Heterogeneous responsiveness to environmental stimuli*

Jerome Cavailles[1], Christoph Kuzmics[2], and Martin Grube[1]

[1]Institute of Biology, University of Graz, Graz, Austria
[2]Institute of Economics, University of Graz, Graz, Austria

March 29, 2024

Keywords: equilibrium, ESS, behavioral ecology, individual differences

Abstract

Individuals of a species cope with environmental variability through behavioral adjustments driven by individuals' responsiveness to environmental stimuli. Three key empirical observations have been made for many animal species: The *coexistence* of different degrees of responsiveness within one species; the *consistency* of an individual's degree of responsiveness across time; and the *correlation* of an individual's degree of responsiveness across contexts. Taking up key elements of existing approaches, we provide one unifying explanation for all three observations, by identifying a unique evolutionarily stable strategy of an appropriately defined game within a stochastic environment that has all three features. Coexistence is explained by a form of negative frequency dependence. Consistency and correlation is explained through potentially small, individual, differences of states animals have and the resulting differential advantages they can get from it. Our results allow us to identify a variety of testable implications.

*We are grateful to Erol Akçay, Adam Clark, Bernat Corominas-Murtra, Niels Dingemanse and his research group, Michael Greinecker, Josef Hofbauer, Laurent Lehmann, John McNamara, Jorge Peña and two anonymous reviewers as well as seminar participants at the Institute for Advanced Studies, Toulouse, the participants at COLIBRI focus workshop on Game Theory in Biology in Graz and the members of the field of excellence initiative COLIBRI at the University of Graz for helpful comments and suggestions.

1 Introduction

Individuals generally cope with frequent changes in their environment [Candolin et al., 2023] by means of behavioral adjustments, see e.g., Coulson et al. [2017], Sih et al. [2012], Wong and Candolin [2015], Sih [2013]. To be able to adjust their behavior to changes in the environment, individuals need to be responsive to environmental stimuli, observable indicators of these changes.[1]

It is well documented that individuals differ in the degree of responsiveness to external stimuli, a phenomenon sometimes referred to as behavioral plasticity, see e.g., Mitchell and Houslay [2021], Laskowski et al. [2022]. More generally, *phenotypic plasticity* – which use the alterations in an organism's behavior, morphology, and physiology as a response to its specific environmental conditions – have been extensively studied. Phenotypic plasticity was once undervalued and largely overlooked as a mechanism and concept of evolution. However, a shift in perspective is occurring due to recent theoretical and empirical studies. These studies highlight the importance of plasticity in fostering novelty and diversity across morphological, physiological, behavioral, and life history traits. For recent overviews, see Pfennig [2021] and Sommer [2020]. The concept of reaction norms originates from Woltereck [1909].

This difference in environmental responsiveness constitutes the main characteristic of personalities, and personalities have been observed in more than 100 species, see e.g., the survey by Nettle and Penke [2010]. In a broader sense, animal personality can be related to plasticity [Dingemanse et al., 2010]. Three key observations have been made for many animal species, as, for instance, highlighted by Wolf et al. [2007], see also Bell et al. [2009]: The *coexistence* of different degrees of responsiveness within one species; the *consistency* of an individual's degree of responsiveness across time; and the consistency, often referred to as *correlation*, of an individual's degree of responsiveness across contexts.[2]

A few theoretical approaches explain one or more of these three observations. The theory of biological sensitivity to context, as in Pluess [2015], Boyce [2016], and Ellis and Boyce [2008] explains the coexistence of different degrees of environmental responsiveness with differences in individuals' experiences in their early development. The theory of differential sensitiv-

[1] Stamps [2016] refers to "[t]he extent to which the phenotype of an agent varies as an immediate response to variation in external stimuli" as "contextual plasticity."

[2] The "suite of correlated behaviors [...] reflecting the individual consistency across [...] situations" has been referred to as a "behavioral syndrome" by Sih et al. [2004].

ity, as in Ellis et al. [2011], interprets the difference of behaviour as a way to hedge future generations against the uncertainty in the environment, recently formalised by Bergstrom [2014] and Frankenhuis et al. [2016].[3] A third approach is built around the idea of "negative frequency dependence:" The more individuals are responsive to environmental stimuli the less the benefits of being responsive. Negative frequency dependence is a cornerstone for explanations of the coexistence of different degrees of environmental responsiveness in the seminal models of Wolf et al. [2008], Wolf et al. [2011], and Wolf and McNamara [2012]. See Dingemanse and Wolf [2010] for a review of earlier models. While negative frequency dependence is able to explain coexistence, consistency and correlation are often explained by an individual's state (e.g., morphology, phenotype, size, etc.), as in Wolf and Weissing [2010] and Dingemanse and Wolf [2010].[4] However, the meta analysis by Niemelä and Dingemanse [2018] shows a weak link between state and personalities (individuals' state can only explain between 3 and 8% of the personality differences).

We study individual responsiveness to environmental stimuli for the specific problem of foraging from multiple food sources, which allows us to additionally build on the existing game-theoretic literature on the *ideal free distribution* of Fretwell and Lucas [1969], see also Křivan et al. [2008]: individuals allocate themselves proportionally to the amount of food available at each food source. For a more general framework about foraging theory see Stephens and Krebs [1986], Pyke [2019]). We enrich these game-theoretic models by embedding them in a stochastic environment, which allows us to incorporate the salient features of the three approaches mentioned above. We do this in a few steps of varying complexity.

The basic model (provided in Section 2) is sketched in Figure 1. To illustrate the model we use fish-feeding birds as an example (inspired from observations of Wakefield et al. [2015], McHuron et al. [2018], Harris et al. [2020], Patrick et al. [2014]), with the simplifying condition that they do not show any social behavior (such as flocking or swarming). Individuals have to choose to forage from one of two food sources, one providing a fixed amount of food, the other a random amount of food. Individuals can, at some small cognitive cost, learn about the state of the food availability at the random

[3] Pluess [2015], also Mitchell and Houslay [2021], argue in favor of an integration of the two theories.

[4] Wolf et al. [2008], Ehlman et al. [2022] explain consistency and correlation with a "positive-feedback mechanism:" responsiveness is less costly for individuals that have been responsive before. Wolf and McNamara [2012] explain consistency and correlation by small variations of individuals' metabolism (which is a form of state).

source. All individuals who go to the same food source are assumed to share the available food there equally. In the basic model the random food source can only have two possible levels of food availability.

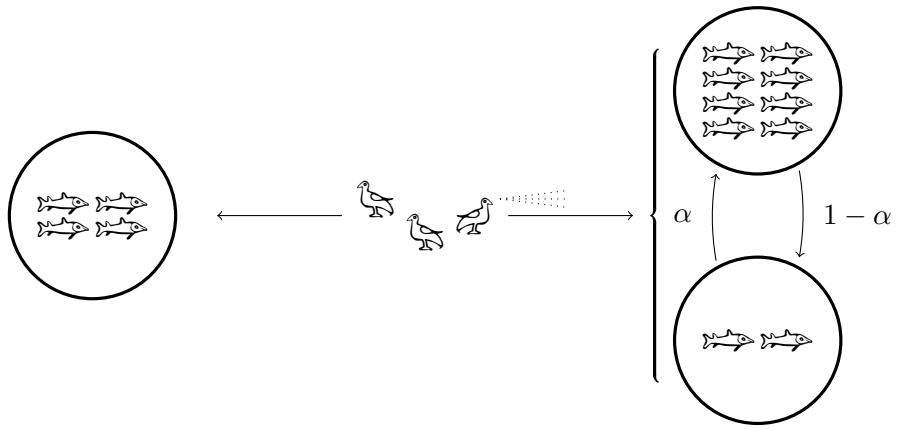

Figure 1: A graphical sketch of the basic model. The birds are the players, who choose which food source to go to: fixed food source A on the left or stochastic food source B on the right, with α the probability of high food availability. The bird's scanning indicates that players can make their decisions based on observing the food availability at the stochastic source. In this concrete example, at the equilibrium, one bird should always go to the left, one bird should always go to the right, and one bird should be responsive to change. They will all eat two fishes when the stochastic source is low, and four fishes when the stochastic source is high.

Our first main result is that the game always has a unique symmetric Nash equilibrium strategy, Nash [1950], and this strategy exhibits coexistence of responsive and non-responsive individuals. This equilibrium is also evolutionarily stable, an ESS, in the sense of Palm [1984], who extended the definition of an ESS by Maynard Smith and Price [1973] to symmetric n-player games, see also Broom et al. [1997]. Moreover, the games are *stable* in the sense of Hofbauer and Sandholm [2009], which implies that the unique symmetric Nash equilibrium is also asymptotically stable under most plausible behavioral adjustment dynamics, including the prominent replicator dynamics of Taylor and Jonker [1978].

The equilibrium is in mixed strategies and individuals are indifferent between being responsive and not being responsive. This raises concerns about the consistency of their behavior (across time and contexts). We then

introduce small perturbations to the model by allowing slight differences in individual-specific food preferences and in the individuals' cost of being responsive. This is the idea of Harsanyi [1973] *purification* and the idea of *threshold decisions* of McNamara and Houston [2005]: The original mixed strategy equilibrium (or ESS) corresponds in the perturbed game to an equilibrium in which every individual has a strict preference to either be or be not responsive. To the extent that these idiosyncratic payoff perturbations are constant over time and across contexts we now also get *consistency* and *correlation* in addition to *coexistence*.

We also provide an analytic expression for the equilibrium pure strategy frequencies as a function of the parameters of the problem. This allows us to derive additional testable predictions of our model. Acknowledging that our model is a highly simplified account of reality, some of these predictions may yet hold beyond the narrower confines of our model. In particular, and perhaps most striking: the unique equilibrium does not change with the stochasticity of the environment, at least when the cost of cognition is negligible. This implies that changes to the stochastic process of the environment may not push behavior out of equilibrium. Individuals' strategies are already sufficiently complex to allow essentially immediate and automatic adaptation to such changes. Of course, this does not imply that individuals are not affected if, for instance, high food availability becomes rarer. Only their strategy is unaffected, not necessarily the amount of food they can consume.

Another testable prediction can be derived from an extension of the basic model, in which individuals only receive (private) noisy information about the state of food availability. At least for small amounts of noise, the equilibrium responsiveness increases when the noise increases. More noise, in some sense, forces individuals to overreact to environmental stimuli. A final testable prediction, again derived from an extension of our basic model, is that for more general distributions of food availability, one would expect to see a continuum of degrees of responsiveness to environmental stimuli. See e.g., Lionetti et al. [2018] for empirical support for this finding.

The rest of the paper proceeds as follows. Section 2 provides the basic game-theoretic model. Section 3 derives the unique symmetric Nash equilibrium of the relevant games. This equilibrium exhibits "coexistence" of different degrees of environmental responsiveness. Section 4 provides a further discussion of the implications of equilibrium behavior. Section 5 then extends the model slightly to accommodate both "consistency" and "correlation" of individuals' degree of responsiveness, across time and contexts, respectively. Section 6 considers the case of noisy information, and Section 7

provides a generalisation of the basic model to general distributions of food availability at the stochastic source. Section 8 finally concludes with a further discussion of the results and further related literature. The Appendix provides some of the more technical arguments behind some of the results in the paper.

2 Methods

In this section, we present the simplest possible model of interest for our problem. Each of $n < \infty$ individuals can go to one of two food sources A or B. Food source A has a fixed amount of food normalized to n units of nutrition to facilitate an easier comparison when we vary the number of individuals n. Food source B has a stochastic amount of food $n \cdot X$, with X drawn from a Bernoulli distribution with $X = \eta$ with probability α and $X = \lambda$ with probability $1 - \alpha$ with λ denoting a low and η a high food-availability at the stochastic food source B. We assume that $\frac{1}{n} < \lambda < \eta < n$. If the number of individuals n is large, the assumption represents barely a restriction. If n is not very large, the assumption represents the most interesting case. Suppose, for instance, that both $\lambda, \eta < \frac{1}{n}$. This would imply that, the food-availability at the stochastic source B is always less than 1, while the food-availability at source A is n. Thus, as we are assuming equal food-sharing (see below), no individual would ever benefit from going to the stochastic food source. Even if all individuals go to food-source A an individual who switches to the stochastic source B gets a food-share of 1 only. Suppose, as another extreme, that both $\lambda, \eta > n$. Then the food-availability at the stochastic source B always exceeds n^2 and, therefore, no individual would ever benefit from going to the fixed source A. They get more than n units of nutrition at food source B and only at most n at food source A.

