Multi-Karma Economies

Student Paper

Author(s):
Oberlin, Patrick

Publication date:
2024-03

Permanent link:
https://doi.org/10.3929/ethz-b-000670203

ETH zürich

Automatic Control Laboratory

Semester Thesis

Multi-karma Economies

Patrick Oberlin

March 15, 2024

Advisors

Ezzat Elokda

Prof. Florian Dörfler

Abstract

Limited public resources include fast lanes on highways, parking spots at high-interest locations near a city center, and access to EV charging stations during peak hours. Elokda et al. proposed karma economies as a non-monetary approach to an effective allocation of such resources and demonstrated an improvement of discomfort cost for all participants with respect to classic solutions. However, so far only single-resource economies have been investigated. This thesis extends the model by allowing agents to use their karma points for other resources than they were obtained from. For this new system, two design instruments are considered, namely non-unit exchange rates between different resources, and karma redistribution schemes. In a numerical analysis, we show that an improvement is achieved by coupling economies. We go on to show that the improvement is largely robust to the specifics of the design. In particular, non-unit exchange rates have negligible impact, whereas the redistribution scheme has shown to be more suitable to affect agents' policies and to implement a policy maker's notion of fairness.

Contents

Chapter 1

Introduction

This thesis investigates the coupled allocation of shared resources among competing agents. Examples of such resources include parking spaces near points of high interest, fast lane access along a congested highway and power consumption or internet bandwidth usage of factories and households. The unregulated usage of these resources leaves no possibility for agents to distinguish their value of time, a private urgency value each agent associates with their necessity to access a specific resource. As a solution to the single-resource allocation problem, karma economies were introduced by Elokda et al. [2], [3].
However, the proposed model only allocates a single resource to the agents. Since subsets of these resources are thematically related, it is natural to ask if the allocation can be pooled. As an example, consider the combination of EV charging and congestion management, where EV charging during rush hours can be incentivized to reduce the peak traffic load. Therefore, the model is extended in this thesis to combine the allocation of multiple resources so that agents can use karma points obtained by yielding one resource to gain access to another. This allows agents who have a higher need for one resource to adjust their bidding policy accordingly. It is then investigated how the combination affects the discomfort costs of agents with heterogeneous needs, and what the implications of different design instruments are.

1.1 Related Literature

Since the number of agents in a typical use case is large, the game can be modeled as a dynamic population game (DPG [1]). In DPGs, a Stationary Nash Equilibrium (SNE) is guaranteed to exist, and an evolutionary dynamics-based algorithm to compute the SNE has been developed [1]. In both [2] and [3], it was shown that karma economies enable an efficient and fair resource allocation in the infinitely repeated game. We will adopt the tools developed in these papers, but extend the model to integrate multiple resources.

1.2 Outline and Contributions

In Chapter 2, the mathematical model for coupled karma economies is presented. Firstly, the state of the agents is extended by the resource they are currently competing for. The state transition function is adapted such that players keep their karma balance when advancing to the next resource. Secondly, an exchange rate is introduced to give the karma points game-dependent values.
In Chapter 3, the results of numerical simulations are discussed. The simulations explore a possible application of the multi-karma economy and highlights how the coupling improves the

reward for all agent types considered. Further, the effects of the exchange rate and redistribution schemes are analyzed.

Chapter 2

Model

In this chapter, the mathematical model of karma economies is presented. It starts with the definition of the terminology, and in the following sections, the different aspects of the model are highlighted.

2.1 Karma Economy

In a karma economy, shared resources are allocated to competing agents. In the example of a congested highway, a policy maker designates part of the highway (e.g., one lane) as the *fast lane*. The access to this lane is restricted to a number of travellers so that it will not congest. The remaining lanes are not regulated and can be used by any agent who wishes. All agents have a karma balance as a fictional currency to bid for the regulated resources. The highest bidding agents will be granted access to the resource and they pay their bid. The accumulated payments of the winners will be redistributed according to some method, increasing yielding players' chances to win future competitions. In a multi karma economy, each resource competition constitutes a *game*, denoted by the letter e, and there are n_e games in an economy that are played each round.

2.2 Type and State

Each agent belongs to a type $\tau \in \mathcal{T} = \{\tau_1, \tau_2, ..., \tau_{n_\tau}\}$. g_{τ_i} denotes the fraction of agents of type i, thus

$$\begin{aligned} 0 \leq g_\tau \leq & 1, \quad \forall \tau, \\ \sum_{i=1}^{n_\tau} g_{\tau_i} = & 1. \end{aligned}$$

The state of an agent is described by the 3-tuple (e, u, k), where e denotes the game the agent is playing, u the urgency to access the resource, and k indicates the current karma balance that is available to an agent to place a bid for the resource. Formally,

$$x = (e, u, k) \in \mathcal{X} = \mathcal{E} \times \mathcal{U} \times \mathbb{N},$$

where $\mathcal{X}$ is the set of possible states, $\mathcal{E} = \{e_1, e_2, ..., e_{n_e}\}$ the set of games in the economy and $\mathcal{U} = \{u_1, u_2, ..., u_{n_u}\}$ the set of urgency states. The urgency state represents an agent's private value for the resource. The urgency in the next game is given by an exogenous Markov process

$$\phi_\tau[e^+, u^+ | e, u] : \mathcal{E} \times \mathcal{U} \to \mathcal{E} \times \mathcal{U}, \tag{2.1}$$

where (e^+, u^+) is the partial state at the next time step. In this thesis, we impose the additional constraint

$$\sum_{u^+ \in \mathcal{U}} \phi_\tau[e_{i+1}, u^+ | e_i, u] = 1, \quad \forall i : 1 \leq i \leq n_e, \quad \forall u, \quad e_{n_e+1} = e_1, \tag{2.2}$$

that is, the games are played consecutively in a given order. Additionally, all players play the same game at any time step. To motivate this modelling choice, consider, for example, the application where commuters travel to the city on a highway and are looking for a good parking spot in the city. Then, all agents compete for the highway first, and only compete for the parking spots afterwards.

