation, we could better explain the intensity of short-term winner effects.

3. Bayesian updating explains variation in the persistence of the winner effect

As we stated in the introduction, there is some discrepancy in whether the winner effect is short-lived or permanent. Our model can account for some of this, based simply on the information content of contests and how behavior is measured, however our model lacks any explicit timescale. Agents are held in stasis between contests. Time could be incorporated to our model by providing additional non-contest information (perhaps each passing moment flattens the prior somewhat), or by having both an “average” prior and a “current” prior that are modified separately. This could help to identify how changes in self-assessment might operate over different timescales.

It is important to note that the updates on self-assessment must occur alongside multiple other mechanisms that can act as winner or loser effects. In our model we have deliberately avoided any impact of contest outcomes on real intrinsic ability in order to focus on self-assessment, however in reality, in the very short term, molecular cascades lead to differences in metabolism that can promote or inhibit contest behavior, not just influencing motivation but that actual ability to convert size and effort into victory. In the longer term, winning a contest may provide access to some resource, providing additional energy to the winner while leaving the loser weaker, allowing the winner to invest more effort in future fights. Over time, the increased resources will compound to drive differences in size. Our model could be expanded to include all of these components (indeed some exist as unused options in the code). Doing so could identify the relative strength of these effects, and how they disrupt or complement one another.

Conclusion

Bayesian updating provides both a *how* and a *why* for the winner effect, while agent based modeling provides an opportunity to incorporate multiple elements to predict individual and network behavior. We hope that our model will help stimulate future theory and drive empirical work to test for a Bayesian winner effect. There are a great deal more questions that can be explored within this framework, many of which require only minor modifications to our existing code. The empirical work is perhaps even more exciting: The winner effect is widely observed across species, however there is considerable variation in the details across systems. Our modeling shows that the winner effect is a natural outcome of Bayesian updating for self-assessment, which is likely favored in many systems. However the specific implementation of shifting behavior in response to a fight could easily vary based on the specific parameters of that system. In many ways, Bayesian updating can act as a null hypothesis, and identifying systems that do not meet the predictions of our model may provide greater insight into how the winner effect function in different systems and across different time scales.

Supplemental PDF

Fight parameter intuition

As described in the methods are settled by based on relative wager (Eq) where wager is a function of relative size and effort, each raised to an exponent between 0 and infinity. In order to explore the parameter space more evenly, we use a scaling function

Supplemental Equation S1 $$l = \cot\left(\frac{(L+1)}{2} * \frac{\pi}{2}\right),\ -1 \leq L \leq 1$$

This effectively maps $L = 1$ to $l = 0$, $L = 0$ to $l = 1$, and $L = -1$ to $l = \infty$, as shown in Figure S1a. We can think of these parameters as magnifying or flattening differences in the different aspects of wager, which varies the importance of these aspects. To help illustrate this, we simulated a fight between two agents, one larger ($x_i = 60$) but low effort ($\phi_i = 0.3$), another smaller ($x_j = 50$) but high effort ($\phi_j = 0.6$). When s is large, the larger agent tends to win, when f is large, the higher investing agent tends to win, and as l increases, the bias against the "underdog" increases.