Before making their choice of food source, individuals can, in principle, inform themselves about the state of food source B; individuals can choose to learn whether $X = \lambda$ or $X = \eta$. If uninformed they can then choose to go to food source A or B. If informed they can react to the information in one of two ways. [5] They can be, what we term, *responsive* by going to A when

[5] In principle, they could also become informed and then ignore their information, by going to A or B regardless of the information they have received. But we assume that the choice of becoming informed bears an arbitrarily small cost $c > 0$. The assumption of positive costs implies that the two strategies of getting informed and then ignoring the information are strictly dominated by the strategy of not getting informed and going to the same food source.

$X = \lambda$ and going to B when $X = \eta$, or, what we term *counter-responsive* by doing the opposite. We denote the set of strategies by $S = \{A, B, R, C\}$, for always going to food source A, always to B, being responsive, and being counter-responsive, respectively.

Symbol	Description	Range
	parameters	
n	number of individuals	$\mathbf{N}$
λ	lower amount of food at the variable food source	$\mathbf{R}^{*+}$
η	higher amount of food at the variable food source	$\mathbf{R}^{*+}$
α	probability to have a high amount of food at the source B	[0, 1]
c	cost for being responsive (or counter-responsive)	$\mathbf{R}^{+}$
X	Stochastic variable of the amount of food at the source B	$\{\lambda, \eta\}$
	strategies	
S	Set of strategies	
A	Strategy *always go to food source A*	
B	Strategy *always go to food source B*	
R	Strategy *being responsive*	
C	Strategy *being counter-responsive*	

Table 1: Overview of key notation

Let $\Delta(S)$ denote the set of all mixed strategies, that is the set of all probability distributions over S. There are two ways to interpret a mixed strategy $\sigma \in \Delta(S)$, both go back at least to Maynard Smith and Price [1973] who speak both of the *probability* and the *frequency* of certain actions. One can think of a mixed strategy as a way an individual randomizes over pure strategies. These pure strategies are then played with the probability specified by σ. But one can also think of a mixed strategy as resulting from a random draw of a large (essentially infinite) population of individuals with given proportions of them playing the various pure strategies. These pure strategies are then present in the population with frequencies given by σ. While mathematically equivalent, this distinction matters for how our results are able to explain both the consistency and correlation of behaviour, which we investigate in the respective section below.

We assume throughout the paper that all individuals, who go to the same food source, share the available food at this source equitably. The payoff to an individual then only depends on the number of other individuals, k, that go to food source A (which implies that $n - 1 - k$ others go to food source B). The payoff to an individual who goes to food source A is given by $\frac{n}{k+1}$; the payoff to an individual who goes to B is given by $\frac{nX}{n-k}$.

Given that individuals can choose mixed strategies, we need to compute individuals' expected payoffs. To do so, consider an arbitrary individual who is facing that all other $n-1$ individuals choose a given (mixed) strategy $\sigma \in \Delta(S)$ with $\sigma(s)$ the probability that pure strategy s is chosen.

Denote by N_s the random variable that is the number of opponents who end up choosing pure strategy $s \in S$ (given the probability of choosing s is $\sigma(s)$). The tuple (N_A, N_B, N_R, N_C) is then multinomially distributed with the probability vector $(\sigma(A), \sigma(B), \sigma(R), \sigma(C))$.

Given σ, let $R_{A\lambda}$ denote the random variable that is the food share available at food source A in state λ. Let food shares $R_{A\eta}$, $R_{B\lambda}$, $R_{B\eta}$, be defined analogously. These are given by

$$\begin{array}{rclrcl} R_{A\lambda} & = & \frac{n}{N_A+N_R+1} & R_{A\eta} & = & \frac{n}{N_A+N_C+1} \\ R_{B\lambda} & = & \frac{\lambda \cdot n}{N_B+N_C+1} & R_{B\eta} & = & \frac{\eta \cdot n}{N_B+N_R+1}. \end{array}$$

We can, then, express an individual's expected payoffs from choosing pure strategy $s \in S$, when all others use mixed strategy σ, as follows.

$$\begin{array}{rclcclcl} u(A,\sigma) & = & \alpha \mathbf{E}\left[R_{A\eta}\right] & + & (1-\alpha)\mathbf{E}\left[R_{A\lambda}\right] & & \\ u(B,\sigma) & = & \alpha \mathbf{E}\left[R_{B\eta}\right] & + & (1-\alpha)\mathbf{E}\left[R_{B\lambda}\right] & & \\ u(R,\sigma) & = & \alpha \mathbf{E}\left[R_{B\eta}\right] & + & (1-\alpha)\mathbf{E}\left[R_{A\lambda}\right] & - & c \\ u(C,\sigma) & = & \alpha \mathbf{E}\left[R_{A\eta}\right] & + & (1-\alpha)\mathbf{E}\left[R_{B\lambda}\right] & - & c, \end{array} \tag{1}$$

where $\mathbf{E}$ denotes the expectation with respect to the randomness created by mixed strategy σ.[6] We extend an individual's payoff function to mixed strategies by taking expectations:

$$u(\sigma', \sigma) = \sum_{s \in S} u(s, \sigma)\sigma'(s).$$

Strategy $\sigma \in \Delta(S)$ is a *symmetric equilibrium* [Nash, 1950] if $u(\sigma, \sigma) \geq u(\sigma', \sigma)$ for all $\sigma' \in \Delta(S)$.

3 Equilibrium

We first show that pure strategies A, B, and R, cannot be equilibria for any $c > 0$ and for any n given the assumption that $\frac{1}{n} < \lambda < \eta < n$. Suppose everyone plays A. Then everyone gets a payoff of $\frac{1}{n}$. An individual deviating to A would at least get λ and since $\lambda > \frac{1}{n}$ this is a profitable deviation.

[6] For instance, $\mathbf{E}\left[R_{A\lambda}\right] = \sum_{k_A=0}^{n-1} \sum_{k_R=0}^{n-1-k_A} \frac{n}{k_A+k_R+1} \frac{(n-1)!}{k_A!k_R!(n-k_A-k_R)!} \sigma(A)^{k_A} \sigma(R)^{k_R} (1-\sigma(A)-\sigma(R))^{n-k_A-k_R}$.

Suppose everyone plays B. Then everyone gets at most $\eta < n$ (if $X = \eta$ and less if $X = \lambda$). Deviating to B would yield a payoff of n. Suppose, finally, that everyone plays R. Then everyone gets a payoff of either 1 (if $X = \lambda$ and everyone goes to food source A) or η (if $X = \eta$ and everyone goes to food source B), minus c, and then a deviation to B gives a higher payoff of $\lambda n > 1$ when $X = \lambda$ (this individual is alone at food source B) and a payoff of $n > \eta$ if $X = \eta$ (this individual is alone at food source A), and thus, a strictly higher expected payoff.

Next we show that there cannot be a symmetric Nash equilibrium that attaches non-zero probability to both pure strategies R and C. If this were the case, then necessarily $u(C, \sigma) = u(R, \sigma)$ and, by the fact that $u(A, \sigma) + u(B, \sigma) = u(R, \sigma) + u(C, \sigma) - 2c$ for all $\sigma \in \Delta(S)$, either pure strategy A or B (or both) provide a strictly higher payoff than both R and C.

There cannot be a symmetric Nash equilibrium in which strategy C is used with non-zero probability. To see this, we show that whenever C is used R must be used with non-zero probability. As we just showed that there cannot be an equilibrium with both C and R, this then proves that there is no equilibrium in which C is used.

We prove this by contradiction. Suppose, that $\sigma(R) = 0$ and $\sigma(C) > 0$ for an equilibrium $\sigma \in \Delta(S)$. We will show that, in this case, either strategy A or B will provide a strictly better payoff than C does. Noting that $N_R = 0$, we can write the expected payoffs as

$$\begin{array}{rcccccc} u(A,\sigma) & = & \alpha\mathbf{E}\left[\frac{n}{N_A+N_C+1}\right] & + & (1-\alpha)\mathbf{E}\left[\frac{n}{N_A+1}\right] & & \\ u(B,\sigma) & = & \alpha\mathbf{E}\left[\frac{\eta\cdot n}{N_B+1}\right] & + & (1-\alpha)\mathbf{E}\left[\frac{\lambda\cdot n}{N_B+N_C+1}\right] & & \\ u(C,\sigma) & = & \alpha\mathbf{E}\left[\frac{n}{N_A+N_C+1}\right] & + & (1-\alpha)\mathbf{E}\left[\frac{\lambda\cdot n}{N_B+N_C+1}\right] & - & c. \end{array}$$

As σ is an equilibrium we must have $u(A, \sigma), u(B, \sigma) \leq u(C, \sigma)$. This implies that

$$\mathbf{E}\left[\frac{\lambda}{N_B+N_C+1}\right] \geq \mathbf{E}\left[\frac{1}{N_A+1}\right]$$

$$\mathbf{E}\left[\frac{1}{N_A+N_C+1}\right] \geq \mathbf{E}\left[\frac{\eta}{N_B+1}\right]$$

But, since we have $\mathbf{E}[N_C] > 0$, we have

$$\mathbf{E}\left[\frac{1}{N_A+1}\right] > \mathbf{E}\left[\frac{1}{N_A+N_C+1}\right] \text{ and } \mathbf{E}\left[\frac{\lambda}{N_B+1}\right] > \mathbf{E}\left[\frac{\lambda}{N_B+N_C+1}\right],$$

which, in combination with the previous two inequalities yields

$$\begin{array}{rcccc} \mathbf{E}\left[\frac{\lambda}{N_B+1}\right] & > & \mathbf{E}\left[\frac{\lambda}{N_B+N_C+1}\right] & \geq & \mathbf{E}\left[\frac{1}{N_A+1}\right] \\ & > & \mathbf{E}\left[\frac{1}{N_A+N_C+1}\right] & \geq & \mathbf{E}\left[\frac{\eta}{N_B+1}\right], \end{array}$$

and, thus,

$$\mathbf{E}\left[\frac{\lambda}{N_B+1}\right] > \mathbf{E}\left[\frac{\eta}{N_B+1}\right],$$

which, finally, contradicts our assumption that $\eta > \lambda$.

Next we show that, given the assumption that $\frac{1}{n} < \lambda < \eta < n$, there cannot be any mixed strategy equilibrium with only two pure strategies in its support. We have already shown that there cannot be an equilibrium with C in its support. Suppose there is an equilibrium σ that has a support of $\{A, R\}$. Then $u(A, \sigma) = u(R, \sigma)$. However, $u(B, \sigma) > u(R, \sigma)$, as B provides the same payoff as R when $X = \eta$ and obtains λn when $X = \lambda$ (as everyone else is at source A), while R obtains $1 < \lambda n$. Suppose there is an equilibrium σ that has a support of $\{B, R\}$. Then $u(B, \sigma) = u(R, \sigma)$. However, $u(A, \sigma) > u(R, \sigma)$, as A provides the same payoff as R when $X = \lambda$ and obtains n when $X = \eta$ (as everyone else is at source B), while R obtains $\eta < n$.

Suppose there is an equilibrium σ that has a support of $\{A, B\}$. Then $u(A, \sigma) = u(B, \sigma)$. As $\eta > \lambda$, then for sufficiently small costs c, R provides a higher payoff than both A and B. In expectation both food sources A and B have the same food share, but R provides a higher payoff by going to B if and only if $X = \eta$.