2.2.1 No-Play Urgency

In [3], only single-karma economies were considered ($n_e = 1$), and all players competed in every round. Since, in this thesis, several games are played consecutively, not necessarily every type of agent will play every game in each round. To model this fact, we designate $0 \in \mathcal{U}$ to represent the no-play urgency state. Having urgency $u = 0$ means that no cost is incurred for not gaining access to the contested resource.

2.3 Social State

The social state (d, π) is defined in a similar fashion as in [2], but with the agents' states extended by the game e. In particular, the joint type-state distribution d is defined as follows.

$$d \in \mathcal{D} := \left\{ d \in \mathbb{R}_+^{n_\tau \times n_e \times |\mathcal{U}| \times \infty} : \forall \tau \in \mathcal{T}, \sum_{e,u,k} d_\tau[e, u, k] = g_\tau \right\}, \tag{2.3}$$

where $d_\tau[e, u, k]$ is the fraction of agents that belong to type τ and are in the state (e, u, k). Since we are interested in the SNE, d can be assumed to be constant in time, given e. The condition on e is necessary because the type-state distribution can differ between games depending on the markov chain (2.1) of the agents.

The action space is made up of the possible bids of an agent, that is, the integers up to the karma balance of the agent,

$$b \in \mathcal{B}^k := \{b \in \mathbb{N} : b \leq k\}. \tag{2.4}$$

For an agent of type τ, the policy, which maps the agent's state (e, u, k) to a probability distribution over the bids b, is defined by

$$\pi_\tau : \mathcal{X} \to \left\{ \sigma \in \mathbb{R}_+^{k+1} : \sum_b \sigma[b] = 1 \right\}. \tag{2.5}$$

Correspondingly, $\pi_\tau[b|e, u, k]$ denotes the probability that an agent of type τ in state (e, u, k) bids b. Furthermore, let's define by $\pi := (\pi_{\tau_1}, \pi_{\tau_2}, ..., \pi_{\tau_{n_\tau}}) \in \Pi$ the *policy* of all agents.

2.4 Immediate Reward

Let $o \in \mathcal{O} = \{0, 1\}$ be the competition outcome of an agent, where $o = 0$ means that the agent was granted access to the regulated resource, and $o = 1$ means that the agent has to use the public resource. Further, let s_e^0 be the capacity of the regulated part of the resource in game e, and s_e^1 the capacity of the public part, both expressed as a fraction of the total number of agents.

Moreover, let $\nu_{\text{loss}}[e](d, \pi)$ be the mass of agents that lost the bid but still uses the resource, i.e., not counting agents that don't bid at all (i.e., $b = 0$) due to no urgency,

$$\nu_{\text{loss}}[e](d, \pi) = \sum_{u>0,\tau,k} \Big(d_\tau[e, u, k]\Big) - s_e^0. \tag{2.6}$$

It is important to note that this mass depends on e, since the mass of players actively requesting different resources is not necessarily the same. Moreover, $\nu_{\text{loss}}[e](d, \pi)$ is constant at the SNE under the DPG model [1].
The cost an agent incurs depends on the outcome and on the mass of players that share that outcome. Namely,

$$c[o, e](d, \pi) = \begin{cases} 0, & o = 0, \\ \max\{0, \frac{\nu_{\text{loss}}[e](d,\pi)-s_e^1}{s_e^1}\}, & o = 1. \end{cases} \tag{2.7}$$

This cost model corresponds to a simple queuing model: the higher the demand-to-capacity ratio of a resource is, the longer an agent has to wait for completed access. Accessing the regulated resource is associated with no cost since by the design of karma economies, there is no congestion. Furthermore, the resource assignment function $\psi[o|e, b](d, \pi)$ is defined as follows. In a manner similar to [3], ψ is dependent on a threshold bid $b^*[e](d, \pi)$ with the conditions,

- if $b > b^*[e](d, \pi)$, then any agent bidding b may access the resource, i.e., $o = 0$,
- if $b < b^*[e](d, \pi)$, then any agent bidding b will not gain access to the resource, i.e., $o = 1$,
- if $b = b^*[e](d, \pi)$, then an agent gains access with a probability that distributes the remaining capacity of the resource among all agents bidding b^* uniformly at random.

The mass of agents playing action b (bidding b) in game e can be defined by

$$\nu[e, b](d, \pi) = \sum_{\tau\in\mathcal{T}} \sum_{u\in\mathcal{U}} \sum_{k\in\mathbb{N}} d_\tau[e, u, k]\pi_\tau[b|e, u, k]. \tag{2.8}$$

Then, the threshold bid in each game $b^*[e](d, \pi)$ is given by

$$b^*[e](d, \pi) = max\Big\{b \in \mathbb{N}| \sum_{b'\geq b} \nu[e, b'](d, \pi) \geq s_e^0\Big\}. \tag{2.9}$$

Hence, the probability that an agent bidding b can access the resource is

$$\psi[o = 0|e, b](d, \pi) = \begin{cases} 1, & b > b^*[e](d, \pi) \\ 0, & b < b^*[e](d, \pi) \\ \frac{s^0-\sum_{b'>b}\nu[e,b'](d,\pi)}{\nu[e,b](d,\pi)}, & b = b^*[e](d, \pi) \end{cases} \tag{2.10}$$

and the probability that an agent cannot access it is $\psi[o = 1|e, b](d, \pi) = 1 - \psi[o = 0|e, b](d, \pi)$.

The immediate reward function is denoted by $\zeta_\tau[e, u, b](d, \pi)$ and gives the expected reward that an agent of type τ, playing game e with urgency u is granted access to the regulated resource in the current time step when bidding b. Note that the immediate reward is a function of the social state (d, π) since an agent competes for access against the whole population, which in turn is represented by this distribution and policy. The reward is defined as the negative cost that is incurred if no access is gained, scaled with the agent's urgency value:

$$\zeta_\tau[e, u, b](d, \pi) = -u \sum_{o\in\mathcal{O}} \psi[o|e, b](d, \pi) \cdot c[e, o](d, \pi). \tag{2.11}$$

2.5 Redistribution Mechanism

After each game is played, the winners pay their bid. The accumulated karma of this payment is then redistributed. The following two mechanisms were implemented and compared in this thesis.