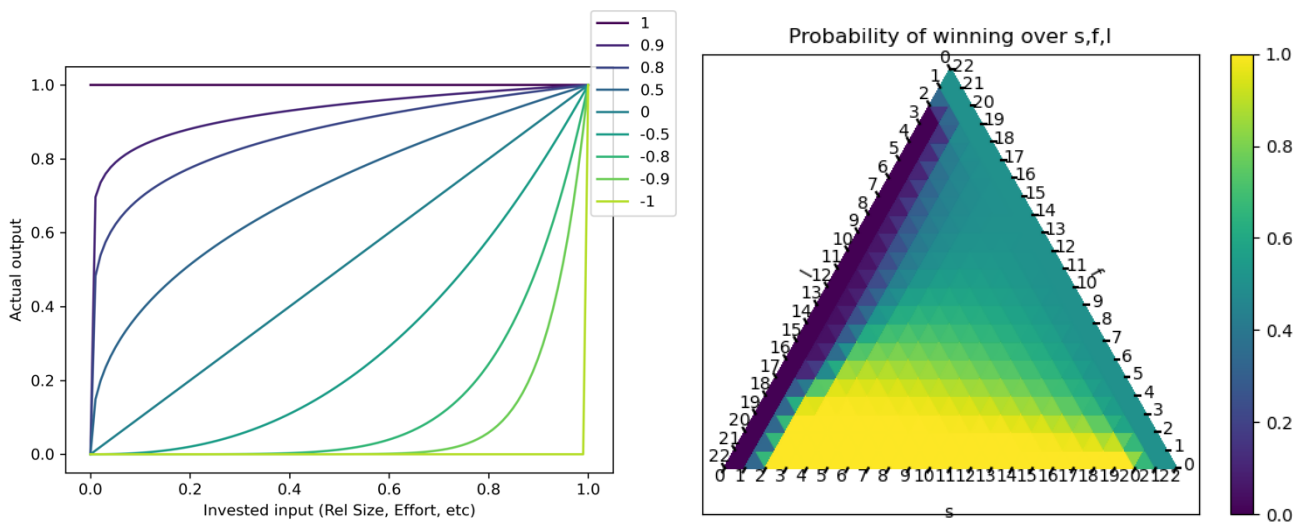

Supplemental Figure 1: Fight parameters. Panel A shows how relative size, effort, and relative wagers are scaled based on the value of s,f or l, respectively. Here again we use Equation S1 to scale parameters, such that 1 represents s^0 while -1 represents s^∞. Panel B shows shows the probability of a larger agent (size=60, effort=0.3) winning against a smaller, smaller competitor investing greater effort (size=50,effort = 0.6). Axes move evenly from -1 to 1 in the scaled parameter space (Supplemental Equation S1).

For results below, we explore this parameter space to observe how our results vary based on the fight parameters.

Linearity persists or increases across parameters

To test whether the stabilization effect that we observed based on bayesian updating was consistent across parameters, we simulated 100 dominance hierarchies across the range of possible parameters. Using the scaling function above, we tested every combination from -1 to 1 in intervals of 0.1. For any dominance hierarchy, the random expectation of linearity is 0.75 (i.e. 3 of 4 triads should be linear just by chance). Dominance hierchies tended to begin more linear when size was a larger component (Figure S2a). We ran these networks for 20 rounds, and again measured linearity (Figure S2b). By taking the difference between this initial linearity, and the end-state linearity, we observed that in general, linearity either increased or stayed the same (Figure S2c) This increase was most dramatic where size had little to no effect on contest outcome, and where effort was a strong factor. Networks that started out strictly linear obviously did not increase.