This, finally, leaves as the only possible equilibria those with support given by $\{A, B, R\}$. Let $r_{yx} = \mathbf{E}[R_{yx}]$ denote the expected food share at source $y \in \{A, B\}$ in state $x \in \{\lambda, \eta\}$. Consider $r_{A\lambda}$. All other expected food shares are calculated analogously. We can write $r_{A\lambda} = n\mathbf{E}[\frac{1}{Y+1}]$, where $Y = N_A + N_R$, and Y, thus, binomially distributed with $n-1$ tries and success probability $p = \sigma(A) + \sigma(R)$. The first negative moment of such a binomial distribution, $\mathbf{E}[\frac{1}{Y+1}]$, is provided in [Chao and Strawderman, 1972, Equation 3.4] and given by $\frac{1-(1-p)^n}{np}$.

We, thus, obtain

$$\begin{aligned}
r_{A\lambda} &= \tfrac{1-(1-\sigma(A)-\sigma(R))^n}{\sigma(A)+\sigma(R)}, \\
&\text{and analogously} \\
r_{A\eta} &= \tfrac{1-(1-\sigma(A)-\sigma(C))^n}{\sigma(A)+\sigma(C)} \\
r_{B\lambda} &= \lambda \cdot \tfrac{1-(1-\sigma(B)-\sigma(C))^n}{\sigma(B)+\sigma(C)} \\
r_{B\eta} &= \eta \cdot \tfrac{1-(1-\sigma(B)-\sigma(R))^n}{\sigma(B)+\sigma(R)}.
\end{aligned} \tag{2}$$

We need to show that there is a unique equilibrium with support $\{A, B, R\}$. To be indifferent between A and R and B and R requires that

$$r_{A\eta} = r_{B\eta} - \frac{c}{\alpha} \text{ and } r_{B\lambda} = r_{A\lambda} - \frac{c}{1-\alpha}.$$

Plugging in the above expressions, and noting that $\sigma(C) = 0$, these two equations lead to

$$\frac{1-(1-\sigma(A))^n}{\sigma(A)} = \eta\frac{1-\sigma(A)^n}{1-\sigma(A)} - \frac{c}{\alpha}, \tag{3}$$

and

$$\lambda\frac{1-(1-\sigma(B))^n}{\sigma(B)} = \frac{1-\sigma(B)^n}{1-\sigma(B)} - \frac{c}{1-\alpha}, \tag{4}$$

respectively.

In general, we cannot provide an analytic solution to this system of equations. If, however, we consider the limit of the number of individuals going to infinity, $n \to \infty$, the system of equations simplifies somewhat to

$$\frac{1}{\sigma(A)} = \eta\frac{1}{1-\sigma(A)} - \frac{c}{\alpha}, \tag{5}$$

and

$$\lambda\frac{1}{\sigma(B)} = \frac{1}{1-\sigma(B)} - \frac{c}{1-\alpha}, \tag{6}$$

respectively. These are quadratic equations that have the following unique closed form solution (in the relevant range between 0 and 1).

$$\begin{aligned}
\sigma(A) &= \tfrac{1}{2c}\left[\sqrt{(\alpha(1+\eta)-c)^2+4\alpha c} - (\alpha(1+\eta)-c)\right] \\
\sigma(B) &= \tfrac{1}{2c}\left[\sqrt{((1-\alpha)(1+\lambda)-c)^2+4(1-\alpha)\lambda c} - ((1-\alpha)(1+\lambda)-c)\right] \\
\sigma(R) &= \tfrac{1}{2c}\left[1+\alpha\eta+(1-\alpha)\lambda - \sqrt{(\alpha(1+\eta)-c)^2+4\alpha c)}\right. \\
&\quad \left. - \sqrt{((1-\alpha)(1+\lambda)-c)^2+4\lambda(1-\alpha)c)}\right]
\end{aligned}$$

If, moreover, we consider the limit case of cognition costs going to zero, $c \to 0$, the system of equations further simplifies to

$$\frac{1}{\sigma(A)} = \eta\frac{1}{1-\sigma(A)} \text{ and } \lambda\frac{1}{\sigma(B)} = \frac{1}{1-\sigma(B)}, \text{ respectively.}$$

The unique solution is then $\sigma(A) = \frac{1}{1+\eta}$ and $\sigma(B) = \frac{\lambda}{1+\lambda}$. This is the only candidate for an equilibrium of this form. As $\sigma(R) = 1 - \sigma(A) - \sigma(B)$ we obtain $\sigma(R) = \frac{\eta-\lambda}{(1+\eta)(1+\lambda)}$. Note, finally, that $\sigma(A), \sigma(B), \sigma(R) \geq 0$, and, thus σ is indeed the unique (mixed strategy) equilibrium.

4 Equilibrium Implications

Note first that the equilibrium frequency $\sigma(A)$ of always going to the constant food source A does not depend on λ, the low state of food-availability in the stochastic source: $\sigma(A)$ does not change when λ changes. This is not only true in the case of $c \to 0$ and $n \to \infty$, but generally true. This can be seen by the fact that Equation 3 is an equation just in $\sigma(A)$ (not also in $\sigma(B)$ or $\sigma(R)$) and the parameter λ does not appear in this equation. Similarly, the equilibrium frequency $\sigma(A)$ of always going to the stochastic food source B does not depend on η, the high state of food-availability in the stochastic source. Only the equilibrium frequency, $\sigma(R)$ of being responsive changes with both levels of food-availability of the stochastic source. For a fixed high level of food-availability η at the stochastic source, increasing λ leaves $\sigma(A)$ constant, increases $\sigma(B)$, and decreases $\sigma(R)$, the frequency of responsive individuals. Analogously, for a fixed low level of food-availability λ at the stochastic source, increasing η decreases $\sigma(A)$, leaves $\sigma(B)$ constant, and increases $\sigma(R)$. All these findings are illustrated in Figure 2. If both η increases and λ decreases (i.e., the stochastic food source becomes more extreme) then $\sigma(A)$ and $\sigma(B)$ decrease and, therefore, the equilibrium frequency of responsive individuals, $\sigma(R)$ increases.

The second thing to note is that in the limit that the cost of cognition becomes negligible (relative to the other parameters), that is when $c \to 0$ (for any n), the equilibrium frequencies do not depend on the stochastic nature of the environment, i.e., do not depend on α. This can be seen by setting $c = 0$ in equations 3 and 4, with the result that α drops out of these equations.

If the cost of cognition, c, is non-negligible, then the equilibrium frequencies depends on this cost and on α, the stochastic nature of the environment, as illustrated in Figure 3.

In Appendix C we show that the unique symmetric Nash equilibrium in this model is also an evolutionarily stable strategy, in the sense of Palm [1984], see also Broom et al. [1997], who have extended the definition of an ESS of Maynard Smith and Price [1973] to symmetric n-player games. We know generally, see e.g., Nachbar [1990], that Nash equilibria are the only candidates for asymptotically stable rest points under most deterministic behavioral adjustment (or evolutionary) dynamics, with the replicator dynamics of Taylor and Jonker [1978] the first and most prominent example. In Appendix D we show that the games we here study are stable games in the sense of Hofbauer and Sandholm [2009], with the implication that the unique symmetric equilibrium of our model is asymptotically stable under

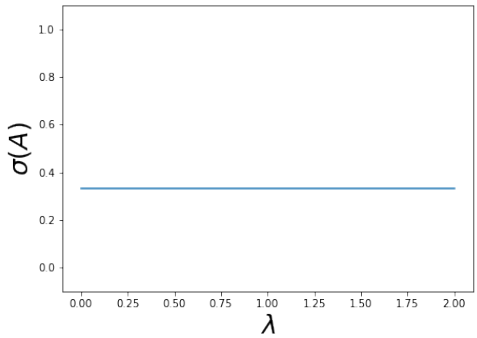

(a) $\sigma(A)$ as a function of λ (with $\eta = 2$, in the limit case $c \to 0$ and $n \to \infty$).

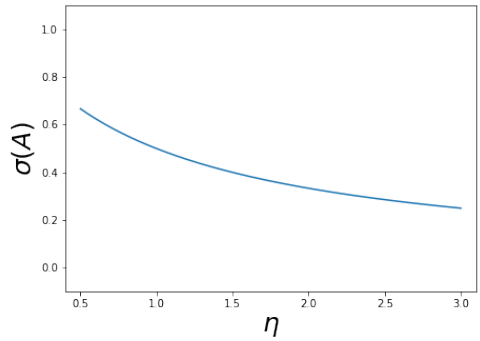

(b) $\sigma(A)$ as a function of η (with $\lambda = 0.5$, in the limit case $c \to 0$ and $n \to \infty$).

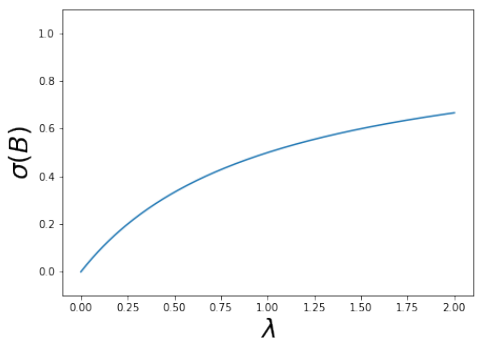

(c) $\sigma(B)$ as a function of λ (with $\eta = 2$, in the limit case $c \to 0$ and $n \to \infty$).

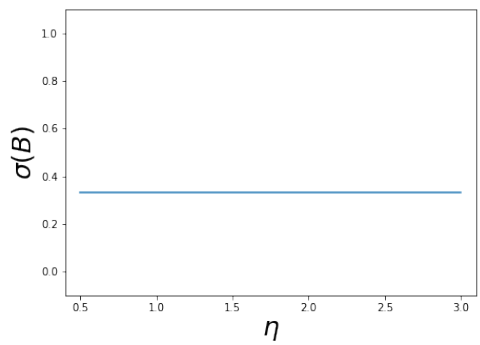

(d) $\sigma(B)$ as a function of η (with $\lambda = 0.5$, in the limit case $c \to 0$ and $n \to \infty$).

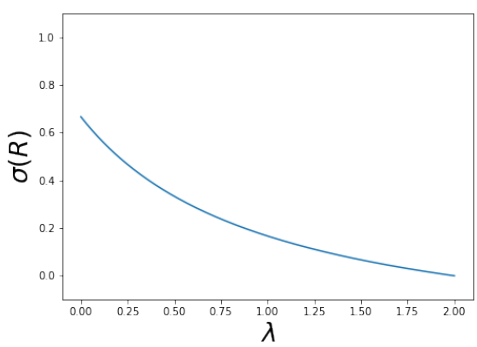

(e) $\sigma(R)$ as a function of λ (with $\eta = 2$, in the limit case $c \to 0$ and $n \to \infty$).

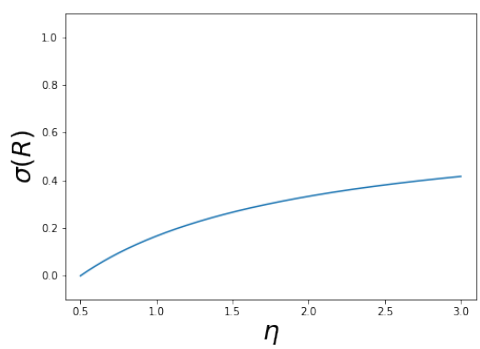

(f) $\sigma(R)$ as a function of η (with $\lambda = 0.5$, in the limit case $c \to 0$ and $n \to \infty$).