1. Uniform redistribution,

2. Redistribution only to players with strictly positive urgency.

1. puts agents that don't play every game at an advantage since they receive karma even after rounds where they don't have a risk to incur any cost. 2. rids this advantage at the price of losing the private information property of the urgency as this value has to be known (at least whether or not it is 0) in order to redistribute karma only to the eligible agents.

2.6 Exchange Rate

Since multiple games are combined into one, one can ask the question if a karma point should always be of the same value in all games. To model differently valued karma points depending on the game, an exchange rate ξ is introduced in this section. Let a karma point in game e_i have a value of ξ_{e_i,e_j} in game e_j, where $\xi_{e_i,e_j} \in \mathbb{R}$ denotes the *exchange rate* between the two games. Then, the exchange rate of the whole economy can be written as a matrix

$$\Xi' = \begin{bmatrix} 1 & \xi_{e_1,e_2} & \xi_{e_1,e_3} & \cdots & \xi_{e_1,e_{n_e}} \\ \xi_{e_2,e_1} & 1 & \xi_{e_2,e_3} & \cdots & \xi_{e_2,e_{n_e}} \\ \xi_{e_3,e_1} & \xi_{e_3,e_2} & 1 & \cdots & \xi_{e_3,e_{n_e}} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \xi_{e_{n_e},e_1} & \xi_{e_{n_e},e_2} & \xi_{e_{n_e},e_3} & \cdots & 1 \end{bmatrix} \in \mathbb{R}^{n_e \times n_e} \tag{2.12}$$

with the karma-preserving properties

$$\xi_{e_i,e_j} \cdot \xi_{e_j,e_k} = \xi_{e_i,e_k}, \quad \forall i,j,k \in \{1,2,...,n_e\}, \tag{2.13}$$

$$\xi_{e_i,e_j} = \frac{1}{\xi_{e_j,e_i}}. \tag{2.14}$$

Note that (2.14) follows directly from (2.13) and $\xi_{e_i,e_i} = 1$. From (2.13), it can also be seen that it suffices to define the consecutive exchange rates, i.e., $\xi_{e_1,e_2}, \xi_{e_2,e_3}, ..., \xi_{e_{n_e-1},e_{n_e}}, \xi_{e_{n_e},e_1}$ as the other ξ can be derived from them. Additionally, considering constraint (2.2), this is indeed all that is needed for the implementation of the model. This allows us to simplify the dependency on the games e to a single subscript,

$$\xi_{e_i} := \xi_{e_i,e_{i+1}}, \quad \forall i,\ e_{n_e+1} = e_1 \tag{2.15}$$

Thus, we can write the exchange rates as

$$\Xi = \begin{bmatrix} \xi_{e_1} & \xi_{e_2} & \cdots & \xi_{e_{n_e-1}} & \xi_{e_{n_e}} \end{bmatrix}^T, \tag{2.16}$$

where $\xi_{e_{n_e}}$ represents the exchange rate from e_{n_e} to e_1.
Further, let us define the exchange function $\chi_{e_i}[k^+|k']$ that assigns a probability to reach a karma balance k^+ in the game e_{i+1} given the karma value k' in game e_i after the redistribution. Since the karma balance is restricted to integers, any fractional values are rounded up or down in a karma-preserving manner. More precisely, if ξ is the exchange rate, a fraction p_{high} of the agents

with balance k' will go to the state $\lceil \xi \cdot k' \rceil$ while the fraction $p_{\text{low}} = 1 - p_{\text{high}}$ will go to state $\lfloor \xi \cdot k' \rfloor$. These probabilities are given as follows.

$$p_{\text{high}} = \xi_{e_i} \cdot k' - \lfloor \xi_{e_i} \cdot k' \rfloor,$$
$$p_{\text{low}} = 1 - p_{\text{high}}.$$

From the perspective of an ego-agent, this means that they will reach the lower (higher) karma state with probability p_{low} (p_{high}).
This leads to the exchange function

$$\chi_{e_i}[k^+|k'] = \begin{bmatrix} 0 & \cdots & p_{\text{low}} & p_{\text{high}} & \cdots & 0 \end{bmatrix}^T, \tag{2.17}$$

where the probabilities p_{low} and p_{high} are at the location of the values $\lfloor \xi_{e_i} \cdot k' \rfloor$ and $\lceil \xi_{e_i} \cdot k' \rceil$, respectively.

2.7 State Transition Function

The game and urgency of an agent in the next competition follows the exogenous process from Equation (2.1). The karma balance available for the resource of the current game, however, depends on values of the current game, namely their bid, the outcome and the redistribution mechanism. Let us denote the karma transition function right after the redistribution similar to [2], Equation (7), by

$$\kappa_r[k'|e, k, b, o](d, \pi). \tag{2.18}$$

Then we can define the complete karma transition function, which includes the effect of the exchange rate by

$$\kappa[k^+|e, k, b, o](d, \pi) = \sum_{k' \in \mathbb{N}} \chi_e[k^+|k'] \cdot \kappa_r[k'|e, k, b, o](d, \pi), \tag{2.19}$$

which leads to the state transition function

$$\rho_\tau[e^+, u^+, k^+|e, u, k, b](d, \pi) = \phi_\tau[e^+, u^+|e, u] \sum_{o \in \mathcal{O}} \psi[o|e, b](d, \pi) \cdot \kappa[k^+|e, k, b, o](d, \pi) \tag{2.20}$$

2.8 Stationary Nash Equilibrium

Elokda et al. reduce the existence of a NE to two conditions in [2],

1. Continuity of the state transition function in (d, π),
2. Karma preservation in expectation.

With the addition of a finite and discrete dimension $\mathcal{E}$, the continuity in the social state (d, π) is not broken.
Karma preservation is given by the nature of the redistribution mechanism, which redistributes all payments to (parts of) the society after every game is played, and Equation (2.14). This ensures that every specific game e is played with the same amount of karma distributed among the agents (though a game $e' \neq e$ might have a different amount due to the exchange rate, but still the always same amount for every unique game). Thus, a Nash Equilibrium is guaranteed for the described model.
The SNE is found by computing the infinite horizon reward of the repeated game.

