Theoretical Computer Science 411 (2010) 1054–1074

Contents lists available at ScienceDirect

Theoretical Computer Science

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

Evolutionary equilibrium in Bayesian routing games: Specialization and niche formation

Petra Berenbrink *, Oliver Schulte

School of Computing Science, Simon Fraser University, Vancouver-Burnaby, B.C., V5A 1S6, Canada

ARTICLE INFO

Article history:
Received 20 February 2008
Received in revised form 8 September 2009
Accepted 16 November 2009
Communicated by M. Mavronicolas

Keywords:
Bayesian Nash Equilibrium
Congestion game
Game theory

ABSTRACT

In this paper we consider Nash equilibria for the selfish task allocation game proposed in Koutsoupias, Papadimitriou (1999) [26], where a set of n users with unsplittable tasks of different size try to access m parallel links with different speeds. In this game, a player can use a mixed strategy (where he uses different links with a positive probability); then he is indifferent between the different link choices. This means that the player may well deviate to a different strategy over time. We propose the concept of evolutionary stable strategies (ESS) as a criterion for stable Nash equilibria, i.e. equilibria where no player is likely to deviate from his strategy. An ESS is a steady state that can be reached by a user community via evolutionary processes in which more successful strategies spread over time. The concept has been used widely in biology and economics to analyze the dynamics of strategic interactions.

We first define a symmetric version of a Bayesian parallel links game where every player is not assigned a task of a fixed size but instead is assigned a task drawn from a distribution, which is the same for all players. We establish that the ESS is *uniquely determined* for a given symmetric Bayesian parallel links game (when it exists). Thus evolutionary stability places strong constraints on the assignment of tasks to links.

We characterize ESS for the Bayesian parallel links game, and investigate the structure of evolutionarily stable equilibria: In an ESS, links acquire niches, meaning that there is minimal overlap in the tasks served by different links. Furthermore, all links with the same speed are interchangeable for every task with weight w: Every player must place a task with weight w on links having the same speed with the same probability. Also, bigger tasks must be assigned to faster links and faster links must have a bigger load. Finally, we introduce a clustering condition – roughly, distinct links must serve distinct tasks – that is sufficient for evolutionary stability, and can be used to find an ESS in many models.

Published by Elsevier B.V.

1. Introduction

We consider the *selfish task allocation game* proposed in [26], where users try to access a set of parallel links. We assume that the users have unsplittable tasks with different sizes (weight) and that the links have different speeds. This scenario gives rise to a strategic interaction between users that combines aspects of both competition, in that users compete for the fastest links, and coordination, in that users want to avoid overloaded links. Koutsoupias and Papadimitriou suggested studying the model in a game-theoretic framework [26]. They compare the cost of the worst case Nash equilibrium with the cost of an optimal solution; this ratio was called *price of anarchy*. Depending on the cost function that is used to assess the optimal solution, the fraction between Nash equilibria (see [32]) and optimal solutions can vary greatly. For example, the

* Corresponding author. Tel.: +1 49524246951.
E-mail addresses: petra@cs.sfu.ca (P. Berenbrink), oschulte@cs.sfu.ca (O. Schulte).

0304-3975/$ – see front matter. Published by Elsevier B.V.
doi:10.1016/j.tcs.2009.11.015

cost of the worst case Nash equilibrium can be similar to the cost of the optimal solution (min–max function considered in [7]), or the cost for every Nash equilibrium can be far away from that of the optimal solution [3].

It is an elementary fact that if a player plays a *mixed Nash strategy*, then he is indifferent between the choices that carry positive probability. So, it is not easy to see what keeps the players from deviating to a different strategy with different probabilities. As an example consider the following instance consisting of n users with uniform tasks and m links with the same speed. One possible mixed Nash equilibrium is the state where every player chooses every link with a probability of $1/m$. Now consider the game from the viewpoint of one fixed user. For him, all the links look identical. Hence, he can use any probability distribution to choose between the m links (such a Nash equilibrium is called weak). Now, if he decides to choose, say, the first link with probability one, the resulting state is not optimal any more from the viewpoint of the other players. Such a Nash equilibrium, having a sequence of single-player strategy changes that do not alter their own payoffs but finally lead to a non-equilibrium position, is called transient (see [12]).

The above example shows that games can have several non-stable and transient Nash equilibria, and it is unlikely that a system will end up in one of these. Hence, instead of calculating the price of anarchy, it might be interesting to answer first the question which Nash equilibria are stable, and then to compare the cost of stable equilibria to the cost of the optimal solution (see [12]). Several stability models were suggested in the literature [40]. One of the most important models is Maynard Smith's concept of an *evolutionarily stable strategy*, abbreviated ESS [30]. The criterion proposed by Maynard Smith is as follows: An equilibrium E is evolutionarily stable if there is a threshold $\overline{\varepsilon} < 1$ for the fraction of players deviating. If the fraction of the deviating players falls below $\overline{\varepsilon}$, then the players following the equilibrium E always do better than the deviants.

The concept of an ESS has had a tremendous impact on evolutionary biology [20]. Economists have applied the concept frequently to analyze strategic interactions between selfish agents [41]. For instance, evolutionary analysis has been applied to analyzing road traffic patterns, which share many structural similarities with computer networks [38, Section 8]. One of the strengths of the concept of evolutionary equilibrium is that it connects with many plausible proposals about the dynamics of social systems, like a user community. Dynamical models specify how the frequency with which strategies are used in a population changes over time. Many evolutionary dynamics have been proposed [41,1,40,20]; the concept of evolutionary stability is robust in the sense that an ESS can be proven to be a stable state, or fixed point, for virtually all of them. For instance, an ESS is an asymptotically stable steady state for the well-known replicator dynamics (the converse also holds under various conditions though not in general) [41,33].

1.1. Previous work

Our work combines three different parts of game theory: task allocation games, games of incomplete information, and evolutionary stability. To our knowledge, this combination of topics is new. *Subsets* of this combination have been studied previously; we organize our review around them.

Parallel links model and congestion games. The Parallel Links Game was introduced by Koutsoupias and Papadimitriou [26] (the KP model), who initiated the study of coordination ratios. In the model of [26], the cost of a collection of strategies is the (expected) maximum load of a link (maximized over all links). The coordination ratio is defined as the ratio between the maximum cost (maximized over all Nash equilibria) divided by the cost of the optimal solution. Koutsoupias and Papadimitriou give bounds on the coordination ratio. These bounds are improved by Mavronicolas and Spirakis [29], and by Czumaj and Vöcking [8] who gave an asymptotically tight bound. Since then several papers considered the problem using different cost functions, using different link models [9,3,27,2], or studying the algorithmic complexity and efficiency of computing equilibria [15,13,17]. The KP model is related to congestion games as defined by Rosenthal [36] where players try to access subsets of the resources instead of single links. In the original definition of a congestion game a pure strategy consists of a set of resources, and the payoff function is the same for all players. This is in contrast to parallel links games where players can have different weights and, therefore, have different payoff functions. For a generalization of congestion games to payoff functions that are specific for each player see [31,28].

Parallel links model and games of incomplete information. Harsanyi [21] introduced the notion of a Bayesian game to analyze games with *incomplete information* where players are uncertain about some aspect of the game such as what preferences or options the other players have. Bayesian games have found many applications in economics; eventually Harsanyi's work earned him the Nobel Prize. In [16] Gairing et al. introduce a Bayesian version of the selfish task allocation game. Following Harsanyi's approach [21], each user can have a set of possible types. Their paper presents a comprehensive collection of results for the Bayesian task allocation game. Note that their model is more general than ours since they allow different types for different users, whereas our users all have the same type space. In our application of the Harsanyi framework, the type space models the uncertainty that players have about what tasks have to be processed. The paper [18] uses the type space formalism to model the uncertainty of players about the link capacities. They show that this kind of Bayesian selfish routing model can be reduced to a game of complete information with player-specific payoff functions.

Stability of mixed Nash equilibria. The potential instability of mixed strategy equilibria has long been recognized and much discussed in game theory; see [34, Chapter 3.2] for a concise summary of the debate. In [40] van Damme surveys a number of different ways of defining stability for Nash equilibria. Harsanyi's celebrated purification theorem provides an interpretation of mixed equilibria without the need for any individual to randomize [22,34,19]. Harsanyi considers a matrix game with perfectly known payoffs as an idealization of a Bayesian game in which each player's payoffs are perturbed by random fluctuations. For example, in a selfish task allocation game, the latency of player *i*'s message may be perturbed by a random noise term ε_i. Supposing that the payoff perturbations of the players are independent, and with mild assumptions about the distribution of the perturbations (e.g., absolute continuity with Lebesgue measure), Harsanyi shows that the perturbation of the original game has a Bayesian equilibrium in which every player chooses a *deterministic* or pure strategy given their private information about their payoff. Since the payoffs randomly fluctuate, every such Bayesian equilibrium induces a distribution over pure strategies; the purification theorem states that for almost every matrix game, all mixed equilibria in the game are the limits of Bayesian equilibria in the perturbed game as the perturbations become arbitrarily small.

Evolutionarily stable strategies. One of the most important criteria for distinguishing stable from unstable mixed equilibria is evolutionary stability. The concept of evolutionary stability is fundamental in evolutionary game theory, which has many applications in theoretical biology and economics. The seminal presentation of the concept of an ESS is due to Maynard Smith [30]. Since then, the concept has played a central role in evolutionary biology and has been used in thousands of studies. Economists have also applied the concept to analyze many economic and social interactions, from currency markets to traffic patterns. Kontogiannis and Spirakis provide an introduction to and motivation for evolutionary analysis from a computer science perspective [25, Section 3]. Kearns and Suri examine evolutionary stability in graphical games [24]. Evolutionary stability and Harsanyi's perturbation concept are similar in that both consider arbitrarily small deviations from an equilibrium point. However, in Harsanyi's model, the payoffs for various strategies fluctuate, whereas the mutations considered in an ESS involve only changes in the players' strategies, while the payoff matrix remains fixed. The papers [23,39] study evolutionary stability in a population interpretation of Harsanyi's perturbed game, where different members of a very large population have different payoff matrices corresponding to the payoff perturbations.

Evolutionary stability and routing/traffic models. In [14] Fischer and Vöcking adopt an evolutionary approach to a related task allocation problem (see [37] for a definition). Sandholm proposes a pricing scheme based on evolutionary stability for minimizing traffic congestion; he notes the potential applicability of his models to computer networks [38, Section 8]. His approach does not apply the concept of evolutionarily stable strategy. The theory of evolution in Bayesian games is developed in [10], based on the Bayesian best response dynamic rather than ESS. To our knowledge, our combination of congestion game + Bayesian incomplete information + ESS is new in the literature. (Ely and Sandholm remark that "nearly all work in evolutionary game theory has considered games of complete information" [10, p.84].)

1.2. New results

In this paper we study evolutionarily stable equilibria for selfish task allocation in Koutsoupias and Papadimitriou's *parallel links model* [26] where the users' tasks cannot be split. See [4] for a preliminary version of this paper. We first define a symmetric version of a Bayesian parallel links game where every player is not assigned a task of a fixed size but, instead, is randomly assigned a task drawn from a distribution (Section 2.1). Then we argue that every ESS in this game is a *symmetric* Bayesian Nash equilibrium, where every player uses the same strategy.

Link group uniqueness. In Section 3 we show that the symmetric Bayesian Nash equilibrium is unique for *link groups*. By *link group uniqueness* we mean the following. Assume that all links with the same speed are grouped together into so-called *link groups*. Then, in every symmetric Bayesian Nash equilibrium, the total probability that tasks of a certain size are sent to a link group is unique. This implies that the only flexibility in a symmetric Bayesian Nash equilibrium is the probability distribution over links from the same link group, not over different link groups. Then we show that in a symmetric equilibrium two links with different speeds cannot both be used by two or more tasks with different weights. In fact, we show an even stronger result: If link ℓ is used for task w and ℓ' for $w' \neq w$, then at least one of the links will not be optimal for the other link's task. We also show that tasks with larger weight must be assigned to links with larger speed.

Uniqueness of ESS. In Section 4 we characterize ESS for the symmetric Bayesian parallel links game. We show that every ESS is a Bayesian Nash equilibrium, and we show that, to evaluate evolutionary stability, we have to consider only best replies to the current strategy. Then we establish that in an ESS, we not only have link group uniqueness, but also the probability distribution with which links of the same group are chosen by tasks has to be unique. In fact, an ESS requires treating two links with equal speed exactly the same. This result establishes the *uniqueness of ESS*.

Specialization. We show that in an ESS even two links with the same speed cannot both be used by two or more tasks with different weights. This implies that in an ESS links acquire niches, meaning that there is minimal overlap in the tasks served by different links. We call this *specialization* in the following. We also show that, unfortunately, the specialization condition is necessary for an ESS, but not sufficient.

Clustering. We introduce a sufficient condition called *clustering* – roughly, links must form disjoint niches – and show that every clustered Bayesian Nash equilibrium is an ESS. Unfortunately, we also show that there exists a game that does not have a clustered ESS, but it has an unclustered ESS, so clustering is not a necessary condition.

In general, the problem of calculating an ESS is very hard; it is contained in Σ_2^P (second level of the polynomial-time hierarchy) and is both NP-hard and coNP-hard [11]. We expect that our uniqueness results and the structural properties of ESS for our game will help to develop algorithms that compute an ESS.

Table 1
Table with frequently used notation.

Symbol	Meaning
$N = [n]$	Set of users
i	Variables for users
$L = \{\ell_1, \ldots, \ell_m\}$	Set of links
ℓ, ℓ', ℓ''	Variables for links
$\mathcal{L}, \mathcal{L}', \mathcal{L}''$	Variables for link groups. A link group is defined as a maximal Set of links with the same speed.
c_ℓ ($c_\mathcal{L}$)	Speed of link ℓ (links in link group $\mathcal{L}$)
$W = \{w_1, \ldots, w_k\}$,	A finite set of task weights or sizes
w, w', w'	Variables for task weights
$w(i)$	Task weight assigned to user i
$\mu(w)$	Probability that a task of weight $w \in W$ is assigned To a user
$\sigma, \sigma', \sigma''$	Variables for strategies
σ_i	Strategy for user i
σ^*	Equilibrium strategy
(σ_i, σ_{-i})	Strategy profile where user i follows strategy σ_i and the other Players' strategies are given by $\sigma_{-i} = \sigma_1, \ldots, \sigma_{i-1}, \sigma_{i+1}, \ldots, \sigma_n$
$\sigma^{(k)}$	Vector with strategy σ repeated k times (for $k \le n$ users)
$\sigma(\ell\|w)$ ($\sigma(\mathcal{L}\|w)$)	Probability that a strategy σ uses link ℓ (some link $\ell \in \mathcal{L}$) for Task size w
$\sigma(\ell\|W')$	Probability that a fixed strategy uses link ℓ for some task Size $w \in W' \subset W$
$\text{load}(\ell\|\sigma_1, \ldots, \sigma_n; w(1), \ldots, w(n))$	Load on link ℓ given that user i ($1 \le i \le n$) has been assigned task $w(i)$, and follows strategy σ_i
$\text{load}(\ell\|\sigma_1, \ldots, \sigma_n)$ ($\text{load}(\mathcal{L}\|\sigma_1, \ldots, \sigma_n)$)	Expected load on link ℓ (sum over expected loads on link $\ell \in \mathcal{L}$) Given that user i follows strategy σ_i
$u(\sigma; \sigma_1, \ldots, \sigma_{n-1})$	Payoff to a user following strategy σ when other users follow $\sigma_1, \ldots, \sigma_{n-1}$
$w \in \text{opt}(\ell\|\sigma)$	Link ℓ is optimal for task w given strategy σ
$w \in \text{support}(\ell\|\sigma)$	Strategy σ uses link ℓ for task w

2. Basic models and concepts

In Section 2.1 introduce Bayesian Parallel Links Games and show some simple observations concerning link load and utilities. In Section 2.2 we introduce population games and define evolutionary stable strategies (ESS).