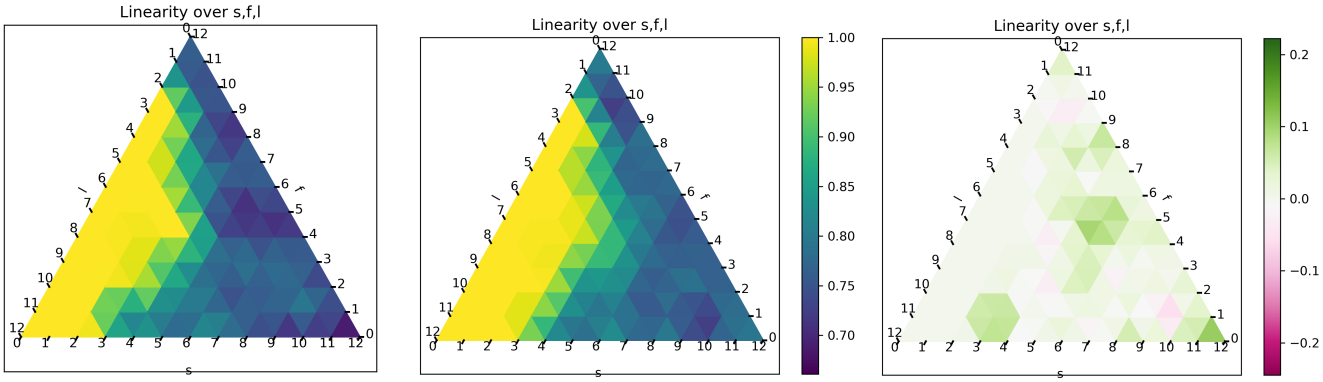

Supplemental Figure S2: Linearity. These plots show the results of simulating dominance hierarchies across various parameters. For each parameter set, we repeated the simulation 100 times. Panel A shows the linearity in the first three rounds of fights. Panel B shows the linearity after 20 rounds. Panel C shows the difference.

How is effort determined?

Agents invested in a fight based on their perceived probability of winning, which was calculated according to (EQUATION), as their probability of winning given their estimate, multiplied by the (rough) probability of their estimate being correct. To demonstrate this, we simulated a contest between a smaller agent and a bigger agent, with various awareness and acuity parameters (Figure S3). As an agent's awareness decreased (i.e., the confidence of its self assessment) so did effort. As acuity increased, effort decreased, since the smaller agents correctly identified that their probability of winning was low.

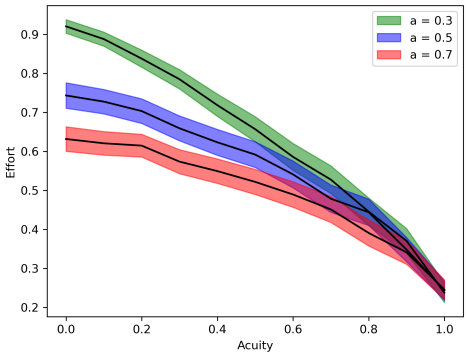

Supplemental Figure S3 The effects of acuity and awareness on effort. This plot shows the mean effort (from 1000 simulated fights) of a smaller agent ($x_i = 40$) against a bigger agent ($x_j = 60$) as a function of acuity. The different lines represent agent of different awareness.

Because they change both the magnitude and variance of effort, acuity and awareness also had an impact on the linearity, and the stabilization of dominance networks. As before, we simulated many networks across the range of acuity and awareness parameters (Figure S4), measuring linearity before (S4a) and after (S4b) repeated interactions. Here the most dramatic stabilization was observed when acuity was low and there was moderate awareness (S4c). Where agents had strong self-assessment, networks were linear from the start, high acuity tended to increase stabilization, since it allows accurate self-assessments to be used effectively.

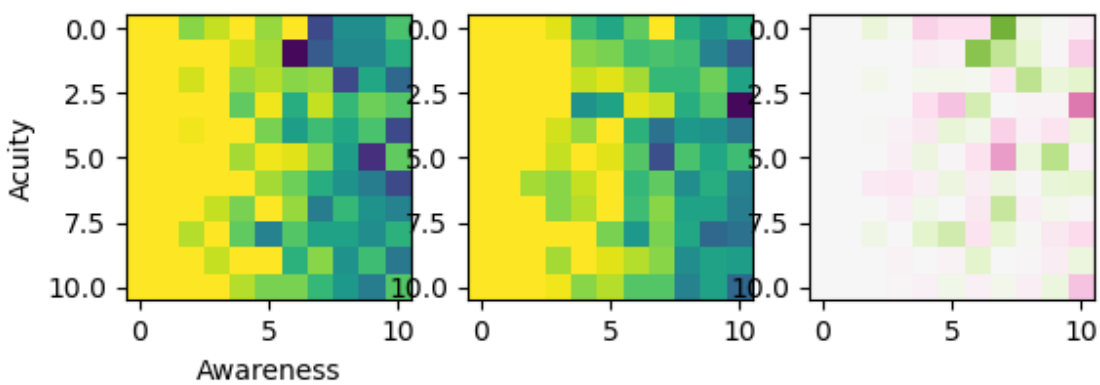

Supplemental Figure S4: Linearity across possible parameters acuity and awareness. As in Figure S2, these panels show the measured linearity before (A) and after (B) 20 rounds of contests. Here, the axes move exponentially from 0 to infinity, and represent the standard deviation around the mean when assessing opponents. Acuity 10 (i.e. $c = \infty$) is the same as pure self assessment, since agents simply assume the opponent is average size.