Chapter 3

Results

In this chapter, we analyze our new model of multi-karma economies using a specific example of two coupled resources, and present the observed properties at the SNE. As shown in previous work, the social state (d, π) can be assumed to be constant in time given e (as seen in Section 2.3) [1]. The objective is to demonstrate that there is a benefit in combining multiple economies.

3.1 Setting

For the numerical analysis of the multi-karma economy, a combination of two games was studied. In the first game e_1, the regulated resource is a fast lane along a highway that leads to a city. The non-regulated lanes can be used by any number of agents, but congestion scales with the load. In the second game e_2, agents bid for a parking spot in a parking lot in the city center. An agent that is not allotted a spot in this parking lot has to find a free parking spot elsewhere in the city, which becomes increasingly difficult the more other agents have joined the search.
In both games, the capacity of the regulated resource is $0.1N$, where N is the number of agents in the game.

3.1.1 Urgency Process

There are two types of agents, one is called the *commuter* and the other the *city worker*. The commuter type lives in the countryside and, commuting to work by car, uses both the highway and a parking spot every day. On the other hand, the city worker lives within the city's suburbs and still travels to work by car, and thus always needs a parking space. Additionally, the city worker makes the occasional factory visit outside the city where they will stay overnight and use the highway along with the commuters in the next morning. Both types sometimes have an important meeting early in the morning, which considerably increases their need to obtain access to the fast lane and the central parking lot. To express this more formally, the set of urgency values for both types is

$$\mathcal{U} = \{0, 1, 6\}, \tag{3.1}$$

where $u = 0$ represents a *no-play state*, meaning that an agent has no necessity for the resource in this round, i.e., the case when the city worker does not visit the factory.
The urgency markov chain for the commuter, ϕ_{commuter}, is depicted in Figure 3.1. The transition probabilities with value 1 make the urgency of the parking game the same as in the highway game of the same round, simulating the important meeting in the morning. Therefore, the two games combined represent one day for the agents. The probability to have high urgency on the next day is $\phi_{\text{commuter}}[e_1, 6|e_2, u] = p_{high} = 0.2$, independently of the urgency u on the current day. Accordingly, the probability to have low urgency is $\phi_{\text{commuter}}[e_1, 1|e_2, u] = p_{low} = 1 - p_{high} = 0.8$.