Figure 2

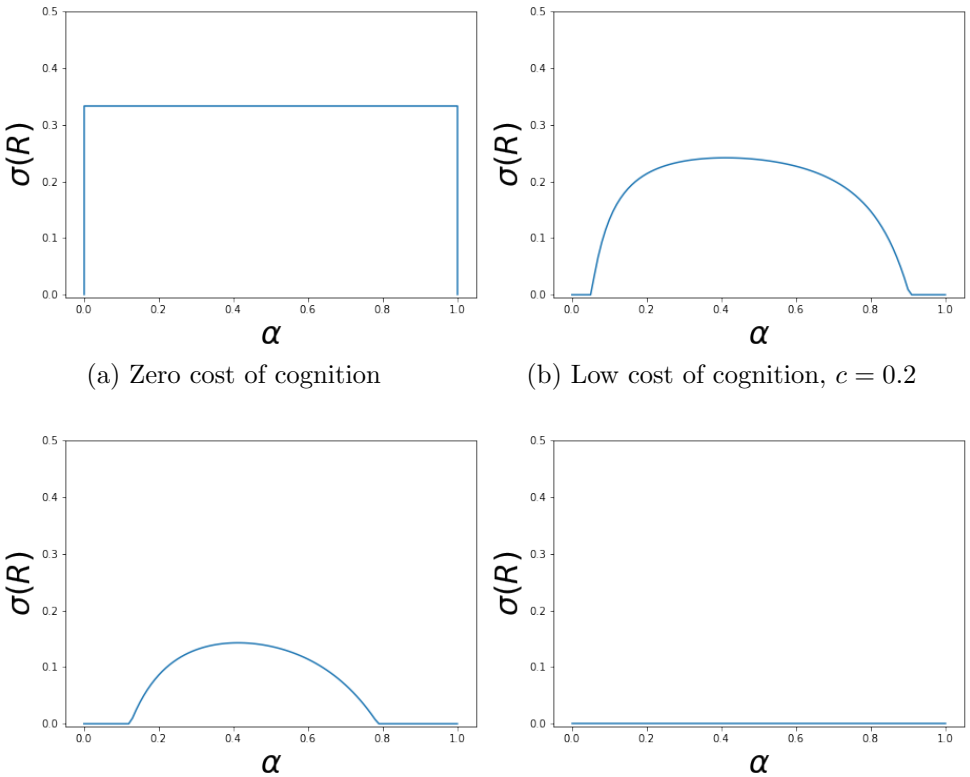

(a) Zero cost of cognition

(b) Low cost of cognition, $c = 0.2$

(c) Higher cost of cognition, $c = 0.4$

(d) Prohibitive cost of cognition

Figure 3: Probability or frequency of being responsive, $\sigma(R)$ as a function of the probability, α of the high food availability at the stochastic food source, for different levels of the cost of cognition, c ($\lambda = 0.5, \eta = 2$).

many different dynamics that includes the replicator dynamics. In Figure 4 we illustrate this finding by sketching the phase diagram of the replicator dynamics for different parameter configurations of our model.

5 Consistency and Correlation

Suppose that the n individuals play the same game given in our basic model over and over again for many periods of time. Suppose that, at every point in time, they play the unique symmetric equilibrium given earlier. An outside observer would note that when the amount of food available at B is high ($X = \eta$) more individuals are to be found at source B than when it is low ($X = \lambda$). They would also observe that the food share each individual receives is the same regardless of which source the individuals go to. The outside observer would conclude that some individuals must be responsive to the stochastic food availability at source B. However, there are, so far, no incentives for individuals to choose the same strategy across time and context.

Recall the two interpretations we can give a mixed strategy equilibrium, as we identified in this model. Suppose that we interpret it as the various individuals involved actually randomizing between the three pure strategies. In that case, the outside observer would note that the identity of responsive individuals varies over time. This, however, is inconsistent with empirical findings (of consistency and correlation as highlighted in the introduction). Alternatively we can interpret the mixed strategy equilibrium as there being a large population of individuals, with certain fractions of these individuals playing the various pure strategies. In this latter case, it is more plausible that if one individual is responsive at one point in time it is also responsive at another point in time. Note, however, that all individuals only have a very weak incentive to stay with their pure strategy choice, as (in any mixed equilibrium) all pure strategies provide the same expected benefit (food share in our model). In principle, the various individuals could switch pure strategies around, as long as the aggregate frequencies remain the same.

A slight change to the basic model, however, provides essentially the same prediction as the basic model, but now with every individual having a strict preference to play a pure strategy with frequencies very close to the mixed strategy equilibrium of the basic model.

This modification is based on the idea of *purification* of Harsanyi [1973], which is very similar to the idea of *threshold decisions* as provided in McNamara and Houston [2005]. The idea is that individuals differ a little bit

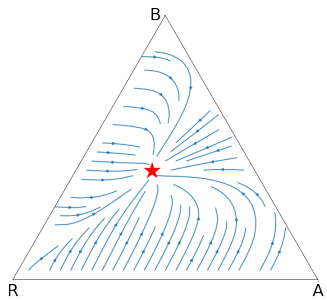

(a) $\alpha = 0.7, \lambda = 0.7, \eta = 3$

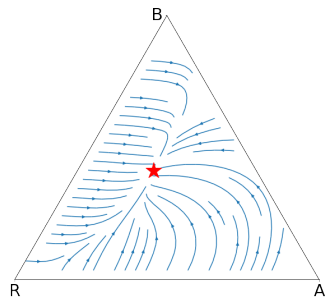

(b) $\alpha = 0.4, \lambda = 0.7, \eta = 3$

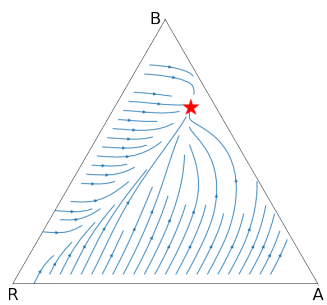

(c) $\alpha = 0.4, \lambda = 2, \eta = 3$

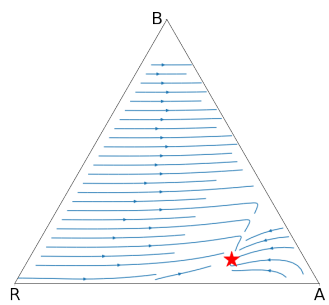

(d) $\alpha = 0.1, \lambda = 0.1, \eta = 0.5$

Figure 4: Phase diagram of the replicator dynamics for $c = 0$, under a variety of different parameter specifications. The only difference between figures a and b is that the stochasticity parameter α changes. This has no effect on the equilibrium itself but does affect somewhat the out of equilibrium dynamics. It can also be seen that not only is the unique equilibrium asymptotically stable, but also in fact a global attractor under the replicator dynamics: All solution paths eventually converge to the equilibrium.

in terms of their personal preferences and actually make a pure strategy choice that is, however, dependent on their own personal preferences that only they themselves know. As a consequence, while the equilibrium looks mixed to other individuals, each individual actually plays a pure strategy.

We adapt the model by replacing the payoff function u of the original model with a slightly perturbed payoff function v_θ that is essentially equal to u plus a small idiosyncratic (individual-specific) preference or perturbation term:

$$\begin{array}{rcccc} v_\theta(A,\sigma) & = & u(A,\sigma) & + & \theta_A \\ v_\theta(B,\sigma) & = & u(B,\sigma) & + & \theta_B \\ v_\theta(R,\sigma) & = & u(R,\sigma) & + & \alpha\theta_B + (1-\alpha)\theta_A + \theta_R \\ v_\theta(C,\sigma) & = & u(C,\sigma) & + & \alpha\theta_A + (1-\alpha)\theta_B + \theta_R, \end{array}$$

where the vector $\theta = (\theta_A, \theta_B, \theta_R)$ is i.i.d. drawn from some arbitrary full support continuous joint distribution F (with density f) over $\Theta = [-\epsilon, \epsilon]^3$, for a small $\epsilon > 0$. It is assumed that an individual's realized θ is that individual's private information, which means that only this individual knows it; it is unknown to other individuals. We have deliberately chosen the same preference perturbation θ_R for pure strategies C and R, as it seems more natural to have an idiosyncratic perturbation of the cost of being responsive rather than for how one is responsive. However, it does not matter what we assume for pure strategy C as long as the payoff perturbation is small, as pure strategy C provides a strictly lower payoff than the other three strategies in the equilibrium given earlier, and small payoff perturbations cannot change that.

In the perturbed game, a strategy is a function $\rho : \Theta \to \Delta(S)$, where ρ attaches to each possible preference type θ a probability distribution over S. Every such strategy ρ induces a strategy $\sigma^\rho \in \Delta(S)$ by $\sigma^\rho(s) = \int_\Theta \rho(\theta)(s) f(\theta) d\theta$ for all $s \in S$, where $\rho(\theta)(s)$ is the probability that $\rho(\theta)$ attaches to s. An equilibrium strategy ρ is then such that $\rho(\theta)(A) = 1$ if

$$\begin{array}{rcl} u(A,\sigma^\rho) - u(B,\sigma^\rho) & > & \theta_B - \theta_A, \text{ and} \\ u(A,\sigma^\rho) - u(R,\sigma^\rho) & > & \alpha(\theta_B - \theta_A) + \theta_R; \end{array}$$

$\rho(\theta)(B) = 1$ if

$$\begin{array}{rcl} u(A,\sigma^\rho) - u(B,\sigma^\rho) & < & \theta_B - \theta_A, \text{ and} \\ u(B,\sigma^\rho) - u(R,\sigma^\rho) & > & (1-\alpha)(\theta_A - \theta_B) + \theta_R; \end{array}$$

and $\rho(\theta)(R) = 1$ if

$$\begin{array}{rcl} u(A,\sigma^\rho) - u(R,\sigma^\rho) & < & \alpha(\theta_A - \theta_B) + \theta_R, \text{ and} \\ u(B,\sigma^\rho) - u(R,\sigma^\rho) & < & (1-\alpha)(\theta_A - \theta_B) + \theta_R. \end{array}$$

Harsanyi [1973] has shown that almost any equilibrium of a complete information game, such as our basic game, is such that any nearby incomplete information game, with payoff perturbations given by the joint distribution F, has a nearby equilibrium and that this nearby equilibrium is essentially in pure strategies. In such a nearby equilibrium there is a parameter region for $\theta \in \Theta$ for which an individual strictly prefers to play A, another region for which an individual strictly prefers to play B, and a final region in which an individual strictly prefers to play R. The set of θ's for which an individual is indifferent between two or three of the three strategies has measure zero (it is statistically impossible). Finally, such purified equilibria can also shown to be dynamically stable under a suitably defined behavioral adjustment dynamics as in Ely and Sandholm [2005], see Sandholm [2007]. In other words, the result of Harsanyi [1973] implies that each individual uses a pure strategy, which they strictly prefer given their own private preference, but the frequency of each strategy at the population level is essentially the same as without the preference perturbation.

Suppose now that the n individuals play the resulting equilibrium of the same slightly perturbed game repeatedly over many time periods. It is then a question of whether the perturbation parameters θ remain the same for each individual over time or not. If they do, it will be the same individuals who always go to food source A, the same individuals who always go to food source B, and the same who are responsive.

As an example of why the perturbed model may be appropriate for our purposes, consider birds who every day have to decide to go to food source A or B from their nesting place. Then the location of their nesting place gives rise to their θ. An approach could be that θ_A and θ_B are proportional to the distance that the bird's nest is from the two food sources, respectively, while θ_R could be more of a personal characteristic of the bird, measuring how much/less cognitively able this bird is relative to other birds.

This model is then also flexible enough to generate a strong consistency over time and a weaker, but some, consistency across contexts, depending on how these consistencies are interpreted. Consider the bird example again. One could imagine that θ_R is an individual bird's specific parameter that does not change over time nor across contexts. On the other hand, the parameters θ_A and θ_B might be constant for one season, but could be different in another season, when the bird's nest location (or the location of the food sources) changes.

Equilibrium purification could even be obtained by introducing a payoff-irrelevant personal and privately know characteristic, such as an individual's prior experiences in life, with individuals playing different pure strategies

depending on their personal prior experiences. This means that, as pointed out e.g., in Wolf et al. [2007], Wolf and Weissing [2010] and Dingemanse and Wolf [2010], the purification threshold could also be based on an individual's life history.