2.1. Bayesian parallel links games

In this section we examine an extension of the original task allocation game called *Bayesian parallel links game*. Our definition below is a special symmetric case of the definition in [16]. The standard parallel links game is not symmetric since the payoff of a user i depends on the task $w(i)$. We summarise our notation in Table 1.

In a Bayesian parallel links game, the uncertainty among the players concerns the task size of the opponents. An agent knows the size of her own message, but not the size of the messages being sent by other users. The Bayesian game of [16] models this uncertainty by a distribution that specifies the probability that w is the task of user i. In our symmetric Bayesian task allocation game, this distribution is the same for all agents. A natural interpretation of this assumption is that agents are assigned tasks drawn from a common distribution.

A game is *symmetric* if (1) all players have the same set of strategy options, and (2) the payoffs only depend on what strategies are chosen and how often they are chosen. The payoffs do not depend on which player is choosing a certain strategy. Our Bayesian version of the game is symmetric, whereas the parallel links game is symmetric for uniform users only.

The formal definition of a symmetric Bayesian task allocation model is as follows.

Definition 1. A **symmetric Bayesian task allocation model** is a tuple $\langle N, W, \mu, L\rangle$ where

1. $N = [n]$ is the set of users.
2. W is a finite set of k task weights, and $\mu : W \to (0, 1]$ is a probability distribution over the weights W. The distribution μ is used to assign weights i.u.r. (independent and uniformly at random) to players $1, \ldots, n$.
3. $L = [m]$ is the set of links. For $\ell \in [m]$, link ℓ has speed c_ℓ.
4. For a fixed user i, a mixed strategy σ_i is a $(k \times n)$ array with one row for every weight $w \in W$. If row r corresponds to weight w, then entry $\sigma_i(r, \ell)$ is the probability that the user assigns a task with weight w to link ℓ. A strategy profile $\sigma_1, \ldots, \sigma_n$ assigns a strategy $\sigma_i \in P$ to each player i.

Now fix a task allocation model with strategy profile $\sigma_1, \ldots, \sigma_n$. In the following $\sigma_i(\cdot|w)$ is the row of the strategy array of user i that corresponds to weight w. We use $\sigma_i(\ell|w)$ for the probability that, in strategy $\sigma_i \in P$, user i assigns a task with weight w to link ℓ. For $1 \leq i \leq n$, the quantity $w(i)$ is the weight assigned to user i. As usual, (σ_i, σ_{-i}) denotes a strategy profile where user i follows strategy σ_i and the other players' strategies are given by $\sigma_{-i} = \sigma_1, \ldots, \sigma_{i-1}, \sigma_{i+1}, \ldots, \sigma_n$. Similarly, $(w(i), w(-i))$ denotes a weight vector where user i is assigned task size $w(i)$ and the other players' weights are given by the vector $w(-i) = w(1), \ldots, w(i-1), w(i+1), \ldots, w(n)$.

The concept of a mixed strategy in a symmetric Bayesian task allocation model may be interpreted as follows. Each player chooses a strategy before the game is played. Then tasks $w(1), w(2), \ldots, w(n)$ are assigned i.u.r. to users 1 through n according to the distribution μ. Each user learns their own task but not that of the others. Next for each user i we "execute" the strategy σ_i given task $w(i)$, such that task $w(i)$ is sent to link ℓ with probability $\sigma_i(\ell|w(i))$. Thus, strategies have a natural interpretation as programs that take as input a task w and output a link for w or a probability distribution over links for w.

Our definition of a mixed strategy is standard in the theory of Bayesian games, but differs from [16] in that we do not define a mixed strategy to be a probability distribution over pure strategies. However, it is easy to see that the two concepts are equivalent: given any probability distribution over pure strategies, there is an equivalent mixed strategy in our sense, and vice versa (cf. [16]).

Like Koutsoupias and Papadimitriou [26], we assume that the latency of a link depends linearily on the load of a link. Thus we have the following definition of the load on a link.

Definition 2. Let $B = \langle N, W, \mu, L\rangle$ be a symmetric Bayesian task allocation model.

1. For fixed $w(1), w(2), \ldots, w(n)$, the **conditional expected load** on link ℓ is

$$\mathrm{load}(\ell|\sigma_1, \ldots, \sigma_n; w(1), \ldots, w(n)) = \frac{1}{c_\ell} \sum_{i \in N} w(i) \cdot \sigma_i(\ell|w(i)).$$

2. The **expected load** on link ℓ is

$$\mathrm{load}(\ell|\sigma_1, \ldots, \sigma_n) = \sum_{w(1),\ldots,w(n) \in W^n} \mathrm{load}(\ell|\sigma_1, \ldots, \sigma_n; w(1), \ldots, w(n)) \cdot \prod_{i \in N} \mu(w(i)).$$

where W^n denotes the n-fold Cartesian product of W.

The next observation shows that the load function is additive in the sense that the total load on link ℓ due to n users is just the sum of the loads due to the individual users.

Observation 2.1. *Let $B = \langle N, W, \mu, L\rangle$ be a symmetric Bayesian task allocation model. Then for any user i we have*

$$\mathrm{load}(\ell|\sigma_1, \ldots, \sigma_n) = \mathrm{load}(\ell|\sigma_{-i}) + \mathrm{load}(\ell|\sigma_i).$$

Therefore $\mathrm{load}(\ell|\sigma_1, \ldots, \sigma_n) = \sum_{i \in N} \mathrm{load}(\ell|\sigma_i)$.

Proof. Without loss of generality we assume $i = 1$.

$$\begin{aligned}
\mathrm{load}(\ell|\sigma_1, \ldots, \sigma_n) &= \sum_{w(1),\ldots,w(n) \in W^n} \mathrm{load}(\ell|\sigma_1, \ldots, \sigma_n; w(1), \ldots, w(n)) \cdot \prod_{i \in N} \mu(w(i)) \\
&= \sum_{w(1),\ldots,w(n) \in W^n} \frac{1}{c_\ell} \cdot \left(\sum_{i \in N} w(i) \cdot \sigma_i(\ell|w(i)) \right) \cdot \prod_{i \in N} \mu(w(i)) \\
&= \sum_{w(1),\ldots,w(n) \in W^n} \frac{1}{c_\ell} \cdot \prod_{i \in N} \mu(w(i)) \cdot w(1) \cdot \sigma_1(\ell|w(1)) \\
&\quad + \sum_{w(1),\ldots,w(n) \in W^n} \frac{1}{c_\ell} \cdot \prod_{i \in N} \mu(w(i)) \cdot \left(\sum_{i \in N-\{1\}} w(i) \cdot \sigma_i(\ell|w(i)) \right) \\
&= \mathrm{load}(\ell|\sigma_1) + \mathrm{load}(\ell|\sigma_{-1}). \quad \square
\end{aligned}$$

A symmetric Bayesian task allocation game is a symmetric Bayesian task allocation model where all players have the same utility function u. Additionally, the payoff of each player depends only on what strategies are chosen, and not on which players choose particular strategies. This allows us to drop the index i for the user from time to time and write, for example, $(\sigma; \sigma_1, \ldots, \sigma_{n-1})$ for strategy σ played against $\sigma_1, \ldots, \sigma_{n-1}$. Here we assume that $\sigma_1, \ldots, \sigma_{n-1}$ is simply a list of strategies without σ_i referring to the strategy of user i. To simplify notation for games in which several players follow the same strategy, we write $\sigma^{(k)}$ for a $\sigma, \ldots, \sigma$ with σ repeated k times. If all players in a symmetric game follow the same strategy, then $\sigma^{(n)}$ is the resulting strategy profile.

Definition 3. A **symmetric Bayesian task allocation game** $G = \langle N, W, \mu, L, u\rangle$ is a task allocation model $\langle N, W, \mu, L\rangle$ together with a utility function u. We write $u(\sigma; \sigma_1, \ldots, \sigma_{n-1})$ to denote the payoff of following strategy σ when the other players' strategies are given by $\sigma_1, \ldots, \sigma_{n-1}$. Then the payoff is defined as

$$u(\sigma; \sigma_1, \ldots, \sigma_{n-1}) = -\sum_{w \in W} \sum_{\ell \in L} (w/c_\ell + \mathrm{load}(\ell|\sigma_1, \ldots, \sigma_{n-1})) \cdot \sigma(\ell|w) \cdot \mu(w).$$

Note that maximizing negative utility function is equivalent to minimizing the cost of a strategy.

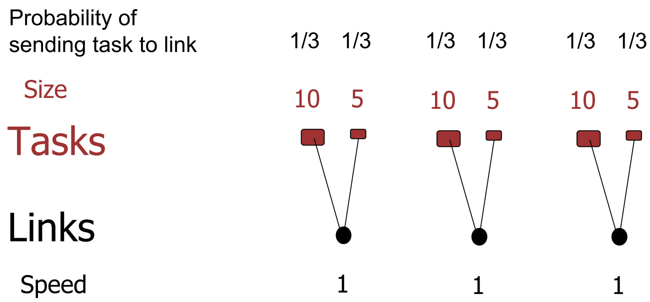

Fig. 1. A Bayesian task allocation game. There are three links, each with capacity or speed 1, and two task sizes, 5 and 10. Each task size occurs equally frequently. The figure shows one possible allocation or mixed strategy that assigns each link to each task size with equal probability.

Example 1. Consider a 2-player Bayesian task allocation game with 3 links 1, 2, 3, each with the same capacity $1 = c_1 = c_2 = c_3$, two task sizes $W = \{5, 10\}$, each occurring equally frequently so $\mu(10) = \mu(5) = 1/2$. Let σ^* be the following mixed strategy.

w	$\sigma^*(1\|w)$	$\sigma^*(2\|w)$	$\sigma^*(3\|w)$
5	1/3	1/3	1/3
10	1/3	1/3	1/3

Fig. 1 illustrates this task allocation game and the mixed strategy σ^*. In this example, the expected load on each link ℓ due to one player following strategy σ^* is the same for all links due to the symmetry of the model; it is given by

$$\text{load}(\ell|\sigma^*) = \sum_{w \in \{5,10\}} w \cdot \sigma^*(\ell|w) \cdot \mu(w) = 5 \cdot \frac{1}{6} + 10 \cdot \frac{1}{6} = \frac{5}{2}.$$

Let (σ^*, σ^*) be the strategy profile in which each player follows strategy σ^*. The expected load on link ℓ due to two players following strategy σ^* is given by

$$\begin{aligned}\text{load}(\ell|\sigma^*, \sigma^*) &= \frac{1}{4} \cdot \sum_{w(1) \in \{5,10\}} \sum_{w(2) \in \{5,10\}} \text{load}(\ell|\sigma^*, \sigma^*; w(1), w(2)) \\ &= \frac{1}{4} \sum_{w(1) \in \{5,10\}} \sum_{w(2) \in \{5,10\}} \left(\frac{w(1)}{3} + \frac{w(2)}{3} \right) \\ &= \frac{1}{4} \cdot \frac{1}{3} \cdot (5 + 5 + 5 + 10 + 10 + 5 + 10 + 10) = \frac{1}{3} \cdot (5 + 10) = 5.\end{aligned}$$

Notice that $\text{load}(\ell|\sigma^*, \sigma^*) = 2 \cdot \text{load}(\ell|\sigma^*)$, which illustrates Observation 2.1. For the expected payoff $u(\sigma^*; \sigma^*)$ for a player following strategy σ^* against another player also using σ^* we have

$$\begin{aligned}u(\sigma^*; \sigma^*) &= - \sum_{w \in \{5,10\}} \sum_{\ell \in [3]} (w + \text{load}(\ell|\sigma^*)) \cdot \frac{1}{3} \cdot \frac{1}{2} \\ &= -\frac{1}{6} \cdot 3 \cdot \left(5 + \frac{5}{2} + 10 + \frac{5}{2} \right) = -10.\end{aligned}$$

The mixed strategy σ_i is a *best reply* to σ_{-i} if for all mixed strategies σ_i' we have

$$u(\sigma_i; \sigma_{-i}) \geq u(\sigma_i'; \sigma_{-i}).$$

A strategy profile $\sigma_1, \ldots, \sigma_n$ is a *Bayesian Nash equilibrium* if every player i plays a best response strategy against σ_{-i}. The strategy profile $\sigma^{(n)}$ is a *symmetric Bayesian Nash equilibrium* if σ is a best reply to $\sigma^{(n-1)}$. Hence, a symmetric Bayesian Nash equilibrium for a symmetric Bayesian task allocation game with n players is a Bayesian Nash equilibrium $(\sigma^{(n)})$ in which each player follows the same strategy. It follows from Nash's existence proof [32] that a symmetric game, such as a symmetric Bayesian Routing Game, has a symmetric Bayesian Nash equilibrium.

In the following we say that link ℓ is **optimal** for task w of player i given σ_{-i} iff ℓ minimizes the function $w/c_\ell + \text{load}(\ell|\sigma_{-i})$. In this case we write $w \in \text{opt}(\ell|\sigma_{-i})$. A mixed strategy σ **uses** link ℓ for task w if $\sigma(\ell|w) > 0$; we write $w \in \text{support}(\ell|\sigma)$.

The next proposition asserts that a best reply σ_i to a strategy profile σ_{-i} uses a link for a task only if the link is optimal for the task given σ_{-i}. So a Bayesian Nash equilibrium requires that each player choose an optimal strategy *for each task* w; a mixed strategy σ_i is optimal just in the case it assigns positive probability only to optimal links for a given task w. This is a variant of the standard characterization of Nash equilibrium according to which all pure strategies in the support of an equilibrium strategy are best replies. The proof can be done similar to the proof of the standard Nash characterization and is omitted.