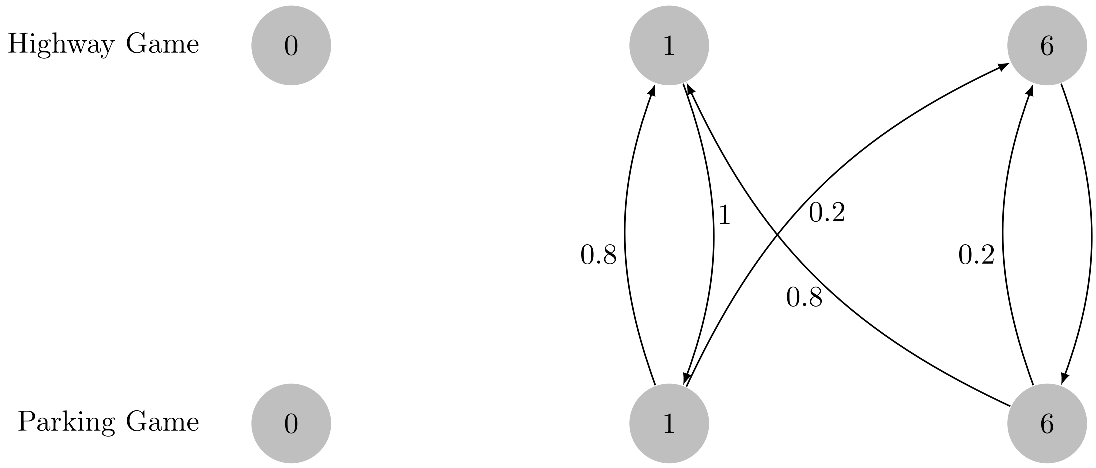

Figure 3.1: The markov chain of the commuter type ϕ_{commute}

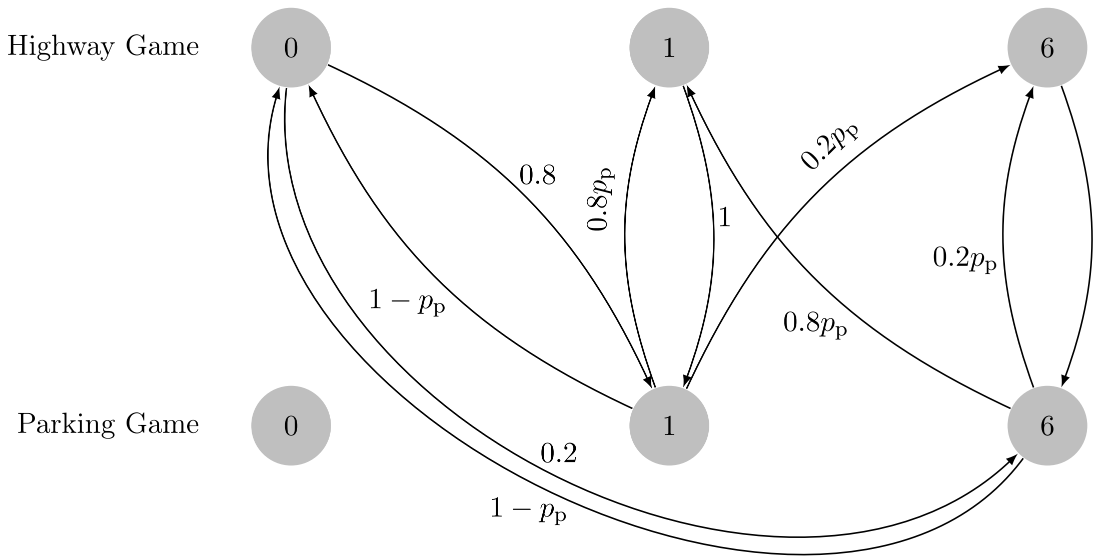

Figure 3.2: The markov chain of the city worker type ϕ_{city}

The urgency markov chain for the city worker, ϕ_{city}, is shown in Figure 3.2. It incorporates the occasional factory visits of the city worker by assigning a non-zero probability to go to the no-play state of the highway game. In particular, the probability that a city worker does not visit the factory and sleeps at home is $\phi_{\text{city}}[e_1, 0|e_2, u] = 1 - p_{\text{p}}$, where p_{p} denotes the participation probability for the highway game. The other probabilities are scaled accordingly, i.e., the probability to have low urgency on the next day is $\phi_{\text{city}}[e_1, 1|e_2, u] = p_{\text{p}} \cdot p_{low}$ and the probability to have high urgency is $\phi_{\text{city}}[e_1, 6|e_2, u] = p_{\text{p}} \cdot p_{high}$. The participation probability p_{p} will be varied in the first part of the analysis. Later, it will be fixed to $p_{\text{p}} = 0.5$.

3.1.2 Methodology

All data displayed in this chapter is generated from the policies and distributions of the agents at the SNE. The measure for the reward of an agent is the expected reward, which can be defined depending on the type τ by

$$R_\tau = \frac{1}{n'_e[\tau]} \sum_{b,e,u,k} \pi_\tau[b|e,u,k] \cdot \zeta_\tau[e,u,b](d,\pi), \tag{3.2}$$

where $n'_e[\tau]$ is the effective number of games played by an agent of the corresponding type, assuming agents do not participate in the bidding for a resource if their urgency is $u = 0$. Normalizing the reward in this way enables direct comparison of different types' rewards since it mitigates the natural advantage of players with lower general demand (players that visit the no-play state more often). The effective number of games played by the commuter is given by $n'_e[\text{commuter}] = n_e$ and by the city worker by

$$n'_e[\text{city}] = \frac{1 + p_{\text{p}}}{n_e}. \tag{3.3}$$

To evaluate the performance of the model, it is compared to the non-regulated economy on one hand, and to the case of two separate economies on the other. Non-regulated means that there is no policy maker intervention for any resource.
Agents discount the future *days* with a factor of $\alpha' = 0.99$. This means that the discount factor $\alpha = \sqrt[n_e]{0.99}$ is used to discount a single game, ensuring that when the game is carried out the next time, the reward is discounted by α'. It permits direct comparison of the separate and combined games with minimal numerical error while keeping the Bellman Equation solvable.
As a social measure, the average reward of all agents was chosen,

$$\frac{1}{g_\tau} \sum_\tau R_\tau. \tag{3.4}$$

3.2 Results

3.2.1 Unit Exchange Rate

First, the direct combination of the two games is analyzed, namely where karma points can be used in any game without any restrictions. Figure 3.3 exhibits the reward for both types for the non-regulated baseline (dotted lines), and for the combined (solid lines) and separate (dashed lines) games as a function of the participation probability of the city worker in the highway game. It can be seen that for all p_{p} an improvement is achieved by combining the two games. The difference between the commuter and the city worker is explained by the slight advantage the city workers have because they obtain karma from the redistribution even if they do not use the highway.

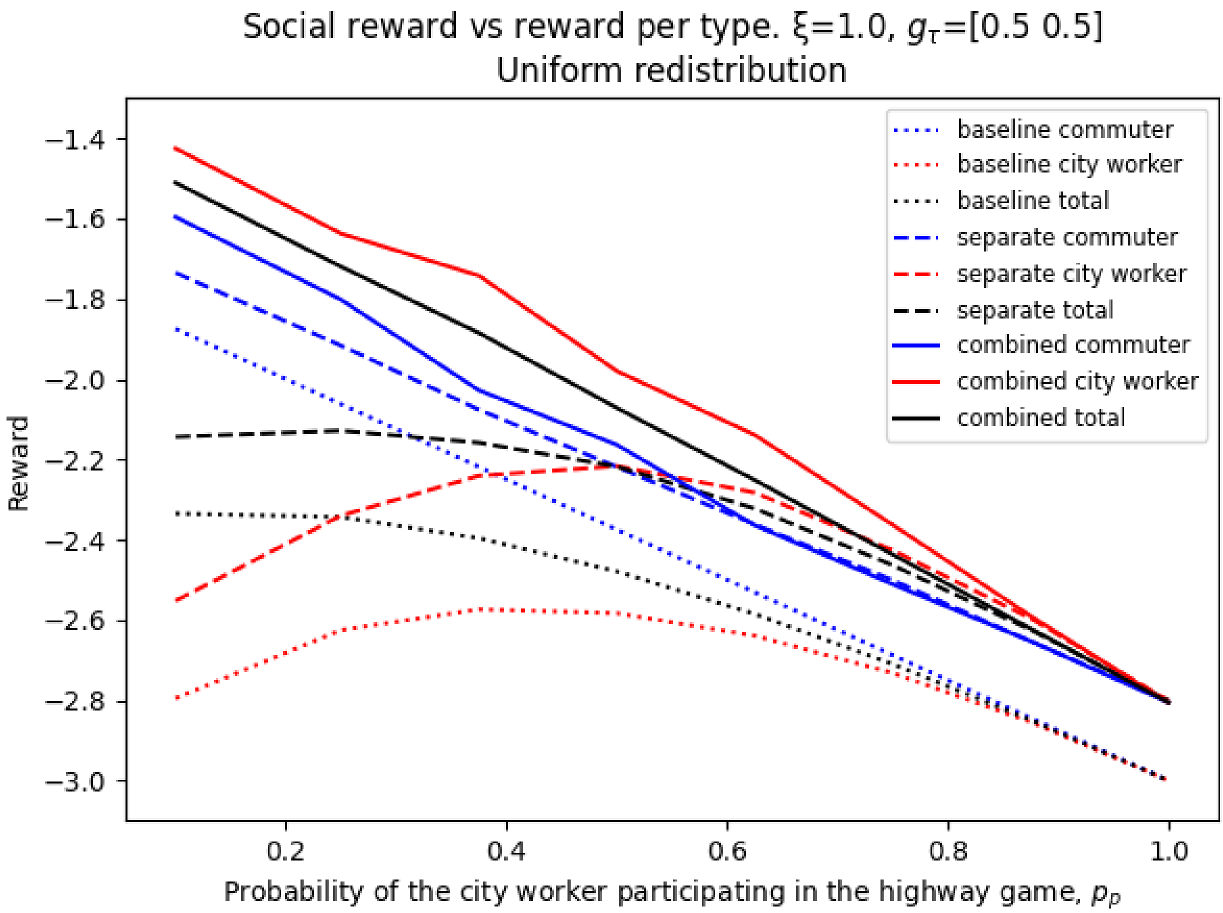

Figure 3.3: Comparison of the rewards per type and in total between the non-regulated baseline, the separate and the combined economies as a function of p_{p}.

