Unique Ergodicity in Feedback Interconnections of Ensembles of Agents

Wynita M. Griggs[a], Ramen Ghosh[b,c], Jakub Mareček[d] and Robert N. Shorten[e]

[a]Monash University, Clayton, Victoria, Australia;
[b]University College Dublin, Dublin, Ireland;
[c]Atlantic Technological University, Sligo, Ireland;
[d]Czech Technical University in Prague, Prague, Czech Republic;
[e]Imperial College London, South Kensington, United Kingdom

ARTICLE HISTORY
Compiled February 28, 2025

ABSTRACT
Recent results (Fioravanti, Mareček, Shorten, Souza, & Wirth, 2019; Kungurtsev, Marecek, Ghosh, & Shorten, 2023; Mareček, Roubalik, Ghosh, Shorten, & Wirth, 2023) on the unique ergodicity of feedback systems, which strive to orchestrate the behaviour of a single ensemble of agents, demonstrate that, under straightforward conditions, convergence in distribution of the feedback loop to a unique invariant measure can be guaranteed. These results were motivated by sharing-economy applications, where one strives to provide guarantees on the allocation of constrained resources under a feedback signal. In this paper, we extend these existing results to the case of larger-scale interconnections containing two or more ensembles of agents. This direction may be construed as an analogue of the classical theory of interconnected dynamical systems under uncertainty of the systems' models. Our key finding is that traditional small-gain-like results do not apply. In the case of linear filters and controllers, interconnections of such ensembles of agents are uniquely ergodic if and only if the controllers and filters are stable.

KEYWORDS
Randomized methods; multi-agent systems; control of interconnected systems; simulation of stochastic systems; dynamic resource allocation

1. Introduction

A number of societal-scale challenges should be addressed using control-theoretic approaches, as has been suggested by a roadmap (Annaswamy, Johansson, Pappas, et al., 2023) co-sponsored by IFAC and IEEE CSS. Many of these challenges require a model of human behaviour, which in turn requires a probabilistic model of the response to a control signal. For instance, one can consider non-linear probability functions (Fioravanti et al., 2019). At the same time, these challenges require that the closed-loop system exhibits behaviour acceptable both from the perspective of the system operator (e.g., in terms of regulation of the aggregate behaviour of a number of individuals) and the individual perspective (e.g., predictability and fairness of the average allocation of a resource to an individual). Traditional control theory does not seem well-equipped

CONTACT W. M. Griggs. Email: wynita.griggs@monash.edu

to deal with such problems in the control of multi-agent systems.

A number of such novel challenges arise from the so-called sharing economy (Annaswamy et al., 2023, Section 2.D). There, one often encounters two-sided markets (Lobel, 2020; Rochet & Tirole, 2006). For example, in modelling online labour platforms, these enable the interaction between human customers submitting jobs over time and workers performing these jobs. On the one side, a controller might be suggesting prices to the customers; while on the other, a controller matches the jobs to the workers. Both the customers and workers likely exhibit what would be considered "random" behaviours in regard to their acceptance or rejection of prices and jobs, respectively, by an outside observer. That is, one customer might accept a price while another does not. One worker might accept a job while another does not. These random behaviours can be modelled using so-called probability functions (Fioravanti et al., 2019).

This paper aims to demonstrate results on ergodicity in regard to these random behaviours. Informally speaking, we call a feedback system uniquely ergodic when, for every agent i, there exists a limit of a long-term average allocation of a resource to the agent, independent of any initial conditions. In turn, the notion of unique ergodicity underlies a natural notion of fairness, distinct from other popular notions of fairness (Bateni, Chen, Ciocan, & Mirrokni, 2022), where the limit coincides across all agents or market participants. A necessary condition for the feedback system to be uniquely ergodic is to be contractive on average.

Specifically, in this paper, we consider interconnections of populations of agents exhibiting stochastic behaviours, where each population has associated with it a stable controller. We show that, provided that the individual systems in the feedback interconnection are ergodic, then the interconnections of these systems will also be ergodic (under some straightforward assumptions).

The results presented form a natural extension of the ideas on ensemble control presented in Fioravanti et al. (2019); Ghosh, Mareček, Griggs, Souza, and Shorten (2022); Kungurtsev et al. (2023); Mareček et al. (2023). In Fioravanti et al. (2019); Ghosh et al. (2022); Kungurtsev et al. (2023); Mareček et al. (2023), all agents can be thought of as responding to the prices a central organising entity sets or signals it provides. It was shown that a population of agents in a simple feedback interconnection with stable controller and stable filter results in ergodic behaviour. Kungurtsev et al. (2023); Mareček et al. (2023) extended the results to non-linear controllers and filters.

We now extend results presented in Fioravanti et al. (2019) to a feedback interconnection containing two (or more) populations of agents exhibiting certain stochastic behaviours. Our results suggest the following:

- When each population is subject to a stable controller and filtered output, one obtains unique ergodicity. (State converges in distribution to a unique invariant measure.)
- Small-gain-like results do not apply to linear controllers and filters: interconnections of such systems are uniquely ergodic if and only if the controllers and filters are stable.

The remainder of this work is organised as follows. In Section 2, we provide the necessary mathematical preliminaries. In Section 3, the main result for an intercon-

nection of two ensembles of agents is presented. In Section 4, the result is extended to the case where there are multiple ensembles of agents (i.e., more than two). A simulated example demonstrating the main result is presented in Section 5. Discussion on potential applications is provided in Section 6. Section 7 contains the conclusions and possible directions for future work.

2. Mathematical Preliminaries

This paper builds on the work of Fioravanti et al. (2019), and aims to reuse as much of the notation as is feasible, albeit noting that we will consider two (or more) different populations of agents, as opposed to the single population of agents considered in Fioravanti et al. (2019). In general, we will use the superscript to distinguish between the populations of agents. For example, we will let $\Pi^1 \subseteq \mathbb{R}$ and $\Pi^2 \subseteq \mathbb{R}$ denote sets of admissible broadcast control signals for the two populations of agents. Suppose that each agent from the first population responds to a control signal $\pi^1 \in \Pi^1$, and each agent from the second population responds to a control signal $\pi^2 \in \Pi^2$.

As in Fioravanti et al. (2019), these responses contain non-deterministic elements that can be modelled by agent-specific probability distributions. For each agent i in each population $o \in \{1, 2\}$, and for all $k \in \mathbb{Z}^+$, we have that

$$\mathbb{P}\Big(x_i^o(k+1) = \mathcal{W}_{ij}^o(x_i^o(k))\Big) = p_{ij}^o(\pi^o(k)), \tag{1a}$$

$$\mathbb{P}\Big(y_i^o(k) = \mathcal{H}_{il}^o(x_i^o(k))\Big) = p_{il}^{o'}(\pi^o(k)), \tag{1b}$$

where $\mathcal{W}_{ij}^o$, for $j = 1, 2, \ldots, w_i^o \in \mathbb{N}$, denotes all possible state transition maps, and $\mathcal{H}_{il}^o$, for $l = 1, 2, \ldots, h_i^o \in \mathbb{N}$, denotes all possible output maps, for each agent i, similar to Fioravanti et al. (2019, Section 2.2). The $p_{ij}^o : \Pi^o \to [0, 1]$ and $p_{il}^{o'} : \Pi^o \to [0, 1]$ are probability functions. Additionally, for all $\pi^o \in \Pi^o$, $o \in \{1, 2\}$, it holds that

$$\sum_{j=1}^{w_i^o} p_{ij}^o(\pi^o) = \sum_{l=1}^{h_i^o} p_{il}^{o'}(\pi^o) = 1. \tag{1c}$$