6 Imperfect private signals of food availability

In our basic model, individuals can learn the state of food availability at the food sources perfectly. In this section, we study how the results change if this learning is imperfect. To do so, we suppose that each individual receives a noisy signal about the level of food availability at food source B. Individuals i receive conditionally independent (and identically distributed) signals $s_i \in \{l, h\}$ such that $P(s_i = h|X = \eta) = P(s_i = l|X = \lambda) = 1 - \epsilon$, with $\epsilon < \frac{1}{2}$. In words, in the high state η the high signal h is more likely than the low signal l, and in the low state λ the low signal l is more likely than the high signal h. The signal is, thus, informative about the true state.

All the arguments made for the perfect signals of the food availability model (the basic model) go through (see Appendix B). Ultimately, we obtain that, provided the level of noise ϵ is not too large, the game has a unique symmetric Nash equilibrium, denoted σ_ϵ, that is also an ESS (and dynamically stable under a large class of evolutionary adjustment models). In the limiting case when the cost of cognition $c \to 0$ the equilibrium is given by

$$\begin{cases} \sigma_\epsilon(R) &= \frac{\eta-\lambda}{(1-2\epsilon)(1+\eta)(1+\lambda)} \\ \sigma_\epsilon(A) &= \frac{1}{1+\eta} - \epsilon\frac{\eta-\lambda}{(1-2\epsilon)(1+\eta)(1+\lambda)} \\ \sigma_\epsilon(B) &= \frac{\lambda}{1+\lambda} - \epsilon\frac{\eta-\lambda}{(1-2\epsilon)(1+\eta)(1+\lambda)}. \end{cases}$$

There are two interesting implications of this model, both seen also in Figure 5. First, for all parameter values λ, η, the frequency of responsive individuals increases with the level of noise ϵ in the information. The noisier the signal the more individuals are responsive to the possible erroneous information about the stochastic food source. The intuition behind this is, that because of the possible error, more individuals have to respond, because some will get the wrong signal and, thus, respond incorrectly. This would lead to not enough people responding appropriately. Second, the impact that an increase of the level of noise ϵ has on the frequency of responsive individuals is higher the higher the frequency of responsive individuals is without noise. This can be seen in Figure 5, and also formally, if we write

$$\sigma_\epsilon(R) = \frac{1}{1-2\epsilon}\sigma_0(R),$$

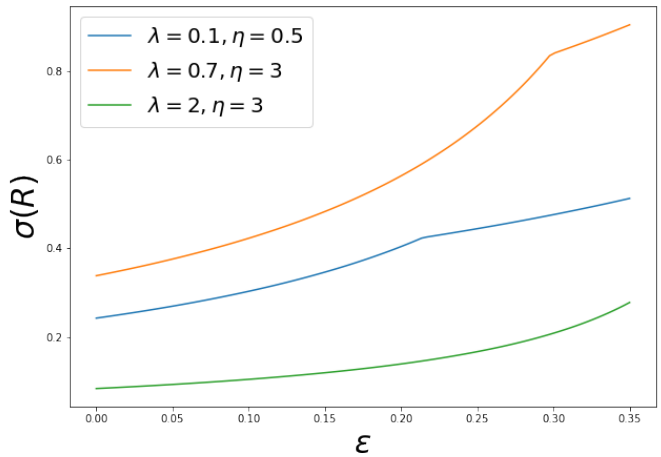

Figure 5: The proportion of responsive individuals increases with the level of noise in the information. The kink in the graph is at the level of noise ϵ at which one of the two equilibrium frequencies $\sigma_\epsilon(A)$ or $\sigma_\epsilon(B)$ become zero.

and, taking the derivative with respect to ϵ obtain

$$\frac{d}{d\epsilon}\sigma_\epsilon(R) = \frac{2}{(1-2\epsilon)^2}\sigma_0(R).$$

Interestingly, these results remain valid even when considering a strictly positive cognition cost.

7 General distributions of food source availability

Recall that in the basic model food source B is assumed to be Bernoulli distributed (i.e., with only two possible levels of available food at that source). In this section we consider an arbitrary distribution for the food availability at food source B.

Let X, the available quantity of food at food source B, be distributed according to some distribution with cdf F with everywhere positive density f on the interval $\chi = [x_L, x_H]$ with $0 \leq x_L < x_H \leq \infty$. To make the analysis tractable we simplify the model in two ways. First, we set the cost of being

responsive, c, to zero.[7] Second, we assume that all individuals learn the value of X, and allow individuals to only use monotone strategies: An individual's strategy can be described by a cutoff value $y \in \chi$ such that the individual goes to food source A if and only if $x < y$. Otherwise the individual goes to food source B. This implies that the strategy space is identical to χ and the set of mixed strategies is the set $\Delta(\chi)$ of all probability distributions over χ. A fully mixed symmetric Nash equilibrium strategy, which can be described by a cdf G on χ must satisfy that any individual is indifferent between using any any cut-off $y \in \chi$.

We then get the following result. In the model of this section, for any n, there is a unique completely mixed symmetric equilibrium. In the limit as n tends to infinity, the equilibrium probability that an individual uses cut-off responsiveness y is given by the cdf $G(y) = \frac{y}{1+y}$, with $G(x_L) = \frac{x_L}{1+x_L}$ the probability of an individual always going to food source B, $1 - G(x_H) = \frac{1}{1+x_H}$ the probability of an individual always going to food source A, and $G(y) - G(x) = \frac{y-x}{(1+x)(1+y)}$ the probability that an individual adopts a degree of responsiveness in the interval $[x, y]$.[8]

This finding is consistent with those in the basic model. For example, the strategy called B in the previous model is similar to choosing the cut-off x_L (since x_L is the minimum possible value of the stochastic source). The equilibrium frequency of this strategy in the basic model is given by $\frac{\lambda}{1+\lambda}$, which is equivalent to $\frac{x_L}{1+x_L}$, as x_L is the smallest possible value of the stochastic source. This generalization to the basic model also delivers a new insight. If there is a wide range of possible levels of food availability at the stochastic food source, then in equilibrium there is a continuum of degrees of responsiveness to environmental stimuli. See e.g., Lionetti et al. [2018] for empirical support for this finding.

8 Discussion

We built a stylized game-theoretic model of foraging behavior in a stochastic environment. For every parameter specification within certain bounds, this model has a unique symmetric Nash equilibrium, that is also the unique

[7] In a setting in which individuals can have different degrees of responsiveness one might want to assume that the cost of responsiveness varies with the degree of responsiveness. We shall not pursue this here, however.

[8] The distribution with cdf $G(y) = \frac{y}{1+y}$ is the distribution of a random variable Y such that its reciprocal (or inverse) $\frac{1}{Y}$ has exactly the same distribution. One could call G the inversion invariant distribution.

ESS and asymptotically stable under a variety of evolutionary dynamics. This equilibrium has the three key features identified in the literature of *coexistence* of differing degrees of environmental responsiveness, *consistency* of individual responsiveness over time, and *correlation* of individual responsiveness across contexts.

By explicitly studying the phenomenon of heterogeneous responsiveness to environmental stimuli in a foraging setting, we are able to identify the push towards the ideal free distribution of Fretwell and Lucas [1969], satisfied in the equilibrium of our game, as a possible driving force of this heterogeneity. See also Křivan et al. [2008] for a game-theoretic treatment (without environmental stochasticity). We derive explicit analytical expressions for the equilibrium frequencies of responsive and non-responsive behavior, at least when the cost of cognition (needed to respond to environmental stimuli) is negligible or at least relatively small. This allows us to study how the equilibrium frequencies change when some of the model parameters change. For instance, we find that, at least when cognition costs are negligible, the exact stochastic nature of the environment does not affect the equilibrium. This finding suggests that changes in the stochastic environmental process would at least not be so disruptive as to push behavior out of equilibrium. Put differently, equilibrium strategies are already complex enough to allow for automatic adaptation to such changes in the stochastic environmental process. In the remainder of this section we discuss some of the limitations of our approach.

Cost of cognition We have mostly explored the case of zero, and by a continuity argument, also of small cost of cognition. For larger cost of cognition, generally, equilibrium behavior will depend on the stochastic nature of the environment, see, for instance, Figure 3, and the equilibrium will not satisfy the ideal free distribution. For most species, however, it is not unreasonable to assume a relatively small cost of cognition, see e.g., [Hendry, 2016, Murren et al., 2015, Auld et al., 2010, Van Buskirk and Steiner, 2009]. Relatedly, the cost of plasticity, as defined, for instance, in DeWitt et al. [1998], see also [Relyea, 2002] is typically small; see, for instance Hendry [2016], the review by Murren et al. [2015], and the meta-analysis by Van Buskirk and Steiner [2009].

Noisy information Another insight that we can derive from an extension of our model is that the higher the noise in the environmental stimuli the more responsive individuals become in equilibrium. This is under the

assumption of individuals receiving private and stochastically independent noisy information about the state of the environment. We have not explored the case of correlated information, such as all individuals receiving the same public information. In such a setting, the ideal free distribution would at best hold in expectation, and there would be a positive variance of food share availability at the random source.

Social Information Another, empirically relevant, informational setting, that we here abstracted away from, is one where not all individuals receive the same quality of information (perhaps not all are equally close to the source of information). One would then expect individuals to infer additional information about the state of the environment from other individuals' behavior. If, for instance, there are many birds flying out to a specific point at sea, another bird might follow them based on the idea that there is information in that behavior. This will certainly be the case for socializing birds, which display behavior of forming flocks and swarms. Such behavior would add another layer of complexity to the game. For a literature review on social information use in a foraging context see Kohles et al. [2022]. According to the "information-sharing theory", as in Clark and Mangel [1984], it has been observed for different species that animals can both search for food on their own or join others who have found food, see e.g., Giraldeau and Beauchamp [1999]. Models, in which individuals choose between the two strategies "discoverers" or "copiers" are referred to as "producer-scrounger games" [Giraldeau and Beauchamp, 1999, Bhattacharya and Vicsek, 2014]. The interest in using social information is notably to avoid paying a cost to acquire private information [Webster and Laland, 2008]. The choice between producer and scrounger is, to some extent, persistent in time [Morand-Ferron et al., 2011]. The connection between *responsiveness* and the producer-scrounger game could be made through the concepts of *personality*. Indeed, the main trait of personality is the axis boldness-shyness, and it has been observed that shyness increases the probability to scrounge [Kurvers et al., 2010].

The number of food sources Our model only has two food sources. This keeps the analysis mathematically tractable, but at the cost of a possible oversimplification. Given our results, however, one would conjecture that in any (evolutionary stable) equilibrium of a game with multiple food sources, the ideal free distribution holds, at least when costs of cognition are negligible: all food sources would have equal food shares, and this would be

true for all states. This would imply that equilibrium behavior would not depend on the exact stochastic nature of the environment. This would also imply that any (evolutionary stable) equilibria would again satisfy coexistence, and for slightly perturbed models, consistency and correlation. One would, however, not necessarily expect a unique equilibrium and it would be harder to characterize these explicitly.

Generalizing We focused here on foraging choice as a concrete setting in which one would expect the coexistence, consistence and correlation of different responsiveness to external stimuli. However, our results could possibly be generalized to different context where there is a resource to share among individuals, with a resource distributing at different point, and some of them stochastic. Such contexts are social interactions, mating behaviour, division of labour [Dall et al., 2012], space-use [Spiegel et al., 2017, Ceia and Ramos, 2015], or niche specialisation [Dall et al., 2012, Montiglio et al., 2013, Bolnick et al., 2003]. Those last studies show similar concepts (state dependence, frequency dependence, social awareness, environmental heterogeneity) applied to niche specialisation. In particular, increasing evidence show a link between specialisation and personality. It is hypothesized that personality implies specialization [Toscano et al., 2016, Harris et al., 2020] or the other way around [Bergmüller and Taborsky, 2010, Montiglio et al., 2013]. Similar to our model, Leimar et al. [2022] used frequency-dependent competition for resources to explain specialization. For a review of individual foraging specialization see Sheppard et al. [2021]. Finally and more speculatively, our results might even be translatable to the issue of stem cell differentiation, which might arise from competition over resources, see Jörg et al. [2019].