Proposition 2.2. *Let G be a Bayesian task allocation game with n players, and let σ_{-i} be a mixed strategy profile. A strategy σ_i is a best reply to σ_{-i} $\iff$ for all tasks w, links ℓ, if strategy σ_i uses link ℓ for task w, then ℓ is an optimal link for w given σ_{-i}.*

Example 2. Consider again the task allocation game and the strategy illustrated in Fig. 1. Since the speed of each link is the same, and since each link carries the same load given the strategy σ^*, it follows that every link is optimal for every task given σ^*. Hence the strategy profile (σ^*, σ^*) is a Bayesian Nash equilibrium.

2.2. Population games and evolutionary stability for the parallel links game

We give a brief introduction to population games and evolutionary stability. A more extended introduction from a computer science point of view is provided in [25, Section 3].

Population equilibria. The standard population game model considers a very large **population** A of agents [41,30]. The agents play a symmetric game like our symmetric Bayesian task allocation game. Every agent in the population follows a strategy σ *fixed before* the game is played. A *match* is a particular instance of the base game that results when we match n i.u.r. chosen agents together to play the base game. Since strategies occur with a certain frequency in the population, the probability that a task with a given size is assigned to a link can be regarded as fixed. *Hence, with a population A we can associate a mixed strategy that we denote by σ_A.* For example, in the task allocation game of Fig. 1, suppose that one third of the population places both tasks of size 5 and size 10 on link 1, one third places both tasks on link 2, and one third places both tasks on link 3. The mixed strategy that describes the aggregate behavior of this population is the one illustrated in Fig. 1, defined by $\sigma^*(\ell|w) = 1/3$ for all links $\ell \in [3]$ and task sizes $w \in \{5, 10\}$. This example illustrates an interesting feature of population models: even if each individual agent chooses a deterministic strategy, mixed strategies can be interpreted as describing the aggregate behavior of the population (cf. [34, Chapter 3.2]).

Consider now the expected payoff that an agent using strategy σ receives in a match with $n-1$ opponents that are randomly selected from a fixed population A. This payoff is the same as the payoff that results from playing strategy σ against $n-1$ opponents whose choices are determined by *the same distribution*, namely the population distribution σ_A. In other words, the expected payoff is given by $u(\sigma; (\sigma_A)^{(n-1)})$, the payoff of using strategy σ when the other $n-1$ players follow mixed strategy σ_A. A population is *in equilibrium* if no agent benefits from changing her strategy unilaterally given the state of the population. Formally, a population A with associated mixed strategy σ_A is in equilibrium if every mixed strategy σ that occurs with frequency greater than zero in the population is a best reply to $(\sigma_A)^{(n-1)}$. It is easy to see that this is the case if and only if the symmetric strategy profile $(\sigma, \sigma, \dots, \sigma)$ is a Bayesian Nash equilibrium. So *population equilibria correspond exactly to symmetric Bayesian Nash equilibria.* While restricting attention to symmetric Bayesian Nash equilibria may seem like an artificial restriction for non-population models, in large population models symmetric Bayesian Nash equilibria characterize the natural equilibrium concept for a population.

Evolutionarily stable population equilibria. The main idea in evolutionary game theory is Maynard Smith's concept of *stability against mutations*, which is a criterion for distinguishing stable from unstable population equilibria. Intuitively, a population is evolutionarily stable if a small group of mutants cannot invade the population. In the context of evolutionary analysis, we refer to the base population as the *incumbents.* Consider an incumbent population A that encounters a group M of mutants. Then the mixed population is $A \cup M$. Suppose that in this mixed population the proportion of mutants is ε. The distribution for the mixed population is the probabilistic mixture $(1-\varepsilon)\sigma_A + \varepsilon\sigma_M$.

We may view a mutation M as successful if the average payoff for invaders *in the mixed population* is at least as great as the average payoff for incumbents in the mixed population. If a sufficiently small mutation is successful, the population is considered unstable. The expected payoff for a strategy σ in the mixed population $A \cup M$ is given by $u(\sigma; [(1-\varepsilon)\sigma_A + \varepsilon\sigma_M]^{(n-1)})$. So the average payoff for the incumbents is $u(\sigma_A; [(1-\varepsilon)\sigma_A + \varepsilon\sigma_M]^{(n-1)})$ and for the mutants it is $u(\sigma_M; [(1-\varepsilon)\sigma_A + \varepsilon\sigma_M]^{(n-1)})$. The next example illustrates the key concepts of population equilibrium and successful mutations in our task allocation game. For readers who are new to the concept of evolutionary stability, we provide a simpler example in the Appendix, using the standard Hawk–Dove game.

Example 3. Consider again the task allocation game of Fig. 1 and the uniform distribution mixed strategy σ^* defined by $\sigma^*(\ell|w) = 1/3$ for all links $\ell \in [3]$ and task sizes $w \in \{5, 10\}$. A possible mutant population M may be described by the following mixed strategy σ.

w	$\sigma(1\|w)$	$\sigma(2\|w)$	$\sigma(3\|w)$
5	29/60	11/60	1/3
10	7/30	13/30	1/3

Suppose that the relative sizes of the current and mutant populations are such that the mixed population $A \cup M$ is described by the mixed strategy $(1\%)p + (99\%)\sigma^*$. Then in the mixed population, the expected payoff for the incumbents is given by

$$u(\sigma^*; (1\%)\sigma + (99\%)\sigma^*) = (1\%) \cdot u(\sigma^*; \sigma) + (99\%) \cdot u(\sigma^*; \sigma^*).$$

As calculated in the preceding section, $u(\sigma^*; \sigma^*) = -10$. We find that the loads due to σ are as follows:

$$\text{load}(1|\sigma) = 10 \cdot \frac{7}{30} \cdot \frac{1}{2} + 5 \cdot \frac{29}{60} \cdot \frac{1}{2} = \frac{19}{8}$$
$$\text{load}(2|\sigma) = 10 \cdot \frac{13}{30} \cdot \frac{1}{2} + 5 \cdot \frac{11}{60} \cdot \frac{1}{2} = \frac{21}{8}$$
$$\text{load}(3|\sigma) = \frac{5}{2}.$$

Then we get

$$u(\sigma^*; \sigma) = -1/6 \cdot (37.5 + 22.5) = -10,$$

so overall

$$u\left(\sigma^*; (1\%)\sigma + (99\%)\sigma^*\right) = -10.$$

In the mixed population, the expected payoff for the mutants is given by

$$u(\sigma; (1\%)\sigma + (99\%)\sigma^*) = (1\%) \cdot u(\sigma; \sigma) + (99\%) \cdot u(\sigma; \sigma^*).$$

Since all links have equal latency given the equilibrium strategy σ^*, we have

$$u(\sigma; \sigma^*) = u(\sigma^*; \sigma^*) = -10.$$

We also find that

$$u(\sigma; \sigma) = -\frac{1}{2} \cdot \left(\frac{501}{40} + \frac{597}{80}\right) = -\frac{1}{2} \cdot \left(\frac{1599}{80}\right) = -9.99375.$$

So all told,

$$u(\sigma; (1\%)\sigma + (99\%)\sigma^*) = (1\%) \cdot (-9.99375) - (99\%) \cdot 10 \geq -10 = u(\sigma^*; (1\%)\sigma + (99\%)\sigma^*).$$

So in the mixed population, the mutants do better on average than the incumbents.

As our Bayesian task allocation game is a symmetric game, we can generalize the standard definition of an ESS [41] for 2-player games to n-player games following [6]. Note that in the following definition ε can depend on n.

Definition 4 (*ESS*). Let G be a symmetric Bayesian task allocation game with n players. A mixed strategy σ^* is an **evolutionarily stable strategy** (*ESS*) $\iff$ there is an $\overline{\varepsilon} > 0$ such that for all $0 < \varepsilon < \overline{\varepsilon}$ and mixed strategies $\sigma \neq \sigma^*$ we have $u(\sigma^*; [\varepsilon\sigma + (1-\varepsilon)\sigma^*]^{(n-1)}) > u(\sigma; [\varepsilon\sigma + (1-\varepsilon)\sigma^*]^{(n-1)})$.

The payoff of every mixed strategy σ' in a mixed population with distribution $\varepsilon\sigma + (1-\varepsilon)\sigma^*$ and a base game with n players can be computed by summing over the payoffs to σ' when faced with $k = 0, 1, \ldots, n-1$ mutants σ and $n-1-k$ incumbents σ^*, weighted by the probability of encountering exactly k mutants. We denote the payoff from playing mixed strategy σ' against k mutants σ by $u(\sigma'; (\sigma^*)^{(n-1-k)}, \sigma^{(k)})$. This payoff is given by

$$u(\sigma'; (\sigma^*)^{(n-1-k)}, \sigma^{(k)}) = -\sum_{w \in W} \sum_{\ell \in L} [w/c_\ell + (n-1-k)\text{load}(\ell|\sigma^*) + k \cdot \text{load}(\ell|\sigma)] \cdot \sigma'(\ell|w) \cdot \mu(w).$$

The probability of encountering exactly k mutants is

$$\binom{n-1}{k} \cdot \varepsilon^k \cdot (1-\varepsilon)^{n-1-k}.$$

Overall, the payoff of using strategy σ in the mixed population can be computed as

$$u(\sigma'; [\varepsilon\sigma + (1-\varepsilon)\sigma^*]^{n-1}) = \sum_{k=0}^{(n-1)} u(\sigma'; (\sigma^*)^{(n-1-k)}, \sigma^{(k)}) \cdot \binom{n-1}{k} \cdot \varepsilon^k \cdot (1-\varepsilon)^{n-1-k}.$$

3. Link group uniqueness of symmetric Bayesian Nash equilibria

This section investigates the structure of symmetric Bayesian Nash equilibria and establishes that symmetric equilibria are uniquely determined up to the distribution of tasks within link groups. Note that, in a large population model, a symmetric equilibrium $\sigma^{(n)}$ represents an equilibrium state of the population: given that the aggregate allocation of tasks to links corresponds to mixed strategy σ, no single member of the population can improve his payoff by unilaterally changing his strategy. Hence, an ESS can be considered a special case of a symmetric Bayesian Nash equilibrium that satisfies a stability condition (see [6] and Section 4). So an ESS inherits the mathematical properties of symmetric Bayesian Nash equilibria.

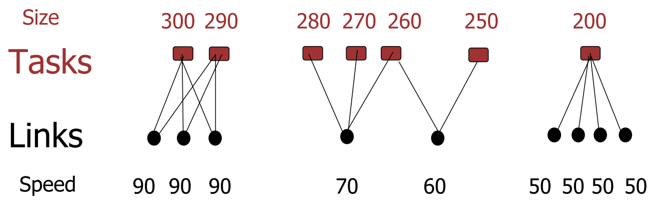

Fig. 2. A typical allocation of tasks to links in a *symmetric Bayesian Nash equilibrium*, as characterized by Lemma 3.1. Links in the same link group (e.g., with speed 90 in the figure) may share two tasks (of size 300 and 290 in the figure). In contrast, the links with speeds 70 and 60 may share at most one task (of size 260 in the figure).

A *link group* $\mathcal{L}$ in a symmetric Bayesian task allocation game G is a maximal set of links with the same speed, that is, $c_\ell = c_{\ell'}$ for all $\ell, \ell' \in \mathcal{L}$. Then, for any mixed strategy σ, the probability that σ sends task w to a link in link group $\mathcal{L}$ is given by

$$\sigma(\mathcal{L}|w) \equiv \sum_{\ell \in \mathcal{L}} \sigma_\ell(w).$$

The main result of this section is that in any symmetric Bayesian task allocation game the aggregate distribution over groups of links with the same speed is *uniquely determined* for symmetric Bayesian Nash equilibria. In other words, the probabilities $\sigma_{\mathcal{L}}$ are uniquely determined in a symmetric Bayesian Nash equilibrium; if $\sigma^{(n)}$ and $(\sigma')^{(n)}$ are Bayesian Nash equilibria in a task allocation game B, then for every link group $\mathcal{L}$ and every task weight w we have $\sigma(\mathcal{L}|w) = \sigma'(\mathcal{L}|w)$.

The next lemma gives a clear picture of what a symmetric Bayesian Nash equilibrium looks like. Intuitively, this picture is the following. (1) Tasks with bigger weights are placed on faster links. (2) Faster links have a bigger load. (3–5) For every link ℓ there is an "interval" of ordered task weights $w_1 < \cdots < w_k$ such that ℓ is optimal for all and only these weights. (6) Any pair of links with different speeds are optimal for at most one common task weight. Fig. 2 illustrates this structure.

Lemma 3.1. *Let G be a symmetric Bayesian task allocation game with n players and a symmetric Bayesian Nash equilibrium $\sigma^{(n)}$. Fix any two links ℓ and ℓ'.*

1. *If $c_\ell > c_{\ell'}$, strategy σ uses ℓ for w and ℓ' for w', then $w \geq w'$.*
2. *If $c_\ell > c_{\ell'}$, then $\text{load}(\ell|\sigma^{(n)}) > \text{load}(\ell'|\sigma^{(n)})$, or $\text{load}(\ell|\sigma^{(n)}) = \text{load}(\ell'|\sigma^{(n)}) = 0$. If $c_\ell = c_{\ell'}$, then $\text{load}(\ell|\sigma^{(n)}) = \text{load}(\ell'|\sigma^{(n)})$.*
3. *If $c_\ell > c_{\ell'}$, then there cannot exist tasks $w > w'$ such that σ uses both links ℓ and ℓ' for both tasks w and w'.*
4. *If $w \geq w' \geq w''$ and link ℓ is optimal for tasks w, w'' given σ, then ℓ is optimal for w' given σ.*
5. *If $c_\ell > c_{\ell'} > c_{\ell''}$ and links ℓ and ℓ'' are both optimal for w given σ, then link ℓ' is optimal for w given σ.*
6. *If $c_\ell > c_{\ell'}$, then there is at most one task w such that both links ℓ and ℓ' are optimal for w given σ.*

Proof. We begin with the proof of Part 1.

Part 1. We show that if $c_\ell > c_{\ell'}$, link ℓ is optimal for w given σ, and link ℓ' is optimal for w' given σ, then $w \geq w'$. Together with Proposition 2.2, the claim follows.
Assume that $c_\ell > c_{\ell'}$ and σ uses link ℓ for w and link ℓ' for w'. Suppose for contradiction that $w' > w$. Since ℓ is optimal for w we have

$$\frac{w}{c_\ell} + (n-1) \cdot \text{load}(\ell|\sigma) \leq \frac{w}{c_{\ell'}} + (n-1) \cdot \text{load}(\ell'|\sigma).$$

Since $c_\ell > c_{\ell'}$, it follows that for any $x > 0$ we have

$$\frac{w}{c_\ell} + \frac{x}{c_\ell} + (n-1) \cdot \text{load}(\ell|\sigma) < \frac{w}{c_{\ell'}} + \frac{x}{c_{\ell'}} + (n-1) \cdot \text{load}(\ell'|\sigma).$$

In particular we may take $x = w' - w$ which yields

$$\frac{w'}{c_\ell} + (n-1) \cdot \text{load}(\ell|\sigma) < \frac{w'}{c_{\ell'}} + (n-1) \cdot \text{load}(\ell'|\sigma).$$

But then link ℓ' is not optimal for w' given $\sigma^{(n-1)}$, which contradicts the hypothesis that $\sigma^{(n)}$ is a Bayesian Nash equilibrium.