To confirm that our stabilization effect was not an artifact of mutual assessment, we repeated our linearity measures with moderate awareness but minimal acuity. Under these conditions, agents still increased their linearity over time (Figure S5).

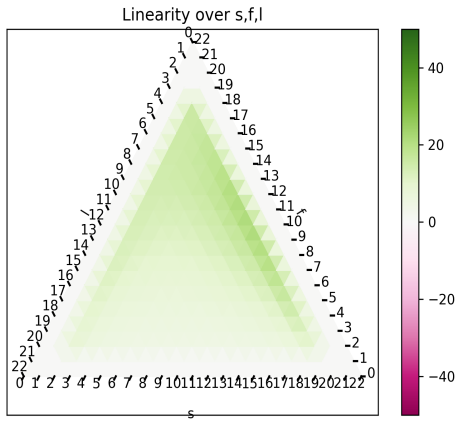

Supplemental Figure S5 : This plot shows the difference in initial and end-state linearity with agents with no opponent assessment (acuity variance = ∞). Except in certain edge cases, linearity always increased.

The number of agents impacts the persistence and extent of linearity

The numbers of agents also impacts linearity. We explored this in a context where all agents were equal size, thus the “correct” dominance hierarchy would be flat and/or random. We simulated many rounds of contests with few (n=5) and many (n=50) individuals (Figure S6). In both cases, linearity started above random chance and initially increased, however with many individuals, linearity eventually fell back to baseline. This is likely due to the large amount of information individuals received causing them to all (correctly) conclude that they were the same size, where in the low number, low information environment, it was easy to establish relative assessments to enforce linearity.

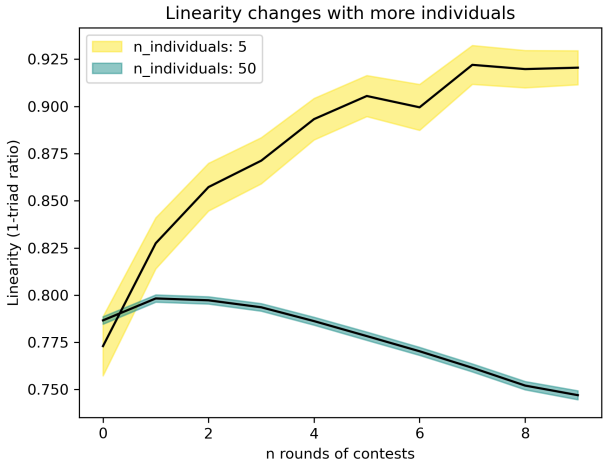

Supplemental Figure S6: This plot shows mean linearity for two simulations, one with 5 agents, a second with 50 agents. Each simulation was repeated 1000 times to calculate error ranges.

All equations

Size:

$$x_i \sim N(\mu_x, \sigma_x)$$

Assessment:

$$\bar{x}_i \sim N(x_i, \sigma_a)$$

$$\mathrm{P_r}(\theta) = N(\bar{x}_i, \sigma_a)$$

$$\bar{x}_o \sim N(x_o, \sigma_c)$$

<u>Effort</u>

Effort calculation

$$\varphi_i = P(win\,|\bar{x}_i, \bar{x}_o) * P(\bar{x}_i \geq 0.9\bar{x}_o)$$

Specific cases of effort

$$\varphi_i = \left(\frac{\bar{x}_i}{\bar{x}_o}\right)^{sl} * P(\bar{x}_i \geq 0.9\bar{x}_o), \qquad if\ \bar{x}_i \leq \bar{x}_o$$

$$\varphi_i = \left[1 - \left(\frac{\bar{x}_o}{\bar{x}_i}\right)^{sl}\right] * P(\bar{x}_i \geq 0.9\bar{x}_o), \qquad if\ \bar{x}_i > \bar{x}_o$$