A General distributions of food source availability (proof)

Let $u(A, G, x)$ and $u(B, G, x)$ denote the expected payoff to an individual who goes to food source A and B, respectively, when all others use mixed strategy G and when the realized state of food availability at food source B is given by x. As $G(x) = P(y \leq x)$ is the probability that a random other individual goes to food source B, using again [Chao and Strawderman, 1972, Equation 3.4], we have

$$u(A, G, x) = \frac{1 - G(x)^n}{1 - G(x)}$$

and

$$u(B, G, x) = x\frac{1-(1-G(x))^n}{G(x)}.$$

The expected payoff for an individual using cut-off $y \in \chi$, when all others use mixed strategy G, is then given by

$$u(y, G) = \int_{x=x_L}^{y} \frac{1-G(x)^n}{1-G(x)} f(x)dx + \int_{x=y}^{x_H} x\frac{1-(1-G(x))^n}{G(x)} f(x)dx.$$

For a completely mixed strategy G to be a symmetric Nash equilibrium we need that $u(y, G)$ is constant for all $y \in \chi$. If $u(y, G)$ is a constant function in y then its derivative with respect to y must be constant as well (and equal to zero). This implies that

$$\frac{1-G(y)^n}{1-G(y)} f(y) - y\frac{1-(1-G(y))^n}{G(y)} f(y) = 0,$$

for all $y \in \chi$. Note that $f(y)$ cancels out and the resulting equilibrium distribution G does again not depend on the stochasticity of the food source, and we obtain

$$\frac{1-G(y)^n}{1-G(y)} = y\frac{1-(1-G(y))^n}{G(y)}.$$

Note that the left hand side of this equation is increasing in $G(y)$ and the right hand side is decreasing in $G(y)$. Note also that the left hand side is strictly less than the right hand side when both are evaluated at $G(y) = 0$ and the inequality is reversed when both are evaluated at $G(y) = 1$. Thus, for any $y \in \chi$, this equation has a unique solution $G(y)$. Note that this solution $G(y)$ must necessarily be higher for higher y. Note also that, necessarily, $G(0) = 0$ and $\lim_{y\to\infty} G(y) = 1$. G is, thus, a cumulative distribution function.

For finite n this solution is not easily expressed analytically. In the limit as $n \to \infty$ we obtain

$$\frac{1}{1-G(y)} - y\frac{1}{G(y)} = 0,$$

for all $y \in \chi$, and, thus, $G(y) = \frac{y}{1+y}$ in this limiting case.

B Imperfect signals

In this case, (random) food shares are given by

$$\begin{array}{rcl} R_{A\lambda} & = & \frac{n}{N_{A\lambda}+1} \\ R_{A\eta} & = & \frac{n}{N_{A\eta}+1} \\ R_{B\lambda} & = & \frac{\lambda \cdot n}{N_{B\lambda}+1} \\ R_{B\eta} & = & \frac{\eta \cdot n}{N_{B\eta}+1}, \end{array}$$

where, given strategy σ (used by all opponents), $N_{A\lambda}$ follows a binomial distribution with $n-1$ trials and success probability $\sigma(A) + (1-\epsilon)\sigma(R) + \epsilon\sigma(C)$. Similarly, $N_{A\eta}, N_{B\lambda}, N_{B\eta}$ are also binomially distributed with $n-1$ trials and success probabilities $\sigma(A) + (1-\epsilon)\sigma(C) + \epsilon\sigma(R)$, $\sigma(B) + (1-\epsilon)\sigma(R) + \epsilon\sigma(C)$, and $\sigma(B) + (1-\epsilon)\sigma(C) + \epsilon\sigma(R)$, respectively.

We, thus, obtain

$$\begin{array}{lcl} r_{A\lambda} & = & \frac{1-(1-\sigma(A)-(1-\epsilon)\sigma(R)-\epsilon\sigma(C))^n}{\sigma(A)+(1-\epsilon)\sigma(R)+\epsilon\sigma(C)}, \\ & \text{and analogously} & \\ r_{A\eta} & = & \frac{1-(1-\sigma(A)-(1-\epsilon)\sigma(C)-\epsilon\sigma(R))^n}{\sigma(A)+(1-\epsilon)\sigma(C)+\epsilon\sigma(R)} \\ r_{B\lambda} & = & \lambda \cdot \frac{1-(1-\sigma(B)-(1-\epsilon)\sigma(C)-\epsilon\sigma(R))^n}{\sigma(B)+(1-\epsilon)\sigma(C)+\epsilon\sigma(R)} \\ r_{B\eta} & = & \eta \cdot \frac{1-(1-\sigma(B)-(1-\epsilon)\sigma(R)-\epsilon\sigma(C))^n}{\sigma(B)+(1-\epsilon)\sigma(R)+\epsilon\sigma(C)}. \end{array} \tag{7}$$

We need to show that there is a unique equilibrium with support $\{A, B, R\}$. To be indifferent between A and R requires that

$$r_{A\eta} = r_{B\eta} - \frac{c}{\alpha},$$

to be indifferent between B and R requires

$$r_{B\lambda} = r_{A\lambda} - \frac{c}{1-\alpha}.$$

Plugging in the above expressions, and noting that $\sigma(C) = 0$, these two equations lead to

$$\frac{1-(1-\sigma(A)-\epsilon\sigma(R))^n}{\sigma(A)+\epsilon\sigma(R)} = \eta\frac{1-(\sigma(A)+\epsilon\sigma(R))^n}{1-\sigma(A)-\epsilon\sigma(R)} - \frac{c}{\alpha}, \tag{8}$$

and

$$\lambda\frac{1-(1-\sigma(B)-\epsilon\sigma(R))^n}{\sigma(B)+\epsilon\sigma(R)} = \frac{1-(\sigma(B)+\epsilon\sigma(R))^n}{1-\sigma(B)-\epsilon\sigma(R)} - \frac{c}{1-\alpha}, \tag{9}$$

respectively. We are interested in the limit as $n \to \infty$ and $c \to 0$ and obtain

$$\frac{1}{\sigma(A) + \epsilon\sigma(R)} = \eta \frac{1}{1 - \sigma(A) - \epsilon\sigma(R)},$$

and

$$\lambda \frac{1}{\sigma(B) + \epsilon\sigma(R)} = \frac{1}{1 - \sigma(B) - \epsilon\sigma(R)},$$

respectively. Solving these two equations yields unique solutions $\sigma(A) + \epsilon\sigma(R) = \frac{1}{1+\eta}$ and $\sigma(B) + \epsilon\sigma(R) = \frac{\lambda}{1+\lambda}$. This is the only candidate for an equilibrium of this form. As $\sigma(R) = 1 - \sigma(A) - \sigma(B)$ we obtain $\sigma(R) = \frac{\eta-\lambda}{(1-2\epsilon)(1+\eta)(1+\lambda)}$ and then $\sigma(A) = \frac{1}{1+\eta} - \epsilon\frac{\eta-\lambda}{(1-2\epsilon)(1+\eta)(1+\lambda)}, \sigma(B) = \frac{\lambda}{1+\lambda} - \epsilon\frac{\eta-\lambda}{(1-2\epsilon)(1+\eta)(1+\lambda)}$. Note, finally, that $\sigma(A), \sigma(B), \sigma(R) \geq 0$, and, thus σ is indeed the unique (mixed strategy) equilibrium.

When ϵ gets large enough, one of the nonreactive strategies A or B can vanish. Then, the equilibrium changes. For example, if B vanishes. Then, the only equation is $\sigma(A) + \epsilon\sigma(R) = \frac{1}{1+\eta}$ with $\sigma(A) + \sigma(R) = 1$, which give the equilibrium $\sigma(R) = \frac{\eta}{(1-\epsilon)(1+\eta)}, \sigma(A) = \frac{1}{1+\eta} - \epsilon\frac{\eta}{(1-\epsilon)(1+\eta)}$

C Evolutionary Stable Strategy

We here show that the unique symmetric equilibrium of the basic game is evolutionarily stable as defined, for symmetric n player games, by Palm [1984], see also Broom et al. [1997]. By [Palm, 1984, Proposition 4] a sufficient condition for a (mixed) strategy σ of a symmetric n-player normal form game is an ESS is that

1. $u(\sigma, \sigma) \geq u(\sigma', \sigma)$ for all $\sigma' \in \Delta(S)$ and
2. if $u(\sigma', \sigma) = u(\sigma, \sigma)$ for some $\sigma' \neq \sigma$ then $u(\sigma', \sigma') < u(\sigma, \sigma')$.

The first condition states that only symmetric Nash equilibria can be an ESS. Thus, the strategy we identified as the unique symmetric Nash equilibrium of the game is the only candidate for an ESS. To show that it is indeed an ESS we need to show that it also satisfies the second condition.

For strictly positive costs pure strategy C yields a strictly lower payoff than the equilibrium payoff. However, all strategies $\sigma' \in \Delta(S)$ with $\sigma'(C) = 0$ satisfy $u(\sigma', \sigma) = u(\sigma, \sigma)$. We need to show that for all such $\sigma' \neq \sigma$ we

have $u(\sigma,\sigma') - u(\sigma',\sigma') > 0$. As $\sigma'(C) = \sigma(C) = 0$. We have that,

$$\begin{array}{ll} u(\sigma,\sigma') - u(\sigma',\sigma') & = \\ (\sigma(A) - \sigma'(A)) \left(\alpha \frac{1-(1-\sigma'(A))^n}{\sigma'(A)} + (1-\alpha)\frac{1-\sigma'(B)^n}{1-\sigma'(B)}\right) & + \\ (\sigma(B) - \sigma'(B)) \left(\alpha\eta \frac{1-\sigma'(A)^n}{1-\sigma'(A)} + (1-\alpha)\lambda\frac{1-(1-\sigma'(B))^n}{\sigma'(B)}\right) & + \\ (\sigma(R) - \sigma'(R)) \left(\alpha\eta \frac{1-\sigma'(A)^n}{1-\sigma'(A)} + (1-\alpha)\frac{1-\sigma'(B)^n}{1-\sigma'(B)} - c\right). & \end{array}$$

As $\sigma(R) = 1 - \sigma(A) - \sigma(B)$ (and the same for σ'), this is equivalent to

$$\begin{array}{ll} u(\sigma,\sigma') - u(\sigma',\sigma') & = \\ \alpha\left(\sigma'(A) - \sigma(A)\right) \left(\frac{1-\sigma'(A)^n}{1-\sigma'(A)}\eta - \frac{1-(1-\sigma'(A))^n}{\sigma'(A)} - \frac{c}{\alpha}\right) & + \\ (1-\alpha)\left(\sigma'(B) - \sigma(B)\right) \left(\frac{1-\sigma'(B)^n}{1-\sigma'(B)} - \lambda\frac{1-(1-\sigma'(B))^n}{\sigma'(B)} - \frac{c}{1-\alpha}\right). & \end{array}$$

Note that last term in brackets in the first line is the equilibrium equation 5 (here for $\sigma'(A)$) and similarly the last term in brackets in the second line is the equilibrium equation 6 (here for $\sigma'(B)$). These terms are, therefore zero, for $\sigma' = \sigma$ and they are, for all $n \geq 2$, increasing functions in $\sigma'(A)$ and $\sigma'(B)$, respectively. To see this for the first of these two expressions,

$$\frac{1-\sigma'(A)^n}{1-\sigma'(A)}\eta - \frac{1-(1-\sigma'(A))^n}{\sigma'(A)} - \frac{c}{\alpha}$$

note that the first term in this expression is

$$r_{B\eta} = \frac{1-\sigma'(A)^n}{1-\sigma'(A)} = \sum_{k=0}^{n-1} \sigma'(A)^k,$$

and, thus, clearly increasing in $\sigma'(A)$. Analogous arguments can be made for the second term in this first expression, as well as for both terms in the second of the two expressions. This means that the first term in brackets is negative exactly when $\sigma'(A) - \sigma(A)$ is negative and is positive exactly when the latter is positive. The same argument is true for the second term in brackets. Therefore, for any $\sigma' \neq \sigma$ with $\sigma'(C) = 0$ we, indeed, have $u(\sigma,\sigma') - u(\sigma',\sigma') > 0$. Therefore, the equilibrium σ is an ESS.