3. Interconnections of Two Ensembles

Consider the system depicted in Figure 1, where $\mathcal{P}^1$ and $\mathcal{P}^2$ denote two different populations of agents of sizes $N^1 \in \mathbb{Z}^+$ and $N^2 \in \mathbb{Z}^+$, respectively. The individual agents' states are described by finite-dimensional, linear, shift-invariant systems, as follows:

$$\mathcal{P}_i^1 \ : \ \begin{cases} x_i^1(k+1) &= A_i^1 x_i^1(k) + b_i^1, \\ y_i^1(k) &= c_i^{1T} x_i^1(k) + d_i^1, \end{cases} \tag{2}$$

for $i = 1, 2, \ldots, N^1$; and

$$\mathcal{P}_i^2 \ : \ \begin{cases} x_i^2(k+1) &= A_i^2 x_i^2(k) + b_i^2, \\ y_i^2(k) &= c_i^{2T} x_i^2(k) + d_i^2, \end{cases} \tag{3}$$

for $i = 1, 2, \ldots, N^2$. It is assumed that all of the $A_i^1 \in \mathbb{R}^{n_i^1 \times n_i^1}$ and $A_i^2 \in \mathbb{R}^{n_i^2 \times n_i^2}$ are Schur matrices (i.e., the spectral radii of the matrices are less than one), where $n_i^1 \in \mathbb{Z}^+$ and $n_i^2 \in \mathbb{Z}^+$; $x_i^1, c_i^1 \in \mathbb{R}^{n_i^1}$ and $x_i^2, c_i^2 \in \mathbb{R}^{n_i^2}$; and the inputs b_i^1, d_i^1, b_i^2 and d_i^2 are random variables that take values in $\mathbb{R}^{n_i^1}$, $\mathbb{R}$, $\mathbb{R}^{n_i^2}$ and $\mathbb{R}$, respectively, with $\mathbb{P}(b_i^1 = b_{ij}^1) = p_{ij}^1(\pi^1)$, $\mathbb{P}(d_i^1 = d_{il}^1) = p_{il}^{1'}(\pi^1)$, $\mathbb{P}(b_i^2 = b_{ij}^2) = p_{ij}^2(\pi^2)$ and $\mathbb{P}(d_i^2 = d_{il}^2) = p_{il}^{2'}(\pi^2)$, where π^1 and π^2 are control signals as described in detail below.

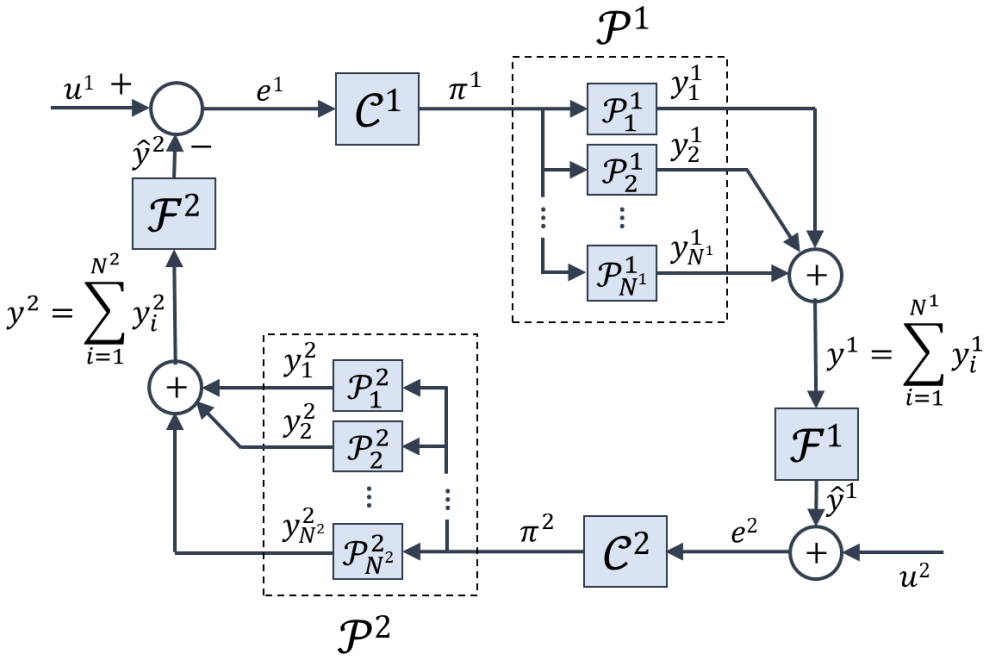

Figure 1. Interconnection of two different populations of agents.

Let $y^1(k) = \sum_{i=1}^{N^1} y_i^1(k)$ and $y^2(k) = \sum_{i=1}^{N^2} y_i^2(k)$, noting that y^1 and y^2 are also random variables. The agents comprising $\mathcal{P}^1$ and $\mathcal{P}^2$ respond to the control signals π^1 and π^2, respectively, which are produced by the controllers

$$\mathcal{C}^1 \; : \; \left\{ \begin{array}{rcl} x_{\mathcal{C}^1}(k+1) & = & A_{\mathcal{C}^1} x_{\mathcal{C}^1}(k) + B_{\mathcal{C}^1} e^1(k), \\ \pi^1(k) & = & C_{\mathcal{C}^1} x_{\mathcal{C}^1}(k) + D_{\mathcal{C}^1} e^1(k), \end{array} \right. \tag{4}$$

and

$$\mathcal{C}^2 \; : \; \left\{ \begin{array}{rcl} x_{\mathcal{C}^2}(k+1) & = & A_{\mathcal{C}^2} x_{\mathcal{C}^2}(k) + B_{\mathcal{C}^2} e^2(k), \\ \pi^2(k) & = & C_{\mathcal{C}^2} x_{\mathcal{C}^2}(k) + D_{\mathcal{C}^2} e^2(k), \end{array} \right. \tag{5}$$

where $x_{\mathcal{C}^1} \in \mathbb{R}^{n_{\mathcal{C}^1}}$, $x_{\mathcal{C}^2} \in \mathbb{R}^{n_{\mathcal{C}^2}}$, $n_{\mathcal{C}^1} \in \mathbb{Z}^+$, $n_{\mathcal{C}^2} \in \mathbb{Z}^+$ and $e^1(k) = u^1(k) - \hat{y}^2(\mathrm{k})$, $e^2(k) = u^2(k) + \hat{y}^1(k)$. Moreover, $\hat{y}^1$ and $\hat{y}^2$ are the outputs of the filters

$$\mathcal{F}^1 \; : \; \left\{ \begin{array}{rcl} x_{\mathcal{F}^1}(k+1) & = & A_{\mathcal{F}^1} x_{\mathcal{F}^1}(k) + B_{\mathcal{F}^1} y^1(k), \\ \hat{y}^1(k) & = & C_{\mathcal{F}^1} x_{\mathcal{F}^1}(k), \end{array} \right. \tag{6}$$

and

$$\mathcal{F}^2 \; : \; \left\{ \begin{array}{rcl} x_{\mathcal{F}^2}(k+1) & = & A_{\mathcal{F}^2} x_{\mathcal{F}^2}(k) + B_{\mathcal{F}^2} y^2(k), \\ \hat{y}^2(k) & = & C_{\mathcal{F}^2} x_{\mathcal{F}^2}(k), \end{array} \right. \tag{7}$$

respectively, where $x_{\mathcal{F}^1} \in \mathbb{R}^{n_{\mathcal{F}^1}}$, $x_{\mathcal{F}^2} \in \mathbb{R}^{n_{\mathcal{F}^2}}$ and $n_{\mathcal{F}^1} \in \mathbb{Z}^+$, $n_{\mathcal{F}^2} \in \mathbb{Z}^+$. The main result of the paper follows.