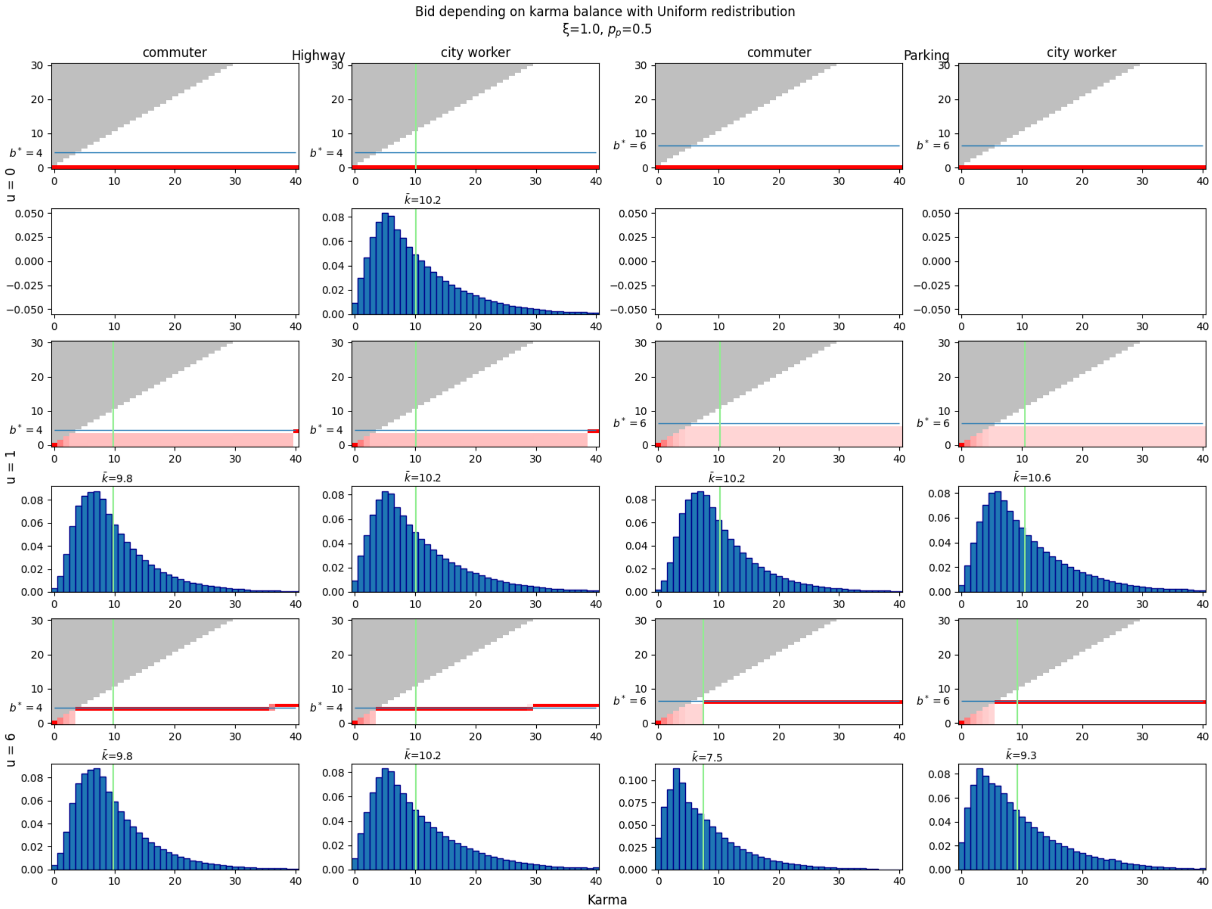

Figure 3.4: Policy and karma distribution of each type in each game at the SNE. The darker the red, the higher the probability that the action is chosen. The gray area covers invalid actions since no bid can be larger than the agent's karma balance. The horizontal blue line indicates the threshold bid $b^*[e](d,\pi)$, and the vertical green line indicates the average karma balance $\bar{k}$ of the agents in the corresponding state.

The higher cost for higher p_p is explained by the additional number of agents playing the highway game. Even though this effect is normalized, the number of agents grows, increasing overall congestion. It is obvious that with city workers at highway participation probability 1, all three lines coincide since the two types are equivalent.

Figure 3.4 displays the policies of both types depending on their current urgency and which game they are currently playing in the case of uniform redistribution and $p_\text{p} = 0.5$. Additionally, the corresponding karma distribution is shown. We see that, at low urgency, both types only place bids below b^*. Any bid $b < b^*$ has equal probability as the exact value b has no influence on the state transition since no payment is due. Agents at a high urgency state bid the threshold value or one point above since now the cost incurred by not getting access to the fast lane or a good parking spot would be high.

Since the number of participants is larger in the parking game, it is more fiercely contested, leading to a higher threshold bid $b^*_\text{Parking} = 6$.