D Dynamic stability

In this section we show that the class of games we here study are all stable games, in the sense of Hofbauer and Sandholm [2009]. Results in Hofbauer

and Sandholm [2009] then imply that the unique symmetric Nash equilibrium that we have identified above is asymptotically stable under a large class of evolutionary (or behavioral adjustment) dynamics, that includes, for instance, also the well-known replicator dynamics of Taylor and Jonker [1978].

In order to prove that our games are stable games, we need to introduce some of the more general notation introduced by Hofbauer and Sandholm [2009]. As we have just a single populations of players, in their notation we have $p = 1$. In their notation X denotes the space of players' mixed strategy space, which in our setting is, therefore, equal to $\Delta(S)$. The payoff function in their setting is denoted by $F : X \to \mathbf{R}^k$, where k is the number of pure strategies, i.e., $k = 4$ in our setting.[9] Translated to our notation, F is, therefore, given by

$$F : \sigma \mapsto \begin{bmatrix} u(A, \sigma) \\ u(B, \sigma) \\ u(R, \sigma) \\ u(C, \sigma) \end{bmatrix}.$$

A population game $F : X \to \mathbf{R}^k$ is a stable game, as defined by Hofbauer and Sandholm [2009], if,

$$(y - x)'(F(y) - F(x)) \leq 0, for\ all\ x, y \in X$$

We now prove that our games are stable games in this sense. We utilize [Hofbauer and Sandholm, 2009, Example 2.9] to show that our games are *negative dominant diagonal games*, which by [Hofbauer and Sandholm, 2009, Proposition 2.7] are stable games.

To do so we need to check two conditions. First, a condition of universal *negative frequency dependence*,

$$\frac{\partial u(s, \sigma)}{\partial \sigma(s)} \leq 0 \text{ for all } s \in \{A, B, R, C\},$$

and second,

$$\left| \frac{\partial u(s, \sigma)}{\partial \sigma(s)} \right| \geq \frac{1}{2} \left(\sum_{s' \neq s} \left| \frac{\partial u(s', \sigma)}{\partial \sigma(s)} \right| + \sum_{s' \neq s} \left| \frac{\partial u(s, \sigma)}{\partial \sigma(s')} \right| \right) \text{ for all } s \in \{A, B, R, C\}.$$

We here only provide the proof that these two conditions are satisfied for $s = A$, the proof for other strategies in $S = \{A, B, R, C\}$ proceeds analogously.

[9] Hofbauer and Sandholm [2009] actually use the letter n instead of k, but for our paper n is reserved to denote the number of players.

Recall that,

$$\begin{bmatrix} u(A,\sigma) \\ u(B,\sigma) \\ u(R,\sigma) \\ u(C,\sigma) \end{bmatrix} = \begin{bmatrix} \alpha r_{A\eta} + (1-\alpha) r_{A\lambda} \\ \alpha r_{B\eta} + (1-\alpha) r_{B\lambda} \\ \alpha r_{B\eta} + (1-\alpha) r_{A\lambda} - c \\ \alpha r_{A\lambda} + (1-\alpha) r_{B\lambda} - c \end{bmatrix},$$

where, by Equation 2,

$$\begin{aligned} r_{A\lambda} &= \tfrac{1-(1-\sigma(A)-\sigma(R))^n}{\sigma(A)+\sigma(R)} \\ r_{A\eta} &= \tfrac{1-(1-\sigma(A)-\sigma(C))^n}{\sigma(A)+\sigma(C)} \\ r_{B\lambda} &= \lambda \cdot \tfrac{1-(1-\sigma(B)-\sigma(C))^n}{\sigma(B)+\sigma(C)} \\ r_{B\eta} &= \eta \cdot \tfrac{1-(1-\sigma(B)-\sigma(R))^n}{\sigma(B)+\sigma(R)}. \end{aligned}$$

We, therefore, have that

$$\frac{\partial u(A,\sigma)}{\partial \sigma(A)} = \alpha \frac{\partial r_{A\eta}(\sigma)}{\partial \sigma(A)} + (1-\alpha) \frac{\partial r_{A\lambda}(\sigma)}{\partial \sigma(A)}.$$

Note that

$$\begin{aligned} r_{A\lambda} &= \tfrac{1-(1-\sigma(A)-\sigma(R))^n}{\sigma(A)+\sigma(R)} \\ &= \textstyle\sum_{j=0}^{n-1} (1-\sigma(A)-\sigma(R))^j, \end{aligned}$$

by the well-known formula for the geometric series. Thus, $\frac{\partial r_{A\lambda}(\sigma)}{\partial \sigma(A)} \leq 0$. By the same argument $\frac{\partial r_{A\eta}(\sigma)}{\partial \sigma(A)}$ is also everywhere non-positive, and, therefore,

$$\frac{\partial u(A,\sigma)}{\partial \sigma(A)} \leq 0,$$

as desired. Analogous arguments yield that the desired first condition is satisfied:

$$\frac{\partial u(s,\sigma)}{\partial \sigma(s)} \leq 0 \text{ for all } s \in \{A, B, R, C\}.$$

Let us now consider the second condition, again for $s = A$, which can be written as

$$\begin{aligned} \left|\tfrac{\partial u(A,\sigma)}{\partial \sigma(A)}\right| &\geq \tfrac{1}{2}\left(\left|\tfrac{\partial u(B,\sigma)}{\partial \sigma(A)}\right| + \left|\tfrac{\partial u(R,\sigma)}{\partial \sigma(A)}\right| + \left|\tfrac{\partial u(C,\sigma)}{\partial \sigma(A)}\right|\right) \\ &+ \tfrac{1}{2}\left(\left|\tfrac{\partial u(A,\sigma)}{\partial \sigma(B)}\right| + \left|\tfrac{\partial u(A,\sigma)}{\partial \sigma(R)}\right| + \left|\tfrac{\partial u(A,\sigma)}{\partial \sigma(C)}\right|\right). \end{aligned}$$

We prove this statement in two steps. Note first that

$$u(A,\sigma) + u(B,\sigma) = u(R,\sigma) + u(C,\sigma) + 2c.$$

Therefore,

$$\frac{\partial u(A,\sigma)}{\partial \sigma(A)} + \frac{\partial u(B,\sigma)}{\partial \sigma(A)} = \frac{\partial u(R,\sigma)}{\partial \sigma(A)} + \frac{\partial u(C,\sigma)}{\partial \sigma(A)}.$$

Since $\frac{\partial u(B,\sigma)}{\partial \sigma(A)} = 0$, we then have

$$\frac{\partial u(A,\sigma)}{\partial \sigma(A)} = \frac{\partial u(B,\sigma)}{\partial \sigma(A)} + \frac{\partial u(R,\sigma)}{\partial \sigma(A)} + \frac{\partial u(C,\sigma)}{\partial \sigma(A)}.$$

As all these derivatives are either negative or equal to zero, we have

$$\left|\frac{\partial u(A,\sigma)}{\partial \sigma(A)}\right| = \left|\frac{\partial u(B,\sigma)}{\partial \sigma(A)}\right| + \left|\frac{\partial u(R,\sigma)}{\partial \sigma(A)}\right| + \left|\frac{\partial u(C,\sigma)}{\partial \sigma(A)}\right|.$$

For the second step, note that $\frac{\partial r_{A\lambda}(\sigma)}{\partial \sigma(A)} = \frac{\partial r_{A\lambda}(\sigma)}{\partial \sigma(C)}$, as well as $\frac{\partial r_{A\eta}(\sigma)}{\partial \sigma(A)} = \frac{\partial r_{A\eta}(\sigma)}{\partial \sigma(R)}$. Moreover, because $\frac{\partial r_{A\lambda}(\sigma)}{\partial \sigma(R)} = \frac{\partial r_{A\eta}(\sigma)}{\partial \sigma(C)} = 0$, we have,

$$\begin{aligned}
\frac{\partial u(A,\sigma)}{\partial \sigma(A)} &= \alpha \frac{\partial r_{A\eta}(\sigma)}{\partial \sigma(A)} + (1-\alpha)\frac{\partial r_{A\lambda}(\sigma)}{\partial \sigma(A)} \\
&= \alpha \frac{\partial r_{A\eta}(\sigma)}{\partial \sigma(R)} + (1-\alpha)\frac{\partial r_{A\lambda}(\sigma)}{\partial \sigma(C)} + (1-\alpha)\frac{\partial r_{A\lambda}(\sigma)}{\partial \sigma(R)} + \alpha \frac{\partial r_{A\eta}(\sigma)}{\partial \sigma(C)} \\
&= \frac{\partial u(A,\sigma)}{\partial \sigma(R)} + \frac{\partial u(A,\sigma)}{\partial \sigma(C)}.
\end{aligned}$$

As $\frac{\partial u(A,\sigma)}{\partial \sigma(B)} = 0$, and again, as all derivatives are either negative or equal to zero we have

$$\left|\frac{\partial u(A,\sigma)}{\partial \sigma(A)}\right| = \left|\frac{\partial u(A,\sigma)}{\partial \sigma(B)}\right| + \left|\frac{\partial u(A,\sigma)}{\partial \sigma(R)}\right| + \left|\frac{\partial u(C,\sigma)}{\partial \sigma(B)}\right|.$$

Finally, by adding the result of the two parts, the second condition follows for $s = A$. For strategies B, R, C the same result can be proven in an analogous manner.

References

Josh R Auld, Anurag A Agrawal, and Rick A Relyea. Re-evaluating the costs and limits of adaptive phenotypic plasticity. *Proceedings of the Royal Society B: Biological Sciences*, 277(1681):503–511, 2010.

Alison M Bell, Shala J Hankison, and Kate L Laskowski. The repeatability of behaviour: a meta-analysis. *Animal Behaviour*, 77(4):771–783, 2009.

Ralph Bergmüller and Michael Taborsky. Animal personality due to social niche specialisation. *Trends in Ecology & Evolution*, 25(9):504–511, 2010.

Theodore C Bergstrom. On the evolution of hoarding, risk-taking, and wealth distribution in nonhuman and human populations. *Proceedings of the National Academy of Sciences*, 111(Supplement 3):10860–10867, 2014.

Kunal Bhattacharya and Tamás Vicsek. Collective foraging in heterogeneous landscapes. *Journal of the Royal Society Interface*, 11(100):20140674, 2014.

Daniel I Bolnick, Richard Svanbäck, James A Fordyce, Louie H Yang, Jeremy M Davis, C Darrin Hulsey, and Matthew L Forister. The ecology of individuals: incidence and implications of individual specialization. *The American Naturalist*, 161(1):1–28, 2003.

W Thomas Boyce. Differential susceptibility of the developing brain to contextual adversity and stress. *Neuropsychopharmacology*, 41(1):142–162, 2016.

M Broom, C Cannings, and GT Vickers. Multi-player matrix games. *Bulletin of Mathematical Biology*, 59(5):931–952, 1997.

Ulrika Candolin, Robert J Fletcher, and Andrea EA Stephens. Animal behaviour in a changing world. *Trends in Ecology & Evolution*, 38(4): 313–315, 2023.

Filipe R Ceia and Jaime A Ramos. Individual specialization in the foraging and feeding strategies of seabirds: a review. *Marine Biology*, 162:1923–1938, 2015.

Min-Te Chao and WE Strawderman. Negative moments of positive random variables. *Journal of the American Statistical Association*, 67(338):429–431, 1972.

Colin W Clark and Marc Mangel. Foraging and flocking strategies: information in an uncertain environment. *The American Naturalist*, 123(5): 626–641, 1984.