Part 2. If $\text{load}(\ell'|\sigma^{(n)}) = 0$, the claim follows immediately. Suppose that $\text{load}(\ell'|\sigma^{(n)}) > 0$; then there is a task w such that σ uses ℓ' for w. For a contradiction, assume $\text{load}(\ell|\sigma^{(n)}) \leq \text{load}(\ell'|\sigma^{(n)})$. We show that then ℓ' is not optimal for w. Using Observation 2.1 we get $\text{load}(\ell|\sigma^{(n)}) = n \cdot \text{load}(\ell|\sigma)$, and it is sufficient to show $(n-1) \cdot \text{load}(\ell|\sigma) \leq (n-1) \cdot \text{load}(\ell'|\sigma)$. Hence, since $c_\ell > c'_{\ell}$,

$$\frac{w}{c_\ell} + (n-1) \cdot \text{load}(\ell|\sigma) < \frac{w}{c_{\ell'}} + (n-1) \cdot \text{load}(\ell'|\sigma),$$

showing that ℓ' is not optimal for w. That contradicts the hypothesis that $\sigma^{(n)}$ is a Bayesian Nash equilibrium.

Part 3. Follows immediately from Part 1.

Part 4. Suppose that it is not the case that $w' \in \text{opt}(\ell|\sigma)$. Then there is some link $\ell' \neq \ell$ such that $w' \in \text{opt}(\ell'|\sigma)$. Note that $c_\ell \neq c_{\ell'}$ for otherwise both ℓ and ℓ' are optimal for w'. First suppose that $c_{\ell'} > c_\ell$. Then, since $w \in \text{opt}(\ell|\sigma)$,

$$\frac{w}{c_{\ell'}} + (n-1) \cdot \text{load}(\ell'|\sigma) \geq \frac{w}{c_\ell} + (n-1) \cdot \text{load}(\ell|\sigma).$$

Let $x = w' - w < 0$. Then

$$\frac{w}{c_{\ell'}} + \frac{x}{c_{\ell'}} + (n-1) \cdot \text{load}(\ell'|\sigma) > \frac{w}{c_\ell} + \frac{x}{c_\ell} + (n-1) \cdot \text{load}(\ell|\sigma),$$

so

$$\frac{w'}{c_{\ell'}} + (n-1) \cdot \text{load}(\ell'|\sigma) > \frac{w'}{c_\ell} + (n-1) \cdot \text{load}(\ell|\sigma),$$

which contradicts the hypothesis that link ℓ' is optimal for w'. The case in which $c_{\ell'} < c_\ell$ is symmetric with w'' instead of w.

Part 5. We show the stronger claim that $\text{opt}(\ell'|\sigma) = \{w\}$.Suppose for contradiction that $w' \in \text{opt}(\ell'|\sigma)$ holds and $w \neq w'$. If $w < w'$, then $w \in \text{opt}(\ell'|\sigma)$ because otherwise we would violate Part 1. Similarly, if $w' < w$, then $w \in \text{opt}(\ell'|\sigma)$ because otherwise we would violate Part 1 (this time with ℓ and ℓ' reversed). This shows that $w = w'$, and so $\text{opt}(\ell'|\sigma) \subseteq \{w\}$. It is also easy to see that there must be some weight on ℓ_2 or else ℓ_3 is not optimal for any weight.

Part 6. Suppose for contradiction that w and w' are each in $\text{opt}(\ell|\sigma) \cap \text{opt}(\ell'|\sigma)$ where $w \neq w'$. Then we have

$$\frac{w}{c_\ell} + (n-1) \cdot \text{load}(\ell|\sigma) = \frac{w}{c_{\ell'}} + (n-1) \cdot \text{load}(\ell'|\sigma).$$

This is equivalent to

$$(n-1) \cdot (\text{load}(\ell|\sigma) - \text{load}(\ell'|\sigma)) = \frac{w}{c_{\ell'}} - \frac{w}{c_\ell}.$$

The same holds for w'. Hence

$$\frac{w}{c_{\ell'}} - \frac{w}{c_\ell} = \frac{w'}{c_{\ell'}} - \frac{w}{c_\ell} \iff \frac{(w-w')}{c_{\ell'}} = \frac{(w-w')}{c_\ell}.$$

Since $w \neq w'$, this implies that $c_\ell = c_{\ell'}$, which is a contradiction. □

We note that Lemma 3.1 holds for Bayesian Nash equilibria in general, not just symmetric ones. Specifically, let σ' be a Bayesian Nash equilibrium for a symmetric Bayesian task allocation game, and fix any player i such that $\sigma' = (\sigma'_i, \sigma'_{-i})$. Then Lemma 3.1 holds if we replace a mixed strategy σ with σ'_i, and $\sigma^{(n-1)}$ with σ'_{-i}, and $\sigma^{(n)}$ with $(\sigma'_i, \sigma'_{-i})$.

We extend our notation for links to link groups $\mathcal{L}$ such that $c_{\mathcal{L}}$ denotes the speed of all links in group $\mathcal{L}$. We also define

$$\text{load}(\mathcal{L}|\sigma^{(n)}) \equiv \sum_{\ell \in \mathcal{L}} \text{load}(\ell|\sigma^{(n)}).$$

The next theorem is the main result of this section. It states that for a user population in equilibrium (corresponding to a symmetric Bayesian Nash equilibrium), the distribution of tasks to *link groups* is uniquely determined. Thus the only way in which population equilibria can differ is by how tasks are allocated within a link group. This result is the first key step for establishing the uniqueness of an ESS for a symmetric Bayesian task allocation game.

Theorem 3.2 (*Link Group Uniqueness*). *Let G be a symmetric Bayesian task allocation game with n players and two symmetric Bayesian Nash equilibria $\sigma^{(n)}$ and $(\sigma')^{(n)}$. Then we have $\sigma(\mathcal{L}|w) = \sigma'(\mathcal{L}|w)$ and $\text{load}(\mathcal{L}|\sigma^{(n)}) = \text{load}(\mathcal{L}|(\sigma')^{(n)})$ for all task sizes w and link groups $\mathcal{L}$ of B.*

Proof. Due to Lemma 3.1 (Parts 1 and 2), for every fixed load distribution there is at most one mixed strategy σ that induces this load distribution over link groups. So it suffices to show $\text{load}(\mathcal{L}|\sigma^{(n)}) = \text{load}(\mathcal{L}|(\sigma')^{(n)})$. Due to Observation 2.1, this is equivalent to showing that $\text{load}(\mathcal{L}|\sigma) = \text{load}(\mathcal{L}|\sigma')$.

For a contradiction, assume there exists a link group $\mathcal{L}$ with $\text{load}(\mathcal{L}|\sigma) \neq \text{load}(\mathcal{L}|\sigma')$. If there are several such link groups choose $\mathcal{L}$ such that $c_{\mathcal{L}} > c_{\mathcal{L}''}$ for all groups $\mathcal{L}''$ with $\text{load}(\mathcal{L}''|\sigma) \neq \text{load}(\mathcal{L}''|\sigma')$. Without loss of generality we can assume that $\text{load}(\mathcal{L}|\sigma) > \text{load}(\mathcal{L}|\sigma')$. Due to Lemma 3.1 (Part 1) and the fact that we have monotone (linear) latency functions, there has to exist a link group $\mathcal{L}'$ with $\text{load}(\mathcal{L}'|\sigma) < \text{load}(\mathcal{L}'|\sigma')$. If there exist several such link groups choose $\mathcal{L}'$ such that $c_{\mathcal{L}'} > c_{\mathcal{L}''}$ for all link group $\mathcal{L}''$ with $\text{load}(\mathcal{L}''|\sigma) < \text{load}(\mathcal{L}''|\sigma')$. This gives us

$$\sum_{\mathcal{L}'' :\, c_{\mathcal{L}''} > c_{\mathcal{L}'}} \text{load}(\mathcal{L}''|\sigma) > \sum_{\mathcal{L}'' :\, c_{\mathcal{L}''} > c_{\mathcal{L}'}} \text{load}(\mathcal{L}''|\sigma'). \tag{1}$$

Due to Lemma 3.1 the task distribution in any Bayesian Nash equilibrium looks as follows. Assume that the link groups are ordered from left to right in non-increasing order of the capacities of their links. Then the tasks are assigned to the links of the link groups in non-increasing order, starting at the leftmost link group. Hence, for $c_\ell > c_{\ell'}$ the tasks assigned to link

ℓ are not smaller than the tasks assigned to ℓ' (Part 1 of Lemma 3.1). Let K be the set of link groups $\mathcal{L}''$ with $c_\ell > c_{\ell'}$ for $\ell \in \mathcal{L}''$ and $\ell' \in \mathcal{L}'$. Then there exists a task w such that w is the smallest task that is assigned to any link in K by σ. Let p (p') be the probability that w is assigned to a task in K by σ (σ'). Then $p > p'$ (see Eq. (1) above). Note that $p' = 0$ is the case the tasks assigned to links in K in σ' is a *proper* subset of the tasks assigned to links in K in σ.

Now let $w' \geq w$ be the largest job used by a task in K such that σ' uses links in $\mathcal{L}'$ for w'. σ uses a group $\mathcal{L}''$ with $c_{\mathcal{L}''} > c_{\mathcal{L}'}$ for some fraction of w (see Lemma 3.1). Note that due to Lemma 3.1 (Part 2) the load of all links in $\mathcal{L}$ is the same. The same holds for the link load of $\mathcal{L}'$ and $\mathcal{L}''$. Now pick arbitrary links $\ell'' \in \mathcal{L}''$, $\ell' \in \mathcal{L}'$, and $\ell \in \mathcal{L}$. Then we get

$$\begin{aligned}\frac{w'}{c_{\ell''}} + \text{load}(\ell''|\sigma) &\leq \frac{w'}{c_{\ell'}} + \text{load}(\ell'|\sigma)\\ &< \frac{w'}{c_{\ell'}} + \text{load}(\ell'|\sigma')\\ &\leq \frac{w'}{c_{\ell''}} + \text{load}(\ell''|\sigma')\\ &\leq \frac{w'}{c_{\ell''}} + \text{load}(\ell''|\sigma).\end{aligned}$$

The first inequality holds since σ is a Bayesian Nash equilibrium. The second strict inequality is due to the choice of $\mathcal{L}'$. The third inequality follows because σ' is a Bayesian Nash equilibrium. The last inequality holds since $c_{\ell''} > c_{\ell'}$ and since $\mathcal{L}'$ is from all link groups with $\text{load}(\mathcal{L}'|\sigma) < \text{load}(\mathcal{L}'|\sigma')$ the one with the largest speed. □

In the task allocation game of Fig. 1, there is only one link group, since all links have the same capacity 1. So Theorem 3.2 implies that there is only symmetric Bayesian Nash equilibrium in this game, which corresponds to the uniform distribution of tasks to links. The next section begins the investigation of evolutionarily stable equilibria; we begin with a characterization of ESS for a symmetric Bayesian task allocation game that provides a simple necessary and sufficient condition for a deviation from a current population state to be successful.

4. Characterization of evolutionary stability

Following the usual convention of evolutionary game theory, we use σ^* to refer to the mixed strategy associated with the incumbent population. In this section we prove a necessary and sufficient condition for a mixed strategy σ^* to be an ESS. The next proposition shows that for sufficiently small sizes of mutations, only best replies to the incumbent distribution σ^* have the potential to do better than the incumbent. The proposition also implies that an ESS corresponds to a symmetric Bayesian Nash equilibrium (Corollary 4.2).

Proposition 4.1. *Let G be a symmetric Bayesian task allocation game with n players, and let σ^* be a mixed strategy. Then there is a threshold $\overline{\varepsilon}$ such that for all ε with $0 < \varepsilon < \overline{\varepsilon}$, for all mixed strategies σ:*

1. *If $u(\sigma^*; (\sigma^*)^{(n-1)}) > u(\sigma; (\sigma^*)^{(n-1)})$, then*
$$u(\sigma^*; [\varepsilon\sigma + (1-\varepsilon)\sigma^*]^{(n-1)}) > u(\sigma; [\varepsilon\sigma + (1-\varepsilon)\sigma^*]^{(n-1)}).$$
2. *If $u(\sigma^*; (\sigma^*)^{(n-1)}) < u(\sigma; (\sigma^*)^{(n-1)})$, then*
$$u(\sigma^*; [\varepsilon\sigma + (1-\varepsilon)\sigma^*]^{(n-1)}) < u(\sigma; [\varepsilon\sigma + (1-\varepsilon)\sigma^*]^{(n-1)}).$$

Proof. The proof requires only standard techniques from evolutionary game theory [41] and is omitted. Intuitively, the result holds because we can choose our threshold $\overline{\varepsilon}$ small enough (as a function of B and σ^*) so that any difference in the case in which the mutant and incumbent face 0 mutants outweighs the differences in their payoffs when they face one or more mutants. □

Proposition 4.1(1) says that if a mutation σ is sufficiently small and a worse reply to the distribution σ^* than σ^* itself, then the mutation does worse in the mixed population than the incumbent. Similarly, the second part of the proposition says that if a mutation σ is sufficiently small and a better reply to the distribution σ^* than σ^* itself, then the mutation does better in the mixed population than the incumbent. Corollary 4.2 shows that an ESS must correspond to a Bayesian Nash equilibrium, and that we need only consider best replies to an incumbent strategy to evaluate evolutionary stability. Then we provide a necessary and sufficient condition for a best reply to be a successful mutation.