Theorem 3.1 (Main Result)**.** *Consider the feedback system illustrated in Figure 1, where $\mathcal{C}^1$ and $\mathcal{C}^2$ are controllers described by* (4) *and* (5)*, respectively, and $\mathcal{F}^1$ and $\mathcal{F}^2$ are filters described by* (6) *and* (7)*, respectively. Assume that each agent i from its respective population has its state dynamics determined by the stochastic difference equations of* (2) *or* (3)*, as appropriate, where all of the A_i^1 and A_i^2 are Schur matrices; and the b_i^1, d_i^1, b_i^2 and d_i^2 are chosen, at each time step, from the sets $\{b_{ij}^1\} \subset \mathbb{R}^{n_i^1}$, $\{d_{il}^1\} \subset \mathbb{R}$, $\{b_{ij}^2\} \subset \mathbb{R}^{n_i^2}$ and $\{d_{il}^2\} \subset \mathbb{R}$, respectively, according to Dini continuous probability functions $p_{ij}^1(\cdot)$, $p_{il}^{1'}(\cdot)$, $p_{ij}^2(\cdot)$, and $p_{il}^{2'}(\cdot)$, respectively, that verify* (1)*. Suppose, further, that the probability functions $p_{ij}^1(\pi^1)$, $p_{il}^{1'}(\pi^1)$, $p_{ij}^2(\pi^2)$, and $p_{il}^{2'}(\pi^2)$ are bounded below by scalars that are strictly greater than zero for all (i,j), or (i,l), and all $\pi^1 \in \Pi^1$, or $\pi^2 \in \Pi^2$, as appropriate. Then, for every pair of stable controllers, $\mathcal{C}^1$ and $\mathcal{C}^2$, and every pair of stable filters, $\mathcal{F}^1$ and $\mathcal{F}^2$, the feedback loop converges in distribution to a unique invariant measure.*

Proof: Let $x^1(k) := [x_1^1(k)^T \; x_2^1(k)^T \; \ldots \; x_{N^1}^1(k)^T]^T$, $x^2(k) := [x_1^2(k)^T \; x_2^2(k)^T \; \ldots \; x_{N^2}^2(k)^T]^T$, $b^1 := [b_1^{1T} \; b_2^{1T} \; \ldots \; b_{N^1}^{1T}]^T$, $b^2 := [b_1^{2T} \; b_2^{2T} \; \ldots \; b_{N^2}^{2T}]^T$, $d^1 := [d_1^1 \; d_2^1 \; \ldots \; d_{N^1}^1]^T$ and $d^2 := [d_1^2 \; d_2^2 \ldots \; d_{N^2}^2]^T$. Define the augmented state $\xi(k) := [x^1(k)^T \; x^2(k)^T \; x_{\mathcal{F}^1}(k)^T \; x_{\mathcal{F}^2}(k)^T \; x_{\mathcal{C}^1}(k)^T \; x_{\mathcal{C}^2}(k)^T]^T$ whose dynamic behaviour is described by the difference equation

$$\xi(k+1) = \mathcal{A}\xi(k) + \beta,$$

where

$$\mathcal{A} := \begin{bmatrix} \hat{A}^1 & 0 & 0 & 0 & 0 & 0 \\ 0 & \hat{A}^2 & 0 & 0 & 0 & 0 \\ B_{\mathcal{F}^1}\hat{C}^1 & 0 & A_{\mathcal{F}^1} & 0 & 0 & 0 \\ 0 & B_{\mathcal{F}^2}\hat{C}^2 & 0 & A_{\mathcal{F}^2} & 0 & 0 \\ 0 & 0 & 0 & -B_{\mathcal{C}^1}C_{\mathcal{F}^2} & A_{\mathcal{C}^1} & 0 \\ 0 & 0 & B_{\mathcal{C}^2}C_{\mathcal{F}^1} & 0 & 0 & A_{\mathcal{C}^2} \end{bmatrix},$$

$\mathbf{1}$ denotes row vectors of appropriate dimensions containing 1's as entries, $\hat{A}^1 := \mathbf{diag}(A_1^1, \ldots, A_{N^1}^1)$, $\hat{A}^2 := \mathbf{diag}(A_1^2, \ldots, A_{N^2}^2)$, $\hat{C}^1 := [c_1^{1T} \; c_2^{1T} \; \ldots \; c_{N^1}^{1T}]$, $\hat{C}^2 := [c_1^{2T} \; c_2^{2T} \; \ldots \; c_{N^2}^{2T}]$ and

$$\beta := \begin{bmatrix} b^1 \\ b^2 \\ B_{\mathcal{F}^1}\mathbf{1}d^1 \\ B_{\mathcal{F}^2}\mathbf{1}d^2 \\ B_{\mathcal{C}^1}u^1(k) \\ B_{\mathcal{C}^2}u^2(k) \end{bmatrix}.$$

The remainder of the proof then follows in a manner similar to the proof of Fioravanti et al. (2019, Theorem 12) upon noting that $\sigma(\mathcal{A}) = \sigma(\hat{A}^1) \cup \sigma(\hat{A}^2) \cup \sigma(A_{\mathcal{F}^1}) \cup \sigma(A_{\mathcal{F}^2}) \cup \sigma(A_{\mathcal{C}^1}) \cup \sigma(A_{\mathcal{C}^2})$, where $\sigma(\cdot)$ denotes the spectrum of a matrix; and that all of the A_i^1 and A_i^2, and $A_{\mathcal{F}^1}$, $A_{\mathcal{F}^2}$, $A_{\mathcal{C}^1}$ and $A_{\mathcal{C}^2}$, are Schur matrices by assumption. ■

4. More Than Two Ensembles

Building on the techniques of Section 3, it is possible to consider interconnections that contain larger numbers of ensembles of agents (i.e., greater than two ensembles), and show that, under certain sufficient conditions, the interconnections converge in distribution to unique invariant measures. Suppose that a large-scale interconnection of ensembles contains Q populations of agents, where $Q \in \mathbb{Z}^+$. Let $\mathcal{P}^q$ denote each of these different populations of agents, where $q = 1, 2, \ldots, Q$. Let the size of the qth population of agents be denoted by $N^q \in \mathbb{Z}^+$.

The individual agents' states in the qth population are described by finite-dimensional, linear, shift-invariant systems, as follows:

$$\mathcal{P}_i^q \; : \; \left\{ \begin{array}{rcl} x_i^q(k+1) & = & A_i^q x_i^q(k) + b_i^q, \\ y_i^q(k) & = & c_i^{qT} x_i^q(k) + d_i^q, \end{array} \right. \tag{8}$$

for $i = 1, 2, \ldots, N^q$. It is assumed that all of the $A_i^q \in \mathbb{R}^{n_i^q \times n_i^q}$ are Schur matrices, where $n_i^q \in \mathbb{Z}^+$; $x_i^q, c_i^q \in \mathbb{R}^{n_i^q}$; and the inputs b_i^q and d_i^q are random variables that take values in $\mathbb{R}^{n_i^q}$ and $\mathbb{R}$, respectively, with $\mathbb{P}(b_i^q = b_{ij}^q) = p_{ij}^q(\pi^q)$ and $\mathbb{P}(d_i^q = d_{il}^q) = p_{il}^{q'}(\pi^q)$, where π^q is a control signal that all agents in the qth population respond to. This control signal is produced by the controller

$$\mathcal{C}^q \; : \; \left\{ \begin{array}{rcl} x_{\mathcal{C}^q}(k+1) & = & A_{\mathcal{C}^q} x_{\mathcal{C}^q}(k) + B_{\mathcal{C}^q} e^q(k), \\ \pi^q(k) & = & C_{\mathcal{C}^q} x_{\mathcal{C}^q}(k) + D_{\mathcal{C}^q} e^q(k), \end{array} \right. \tag{9}$$

where $x_{\mathcal{C}^q} \in \mathbb{R}^{n_{\mathcal{C}^q}}$, $n_{\mathcal{C}^q} \in \mathbb{Z}^+$ and $e^q(k)$ is the input to the controller.