Tim Coulson, Bruce E Kendall, Julia Barthold, Floriane Plard, Susanne Schindler, Arpat Ozgul, and Jean-Michel Gaillard. Modeling adaptive and nonadaptive responses of populations to environmental change. *The American Naturalist*, 190(3):313–336, 2017.

Sasha RX Dall, Alison M Bell, Daniel I Bolnick, and Francis LW Ratnieks. An evolutionary ecology of individual differences. *Ecology Letters*, 15(10): 1189–1198, 2012.

Thomas J DeWitt, Andrew Sih, and David Sloan Wilson. Costs and limits of phenotypic plasticity. *Trends in Ecology & Evolution*, 13(2):77–81, 1998.

Niels J Dingemanse and Max Wolf. Recent models for adaptive personality differences: a review. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 365(1560):3947–3958, 2010.

Niels J Dingemanse, Anahita JN Kazem, Denis Réale, and Jonathan Wright. Behavioural reaction norms: animal personality meets individual plasticity. *Trends in Ecology & Evolution*, 25(2):81–89, 2010.

Sean M Ehlman, Ulrike Scherer, and Max Wolf. Developmental feedbacks and the emergence of individuality. *Royal Society Open Science*, 9(11): 221189, 2022.

Bruce J Ellis and W Thomas Boyce. Biological sensitivity to context. *Current Directions in Psychological Science*, 17(3):183–187, 2008.

Bruce J Ellis, W Thomas Boyce, Jay Belsky, Marian J Bakermans-Kranenburg, and Marinus H Van Ijzendoorn. Differential susceptibility to the environment: an evolutionary–neurodevelopmental theory. *Development and Psychopathology*, 23(1):7–28, 2011.

Jeffrey C Ely and William H Sandholm. Evolution in Bayesian games I: theory. *Games and Economic Behavior*, 53(1):83–109, 2005.

Willem E Frankenhuis, Karthik Panchanathan, and Jay Belsky. A mathematical model of the evolution of individual differences in developmental plasticity arising through parental bet-hedging. *Developmental Science*, 19(2):251–274, 2016.

Stephen Dewitt Fretwell and H. L. Lucas. On territorial behavior and other factors influencing habitat distribution in birds. *Acta Biotheoretica*, 19 (1):16–32, 1969.

Luc-Alain Giraldeau and Guy Beauchamp. Food exploitation: searching for the optimal joining policy. *Trends in Ecology & Evolution*, 14(3):102–106, 1999.

Stephanie M Harris, Sébastien Descamps, Lynne U Sneddon, Philip Bertrand, Olivier Chastel, and Samantha C Patrick. Personality predicts foraging site fidelity and trip repeatability in a marine predator. *Journal of Animal Ecology*, 89(1):68–79, 2020.

John C Harsanyi. Games with randomly disturbed payoffs: a new rationale for mixed-strategy equilibrium points. *International Journal of Game Theory*, 2(1):1–23, 1973.

Andrew P Hendry. Key questions on the role of phenotypic plasticity in eco-evolutionary dynamics. *Journal of Heredity*, 107(1):25–41, 2016.

Josef Hofbauer and William H Sandholm. Stable games and their dynamics. *Journal of Economic Theory*, 144(4):1665–1693, 2009.

David J Jörg, Yu Kitadate, Shosei Yoshida, and Benjamin D Simons. Competition for stem cell fate determinants as a mechanism for tissue homeostasis. arXiv preprint arXiv:1901.03903, 2019.

Jenna E Kohles, M Teague O'Mara, and Dina KN Dechmann. A conceptual framework to predict social information use based on food ephemerality and individual resource requirements. *Biological Reviews*, 97(6):2039–2056, 2022.

Vlastimil Křivan, Ross Cressman, and Candace Schneider. The ideal free distribution: a review and synthesis of the game-theoretic perspective. *Theoretical Population Biology*, 73(3):403–425, 2008.

Ralf HJM Kurvers, Herbert HT Prins, Sipke E van Wieren, Kees van Oers, Bart A Nolet, and Ronald C Ydenberg. The effect of personality on social foraging: shy barnacle geese scrounge more. *Proceedings of the Royal Society B: Biological Sciences*, 277(1681):601–608, 2010.

Kate L Laskowski, Chia-Chen Chang, Kirsten Sheehy, and Jonathan Aguiñaga. Consistent individual behavioral variation: what do we know and where are we going? *Annual Review of Ecology, Evolution, and Systematics*, 53:161–182, 2022.

Olof Leimar, Sasha RX Dall, Alasdair I Houston, and John M McNamara. Behavioural specialization and learning in social networks. *Proceedings of the Royal Society B*, 289(1980):20220954, 2022.

Francesca Lionetti, Arthur Aron, Elaine N Aron, G Leonard Burns, Jadzia Jagiellowicz, and Michael Pluess. Dandelions, tulips and orchids: evidence for the existence of low-sensitive, medium-sensitive and high-sensitive individuals. *Translational Psychiatry*, 8(1):1–11, 2018.

John Maynard Smith and George R Price. The logic of animal conflict. *Nature*, 246(5427):15–18, 1973.

Elizabeth A McHuron, Elliott Hazen, and Daniel P Costa. Constrained by consistency? Repeatability of foraging behavior at multiple timescales for a generalist marine predator. *Marine Biology*, 165:1–13, 2018.

John M McNamara and Alasdair I Houston. If animals know their own fighting ability, the evolutionarily stable level of fighting is reduced. *Journal of Theoretical Biology*, 232(1):1–6, 2005.

David J Mitchell and Thomas M Houslay. Context-dependent trait covariances: how plasticity shapes behavioral syndromes. *Behavioral Ecology*, 32(1):25–29, 2021.

Pierre-Olivier Montiglio, Caterina Ferrari, and Denis Reale. Social niche specialization under constraints: personality, social interactions and environmental heterogeneity. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 368(1618):20120343, 2013.

Julie Morand-Ferron, Gi-Mick Wu, and Luc-Alain Giraldeau. Persistent individual differences in tactic use in a producer–scrounger game are group dependent. *Animal Behaviour*, 82(4):811–816, 2011.

Courtney J Murren, Josh R Auld, H Callahan, Cameron K Ghalambor, Corey A Handelsman, Mary A Heskel, JG Kingsolver, Heidi J Maclean, Joanna Masel, Heather Maughan, et al. Constraints on the evolution of phenotypic plasticity: limits and costs of phenotype and plasticity. *Heredity*, 115(4):293–301, 2015.

John H Nachbar. "Evolutionary" selection dynamics in games: convergence and limit properties. *International Journal of Game Theory*, 19(1):59–89, 1990.

John F Nash. Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences*, 36(1):48–49, 1950.

Daniel Nettle and Lars Penke. Personality: bridging the literatures from human psychology and behavioural ecology. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 365(1560):4043–4050, 2010.

Petri T Niemelä and Niels J Dingemanse. Meta-analysis reveals weak associations between intrinsic state and personality. *Proceedings of the Royal Society B: Biological Sciences*, 285(1873):20172823, 2018.

Günther Palm. Evolutionary stable strategies and game dynamics for n-person games. *Journal of Mathematical Biology*, 19(3):329–334, 1984.

Samantha C Patrick, Stuart Bearhop, David Grémillet, Amélie Lescroël, W James Grecian, Thomas W Bodey, Keith C Hamer, Ewan Wakefield, Mélanie Le Nuz, and Stephen C Votier. Individual differences in searching behaviour and spatial foraging consistency in a central place marine predator. *Oikos*, 123(1):33–40, 2014.

David W Pfennig. *Phenotypic Plasticity & Evolution: Causes, Consequences, Controversies*. Taylor & Francis, 2021.

Michael Pluess. Individual differences in environmental sensitivity. *Child Development Perspectives*, 9(3):138–143, 2015.

Graham Pyke. Optimal foraging theory: an introduction. In *Encyclopedia of Animal Behavior*, pages 111–117. Elsevier Academic Press, 2019.

Rick A Relyea. Costs of phenotypic plasticity. *The American Naturalist*, 159(3):272–282, 2002.

William H Sandholm. Evolution in Bayesian games II: stability of purified equilibria. *Journal of Economic Theory*, 136(1):641–667, 2007.

Catherine E Sheppard, Rachel Heaphy, Michael A Cant, and Harry H Marshall. Individual foraging specialization in group-living species. *Animal Behaviour*, 182:285–294, 2021.

Andrew Sih. Understanding variation in behavioural responses to human-induced rapid environmental change: a conceptual overview. *Animal Behaviour*, 85(5):1077–1088, 2013.

Andrew Sih, Alison Bell, and J Chadwick Johnson. Behavioral syndromes: an ecological and evolutionary overview. *Trends in Ecology & Evolution*, 19(7):372–378, 2004.

Andrew Sih, Julien Cote, Mara Evans, Sean Fogarty, and Jonathan Pruitt. Ecological implications of behavioural syndromes. *Ecology Letters*, 15(3): 278–289, 2012.

Ralf J Sommer. Phenotypic plasticity: from theory and genetics to current and future challenges. *Genetics*, 215(1):1–13, 2020.

Orr Spiegel, Stephan T Leu, C Michael Bull, and Andrew Sih. What's your move? Movement as a link between personality and spatial dynamics in animal populations. *Ecology Letters*, 20(1):3–18, 2017.

Judy A Stamps. Individual differences in behavioural plasticities. *Biological Reviews*, 91(2):534–567, 2016.

David W Stephens and John R Krebs. *Foraging Theory*, volume 1. Princeton University Press, 1986.

Peter D Taylor and Leo B Jonker. Evolutionary stable strategies and game dynamics. *Mathematical Biosciences*, 40(1-2):145–156, 1978.

Benjamin J Toscano, Natasha J Gownaris, Sarah M Heerhartz, and Cristián J Monaco. Personality, foraging behavior and specialization: integrating behavioral and food web ecology at the individual level. *Oecologia*, 182(1):55–69, 2016.

Josh Van Buskirk and U K Steiner. The fitness costs of developmental canalization and plasticity. *Journal of Evolutionary Biology*, 22(4):852–860, 2009.

Ewan D Wakefield, Ian R Cleasby, Stuart Bearhop, Thomas W Bodey, Rachel D Davies, Peter I Miller, Jason Newton, Stephen C Votier, and Keith C Hamer. Long-term individual foraging site fidelity—why some gannets don't change their spots. *Ecology*, 96(11):3058–3074, 2015.

Mike M Webster and KN Laland. Social learning strategies and predation risk: minnows copy only when using private information would be costly. *Proceedings of the Royal Society B: Biological Sciences*, 275(1653):2869–2876, 2008.

Max Wolf and John M McNamara. On the evolution of personalities via frequency-dependent selection. *The American Naturalist*, 179(6):679–692, 2012.

Max Wolf and Franz J Weissing. An explanatory framework for adaptive personality differences. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 365(1560):3959–3968, 2010.

Max Wolf, G Sander Van Doorn, Olof Leimar, and Franz J Weissing. Life-history trade-offs favour the evolution of animal personalities. *Nature*, 447 (7144):581–584, 2007.

Max Wolf, G Sander Van Doorn, and Franz J Weissing. Evolutionary emergence of responsive and unresponsive personalities. *Proceedings of the National Academy of Sciences*, 105(41):15825–15830, 2008.

Max Wolf, G Sander Van Doorn, and Franz J Weissing. On the coevolution of social responsiveness and behavioural consistency. *Proceedings of the Royal Society B: Biological Sciences*, 278(1704):440–448, 2011.

Richard Woltereck. Weitere experimentelle Untersuchungen über Artveränderung, speziell über das Wesen quantitativer Artunterschiede bei Daphniden. *Verh. D. Tsch. Zool. Ges*, 1909:110–172, 1909.

Bob Wong and Ulrika Candolin. Behavioral responses to changing environments. *Behavioral Ecology*, 26(3):665–673, 2015.