Confidence correction

$$P(x_i \geq 0.9x_o) = \Phi_{\mu_\delta, \sigma_\delta}(-0.1 * x_o)$$

Parameters for CDF for confidence

$$\mu_\delta = x_i - x_o$$

$$\sigma_\delta = \sigma_c + \sigma_\theta$$

Edge case:

$$P(x \geq 0.9x_o) = 1, \qquad if\ \sigma_c = 0, and\ \sigma_\theta = 0$$

<u>Fights</u>

Wager calculation:

$$w_i = \left(\frac{x_i}{x_{max}}\right)^s * \varphi_i{}^f$$

Probability calculation:

$$P(upset) = \frac{1}{2} * \left(\frac{w_{min}}{w_{max}}\right)^l$$

Parameter scaling:

$$l = \cot\left(\frac{(L+1)}{2} * \frac{\pi}{2}\right), \qquad -1 \leq L \leq 1$$

or

$$l = \tan\left(\frac{\pi}{2} - \frac{(L+1)}{2} * \frac{\pi}{2}\right), \qquad -1 \leq L \leq 1$$

Updating

$$P_{post}(\theta|outcome) = \frac{P(outcome|\theta) * P_{prior}(\theta)}{P(outcome)}$$

$$P(x_i = x|win) = \frac{P(win|x_i = x) * P(x_i = x)}{P(win)}$$

$$P(win|x_i = x) = P(upset), \qquad if\ x \leq x_o$$

$$P(win|x_i = x) = 1 - P(upset), \qquad if\ x > x_o$$

$$P(win) = \sum_{x=0}^{x_{max}} P(win|x_i = x) * P(x_i = x)$$

Full updating if you didn't know effort and size, (too hard, so I don't do it)

$$P(win) = \sum_{x=0}^{x_{max}} \left[\sum_{y=0}^{x_{max}} \left[\sum_{\phi=0}^{1} P(win|x_i = x, x_o = y, \phi_o = \phi)\, P(\phi_o = \phi) \right] * P(x_o = y) \right] * P(x_i = x)$$

New Estimate & std (for confidence):

$$\bar{x}_i = L(\theta) = \sum_{x=1}^{x_{max}} x * P(x_i = x)$$

$$\sigma_\theta = \sum_{x=1}^{x_{max}} P(x_i = x) * (x - \bar{x}_i)^2$$

Linear updating:

$$\bar{x}_{i_{t+1}} = p\bar{x}_{i_t} + k$$

Stepped updating:

$$p = \frac{x_{max} - \bar{x}_i}{q}, \qquad if\ win$$

$$p = \frac{x_{min} - \bar{x}_i}{q}, \qquad if\ loss$$

Table of Parameters:

Parameter	Default	Range	Explanation
$x: \quad (x, x_i, x_o, \bar{x}_i)$	Random	$1 \leq x \leq 100$	Size of an individual. x_i represents actual size for individual i, while x_o is the actual size of opponent. $\bar{x}$ denotes an estimated size
μ_x	50	$1 \leq x \leq 100$	The mean of the population from which individuals are drawn
σ_x	10*	$0 \leq \sigma_x \leq \infty$	Standard deviation of population from which individuals are drawn. At $\sigma = \infty$, uniform distribution is used.
$\phi: \quad (\phi, \phi_i, \phi_o)$ $v: \quad (v, v_i, v_o)$?	None	$0 \leq \phi \leq 1$	Effort invested in a contest, as with size, subscripts distinguish effort for the focal individual ϕ_i from opponent's, ϕ_o
s	1*	$0 \leq s \leq \infty$	Relative size dampener, acts as an exponent such that as s increases, so does the importance of relative size
f	1*	$0 \leq f \leq \infty$	Effort dampener. Acts as an exponent, such that as f increases, so does marginal benefit of increased effort relative to the opponent
l	13*	$0 \leq l \leq \infty$	Relative wager dampener, essentially the importance of luck, acting as an exponent on relative wager, such that as l increases, the probability of upsets decreases
σ_a	10*	$0 \leq \sigma_a \leq \infty$	*awareness*, the accuracy of individuals' starting guesses for their self-assessment.
σ_c	10*	$0 \leq \sigma_c \leq \infty$	*acuity*, the accuracy of individuals when guessing the size of opponents prior to contests
$P(\theta)$	-	-	Used to denote the prior, an estimated, discrete distribution of the probability of being some size.
Φ	-	-	Used to denote the cumulative distribution function used to calculate confidence
p	1.1	$0 \leq x \leq 100$	Proportional shift after a fight, used in linear updating
k	0	$0 \leq x \leq 100$	Constant shift after a fight, used in linear updating