For the qth population, let $y^q(k) = \sum_{i=1}^{N^q} y_i^q(k)$, noting that y^q is a random variable. Each population has its output filtered. Specifically, for $q = 1, 2, \ldots, Q$, let $\hat{y}^q$ be the output of the filter

$$\mathcal{F}^q \; : \; \left\{ \begin{array}{rcl} x_{\mathcal{F}^q}(k+1) & = & A_{\mathcal{F}^q} x_{\mathcal{F}^q}(k) + B_{\mathcal{F}^q} y^q(k), \\ \hat{y}^q(k) & = & C_{\mathcal{F}^q} x_{\mathcal{F}^q}(k), \end{array} \right. \tag{10}$$

where $x_{\mathcal{F}^q} \in \mathbb{R}^{n_{\mathcal{F}^q}}$ and $n_{\mathcal{F}^q} \in \mathbb{Z}^+$. Next, we describe our large-scale interconnection structure, which is depicted in Figure 2.

For $p = 1, 2, \ldots, Q$, let

$$e^p(k) = u^p(k) - \sum_{q=1}^{Q} H^{pq} \hat{y}^q(k),$$

where each $u^p(k)$ represents an external input, and each H^{pq} represents a matrix with real, constant entries. (Specifically, H^{pq} has the same number of rows as e^p and u^p,

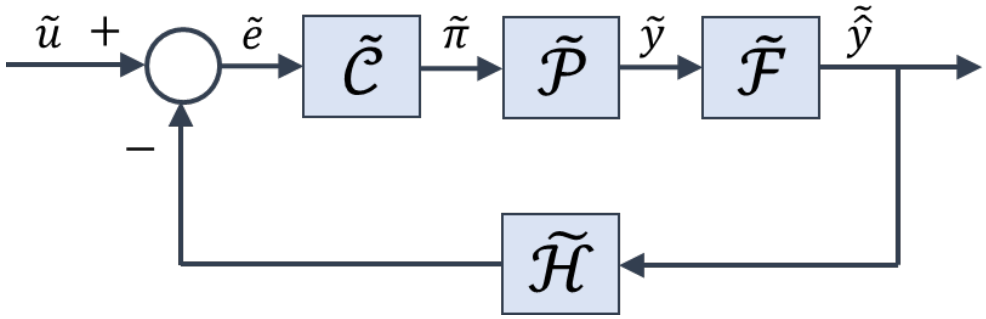

Figure 2. Feedback system containing many different populations of agents.

and the same number of columns as the transpose of $\hat{y}^q$, for all $p = 1, 2, \ldots, Q$, and all $q = 1, 2, \ldots, Q$.) By letting $\tilde{\mathbf{e}}(k) := [e^1(k)^T\ e^2(k)^T\ \ldots\ e^Q(k)^T]^T$, $\tilde{\mathbf{u}}(k) := [u^1(k)^T\ u^2(k)^T\ \ldots\ u^Q(k)^T]^T$ and $\tilde{\hat{\mathbf{y}}}(k) := [\hat{y}^1(k)^T\ \hat{y}^2(k)^T\ \ldots\ \hat{y}^Q(k)^T]^T$, the interconnection can be expressed in the compact form

$$\tilde{\mathbf{e}}(k) = \tilde{\mathbf{u}}(k) - \tilde{\mathcal{H}}\tilde{\hat{\mathbf{y}}}(k),$$

where $\tilde{\mathcal{H}}$ is a block matrix whose elements are the H^{pq}. Let $\tilde{\mathcal{C}} := \mathbf{diag}(\mathcal{C}^1, \ldots, \mathcal{C}^Q)$, $\tilde{\mathcal{P}} := \mathbf{diag}(\mathcal{P}^1, \ldots, \mathcal{P}^Q)$, $\tilde{\mathcal{F}} := \mathbf{diag}(\mathcal{F}^1, \ldots, \mathcal{F}^Q)$, $\tilde{\pi}(k) := [\pi^1(k)\ \pi^2(k)\ \ldots\ \pi^Q(k)]^T$ and $\tilde{y}(k) := [y^1(k)\ y^2(k)\ \ldots\ y^Q(k)]^T$. Then, in the manner of Theorem 3.1, we have the following result.

Theorem 4.1. *Consider the large-scale interconnection depicted in Figure 2, as described in the preamble of Section 4, above. Assume that the ith agent in the qth population of agents, $\mathcal{P}^q$, has its state dynamics determined by the stochastic difference equations of* (8)*, where A_i^q is a Schur matrix; and b_i^q and d_i^q are chosen, at each time step, from the sets $\{b_{ij}^q\} \subset \mathbb{R}^{n_i^q}$ and $\{d_{il}^q\} \subset \mathbb{R}$, respectively, according to Dini continuous probability functions $p_{ij}^q(\cdot)$ and $p_{il}^{q'}(\cdot)$, respectively, that verify* (1)*. Suppose, further, that the probability functions $p_{ij}^q(\cdot)$ and $p_{il}^{q'}(\cdot)$ are bounded below by scalars that are strictly greater than zero for all (i, j) or (i, l), as appropriate, and all $\pi^q \in \Pi^q$. Then, for any Q stable controllers, $\mathcal{C}^1$, $\mathcal{C}^2$, ..., $\mathcal{C}^Q$, as described by* (9)*, and any Q stable filters, $\mathcal{F}^1$, $\mathcal{F}^2$, ..., $\mathcal{F}^Q$, as described by* (10)*, the large-scale interconnection converges in distribution to a unique invariant measure.*

Proof: For all $q = 1, 2, \ldots, Q$, set $x^q(k) := [x_1^q(k)^T\ x_2^q(k)^T\ \ldots\ x_{N^q}^q(k)^T]^T$, $\hat{C}^q := [c_1^{qT}\ c_2^{qT}\ \ldots\ c_{N^q}^{qT}]$, $b^q := [b_1^{qT}\ b_2^{qT}\ \ldots\ b_{N^q}^{qT}]^T$ and $d^q := [d_1^q\ d_2^q\ \ldots\ d_{N^q}^q]^T$. Then, let $\tilde{x}(k) := [x^1(k)^T\ x^2(k)^T\ \ldots\ x^Q(k)^T]^T$, $\tilde{b} := [b^{1T}\ b^{2T}\ \ldots\ b^{QT}]^T$ and $\tilde{d} := [d^{1T}\ d^{2T}\ \ldots\ d^{QT}]^T$. Furthermore, let $\tilde{x}_{\mathcal{F}}(k) := [x_{\mathcal{F}^1}(k)^T\ x_{\mathcal{F}^2}(k)^T\ \ldots\ x_{\mathcal{F}^Q}(k)^T]^T$ and $\tilde{x}_{\mathcal{C}}(k) := [x_{\mathcal{C}^1}(k)^T\ x_{\mathcal{C}^2}(k)^T\ \ldots\ x_{\mathcal{C}^Q}(k)^T]^T$. Define the augmented state $\xi(k) := [\tilde{x}(k)^T\ \tilde{x}_{\mathcal{F}}(k)^T\ \tilde{x}_{\mathcal{C}}(k)^T]^T$ whose dynamic behaviour is described by the difference equation

$$\xi(k+1) = \mathcal{A}\xi(k) + \beta,$$

where

$$\mathcal{A} := \begin{bmatrix} \hat{A} & 0 & 0 \\ \hat{B}_{\mathcal{F}}\hat{C} & \hat{A}_{\mathcal{F}} & 0 \\ 0 & -\hat{B}_{\mathcal{C}}\tilde{\mathcal{H}}\hat{C}_{\mathcal{F}} & \hat{A}_{\mathcal{C}} \end{bmatrix},$$