Corollary 4.2. *Let G be a symmetric Bayesian task allocation game with n players, and let σ^* be an* ESS. *Then $(\sigma^*)^{(n)}$ is also a Bayesian Nash equilibrium.*

Proof. If σ^* is not a best reply to $(\sigma^*)^{(n-1)}$, then there is a mutant σ such that $u(\sigma; (\sigma^*)^{(n-1)}) > u(\sigma^*; (\sigma^*)^{(n-1)})$. Proposition 4.1(2) then implies that σ is a successful mutation no matter how low we choose the positive threshold $\overline{\varepsilon}$. □

The next Lemma 4.3 provides a necessary and sufficient condition for when a best reply is a successful mutation, which is key for our analysis of evolutionarily stable strategies in a given network game. In the following let

$$\sigma(\ell|W) := \sum_{w\in W} \sigma(\ell|w)\cdot\mu(w).$$

Then the condition of the lemma is as follows: consider an equilibrium strategy σ^* and a mutation $\sigma \neq \sigma^*$ that is a best reply to σ^*. Then σ is successful if and only if

$$\sum_{\ell\in L}[\text{load}(\ell|\sigma^*) - \text{load}(\ell|\sigma)]\cdot[\sigma^*(\ell|W) - \sigma(\ell|W)] < 0.$$

The condition of Lemma 4.3 may be interpreted as follows. For a fixed link ℓ, the expression $\text{load}(\ell|\sigma^*) - \text{load}(\ell|\sigma)$ measures the difference between the loads on ℓ caused by the incumbent population σ^* and caused by the mutant population σ. In the first case, $\text{load}(\ell|\sigma) > \text{load}(\ell|\sigma^*)$. Then, the mutation increases the load on link ℓ. Therefore, in the mutant population, link ℓ is slower compared to the link in the incumbent population. Hence, the cost of a strategy σ' with respect to link ℓ increases as the probability $\sigma'(\ell|W)$ increases. In particular, if a) $\sigma^*(\ell|W) > \sigma(\ell|W)$, then the mutation does relatively better on link ℓ. In this case $\text{load}(\ell|\sigma^*) - \text{load}(\ell|\sigma) < 0$ and $\sigma^*(\ell|W) - \sigma(\ell|W) > 0$ so the product $[\text{load}(\ell|\sigma^*) - \text{load}(\ell|\sigma)]\cdot[\sigma^*(\ell|W) - \sigma(\ell|W)]$ is negative. If b) $\sigma^*(\ell|W) < \sigma(\ell|W)$, then $\sigma^*(\ell|W) - \sigma(\ell|W) < 0$ and the product is positive. In the second case $\text{load}(\ell|\sigma^*) > \text{load}(\ell|\sigma)$. Then similarly the product $[\text{load}(\ell|\sigma^*) - \text{load}(\ell|\sigma)]\cdot[\sigma^*(\ell|W) - \sigma(\ell|W)]$ is negative if the decrease benefits the mutants more than the incumbents, and positive otherwise.

Example 4. Consider again the task allocation game of Fig. 1 and the uniform distribution mixed strategy σ^* defined by $\sigma^*(\ell|w) = 1/3$ for all links $l \in [3]$. The task sizes $w \in \{5, 10\}$, each size is generated with probability 1/2. A possible mutant population M may be described by the following mixed strategy σ.

w	$\sigma(1\|w)$	$\sigma(2\|w)$	$\sigma(3\|w)$
5	29/60	11/60	1/3
10	7/30	13/30	1/3

Since given the uniform distribution σ^*, all links are optimal for every task, every mutant distribution is a best reply to σ^*; in particular, we found in Section 2.2 that $u(\sigma^*;\sigma^*) = u(\sigma;\sigma^*) = -10$.

The aggregate probabilities of σ^* are the same for all links:

$$\sigma^*(\ell|W) = \frac{1}{3}\cdot\frac{1}{2} + \frac{1}{3}\cdot\frac{1}{2} = \frac{1}{3}.$$

For σ we have

$$\sigma(1|W) = \frac{7}{30}\cdot\frac{1}{2} + \frac{29}{60}\cdot\frac{1}{2} = \frac{43}{120}$$
$$\sigma(2|W) = \frac{13}{30}\cdot\frac{1}{2} + \frac{11}{60}\cdot\frac{1}{2} = \frac{37}{120}$$
$$\sigma(3|W) = \frac{1}{3} = \sigma(3|W).$$

So the differences in the aggregate probabilities are as follows:

$$\sigma^*(1|W) - \sigma(1|W) = \frac{1}{3} - \frac{43}{120} = \frac{-1}{40}$$
$$\sigma^*(2|W) - \sigma(2|W) = \frac{1}{3} - \frac{37}{120} = \frac{1}{40}$$
$$\sigma^*(3|W) - \sigma(3|W) = \frac{1}{3} - \frac{1}{3} = 0.$$

The load differences due to the incumbent and mutant strategies are as follows (see Example 3):

$$\text{load}(1|\sigma^*) - \text{load}(1|\sigma) = \frac{20}{8} - \frac{19}{8} = \frac{1}{8}$$
$$\text{load}(2|\sigma^*) - \text{load}(2|\sigma) = \frac{20}{8} - \frac{21}{8} = \frac{-1}{8}$$
$$\text{load}(3|\sigma^*) - \text{load}(3|\sigma) = \frac{20}{8} - \frac{20}{8} = 0.$$

Now comparing the aggregate usage, we find that there is no difference with respect to link 3 since both strategies use it in the same way. The mutant strategy uses link 1 more often overall than the incumbent ($\sigma^*(1|W) - \sigma(1|W) < 0$) and places less load on it ($\text{load}(1|\sigma^*) - \text{load}(1|\sigma) > 0$). The mutant strategy uses link 2 less often overall than the incumbent ($\sigma^*(2|W) - \sigma(2|W) > 0$) and places more load on it ($\text{load}(2|\sigma^*) - \text{load}(2|\sigma) < 0$). So the sum

$$\sum_{\ell \in L} [\text{load}(\ell|\sigma^*) - \text{load}(\ell|\sigma)] \cdot [\sigma^*(\ell|W) - \sigma(\ell|W)]$$

is strictly negative, which according to the following lemma is necessary and sufficient for a successful mutation. We verified in Section 2.2 directly that σ describes a successful mutation at a relative size of 1% compared to the incumbent population.

Lemma 4.3. *Let G be a symmetric Bayesian task allocation game with n players. Let $(\sigma^*)^{(n)}$ be a Bayesian Nash equilibrium, and consider any best reply σ to $(\sigma^*)^{(n-1)}$. Then the sum*

$$\sum_{\ell \in L} [\text{load}(\ell|\sigma^*) - \text{load}(\ell|\sigma)] \cdot [\sigma^*(\ell|W) - \sigma(\ell|W)]$$

1. *is negative if and only if for all ε with $0 < \varepsilon < 1$,*
$$u(\sigma^*; [\varepsilon\sigma + (1-\varepsilon)\sigma^*]^{(n-1)}) < u(\sigma; [\varepsilon\sigma + (1-\varepsilon)\sigma^*]^{(n-1)})$$
(i.e., if and only if the mutation does better in the mixed population).
2. *is positive if and only if for all ε with $0 < \varepsilon < 1$,*
$$u(\sigma^*; [\varepsilon\sigma + (1-\varepsilon)\sigma^*]^{(n-1)}) > u(\sigma; [\varepsilon\sigma + (1-\varepsilon)\sigma^*]^{(n-1)})$$
(i.e., if and only if the mutation does worse in the mixed population).

Proof. We begin with two preliminary claims. As discussed in Section 2.2, for deviant strategies that are best replies to the current equilibrium, the success of the mutation depends on how well the mutation and the incumbent strategies do against $k = 1, \ldots, n-1$ mutants respectively. That is, we have to compare $u(\sigma^*; (\sigma^*)^{(n-k-1)}; \sigma^k)$ with $u(\sigma; (\sigma^*)^{(n-k-1)}; \sigma^k)$ for $k = 1, \ldots, n-1$. The next claims helps to compute these differences.

Claim 1. *Let G be a symmetric task allocation game with n players and mixed strategies $\sigma, \sigma^*, \sigma'$. Then for each k with $0 \le k < n$ we have*

$$u(\sigma'; (\sigma^*)^{(n-k-2)}, \sigma^{(k+1)}) - u(\sigma'; (\sigma^*)^{(n-k-1)}, \sigma^{(k)}) = \sum_{w \in W} \sum_{\ell \in L} \left[\text{load}(\ell|\sigma^*) - \text{load}(\ell|\sigma)\right] \cdot \sigma'(\ell|w) \cdot \mu(w).$$

Proof of claim. By definition

$$\begin{aligned}
&u(\sigma'; (\sigma^*)^{(n-k-2)}, \sigma^{(k+1)}) - u(\sigma'; (\sigma^*)^{(n-k-1)}, \sigma^{(k)}) \\
&\quad = -\sum_{w \in W} \sum_{\ell \in L} \left(\frac{w}{c_\ell} + (n-k-2) \cdot \text{load}(\ell|\sigma^*) + (k+1)\text{load}(\ell|\sigma)\right) \cdot \sigma'(\ell|w) \cdot \mu(w) \\
&\qquad + \sum_{w \in W} \sum_{\ell \in L} \left(\frac{w}{c_\ell} + (n-k-1) \cdot \text{load}(\ell|\sigma^*) + k \cdot \text{load}(\ell|\sigma)\right) \cdot \sigma'(\ell|w) \cdot \mu(w) \\
&\quad = \sum_{w \in W} \sum_{\ell \in L} \left[\text{load}(\ell|\sigma^*) - \text{load}(\ell|\sigma)\right] \cdot \sigma'(\ell|w) \cdot \mu(w).
\end{aligned}$$

Claim 2.

$$\begin{aligned}
\sum_{\ell \in L} \left[\text{load}(\ell|\sigma^*) - \text{load}(\ell|\sigma)\right]\left[\sigma^*(\ell|W) - \sigma(\ell|W)\right] &= \left[u(\sigma^*; (\sigma^*)^{(n-k-2)}, \sigma^{(k+1)}) - u(\sigma^*; (\sigma^*)^{(n-k-1)}, \sigma^{(k)})\right] \\
&\quad - \left[u(\sigma; (\sigma^*)^{(n-k-2)}, \sigma^{(k+1)}) - u(\sigma; (\sigma^*)^{(n-k-1)}, \sigma^{(k)})\right].
\end{aligned}$$

Proof of claim. Using Claim 1 we get

$$\begin{aligned}
&\left[u(\sigma^*; (\sigma^*)^{(n-k-2)}, \sigma^{(k+1)}) - u(\sigma^*; (\sigma^*)^{(n-k-1)}, \sigma^{(k)})\right] - \left[u(\sigma; (\sigma^*)^{(n-k-2)}, \sigma^{(k+1)}) - u(\sigma; (\sigma^*)^{(n-k-1)}, \sigma^{(k)})\right] \\
&\quad = \sum_{w \in W} \sum_{\ell \in L} \left[\text{load}(\ell|\sigma^*) - \text{load}(\ell|\sigma)\right] \cdot \sigma^*(\ell|w) \cdot \mu(w) - \sum_{w \in W} \sum_{\ell \in L} \left[\text{load}(\ell|\sigma^*) - \text{load}(\ell|\sigma)\right] \cdot \sigma(\ell|w) \cdot \mu(w) \\
&\quad = \sum_{w \in W} \sum_{\ell \in L} [\text{load}(\ell|\sigma^*) - \text{load}(\ell|\sigma)] \cdot \mu(w) \cdot [\sigma^*(\ell|w) - \sigma(\ell|w)] \\
&\quad = \sum_{\ell \in L} \left[\text{load}(\ell|\sigma^*) - \text{load}(\ell|\sigma)\right] \cdot \sum_{w \in W} \mu(w) \cdot \left(\sigma^*(\ell|w) - \sigma(\ell|w)\right).
\end{aligned}$$

By definition

$$\sum_{w \in W} \mu(w) \cdot \left(\sigma^*(\ell|w) - \sigma(\ell|w)\right) = \sum_{w \in W} \sigma^*(\ell|w) \cdot \mu(w) - \sum_{w \in W} \sigma(\ell|w) \cdot \mu(w) = [\sigma^*(\ell|W) - \sigma(\ell|W)],$$

and the claim follows.

Proof of Lemma 4.3 continued. Now we present the proof of Part (1) of Lemma 4.3. The proofs of Part (2) and (3) can be done in the same way.

Suppose that

$$\sum_{\ell \in L} \left[\text{load}(\ell|\sigma^*) - \text{load}(\ell|\sigma)\right] \cdot \left[\sigma^*(\ell|W) - \sigma(\ell|W)\right] < 0.$$

We show by induction on k that whenever $0 \leq k < n-1$, the difference

$$u(\sigma^*; (\sigma^*)^{(n-k-2)}, \sigma^{(k+1)}) - u(\sigma; (\sigma^*)^{(n-k-2)}, \sigma^{(k+1)}) < 0.$$

Since the payoff difference $u(\sigma^*; [\varepsilon\sigma + (1-\varepsilon)\sigma^*]^{(n-1)}) - u(\sigma; [\varepsilon\sigma + (1-\varepsilon)\sigma^*]^{(n-1)})$ can be computed as the sum of differences over the numbers k of mutants that strategy σ or σ^* may encounter, it follows that this difference is negative, which establishes the lemma.

Base case, $k = 0$. Since both σ^* and σ are best replies to $(\sigma^*)^{(n-1)}$, we have $u(\sigma^*; (\sigma^*)^{(n-1)}) = u(\sigma; (\sigma^*)^{(n-1)})$. So, using Claim 2

$$\begin{aligned} &u(\sigma^*; (\sigma^*)^{(n-2)}, \sigma) - u(\sigma; (\sigma^*)^{(n-2)}, \sigma) \\ &= [u(\sigma^*; (\sigma^*)^{(n-2)}, \sigma) - u(\sigma^*; (\sigma^*)^{(n-1)})] - [u(\sigma; (\sigma^*)^{(n-2)}, \sigma) - u(\sigma; (\sigma^*)^{(n-1)})] \\ &= \sum_{\ell \in L} \left[\text{load}(\ell|\sigma^*) - \text{load}(\ell|\sigma)\right] \cdot \left[\sigma^*(\ell|W) - \sigma(\ell|W)\right] < 0. \end{aligned}$$

Inductive step: Assume the hypothesis for k and consider $k+1 < n$. By inductive hypothesis we get

$$u(\sigma^*; (\sigma^*)^{(n-k-1)}, \sigma^{(k)}) - u(\sigma; (\sigma^*)^{(n-k-1)}, \sigma^{(k)}) < 0.$$

Hence

$$\begin{aligned} &u(\sigma^*; (\sigma^*)^{(n-k-2)}, \sigma^{(k+1)}) - u(\sigma; (\sigma^*)^{(n-k-2)}, \sigma^{(k+1)}) \\ &< u(\sigma^*; (\sigma^*)^{(n-k-2)}, \sigma^{(k+1)}) - u(\sigma; (\sigma^*)^{(n-k-2)}, \sigma^{(k+1)}) + u(\sigma; (\sigma^*)^{(n-k-1)}, \sigma^{(k)}) - u(\sigma^*; (\sigma^*)^{(n-k-1)}, \sigma^{(k)}) \\ &= [u(\sigma^*; (\sigma^*)^{(n-k-2)}, \sigma^{(k+1)}) - u(\sigma^*; (\sigma^*)^{(n-k-1)}, \sigma^{(k)})] - [u(\sigma; (\sigma^*)^{(n-k-2)}, \sigma^{(k+1)}) - u(\sigma; (\sigma^*)^{(n-k-1)}, \sigma^{(k)})] \\ &= \sum_{\ell \in L} \left[\text{load}(\ell|\sigma^*) - \text{load}(\ell|\sigma)\right] \cdot \left[\sigma^*(\ell|W) - \sigma(\ell|W)\right] < 0. \end{aligned}$$

The last equality follows from Claim 2. So overall we have

$$\begin{aligned} &u(\sigma^*; (\sigma^*)^{(n-k-2)}, \sigma^{(k+1)}) - u(\sigma; (\sigma^*)^{(n-k-2)}, \sigma^{(k+1)}) \\ &< \sum_{\ell \in L} \left[\text{load}(\ell|\sigma^*) - \text{load}(\ell|\sigma)\right] \cdot \left[\sigma^*(\ell|W) - \sigma(\ell|W)\right] < 0, \end{aligned}$$

which completes the inductive step. □

The proof of Lemma 4.3 shows that a best reply σ to $(\sigma^*)^{(n-1)}$ that has a negative sum

$$\sum_{\ell \in L} [\text{load}(\ell|\sigma) - \text{load}(\ell|\sigma^*)] \cdot [\sigma^*(\ell|W) - \sigma(\ell|W)]$$

is successful in the following strong sense. The best reply yields a better payoff than the incumbent strategy σ^* *regardless* of the size of ε.