Further, we can observe differences in the karma distributions. In the highway game, the city workers are at a slightly favourable position that stems from the redistribution mechanism that rewards agents not using the highway. For both types, however, the distributions are equivalent for all urgency states, since agents have no knowledge during the previous round in which state they will be. Moreover, in the parking game, the effect of the redistribution is visible as the

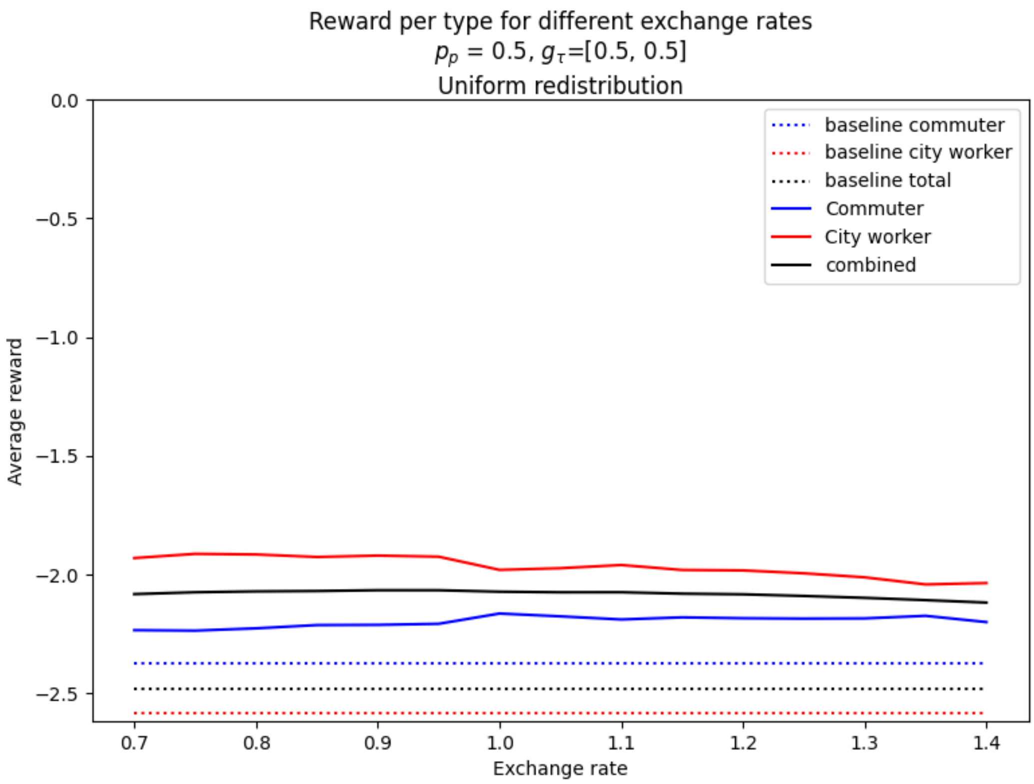

Figure 3.5: The reward per type and in total as a function of the exchange rate. For comparison, the baseline for the non-regulated case is also shown. Since karma points can only be used for the resource they were obtained from, the baseline is constant in the exchange rate.

average karma balance increased for low-urgency agents while it decreased for the high-urgency agents. Since some city workers did not play the highway game at all but now have high urgency for the parking, their average karma is almost 2 karma points higher than the commuters'.

3.2.2 Non-unit Exchange Rate

In the previous section, it became evident that an improvement is achieved by combining the two games. But can the result be further improved by varying the exchange rate between the games? Figure 3.5 depicts the reward as a function of the exchange rate. The participation probability of the city worker is again fixed at $p_{\mathrm{p}} = 0.5$ and the uniform redistribution mechanism is applied.

It is evident that the exchange rate has almost no impact on the rewards. To reason why, let's take a look at Figure 3.6, where a change of the threshold bid can be observed. Due to the exchange rate of $\xi = 0.8$, one karma point in the parking game is worth more than one karma point in the highway game, mitigating the effect of the exchange rate on the threshold bid, as the effective value of the bid is approximately the same as in the unit exchange rate case. Since b^* can only take integer values, rounding has an impact as well and explains the large difference in this specific case. Equivalent reasoning holds for $\xi = 1.2$, where b^* is larger than for the unit exchange rate.

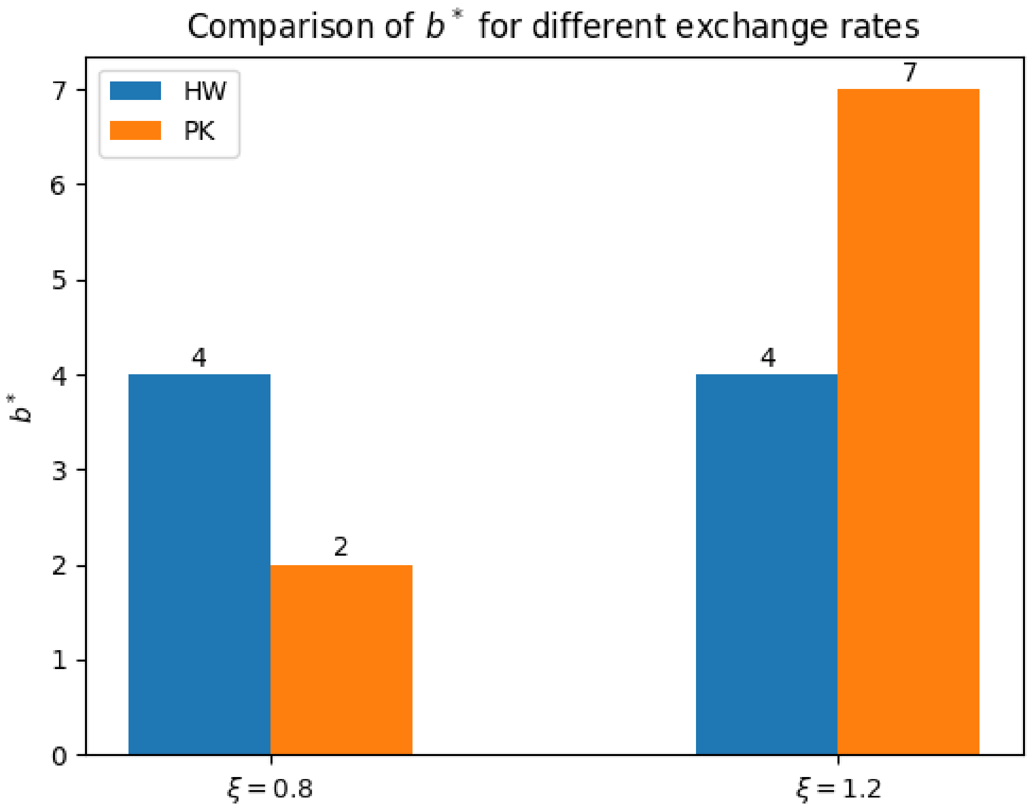

Figure 3.6: A comparison between the threshold bids for different exchange rates.