$\hat{A} := \mathbf{diag}(A_1^1, A_2^1, \ldots, A_{N^1}^1, A_1^2, A_2^2, \ldots, A_{N^2}^2, \ldots, A_1^Q, A_2^Q, \ldots, A_{N^Q}^Q)$, $\hat{C} := \mathbf{diag}(\hat{C}^1, \ldots, \hat{C}^Q)$, $\hat{A}_{\mathcal{F}} := \mathbf{diag}(A_{\mathcal{F}^1}, \ldots, A_{\mathcal{F}^Q})$, $\hat{B}_{\mathcal{F}} := \mathbf{diag}(B_{\mathcal{F}^1}, \ldots, B_{\mathcal{F}^Q})$, $\hat{C}_{\mathcal{F}} := \mathbf{diag}(C_{\mathcal{F}^1}, \ldots, C_{\mathcal{F}^Q})$, $\hat{A}_{\mathcal{C}} := \mathbf{diag}(A_{\mathcal{C}^1}, \ldots, A_{\mathcal{C}^Q})$, $\hat{B}_{\mathcal{C}} := \mathbf{diag}(B_{\mathcal{C}^1}, \ldots, B_{\mathcal{C}^Q})$, $\mathbf{1}$ denotes row vectors of appropriate dimensions containing 1's as entries, $\hat{\mathbf{1}} := \mathbf{diag}(\mathbf{1}, \mathbf{1}, \ldots, \mathbf{1})$ and

$$\beta := \begin{bmatrix} \tilde{b} \\ \hat{B}_{\mathcal{F}}\hat{\mathbf{1}}\tilde{d} \\ \hat{B}_{\mathcal{C}}\tilde{\mathbf{u}}(k) \end{bmatrix}.$$

Observe that $\sigma(\mathcal{A}) = \sigma(\hat{A}) \cup \sigma(\hat{A}_{\mathcal{F}}) \cup \sigma(\hat{A}_{\mathcal{C}})$ and all of the matrices that comprise $\hat{A}$, $\hat{A}_{\mathcal{F}}$ and $\hat{A}_{\mathcal{C}}$ are Schur matrices by assumption. The rest of the proof then follows in the spirit of Theorem 3.1. ■

5. Simulations

To demonstrate Theorem 3.1, namely, the existence of a unique invariant measure, irrespective of the initial conditions of the agents, the following toy experiment was performed.

Populations 1 and 2 were created to consist of agents that could be either "On" or "Off"; that is, each agent's state at time k was permitted to take a value of either 1 (i.e., "On") or 0 (i.e., "Off"). One could think of these populations of agents as forming part of a broader scenario involving a digital labour platform, for example, in which the first population consists of individual customers deciding "Yes" or "No" as to whether to purchase a good or service at the currently offered price, and the second population comprises of workers deciding "Yes" or "No" as to whether to accept jobs. The total sizes of Populations 1 and 2 were set to 100 agents each. The Dini continuous functions used to describe the probabilities of agents' states being "On" according to the controllers' outputs are illustrated in Figure 3. For illustrative purposes, these probability functions were selected such that their inflection points lie on the point $(0, 0.5)$. Note, however, that in a real application, one would typically derive the probability functions from real-world data on human behaviour. Moreover, the probability functions illustrated in Figure 3 were chosen for their plausible sensibility as models depicting the likelihood of agents' responses being "Yes" or "No" based on the cost or benefit to an agent: the more a good or service costs, the less likely a customer is to say "Yes" to the purchase; while, the more a person is offered in terms of financial incentive to work, the more likely they are to accept jobs.

One thousand simulations comprising the experiment were run in total. Each simulation ran for 1,250 time steps, and the external inputs were set to $u^1 = u^2 = 80$ for all simulation runs. At the beginning of each simulation run, the probabilities of agents

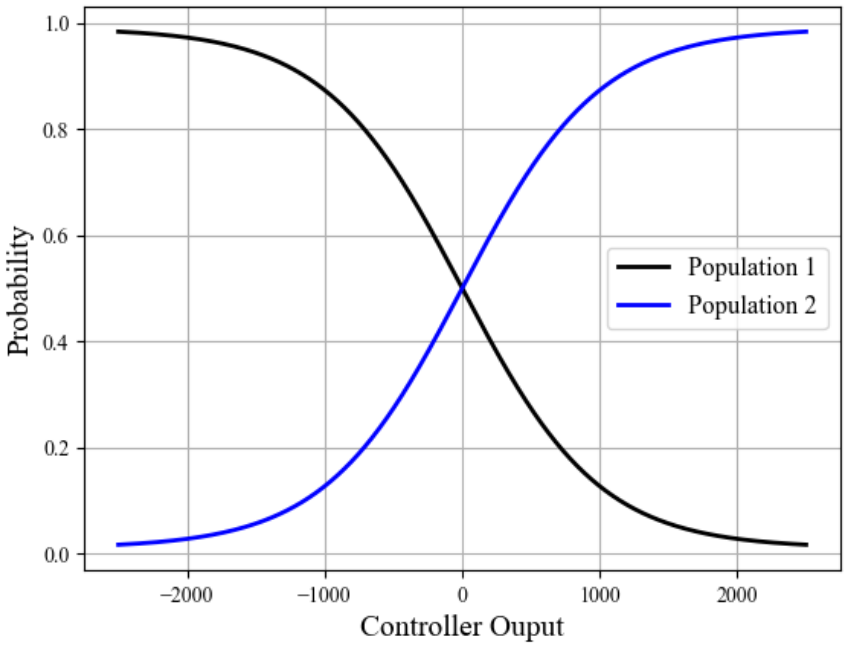

Figure 3. Probability functions. The black line illustrates the probability of agents' states from Population 1 being "On" as a function of π^1. The blue line illustrates the probability of agents' states from Population 2 being "On" as a function of π^2.

from Populations 1 and 2, respectively, being initially "On", were randomly generated to demonstrate different initial conditions. Controllers 1 and 2, respectively, were lag controllers that updated every time step, as described by the difference equations

$$\pi^1(k) = \beta_1 \pi^1(k-1) + \kappa_1[e^1(k) - \alpha_1 e^1(k-1)],$$
$$\pi^2(k) = \beta_2 \pi^2(k-1) + \kappa_2[e^2(k) - \alpha_2 e^2(k-1)],$$

respectively, where $\alpha_1 = -0.1$, $\beta_1 = 0.995$, $\kappa_1 = 0.9$, $\alpha_2 = -12.0$, $\beta_2 = 0.92$, $\kappa_2 = 0.85$ and $k \in \mathbb{N}$. (For simplicity, the controller parameters, α_1, β_1, κ_1, α_2, β_2 and κ_2, were tuned by trial and error. In a real application, more sophisticated techniques, such as those introduced by Nazarov et al. (2024) incorporating artificial intelligence, could be employed.) Filters 1 and 2, respectively, were moving average filters, as described by the equations

$$\hat{y}^1(k) = -0.5[y^1(k) + y^1(k-1)],$$
$$\hat{y}^2(k) = 0.5[y^2(k) + y^2(k-1)],$$

respectively.

Figures 4 and 5 illustrate the average numbers of "On" agents from Populations 1 and 2, respectively, over time, from the 1,000 simulation runs. Figures 6 and 7 depict the mean outputs from Controllers 1 and 2, respectively, over time. Figures 8 and 9 illustrate the mean inputs to Controllers 1 and 2, respectively, over time.

Convergence of the means in all cases is evident. The lag controllers did a reasonable job at achieving small mean steady state errors (see Figures 8 and 9), noting that the use of lag controllers was given preference over PI controllers due to the latter's potential to introduce instabilities. Note that the outputs of Controller 1 can be used in association with the black line in Figure 3, and the outputs of Controller 2 can be used in association with the blue line in Figure 3, to determine the probabilities of agents' states from Populations 1 and 2, respectively, being "On".