It will be convenient to say that a mutation σ **defeats** the incumbent strategy σ^* if the sum is negative. Similarly, we say that a mutation σ **equals** an incumbent σ^* if the sum is zero. In this terminology our results so far yield the following characterization of evolutionary stability. The proof directly follows from Proposition 4.1 and Lemma 4.3.

Corollary 4.4. *Let G be a symmetric Bayesian task allocation game with n players. A mixed strategy σ^* is an ESS for G $\iff$ the strategy profile $(\sigma^*)^{(n)}$ is a Bayesian Nash equilibrium, and no best reply $\sigma \neq \sigma^*$ to $(\sigma^*)^{(n-1)}$ defeats or equals σ^*.*

Lemma 3.1 clarified the structure of user populations in equilibrium. The next section applies the criterion from Corollary 4.4 to establish additional properties of populations in an evolutionarily stable equilibrium. In fact, these properties imply that an evolutionarily stable equilibrium is unique when it exists.

5. Uniqueness and structure of evolutionary stable strategies

We analyze the structure of evolutionary equilibria and show the uniqueness of ESS. For the first point, our focus is on the allocation of tasks to links that are consistent with evolutionary stability. Such results tell us how the structure of the network shapes evolutionary dynamics. They can be helpful for the development of algorithms calculating an ESS for a given system. The next theorem shows that in an evolutionary equilibrium there is minimal overlap in the tasks served by different links, in that two distinct links (even with the same speed) may not be used by tasks with different weights. In fact the result is stronger in that if link ℓ is used for task w and ℓ' for $w' \neq w$, then at least one of the links must not be optimal for the other link's task. This specialization result can be regarded as a stronger version of Lemma 3.1(6), where ℓ and ℓ' can have the

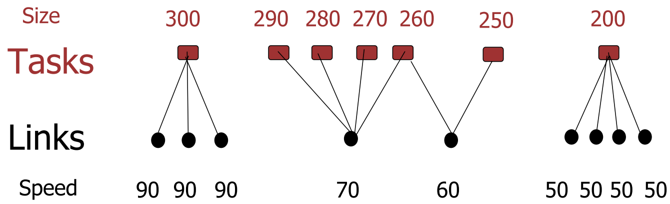

Fig. 3. A typical allocation of tasks to links in an ESS illustrating the necessary condition of Theorem 5.1. Links in the same link group (e.g., with speed 90 in the figure) may be allocated at most one task size (300 in the figure). Links that are the only ones with a given speed may be allocated more than one task size (see the links with speeds 70 and 60 in the figure), but may share at most one task size (260 in the figure).

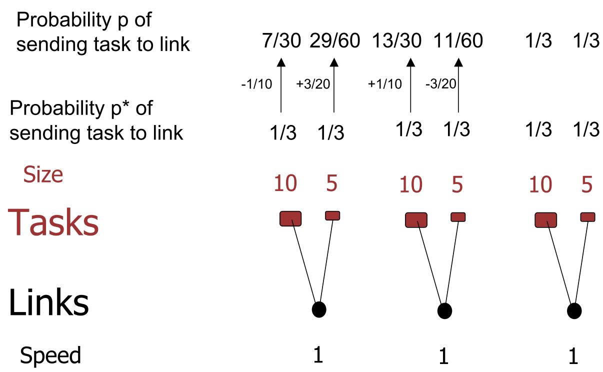

Fig. 4. To illustrate the construction for the proof of Theorem 5.1: Suppose that two links are respectively used by two tasks. Then a mutant distribution can increase the load on one of the two links – 2 in the example – while using it less frequently on average, and it can decrease the load on the other link – 1 in the example – while using it more frequently on average.

same speed. Unfortunately, the specialization condition of the Theorem is necessary but not sufficient, as Observation 5.3 will show. Fig. 3 illustrates the specialization condition.

The idea of the proof of the next theorem is that if two distinct links ℓ and ℓ' are used with a probability > 0 by users with different tasks, it is possible to create a "better" mutant distribution. The mutant distribution increases the load on one of the two links, say ℓ (by putting the task with the bigger weight with a larger probability onto ℓ, and, in turn, by putting the smaller task with smaller probability onto ℓ'), but uses the link overall with a smaller probability. Note that this strategy is possible only if we have different task weights.

Example 5. Consider again the task allocation game and the strategy illustrated in Fig. 1, where the uniform assignment of tasks to links is denoted σ^*. As we saw in Section 4, the mutation σ shifts load between links 1 and 2 to increase the load on 2 while decreasing its overall usage, and to decrease the load on 1 while increasing its overall usage. Fig. 4 diagrams this shift.

Theorem 5.1 (*Specialization*). *Let G be a symmetric Bayesian task allocation game with mixed strategy σ^*. Assume $w \neq w'$, $\ell \neq \ell'$ with $c_\ell \geq c_{\ell'}$, and suppose the following conditions are fulfilled.*

1. $w \in \text{support}(\ell|\sigma^*)$ *and* $w' \in \text{support}(\ell'|\sigma^*)$,
2. $w, w' \in \text{opt}(\ell|\sigma^*)$, *and* $w, w' \in \text{opt}(\ell'|\sigma^*)$.

Then there is a mutation σ that defeats σ^*, and hence σ^* is not evolutionarily stable.

Proof. To show a contradiction suppose that $(\sigma^*)^{(n)}$ is a Bayesian Nash equilibrium and, without loss of generality, assume that $w > w'$. We now want to define a mutant strategy that defeats the original strategy σ^*. For that we adjust the probabilities $\sigma^*(\ell|w)$, $\sigma^*(\ell'|w)$, $\sigma^*(\ell|w')$, and $\sigma^*(\ell'|w')$ by d_1 and d_2. The quantities d_1 and d_2 are chosen such that the following inequalities are fulfilled.

1. $\mu(w) \cdot d_1 < \mu(w') \cdot d_2$, and $d_1 \cdot \mu(w) \cdot w > d_2 \cdot \mu(w') \cdot w'$.
2. $\sigma^*(\ell|w) - d_1 > 0$, $\sigma^*(\ell|w') + d_2 < 1$, $\sigma^*(\ell'|w) + d_1 < 1$, and $\sigma^*(\ell'|w') - d_2 > 0$.

In the following we assume for the time being that these quantities d_1, d_2 exist with $d_1 \neq 0$ and $d_2 \neq 0$; at the end of the proof we show how to compute suitable values for d_1 and d_2. Then we define a mutant σ as follows. First, if $w'' \notin \{w, w'\}$ or $\ell'' \notin \{\ell, \ell'\}$ then $\sigma(\ell''|w'') = \sigma^*(\ell''|w'')$. Thus for all links $\ell'' \notin \{\ell, \ell'\}$ we have $\text{load}(\ell''|\sigma^*) = \text{load}(\ell''|\sigma)$. For the rest of the links and tasks we define σ as follows.

- $\sigma(\ell|w) = \sigma^*(\ell|w) - d_1$, and $\sigma(\ell|w') = \sigma^*(\ell|w') + d_2$,
- $\sigma(\ell'|w) = \sigma^*(\ell'|w) + d_1$, and $\sigma(\ell'|w') = \sigma^*(\ell'|w') - d_2$.

By the hypothesis of the theorem that the links ℓ, ℓ' are optimal for both w and w' it follows that σ is a best reply to $(\sigma^*)^{(n-1)}$. Observation 5.2 implies that

$$\sum_{\ell'' \in L}[\text{load}(\ell''|\sigma^*) - \text{load}(\ell''|\sigma)] \cdot [\sigma^*(\ell''|W) - \sigma(\ell''|W)] = [\text{load}(\ell|\sigma^*) - \text{load}(\ell|\sigma)] \cdot [\sigma^*(\ell|W) - \sigma(\ell|W)]$$
$$+ [\text{load}(\ell'|\sigma^*) - \text{load}(\ell'|\sigma)] \cdot [\sigma^*(\ell'|W) - \sigma(\ell'|W)] < 0.$$

By Lemma 4.3 σ is a successful mutation, contradicting the hypothesis that σ^* is an ESS.

It remains to show the existence of d_1, d_2. It is easy to see that the first condition in the definition of d_1 and d_2 gives us

$$\frac{\mu(w')}{\mu(w)} > \frac{d_1}{d_2} > \frac{\mu(w') \cdot w'}{\mu(w) \cdot w}.$$

Also, the conditions for d_1 are equivalent to $d_1 < \min\{\sigma^*(\ell|w), 1 - \sigma^*(\ell'|w)\}$. Since $1 - \sigma^*(\ell'|w) \geq \sigma^*(\ell|w)$ it is enough to show that we can find d_1 such that $d_1 < \sigma^*(\ell|w)$. Similarly, we can argue that is enough to show that we can find d_2 such that $d_2 < \sigma^*(\ell'|w')$. Altogether we have the following conditions equivalent to the original ones:

$$\frac{\mu(w')}{\mu(w)} > \frac{d_1}{d_2} > \frac{\mu(w') \cdot w'}{\mu(w) \cdot w}, \qquad \sigma^*(\ell|w) > d_1, \quad \text{and} \quad \sigma^*(\ell'|w') > d_2.$$

Since $w > w'$ we have $w'/w < 1$. Therefore

$$\frac{\mu(w')}{\mu(w)} > \frac{\mu(w') \cdot w'}{\mu(w) \cdot w}.$$

Now by the density of rationals there is a rational between any two reals, and so there is a rational a/b such that

$$\frac{\mu(w')}{\mu(w)} > \frac{a}{b} > \frac{\mu(w') \cdot w'}{\mu(w) \cdot w}.$$

Since σ^* uses link ℓ for w we have $\sigma^*(\ell|w) > 0$, and similarly $\sigma^*(\ell'|w') > 0$. Choose an integer k such that $1/(b \cdot k) < \sigma^*(\ell|w)$ and $1/(a \cdot k) < \sigma^*(\ell'|w')$. Now setting $d_1 = 1/(b \cdot k)$ and $d_2 = 1/(a \cdot k)$ satisfies the second two inequalities. Also we have $d_1/d_2 = (a \cdot k)/(b \cdot k) = a/b$, and the first equation of the new specification of d_1 and d_2 is fulfilled. □

Observation 5.2. *Under the conditions of Theorem 5.1 we have*

$$\left(\text{load}(\ell|\sigma^*) - \text{load}(\ell|\sigma)\right) \cdot \left(\sigma^*(\ell|W) - \sigma(\ell|W)\right) < 0$$

and

$$\left(\text{load}(\ell'|\sigma^*) - \text{load}(\ell'|\sigma)\right) \cdot \left(\sigma^*(\ell'|W) - \sigma(\ell'|W)\right) < 0.$$

Proof. To show this result we show the following inequalities.

1. $\text{load}(\ell|\sigma^*) - \text{load}(\ell|\sigma) > 0$.
2. $\sigma^*(\ell|W) - \sigma(\ell|W) < 0$.
3. $\text{load}(\ell'|\sigma^*) - \text{load}(\ell'|\sigma) < 0$.
4. $\sigma^*(\ell'|W) - \sigma(\ell'|W) > 0$.

1. We have

$$\begin{aligned}\text{load}(\ell|\sigma^*) - \text{load}(\ell|\sigma) &= \sum_{w'' \in W} \mu(w'') \cdot \frac{w''}{c_\ell} \cdot \sigma^*(\ell|w'') - \sum_{w'' \in W} \mu(w'') \cdot \frac{w''}{c_\ell} \cdot \sigma(\ell|w'') \\ &= \sum_{w'' \in W} \mu(w'') \cdot \frac{w''}{c_\ell} \cdot (\sigma^*(\ell|w'') - \sigma(\ell|w'')) \\ &= \frac{1}{c_\ell} \cdot \left(\mu(w) \cdot w \cdot d_1 - \mu(w') \cdot w' \cdot d_2\right) > 0.\end{aligned}$$

The last inequality follows from the conditions on d_1, d_2.

2. Consider the difference in the marginal probabilities on the links. It is easy to see that for all $\ell'' \in L$ we have

$$\sigma^*(\ell''|W) - \sigma(\ell''|W) = \mu(w) \cdot (\sigma^*(\ell''|w) - \sigma(\ell''|w)) + \mu(w') \cdot (\sigma^*(\ell''|w') - \sigma(\ell''|w')).$$

By definition of σ we have and the choice of d_1 and d_2 we have

$$\begin{aligned}\sigma^*(\ell|W) - \sigma(\ell|W) &= \mu(w) \cdot \left(\sigma^*(\ell|(w) - (\sigma^*(\ell|w) - d_1)\right) + \mu(w') \cdot (\sigma^*(\ell|w') - (\sigma^*(\ell|w') + d_2)) \\ &= \mu(w) \cdot d_1 - \mu(w') \cdot d_2 < 0.\end{aligned}$$

3. Similar to (1)

$$\begin{aligned}\text{load}(\ell'|\sigma^*) - \text{load}(\ell'|\sigma) &= \sum_{w''\in W} \mu(w'') \cdot \frac{w''}{c_{\ell'}} \cdot (\sigma^*(\ell'|w'') - \sigma(\ell'|w'')) \\ &= \frac{1}{c_{\ell'}} \cdot \left(\mu(w') \cdot w' \cdot d_2 - \mu(w) \cdot w \cdot d_1\right) < 0.\end{aligned}$$

4. Similarly to (2) we have

$$\begin{aligned}\sigma^*(\ell'|W) - \sigma(\ell'|W) &= \mu(w) \cdot (\sigma^*(\ell'|w) - (\sigma^*(\ell'|w) + d_1)) + \mu(w') \cdot (\sigma^*(\ell'|w') - (\sigma^*(\ell'|w') - d_2)) \\ &= \mu(w') \cdot d_2 - \mu(w) \cdot d_1 > 0. \quad \square\end{aligned}$$

Theorem 3.2 implies that there is only one symmetric Bayesian Nash equilibrium for the task allocation game of Fig. 1, which by the construction of Fig. 4 is not an ESS; so there is no ESS in this task allocation game. More generally, Theorems 3.2 and 5.1 imply that there is no ESS in a task allocation game with more than one link and more than one task size where all links have the same speed.