References

Albers, P. C. H., and H. De Vries. 2001. Elo-rating as a tool in the sequential estimation of dominance strengths. Animal Behaviour.

Appleby, M. C. 1983. The probability of linearity in hierarchies. Animal Behaviour 31.

Arnott, G., and R. W. Elwood. 2008. Information gathering and decision making about resource value in animal contests. Animal Behaviour.

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.

Bierbach, D., M. Klein, V. Sassmannshausen, I. Schlupp, R. Riesch, J. Parzefall, and M. Plath. 2012. Divergent Evolution of Male Aggressive Behaviour: Another Reproductive Isolation Barrier in Extremophile Poeciliid Fishes? International Journal of Evolutionary Biology 2012.

Bonabeau, E., G. Theraulaz, and J. L. Deneubourg. 1999. Dominance orders in animal societies: The self-organization hypothesis revisited. Bulletin of Mathematical Biology 61.

Chase, I. D., C. Bartolomeo, and L. A. Dugatkin. 1994. Aggressive interactions and inter-contest interval: How long do winners keep winning? Animal Behaviour 48.

Chase, I. D., D. Coelho, W. Lee, K. Mueller, and J. P. Curley. 2022. Networks never rest: An investigation of network evolution in three species of animals. Social Networks 68.

Chase, I. D., C. Tovey, D. Spangler-Martin, and M. Manfredonia. 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.

Dewsbury, D. A. 1988. Copulatory Behavior as Courtship Communication. Ethology.

Drews, C. 1993. The concept and definition of dominance in animal behaviour. Behaviour 125.

Dugatkin, L. A. 1997. Winner and loser effects and the structure of dominance hierarchies. Behavioral Ecology 8.

Ellis, L. 1995. Dominance and reproductive success among nonhuman animals: A cross-species comparison. Ethology and Sociobiology 16.

Francis, R. C. 1990. Temperament in a Fish: A Longitudinal Study of the Development of Individual Differences in Aggression and Social Rank in the Midas Cichlid. Ethology 86.

Goubault, M., and M. Decuignière. 2012. Previous experience and contest outcome: Winner effects persist in absence of evident loser effects in a parasitoid wasp. American Naturalist 180.

Green, P. A., and S. N. Patek. 2018. Mutual assessment during ritualized fighting in mantis shrimp (Stomatopoda). Proceedings of the Royal Society B: Biological Sciences 285.

Hemelrijk, C. K. 2002. Despotic societies, sexual attraction and the emergence of male "tolerance": An agent-based model. Behaviour 139.

Hickey, J., and J. Davidsen. 2019. Self-organization and time-stability of social hierarchies. PLoS ONE 14.

Hock, K., and R. Huber. 2006. Modeling the acquisition of social rank in crayfish: Winner and loser effects and self-structuring properties. Behaviour 143.

Hsu, Y., R. L. Earley, and L. L. Wolf. 2006. Modulation of aggressive behaviour by fighting experience: Mechanisms and contest outcomes. Biological Reviews of the Cambridge Philosophical Society.