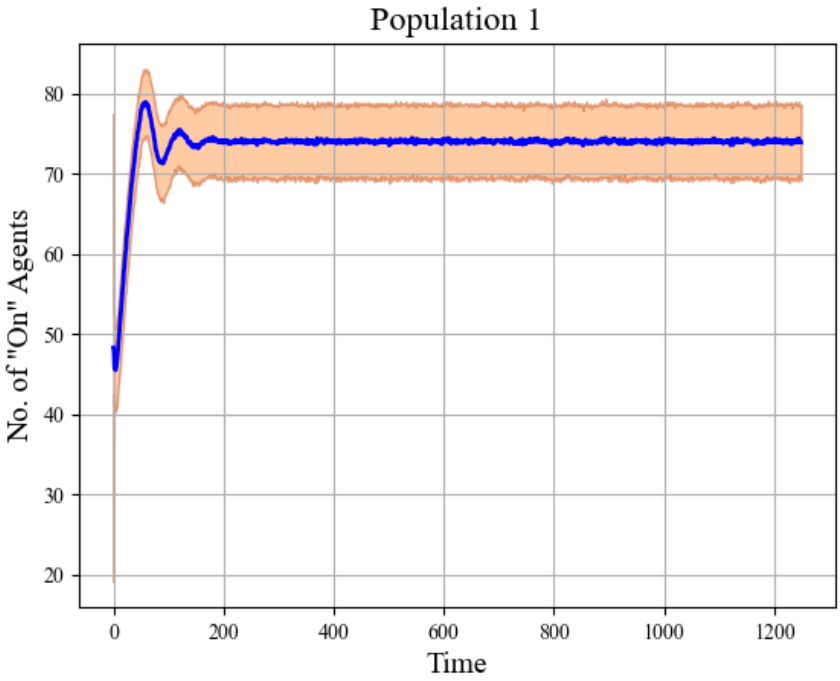

Figure 4. Evolution of "On" agents from Population 1. The mean number of "On" agents is represented by the blue line, while the orange shaded area indicates one standard deviation from the mean.

6. Applications

There has been much recent work on two-sided and multi-sided markets (e.g., by Evans and Schmalensee (2016); Hagiu and Wright (2015); G. Parker, Van Alstyne, and Jiang (2016); Weyl (2010) and most recently by Arnosti, Johari, and Kanoria (2021); Ashlagi, Braverman, Kanoria, and Shi (2020); Benjaafar and Hu (2021); Kanoria, Saban, and Sethuraman (2018); Liu, Zhang, and Zhang (2021); Möhlmann, Zalmanson, Henfridsson, and Gregory (2021)), motivated by the success of the business models of Uber Technologies, Inc. and Upwork Global Inc., following Jean Tirole's pioneering research (Rochet & Tirole, 2003, 2004a, 2004b, 2006). In general, one would like to analyse the long-run properties of such markets, including their unique ergodicity (Fioravanti et al., 2019; Ghosh et al., 2022; Kungurtsev et al., 2023; Mareček et al., 2023) and fairness (Ghosh et al., 2022; Mareček et al., 2023) to individuals.

In the case of Uber Technologies, Inc., our results are reflected in the move from

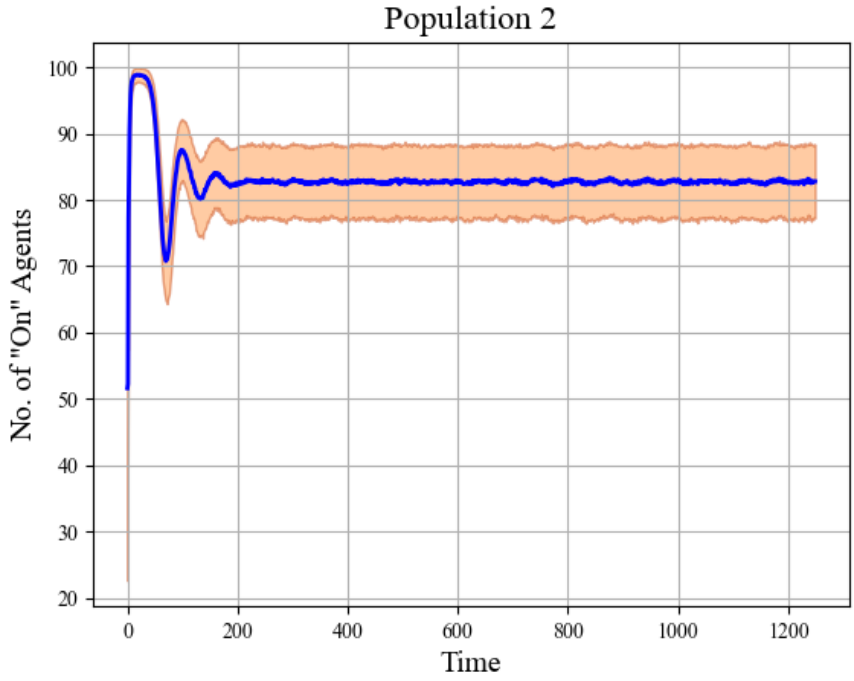

Figure 5. Evolution of "On" agents from Population 2. The mean number of "On" agents is represented by the blue line, while the orange shaded area indicates one standard deviation from the mean.

multiplicative surge pricing to additive surge pricing (Garg & Nazerzadeh, 2020). Multiplicative surge pricing is not a stable controller in the sense of Theorem 3.1, and unique ergodicity cannot be guaranteed. On the other hand, additive surge pricing is a stable controller in the sense of Theorem 3.1 and unique ergodicity can be guaranteed. See, also, Garg and Nazerzadeh (2020) for Uber's own explanation of the move, based on the multiplicative surge pricing not being incentive compatible in a dynamic setting.

Expanding upon Section 5, our simulation results can be interpreted as such a two-sided market (albeit a simplified one): controller $\mathcal{C}^1$ suggests prices π^1 (based on the distance travelled and the so-called driver surge pricing (Cachon, Daniels, & Lobel, 2017; Castillo, 2020; Castillo, Knoepfle, & Weyl, 2017; Chen, 2016; Garg & Nazerzadeh, 2020) in Uber) to customers $\mathcal{P}_i^1$, whose requests $y_i^1(k)$ for jobs (rides) at time k are based on some internal state of each customer i at time k, $x_i^1(k)$, which is not directly observable. A controller $\mathcal{C}^2$ for the other side of the market matches the jobs (rides) to workers (driver-partners) $\mathcal{P}_i^2$ whose states $x_i^2(k)$ at time k may be partially observable (e.g., availability, position) and partially not observable (e.g., appetite for further work that day). Usually (e.g., Aouad and Saritaç (2020); Araman, Calmon, and Fridgeirsdottir (2019); Özkan (2020); Simonetto, Monteil, and Gambella (2019); Yan, Zhu, Korolko, and Woodard (2020)), $\mathcal{C}^2$ is implemented using an on-line matching algorithm. Its matches are provided to workers (drivers), whose total number $y^2(k)$ of accepted matches is then filtered to obtain the proportion of empty cars on the road, which is then the input into the controller $\mathcal{C}^1$ that suggests prices with driver surge pricing implemented, if there are too few empty cars.

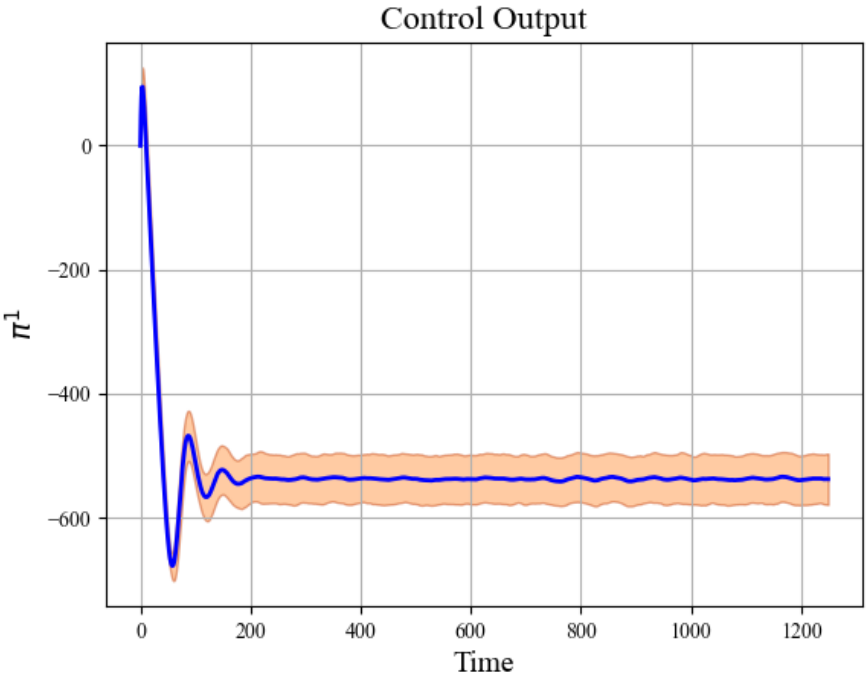

Figure 6. Output of Controller 1 versus time. The mean output is represented by the blue line, while the orange shaded area indicates one standard deviation from the mean.