The next observation gives a counterexample showing that the specialization condition from Theorem 5.1 is necessary but unfortunately not sufficient for an ESS.

Observation 5.3. *There exists a symmetric Bayesian task allocation game G with a strategy σ such that σ meets the specialization condition of Theorem 5.1 for any $w \neq w'$ and $\ell \neq \ell'$, but σ is not an ESS.*

Proof. Assume three resources 1, 2, 3 with speeds $c_1 = 6$, $c_2 = 4$, and $c_3 = 2$ and two task sizes $w = 21$ and $w' = 1$. We define $\mu(21) = 2/3$ and $\mu(1) = 1/3$. The strategy σ is defined as follows.

w	$\sigma(1\|w)$	$\sigma(2\|w)$	$\sigma(3\|w)$
1	0	1/3	2/3
21	19/21	2/21	0

σ defines a symmetric Bayesian Nash equilibrium fulfilling the necessary condition for an ESS established by Theorem 5.1. The next strategy σ' constitutes a successful mutation.

w	$\sigma'(1\|w)$	$\sigma'(2\|w)$	$\sigma'(3\|w)$
1	0	$1/3 - 0.008$	$2/3 + 0.008$
21	$19/21 - 0.001$	$2/21 + 0.001$	0 $\square$

Theorem 5.1 is the last result required to establish the uniqueness of an ESS for symmetric Bayesian task allocation games.

Theorem 5.4 (*Uniqueness*). *Let G be a symmetric Bayesian task allocation game with ESS σ^*.*

1. *Fix any two links $\ell \neq \ell'$ with the same speed, i.e. $c_\ell = c_{\ell'}$. Then for all task weights w we have $\sigma^*(\ell|w) = \sigma^*(\ell'|w)$ and $|\text{support}(\ell|\sigma^*)| \leq 1$.*
2. *The ESS σ^* is the unique ESS for G.*

Proof. Recall that, in a Bayesian Nash equilibrium, two links ℓ and ℓ' with the same speed have always the same load (Lemma 3.1(2)).

Part 1. We show for contradiction that if $c_\ell = c_{\ell'}$ and there is a weight $w \in W$ such that $\sigma^*(\ell|w) \neq \sigma^*(\ell'|w)$, then there is a mutation σ that defeats σ^*. W.l.o.g., suppose that $\sigma^*(\ell|w) > 0$. We know that σ^* is a best reply to $(\sigma^*)^{(n-1)}$. Since $c_\ell = c_{\ell'}$ and the loads on links ℓ and ℓ' are equal, it follows that for all tasks w, either both ℓ and ℓ' are optimal for w, or both are not optimal for w. We have to consider several cases.

1. σ^* uses ℓ and ℓ' for task w but not for any other tasks (i.e., $\text{support}(\ell|\sigma^*) = \text{support}(\ell'|\sigma^*) = \{w\}$). Then since $\text{load}(\ell|\sigma^*) = \text{load}(\ell'|\sigma^*)$ we immediately get $\sigma^*(\ell|w) = \sigma^*(\ell'|w)$, which is a contradiction.
2. There is a task $w' \neq w$ such that σ^* uses link ℓ' for task w'. Then we have two distinct tasks w, w' such that σ^* uses link ℓ for w and link ℓ' for w'. Since for all tasks w, either both ℓ and ℓ' are optimal or both are not optimal, and due to the fact that σ^* is a best reply to $(\sigma^*)^{(n-1)}$, we have ℓ and hence ℓ' is optimal for w. Similarly, ℓ' and ℓ are optimal for w' given $(\sigma^*)^{(n-1)}$. So Theorem 5.1 applies, and there is a mutation σ that defeats σ^*.
3. $\sigma^*(\ell'|w') = 0$ for all $w' \neq w$. Then σ^* must use link ℓ' for task w, since $\text{load}(\ell'|\sigma^*) = \text{load}(\ell|\sigma^*) > 0$. Since Case 1 does not hold, there is a task $w' \neq w$ such that σ^* uses link ℓ for w'. So again we have link ℓ is optimal for w and link ℓ' is optimal for w' given $(\sigma^*)^{(n-1)}$, Theorem 5.1 applies, and there is a mutation σ that defeats σ^*.

This establishes that $\sigma^*(\ell|w) = \sigma^*(\ell'|w)$ for all task sizes w. Thus $\text{support}(\ell|\sigma^*) = \text{support}(\ell'|\sigma^*)$. Now if σ^* uses ℓ and ℓ' for more than one task size, we have two different task sizes both optimal for ℓ and ℓ', Theorem 5.1 applies, and there is a mutation σ that defeats σ^*. Therefore $|\text{support}(\ell|\sigma^*)| \leq 1$ as required.

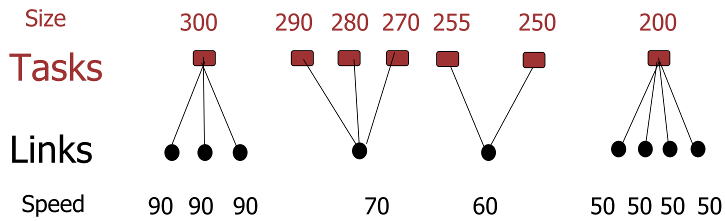

Fig. 5. A typical allocation of tasks to links in a *clustered symmetric Bayesian Nash equilibrium*, which is guaranteed to be an ESS by Theorem 5.5. Links in the same link group (e.g., with speed 90 in the figure) may be allocated at most one task size (300 in the figure). Links with different speeds may be allocated more than one task size, but then tasks with that size are not allocated anywhere else (e.g., in the figure, the link with speed 60 exclusively serves tasks of size 250 and 255).

Part 2. The first part of this theorem implies that two links ℓ, ℓ' with the same speed can share at most one task, that is, $|\text{support}(\ell|\sigma) \cap \text{support}(\ell'|\sigma)| \le 1$. (For two links with different speeds the bound on the shared support follows already from Lemma 3.1(6).) By Theorem 3.2, all symmetric equilibria σ agree on the distribution $\sigma_{\mathcal{L}}$ on a link group $\mathcal{L}$. Now there are two cases.

Case 1*:* $\mathcal{L} = \{\ell\}$. The distribution $\sigma^*(\ell|\cdot)$ is uniquely determined since the distribution $\sigma^*(\mathcal{L}|\cdot)$ is.

Case 2*:* $|\mathcal{L}| > 1$. Then by Theorem 5.4(1), there is at most one task size w such that $\sigma^*(\mathcal{L}|w) > 0$. If there is exactly one task size w with $\sigma^*(\mathcal{L}|w) > 0$ then for all links $\ell \in \mathcal{L}$ we have $\sigma^*(\ell|w) = 1/|\mathcal{L}|$ for any ESS σ^*. Otherwise (no task weight w with $\sigma^*(\mathcal{L}|w) > 0$) for all links $\ell \in \mathcal{L}$ we have $\sigma^*(\ell|w) = 0$ for any ESS σ^*. In either case again, the distribution for each link $\ell \in \mathcal{L}$ is uniquely determined. □

Now we give a structural condition that is sufficient for an ESS. It can be used to construct an ESS in a wide variety of models where the ESS exists. Theorem 5.1 shows that an ESS requires links to "specialize" in tasks where distinct links do not share two distinct tasks. A stronger condition is to require that if a link is optimal for two distinct tasks, then no other link is optimal for either of the tasks. We call such a distribution clustered.

Definition 5. In an n-player symmetric Bayesian task allocation game G, a symmetric strategy profile $\sigma^{(n)}$ is **clustered** if for any two distinct tasks w, w' and any link ℓ, if ℓ is optimal for both w and w' given σ^{n-1}, then no other link is optimal for w or w' given σ^{n-1}.

Fig. 5 illustrates clustering. The next theorem establishes that every clustered symmetric Bayesian Nash equilibrium is an ESS. Intuitively, the reason for this is as follows. Assuming the clustering condition, there are three cases for a given Bayesian Nash equilibrium σ^* and a possible mutation σ that is a best reply to σ^*. First, a link ℓ may be optimal for no task size given σ^*. Then neither σ^* nor σ use ℓ, so this link makes no difference to the relative performance of σ and σ^*. Second, a link ℓ may be optimal for more than one task size given σ^*, say task sizes w, w'. Then clustering requires that ℓ is the *only* optimal link for w and w', so both σ^* and σ use ℓ with probability 1 for w and w', and there is again no difference between them with respect to ℓ. Third, a link ℓ may be optimal for exact one task w_ℓ. Then all links optimal for w are optimal for w only. In that case the cluster of these links forms a Bayesian Nash equilibrium on its own with the single task w; it is easily shown that with just one task size to consider, any equilibrium is an ESS since there is no room to shift links from one task size to another.

Theorem 5.5 (*Clustering*). *Every clustered Bayesian Nash equilibrium is an ESS, but not vice versa. More precisely:*

1. *Let G be a symmetric Bayesian task allocation game. If $(\sigma^*)^{(n)}$ is a clustered Bayesian Nash equilibrium in G, then σ^* is an ESS in G.*
2. *There is a symmetric Bayesian task allocation game G that has a non-clustered ESS and no clustered ESS.*

Proof. We first prove (1), the sufficiency of the clustering condition.

Part 1. To establish that a clustered Bayesian Nash equilibrium is an ESS, let σ be any best reply to $(\sigma^*)^{(n-1)}$. We show that

$$\sum_{\ell \in L} [\text{load}(\ell|\sigma^*) - \text{load}(\ell|\sigma)] \cdot [\sigma^*(\ell|W) - \sigma(\ell|W)] > 0$$

thus by Lemma 4.3 σ^* is an ESS. Our argument decomposes the sum above into two groups; the first group contains links l that are optimal *for exactly one task* given $(\sigma^*)^{(n-1)}$. The second group contains all links that are optimal for more than one task size and links that are optimal for none task size. The set of links in the first group are denoted by $S \subseteq L$. For each link $\ell \in S$, we write w_ℓ for the unique task such that ℓ is optimal for w_ℓ given $(\sigma^*)^{(n-1)}$. The set of links in the second group are denoted by $\bar{S} \subseteq L$.
Group 1. If $\ell \in S$, then clustering implies that all links optimal for task w_ℓ are optimal only for w_ℓ. The application of Proposition 2.2 and the fact that both σ and σ^* are best replies to $(\sigma^*)^{(n-1)}$ shows that

$$\text{support}(\ell|\sigma^*) \subseteq \{w_\ell\}$$

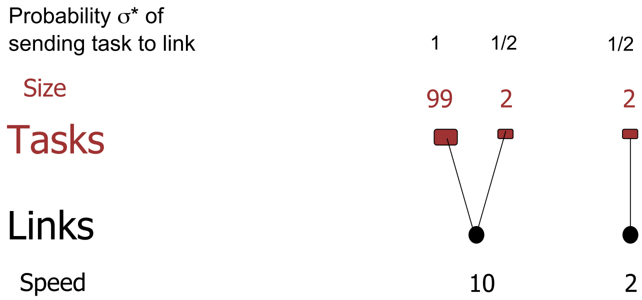

Fig. 6. An example of an ESS that is not clustered, so clustering is sufficient but not necessary for an ESS.

and

$$\text{support}(\ell|\sigma) \subseteq \{w_\ell\}.$$

Therefore the differences in load and frequency of use of link ℓ depend only on the task weight w_ℓ, which implies that

$$\begin{aligned}\left[\text{load}(\ell|\sigma^*) - \text{load}(\ell|\sigma)\right] \cdot \left[\sigma^*(\ell|W) - \sigma(\ell|W)\right] &= \left[\frac{w_\ell}{c_\ell} \cdot \sigma^*(\ell|w_\ell) - \sigma(\ell|w_\ell)\right] \cdot \left[\sigma^*(\ell|w_\ell) - \sigma(\ell|w_\ell)\right] \\ &= \frac{w_\ell}{c_\ell} \cdot \left[\sigma^*(\ell|w_\ell) - \sigma(\ell|w_\ell)\right]^2 \geq 0. \qquad (2)\end{aligned}$$

Group 2. Here we argue that, if a link $\ell \in \bar{S}$, then σ^* and σ agree on the distribution of tasks on ℓ. If a link $\ell \in \bar{S}$ is optimal for no task size given $(\sigma^*)^{(n-1)}$, then by Proposition 2.2, neither σ^* nor σ use ℓ for any task size w, that is, $\sigma^*(\ell|w) = \sigma(\ell|w) = 0$ for all task sizes w.

If a link $\ell \in \bar{S}$ is optimal for more than one task size, then clustering requires that ℓ is uniquely optimal for these task sizes. That is, for all $w \in \text{opt}(\ell|(\sigma^*)^{n-1})$, no other link $\ell' \neq \ell$ is optimal for w given $(\sigma^*)^{n-1}$). With Proposition 2.2, this implies that for all $w \in \text{opt}(\ell|(\sigma^*)^{n-1})$ we have $\sigma^*(\ell|w) = \sigma(\ell|w) = 1$. For all $w \notin \text{opt}(\ell|(\sigma^*)^{n-1})$, we have $\sigma^*(\ell|w) = \sigma(\ell|w) = 0$.