Hsu, Y., and L. L. Wolf. 1999. The winner and loser effect: Integrating multiple experiences. Animal Behaviour 57.
Johnsson, J. I. 1997. Individual recognition affects aggression and dominance relations in rainbow trout, Oncorhynchus mykiss. Ethology 103.
Lan, Y. T., and Y. Hsu. 2011. Prior contest experience exerts a long-term influence on subsequent winner and loser effects. Frontiers in Zoology 8.
Landau, H. G. 1951. On dominance relations and the structure of animal societies: I. Effect of inherent characteristics. The Bulletin of Mathematical Biophysics 13.
Laskowski, K. L., M. Wolf, and D. Bierbach. 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.
Lindquist, W. B., and I. D. Chase. 2009. Data-based analysis of winner-loser models of hierarchy formation in animals. Bulletin of Mathematical Biology.
Luttbeg, B. 1996. A comparative Bayes tactic for mate assessment and choice. Behavioral Ecology 7.
McDonald, A. L., N. W. Heimstra, and D. K. Damkot. 1968. Social modification of agonistic behaviour in fish. Animal Behaviour 16.
McNamara, J. M., R. F. Green, and O. Olsson. 2006. Bayes' theorem and its applications in animal behaviour. Oikos (Vol. 112).
Neumann, C., J. Duboscq, C. Dubuc, A. Ginting, A. M. Irwan, M. Agil, A. Widdig, et al. 2011. Assessing dominance hierarchies: Validation and advantages of progressive evaluation with Elo-rating. Animal Behaviour.
O'Connor, C. M., A. R. Reddon, I. Y. Ligocki, J. K. Hellmann, K. A. Garvy, S. E. Marsh-Rollo, I. M. Hamilton, et al. 2015. Motivation but not body size influences territorial contest dynamics in a wild cichlid fish. Animal Behaviour 107.
Okasha, S. 2013. The evolution of Bayesian updating. Philosophy of Science 80.
Olsson, O. 2006. Decisions under uncertainty: Bayesian foraging - A meeting held at Lund University in August 2003. Oikos (Vol. 112).
Parker, G. A. 1974. Assessment strategy and the evolution of animal conflicts. Journal of theoretical Biology 47:223–243.
Rowell, T. E. 1974. The concept of social dominance. Behavioral Biology 11.
Rutte, C., M. Taborsky, and M. W. G. Brinkhof. 2006. What sets the odds of winning and losing? Trends in Ecology and Evolution 21.
Schlinger, B. A., B. Palter, and G. V. Callard. 1987. A method to quantify aggressiveness in Japanese quail (Coturnix c. japonica). Physiology and Behavior 40.
Searcy, W. A., and S. Nowicki. 2010. The evolution of animal communication. Princeton University Press.
Shizuka, D., and D. B. McDonald. 2015. The network motif architecture of dominance hierarchies. Journal of the Royal Society Interface 12.
Simons, N. D., V. Michopoulos, M. Wilson, L. B. Barreiro, and J. Tung. 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.

Snyder-Mackler, N., J. R. Burger, L. Gaydosh, D. W. Belsky, G. A. Noppert, F. A. Campos, A. Bartolomucci, et al. 2020. Social determinants of health and survival in humans and other animals. Science.
Stevenson, P. A., H. A. Hofmann, K. Schoch, and K. Schildberger. 2000. The fight and flight responses of crickets depleted of biogenic amines. Journal of Neurobiology 43.
Strauss, E. D., A. R. Decasien, G. Galindo, E. A. Hobson, D. Shizuka, and J. P. Curley. 2022. DomArchive: a century of published dominance data. Philosophical Transactions of the Royal Society B: Biological Sciences 377.
Tibbetts, E. A., J. Pardo-Sanchez, and C. Weise. 2022. The establishment and maintenance of dominance hierarchies. Philosophical Transactions of the Royal Society B: Biological Sciences 377.
Van Doorn, G. S., G. M. Hengeveld, and F. J. Weissing. 2003. The evolution of social dominance I: Two-player models. Behaviour 140.
Whitehouse, M. E. A. 1997. Experience influences male-male contests in the spider Argyrodes antipodiana (Theridiidae: Araneae). Animal Behaviour 53.