Further important applications lie within multi-sided platforms (Evans & Schmalensee, 2016; Hagiu & Wright, 2015; Weyl, 2010) and networked markets (G. Parker et al., 2016). To continue our ride-hailing example, one could see the ride-hailing system of Uber as a multi-sided market if one also considers "fleet partners" (who are intermediaries for car manufacturers and car leasing providers) and taxi operators. Indeed, Uber's Vehicle Solutions Program is a platform for fleet partners to offer their vehicles to driver-partners. Drivers, in turn, sometimes happen to work also as licensed taxi drivers.

Similarly, Google's Android ecosystem and Microsoft's Windows ecosystem are sometimes seen (Hagiu & Wright, 2015; G. G. Parker, Van Alstyne, & Choudary, 2016) as three-sided platforms connecting consumers, software providers, and hardware providers. While some of the incentives offered to independent software providers (e.g., free hardware samples, no-cost licenses of development tools) are not being adjusted in real-time, others (e.g., promotions for their apps) are. The details of the control mechanisms have not been made public in this case.

7. Conclusions

In feedback control systems, a demanding area for further study is the control of ensembles of agents. In practice, an example involving multiple ensembles of agents is an online labour platform (Möhlmann et al., 2021). See, also, the work of Scheid, Boursier, Durmus, Moulines, and Jordan (2025). There are two main differences between control

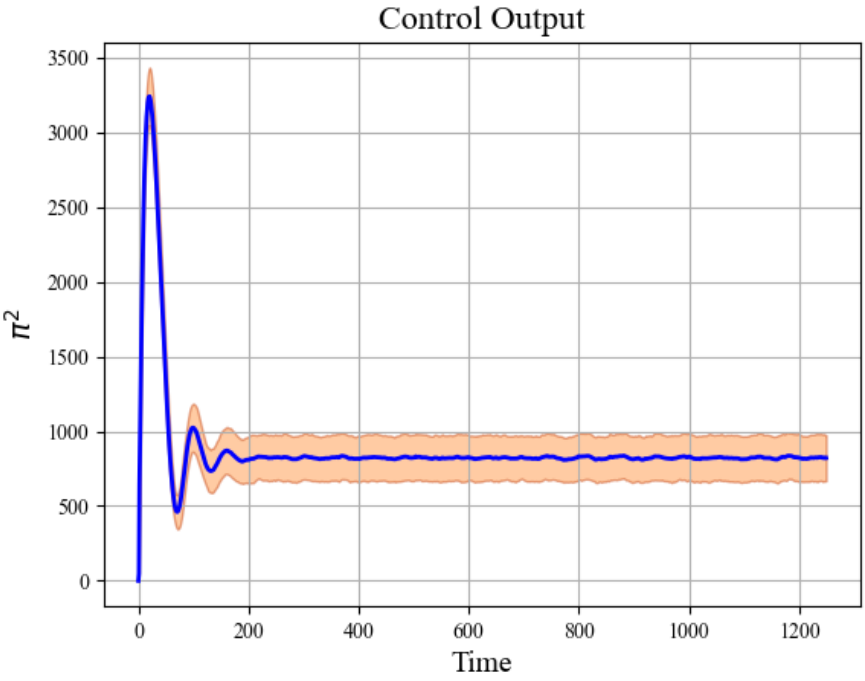

Figure 7. Output of Controller 2 versus time. The mean output is represented by the blue line, while the orange shaded area indicates one standard deviation from the mean.

problems involving ensembles versus classical control problems. First, although these ensembles generally are too large to allow for a microscopic approach, they are not sufficiently large to allow for a meaningful fluid (mean-field) approximation. Second, the regulation problem concerns the ensemble and the individual agents; a certain quality of service should be provided to each agent. We have formulated this problem as an iterated random function (Diaconis & Freedman, 1999; Ghosh & Marecek, 2022) to design an ergodic control which is the key to delivering the expected quality of service to the agents across the network.

Future work in this area can be undertaken in multiple directions. First, more elaborate use cases, compared to the example demonstrated in Section 5, can be explored. Second, of both theoretical and practical interest is a more in-depth study of time delays and any effects of these on the results. Generally speaking, it is well-known that time delays can impact a closed-loop system's stability and performance and, in the classical control sense, strategies for mitigating time delay effects remain an active area of research. Finally, while it was demonstrated in Theorems 3.1 and 4.1 that convergence in distribution to a unique invariant measure is guaranteed in regard to the closed-loop interconnections, it is worth noting that we still know relatively little in terms of further information about this convergence; for instance, how long the mixing time will be. Answers to such questions may help to inform more tailored and intelligent control design. Thus, the three aforementioned areas of future work are the subject of current explorations.

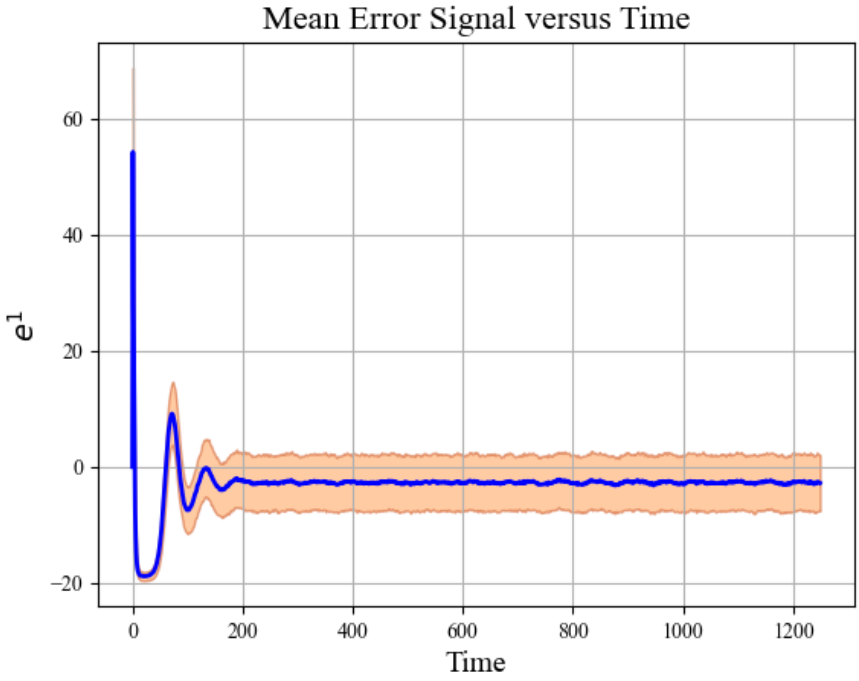

Figure 8. Input to Controller 1 versus time. The mean input is represented by the blue line, while the orange shaded area indicates one standard deviation from the mean.