The above two cases showed that we can assume

$$\sigma^*(\ell|w) = \sigma(\ell|w)$$

and for all links $\ell \in \bar{S}$ and task sizes w we have

$$[\text{load}(\ell|\sigma^*) - \text{load}(\ell|\sigma)] \cdot [\sigma^*(\ell|W) - \sigma(\ell|W)]. \qquad (3)$$

Since $\sigma \neq \sigma^*$, there is a link ℓ_δ and a task w_{ℓ_δ} such that

$$\sigma^*(\ell_\delta|w_{\ell_\delta}) \neq \sigma(\ell_\delta|w_{\ell_\delta}).$$

As the two strategies agree on links optimal for 0 or more than task size (Eq. (3)), we have $\ell \in S$. So Eq. (2) and the fact that both w_{ℓ_δ} and c_{ℓ_δ} are strictly positive imply that

$$\left[\text{load}(\ell_\delta|\sigma^*) - \text{load}(\ell_\delta|\sigma)\right] \cdot \left[\sigma^*(\ell_\delta|W) - \sigma(\ell_\delta|W)\right] = \frac{w_{\ell_\delta}}{c_{\ell_\delta}} \cdot \left[\sigma^*(\ell_\delta|w_{\ell_\delta}) - \sigma(\ell_\delta|w_{\ell_\delta})\right]^2 > 0. \qquad (4)$$

Combining Eqs. (2)–(4) yields the desired result:

$$\begin{aligned}&\sum_{\ell\in L}[\text{load}(\ell|\sigma^*) - \text{load}(\ell|\sigma)] \cdot [\sigma^*(\ell|W) - \sigma(\ell|W)] = \sum_{\ell\in\bar{S}}[\text{load}(\ell|\sigma^*) - \text{load}(\ell|\sigma)] \cdot [\sigma^*(\ell|W) - \sigma(\ell|W)] \\ &\quad + \sum_{\ell\in S-\{\ell_\delta\}} [\text{load}(\ell|\sigma^*) - \text{load}(\ell|\sigma)] \cdot [\sigma^*(\ell|W) - \sigma(\ell|W)] + [\text{load}(\ell_\delta|\sigma^*) - \text{load}(\ell_\delta|\sigma)] \\ &\quad \cdot [\sigma^*(\ell_\delta|W) - \sigma(\ell_\delta|W)] \\ &= \sum_{\ell\in S-\{\ell_\delta\}} [\text{load}(\ell|\sigma^*) - \text{load}(\ell|\sigma)] \cdot [\sigma^*(\ell|W) - \sigma(\ell|W)] + \frac{w_{\ell_\delta}}{c_{\ell_\delta}} \cdot \left[\sigma^*(\ell_\delta|w_{\ell_\delta}) - \sigma(\ell_\delta|w_{\ell_\delta})\right]^2 > 0.\end{aligned}$$

Part 2. Now we show that there exists a task allocation game G that has a non-clustered ESS. The game G has two links 1 and 2 with $c_1 = 10$ and $c_2 = 2$. We have two task sizes 2 and 99, distributed as $\mu(99) = 1/4$ and $\mu(2) = 3/4$. Then the state $\sigma^*(1|99) = 1, \sigma^*(2|99) = 0, \sigma^*(1|2) = 1/2$, and $\sigma^*(2|2) = 1/2$ is an ESS. (Mainly because the only optimal link for task size 99 is link 1.) It is not clustered since link 1 is used for both tasks, and both links are used for task size 2. It is also easy to see that there is no clustered Bayesian Nash equilibrium, and, hence, no clustered ESS in this task allocation game. Fig. 6 illustrates this counterexample. □

6. Conclusions and future work

We proposed the notion of an evolutionarily stable strategy as a refinement of Bayesian Nash equilibrium that (1) provides a criterion for separating stable from unstable equilibria, and (2) corresponds to a steady state of evolutionary

Table 2
A Hawk–Dove game.

	Hawk (H)	Dove (D)
Hawk	$-2, -2$	6, 0
Dove	0, 6	3, 3

process in which more successful strategies spread in a user community over time. We investigated the symmetric Bayesian task allocation game, a selfish task allocation game which is a special case of general Bayesian task allocation games. Our results provide a necessary and sufficient condition under which a deviation (mutation) from a Bayesian Nash equilibrium is evolutionarily successful. If the structure of task–link allocations is consistent with evolutionary stability, then links have to "specialize" in tasks in that they may serve at most one task in common with another link. Finally, we established the uniqueness of evolutionarily stable strategies.

An important open question is to determine the computational complexity of finding evolutionary equilibria in a Bayesian task allocation game. Both our necessary condition (specialization) and our sufficient condition (clustering) for an ESS can be checked in polynomial time for a given candidate ESS. Generating a candidate ESS involves finding a symmetric Bayesian Nash equilibrium. Research into computing Bayesian Nash equilibria shows that symmetric equilibria are significantly easier to find than Bayesian Nash equilibria in general, and highly optimized local iterative search techniques are available [35,5]. So we expect that in many symmetric Bayesian task allocation games evolutionary analysis can proceed efficiently by first finding a symmetric Bayesian Nash equilibrium, and then checking the clustering criterion to establish that the equilibrium is an ESS, or checking the specialization criterion to establish that there is no ESS. The difficult equilibria to analyze are the ones that are specialized but not clustered, because in this case evolutionary stability depends on the exact magnitudes of the parameters of the model, as Observation 5.3 shows.

Another important avenue for future research is to investigate evolutionary dynamics in the symmetric Bayesian task allocation game, for example the replicator dynamics, and to obtain bounds on the speed of convergence to an ESS (cf. [14]).

Acknowledgements

We thank Funda Ergun for helpful discussions and Tom Friedetzky for the example in Observation 5.3. We also thank the anonymous reviewer who simplified the proof of Theorem 3.2. A preliminary version of this paper was presented at the Center for Statistics and Social Sciences, University of Washington; we thank the audience and especially Anna Karlin for valuable comments. Anonymous referees for the ESA 2008 symposium provided helpful suggestions on a previous related paper. This research was supported by discovery grants to the authors from the National Sciences and Engineering Research Council of Canada.

Appendix. Evolutionary stability in the Hawk–Dove game

This section introduces the concept of evolutionary stability in a simple 2×2 matrix game that is often used to illustrate ESS in biological game theory. A 2-player Hawk–Dove game can be represented with the payoff matrix shown in Table 2. The two players have the same strategy set $S = \{H, D\}$. By convention, the row player is player 1, column player 2 and the row player's payoff is given first. Thus

$$u_1(H, H) = -2 = u_2(H, H).$$

One interpretation of this game is that it represents a struggle between two animals over a food source of value 6. Each animal may engage in hawkish or in dovish behavior. A hawk gains all the food against a dove. Two doves share the food equally. Two hawks fight, with a risk of injury to each, leading to an expected payoff of -2.

Let $H_{3/5}$ denote the mixed strategy in which hawk is chosen with probability 3/5. It is easy to see that the only mixed strategy equilibrium for the Hawk–Dove game is the symmetric equilibrium $(H_{3/5}, H_{3/5})$ where each player chooses $H_{3/5}$. A possible population whose distribution corresponds to $H_{3/5}$ is the multiset

$$P^* = [H, H, H, H, H, H, D, D, D, D].$$

The mixed strategy $H_{3/5}$ represents an equilibrium state of this population.

Suppose we have a group of mutants $M = [H, H, D, D]$. Then

$$P \cup M = [H, H, H, H, H, H, D, D, D, D, H, H, D, D].$$

By inspection, the corresponding population distribution assigns frequency $(6 + 2)/(10 + 4) = 4/7$ to Hawk. The relative size of the mutant group is $\varepsilon = 4/14$. Thus the mixture $(1 - \varepsilon)\pi + \varepsilon\pi_{pM}$ assigns $10/14 \cdot 6/10 + 4/14 \cdot 2/4$ to Hawk, which is equal to 4/7. Thus, the payoff from choosing H in the mixed population is

$$u(H, 4/7) = -2 \cdot 4/7 + 6 \cdot 3/7 = 10/7,$$

whereas the payoff from choosing D is

$$u(D, 4/7) = 3 \cdot 3/7 = 9/7.$$

Since $u(H, 4/7) > u(D, 4/7)$, the more probability a mixed strategy assigns to H, the higher payoff it achieves against the distribution 4/7. Now for the original distribution $P(H) = 3/5$, whereas for the mutants $\pi_M(H) = 1/2$, so $\pi_P(H) > \pi_M(H)$ and it follows that

$$u(\pi_P, 4/7) > u(\pi_M, 4/7).$$

So a mutation whose strategy distribution is $\pi_M(H) = 1/2$ fails if its relative size is below $\varepsilon = 2/7$.

References

[1] P. Battigalli, M. Gilli, M.C. Molinari, Learning and Convergence to Equilibrium in Repeated Strategic Interactions: An Introductory Survey, in: Ricerche Economiche, vol. 46, 1992, pp. 335–377.
[2] R. Beier, A. Czumaj, P. Krysta, B. Vöcking, Computing equilibria for congestion games with (im)perfect information, in: Proc. 13th Annual Symposium on Discrete Algorithms, SODA, 2004, pp. 746–755.
[3] P. Berenbrink, L.A. Goldberg, P. Goldberg, R. Martin, Utilitarian resource assignment, Journal of Discrete Algorithms 4 (4) (2006) 567–587.
[4] P. Berenbrink, O. Schulte, Evolutionary equilibrium in Bayesian routing games: Specialization and Niche formation, in: Proc. of 15th European Symposium on Algorithms, ESA, 2007, pp. 29–40.
[5] N.A.R. Bhat, K. Leyton-Brown, Computing Nash equilibria of action-graph games, in: Proc. of 20th Conference on Uncertainty in Artificial Intelligence, UAI, 2004, pp. 35–42.
[6] M. Broom, C. Cannings, G.T. Vickers, Multi-player matrix games, Bulletin of Mathematical Biology 59 (5) (1997) 931–952.
[7] George Christodoulou, E. Koutsoupias, The price of anarchy of finite congestion games, in: Proc. of the 37th Annual Symposium on Theory of Computing, STOC, 2005, pp. 67–73.
[8] A. Czumaj, B. Vöcking, Tight bounds for worst-case equilibria, in: Proc. of 13th Annual Symposium on Discrete Algorithms, SODA, 2002, pp. 413–420.
[9] A. Czumaj, P. Krysta, B. Vöcking, Selfish traffic allocation for server farms, in: Proc. of 34th Annual Symposium on Theory of Computing, STOC, 2002, pp. 287–296.
[10] J.C. Ely, W.H. Sandholm, Evolution in Bayesian games I: Theory, Games and Economic Behavior 53 (1) (2005) 83–109.
[11] K. Etessami, A. Lochbihler, The computational complexity of evolutionarily stable strategies, Electronic Colloquium on Computational Complexity 55 (2004).
[12] A. Fabrikant, A. Luthra, E.N. Maneva, C.H. Papadimitriou, S. Shenker, On a network creation game, in: Proc. of 22nd Symposium on Principles of Distributed Computing, PODC, 2003, pp. 347–351.
[13] R. Feldmann, M. Gairing, T. Lücking, B. Monien, M. Rode, Nashification and the coordination ratio for a selfish routing game, in: Proc.of 30th International Colloquium on Automata, Languages and Programming, ICALP, 2003, pp. 514–526.
[14] S. Fischer, B. Vöcking, On the evolution of selfish routing, in: Proc. of 12th European Symposium on Algorithms, ESA, 2004, pp. 323–334.
[15] D. Fotakis, S. Kontogiannis, E. Koutsoupias, M. Mavronicolas, P. Spirakis, The structure and complexity of Nash equilibria for a selfish routing game, in: Proc. of 29th International Colloquium on Automata, Languages, and Programming, ICALP, 2002, pp. 123–134.
[16] M. Gairing, B. Monien, K. Tiemann, Selfish routing with incomplete information, Theory of Computing Systems 42 (1) (2008) 91–130.
[17] M. Gairing, T. Lücking, M. Mavronicolas, B. Monien, Computing Nash equilibria for scheduling on restricted parallel links, in: Proc. of the 36th Annual Symposium on Theory of Computing, STOC, 2004, pp. 613–622.
[18] C. Georgiou, T. Pavlides, A. Philippou, Selfish routing in the presence of network uncertainty, Parallel Processing Letters 19 (1) (2009) 141–157.
[19] S. Govindan, P.J. Reny, A.J. Robson, A short proof of Harsanyi's purification theorem, Games and Economic Behavior 45 (2) (2003) 369–374.
[20] P. Hammerstein, R. Selten, Game Theory and Evolutionary Biology, in: R.J. Aumann, S. Hart (Eds.), Handbook of Game Theory, vol. 2, Elsevier, 1994, pp. 929–993 (Chapter 2).
[21] J.C. Harsanyi, Games with incomplete information played by 'Bayesian players', Parts I, II, and III, Management Science 14 (1967) 159–182, pp. 320–334, and pp. 486–502.
[22] J.C. Harsanyi, Games with randomly disturbed payoffs: A new rationale for mixed-strategy equilibrium points, International Journal of Game Theory 2 (1973) 1–23.
[23] J. Hofbauer, W.H. Sandholm, Evolution in games with randomly disturbed payoffs, Journal of Economic Theory 132 (2007) 47–69.
[24] M.S. Kearns, S. Suri, Networks preserving evolutionary equilibria and the power of randomization, in: Proc. of the 7th Conference on Electronic Commerce 2006, 2006, pp. 200–207.
[25] S. Kontogiannis, P. Spirakis, The contribution of game theory to complex systems, in: Proc. of 10th Panhellenic Conference on Informatics, PCI, 2005, pp. 101–112.
[26] E. Koutsoupias, C.H. Papadimitriou, Worst-case equilibria, in: Proc. of the 16th Annual Symposium on Theoretical Aspects of Computer Science, STACS, 1999, pp. 404–413.
[27] T. Lücking, M. Mavronicolas, B. Monien, M. Rode, A new model for selfish routing, in: Proc. of the 21st Annual Symposium on Theoretical Aspects of Computer Science, STACS, 2004, pp. 547–558.
[28] M. Mavronicolas, I. Milchtaich, B. Monien, K. Tiemann, Congestion games with player-specific constants, in: Proc. of Mathematical foundations of Computer Science, MFCS, 2007, pp. 633–644.
[29] M. Mavronicolas, P. Spirakis, The price of selfish routing, in: Proc. of the 33rd Annual Symposium on Theory of Computing, STOC, 2001, pp. 510–519.
[30] J. Maynard Smith, Evolution and the Theory of Games, Cambridge University Press, 1982.
[31] I. Milchtaich, Congestion games with player-specific payoff functions, Games and Economic Behavior 13 (1) (1996) 111–124.
[32] J.F. Nash, Equilibrium points in N-person games, Proceedings of the National Academy of Sciences of the United States of America 36 (1950) 48–49.
[33] M. Nowak, Evolutionary Dynamics: Exploring the Equations of Life, Belknap Press, 2006.
[34] M.J. Osborne, A. Rubinstein, A Course in Game Theory, MIT Press, 1994.
[35] C.H. Papadimitriou, T. Roughgarden, Computing equilibria in multi-player games, in: Proc. of the Symposium on Discrete algorithms, SODA, 2005, pp. 82–91.
[36] R.W. Rosenthal, A class of games possessing pure-strategy Nash equilibria, International Journal of Game Theory 2 (1973) 65–67.
[37] T. Roughgarden, É. Tardos, How bad is selfish routing? Journal of the ACM 49 (2) (2002) 236–259.
[38] W.H. Sandholm, Evolutionary implementation and congestion pricing, Review of Economic Studies 69 (2002) 667–689.
[39] W.H. Sandholm, Evolution in Bayesian games II: Stability of purified equilibrium, Journal of Economic Theory 136 (1) (2007) 641–667.
[40] E. van Damme, Stability and Perfection of Nash Equilibria, 2nd edition, Springer-Verlag, Berlin, 1991.
[41] J. Weibull, Evolutionary Game Theory, The MIT Press, Cambridge, MA, 1995.