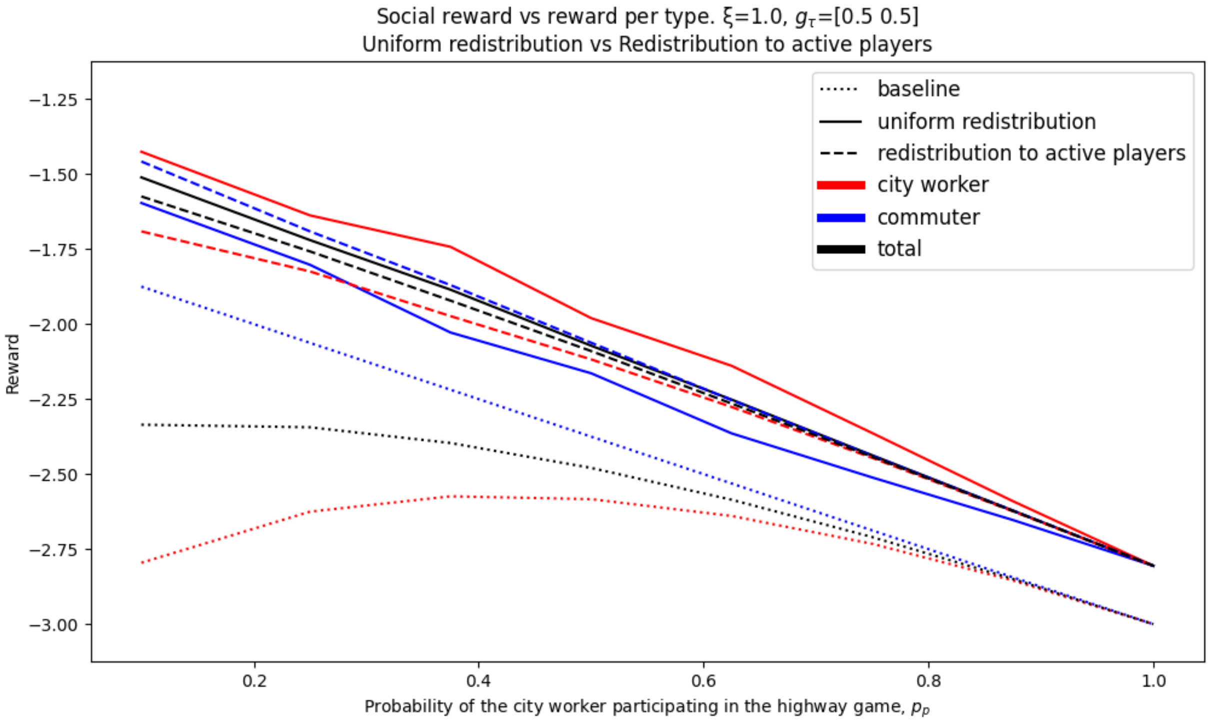

Figure 3.7: A comparison of different redistribution mechanisms. The dotted line shows the non-regulated baseline reward, the solid lines the rewards for uniform redistribution, and the dashed lines the rewards with the mechanism that only redistributes to players that actually used the highway.

3.2.3 Redistribution Mechanisms

In Figure 3.7, a comparison between the two redistribution mechanisms introduced in Chapter 2.5 is depicted. It shows the baseline of the non-regulated case with the dotted lines as a reference. We can observe that the rewards are very similar in terms of the improvement compared to that baseline. However, a big difference is that while in the uniform redistribution the city workers are at an advantage over the commuters, the commuters are favoured when the karma is only redistributed to active players. Since the city workers also obtain karma when they do not use the highway at all, they receive more karma per game effectively played, which explains the gap. Under the redistribution to active players on the other hand, the city workers lose this benefit. Moreover, since both types need a parking spot every day in the scenario, the latter is the more fiercely contested resource. Because the city workers play the parking game at a higher rate (as a fraction of their effective number of games played), they face more competition on average, leading to the slight disadvantage over the commuters.

Conclusion

This thesis extends the mathematical framework of karma economies to couple multiple resource domains. It is shown that, even under simple design decisions such as uniform redistribution, an improvement for heterogeneous types is achieved. The newly introduced exchange rate enables valuing karma differently, depending on the resource that is contested. However, it solely impacts intermediate numerical results, that is the threshold bids, and only shows negligible effect on the agents' rewards. The redistribution mechanism proved to be a more effective tool in altering the outcomes. Nevertheless, neither of the two considered mechanisms attains a strict improvement over the other, but rather a shift in benefits for agents with different behaviour patterns. The decision which design choice should be implemented is to be made by a policy maker.

Future Ideas

Future work could include assessing the benefits of coupling more than two resources, considering more types of agents with different behaviour patterns, as well as investigating a broader spectrum of design choices and their implications for policy makers.

Acknowledgements

First and foremost, I thank Prof. Florian Dörfler and my supervisor Ezzat Elokda for the opportunity to work on this thesis. Further, I thank Ezzat for his insight and guidance throughout the project. Finally, I thank my family and friends for their help in catching errors as well as typos while proof-reading the thesis.

Bibliography

[1] Ezzat Elokda, Andrea Censi, and Saverio Bolognani. "Dynamic population games". In: *arXiv preprint arXiv:2104.14662* (2021).

[2] Ezzat Elokda et al. "A self-contained karma economy for the dynamic allocation of common resources". In: *Dynamic Games and Applications* (2023), pp. 1–33.

[3] Ezzat Elokda et al. "Carma: Fair and efficient bottleneck congestion management with karma". In: *arXiv preprint arXiv:2208.07113* (2022).