Acknowledgements

This work has received funding from the European Union's Horizon Europe research and innovation programme under grant agreement no. 101070568. This work was also supported by Innovate UK under the Horizon Europe Guarantee, UKRI Reference No. 10040569 (Human-Compatible Artificial Intelligence with Guarantees (AutoFair)); Innovate UK under the Horizon Europe Guarantee, UKRI (iCircular) under the European Union's competitive HORIZON-MSCA-2021-DN-01 (Marie Sklodowska-Curie Doctoral Networks) programme under Grant Agreement No. 101073508.

References

Annaswamy, A. M., Johansson, K. H., Pappas, G. J., et al. (2023). Control for societal-scale challenges: Road map 2030. *IEEE Control Systems Society Publication: Piscataway, NJ, USA*.

Aouad, A., & Saritaç, Ö. (2020). Dynamic stochastic matching under limited time. In *Proceedings of the 21st acm conference on economics and computation* (pp. 789–790).

Araman, V. F., Calmon, A., & Fridgeirsdottir, K. (2019). Pricing and job allocation in online labor platforms. *INSEAD Working Paper No. 2019/32/TOM*.

Arnosti, N., Johari, R., & Kanoria, Y. (2021). Managing congestion in matching markets. *Manufacturing & Service Operations Management*, *23*(3), 620-636.

Ashlagi, I., Braverman, M., Kanoria, Y., & Shi, P. (2020). Clearing matching markets efficiently: informative signals and match recommendations. *Management Science*, *66*(5), 2163–2193.

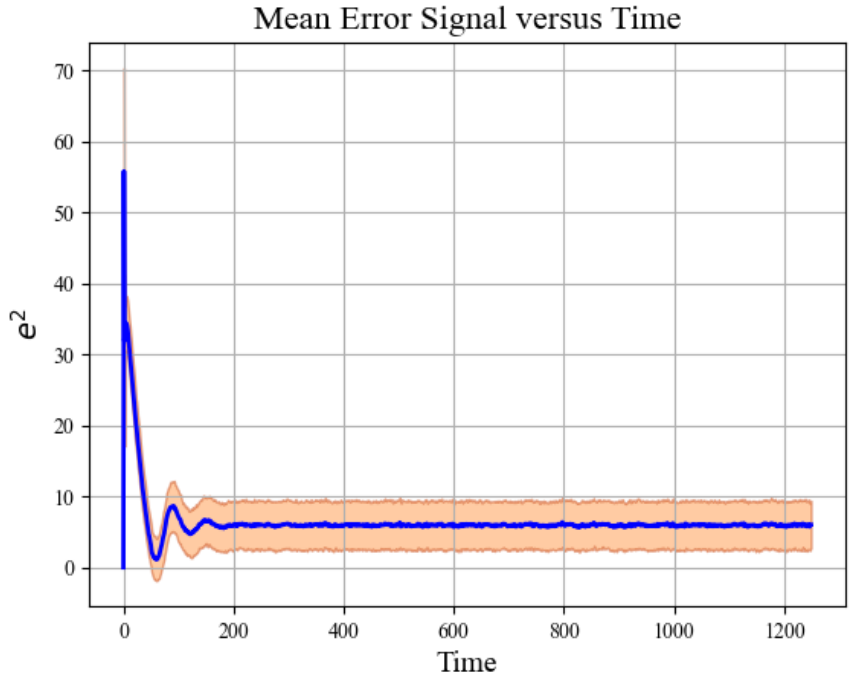

Figure 9. Input to Controller 2 versus time. The mean input is represented by the blue line, while the orange shaded area indicates one standard deviation from the mean.

Bateni, M., Chen, Y., Ciocan, D. F., & Mirrokni, V. (2022). Fair resource allocation in a volatile marketplace. *Operations Research*, *70*(1), 288–308.

Benjaafar, S., & Hu, M. (2021). Introduction to the special issue on sharing economy and innovative marketplaces. *Manufacturing & Service Operations Management*, *23*(3), 549-552. Retrieved from `https://doi.org/10.1287/msom.2021.0998`

Cachon, G. P., Daniels, K. M., & Lobel, R. (2017). The role of surge pricing on a service platform with self-scheduling capacity. *Manufacturing & Service Operations Management*, *19*(3), 368–384.

Castillo, J. C. (2020). Who benefits from surge pricing? *Available at SSRN 3245533*.

Castillo, J. C., Knoepfle, D., & Weyl, G. (2017). Surge pricing solves the wild goose chase. In *Proceedings of the 2017 acm conference on economics and computation* (pp. 241–242).

Chen, M. K. (2016). Dynamic pricing in a labor market: Surge pricing and flexible work on the uber platform. In *Proceedings of the 2016 acm conference on economics and computation* (p. 455). New York, NY, USA: Association for Computing Machinery. Retrieved from `https://doi.org/10.1145/2940716.2940798`

Diaconis, P., & Freedman, D. (1999). Iterated random functions. *SIAM Review*, *41*(1), 45–76. Retrieved from `https://doi.org/10.1137/S0036144598338446`

Evans, D. S., & Schmalensee, R. (2016). *Matchmakers: The new economics of multisided platforms*. Harvard Business Review Press.

Fioravanti, A. R., Mareček, J., Shorten, R. N., Souza, M., & Wirth, F. R. (2019). On the ergodic control of ensembles. *Automatica*, *108*.

Garg, N., & Nazerzadeh, H. (2020). Driver surge pricing. In *Proceedings of the 21st acm conference on economics and computation* (p. 501). New York, NY, USA: Association for Computing Machinery. Retrieved from `https://doi.org/10.1145/3391403.3399476`

Ghosh, R., & Marecek, J. (2022). Iterated function systems: A comprehensive survey. *arXiv preprint arXiv:2211.14661*.

Ghosh, R., Mareček, J., Griggs, W. M., Souza, M., & Shorten, R. N. (2022). Predictability and fairness in social sensing. *IEEE Internet of Things Journal*, *9*(1), 37-54.

Hagiu, A., & Wright, J. (2015). Multi-sided platforms. *International Journal of Industrial Organization*, *43*, 162–174.

Kanoria, Y., Saban, D., & Sethuraman, J. (2018). Convergence of the core in assignment markets. *Operations Research*, *66*(3), 620–636.

Kungurtsev, V., Marecek, J., Ghosh, R., & Shorten, R. (2023). On the ergodic control of ensembles in the presence of non-linear filters. *Automatica*, *152*, 110946.

Liu, Z., Zhang, D. J., & Zhang, F. (2021). Information sharing on retail platforms. *Manufacturing & Service Operations Management*, *23*(3), 606-619. Retrieved from `https://pubsonline.informs.org/doi/abs/10.1287/msom.2020.0915`

Lobel, I. (2020). Revenue management and the rise of the algorithmic economy. *Management Science*, *67*(9), 5389–5398.

Mareček, J., Roubalik, M., Ghosh, R., Shorten, R. N., & Wirth, F. R. (2023). Predictability and fairness in load aggregation and operations of virtual power plants. *Automatica*, *147*, 110743.

Möhlmann, M., Zalmanson, L., Henfridsson, O., & Gregory, R. W. (2021). Algorithmic management of work on online labor platforms: When matching meets control. *MIS Quarterly*, *45*.

Nazarov, R., et al. (2024). *Closed-loop models with a priori fairness guarantees: a toolkit.* GitHub. Retrieved from `https://github.com/humancompatible/interconnect`

Özkan, E. (2020). Joint pricing and matching in ride-sharing systems. *European Journal of Operational Research*, *287*(3), 1149–1160.

Parker, G., Van Alstyne, M., & Jiang, X. (2016). Platform ecosystems: How developers invert the firm. *MIS Quarterly*, *41*(1), 255-266.

Parker, G. G., Van Alstyne, M. W., & Choudary, S. P. (2016). *Platform revolution: How networked markets are transforming the economy and how to make them work for you.* WW Norton & Company.

Rochet, J.-C., & Tirole, J. (2003, 06). Platform Competition in Two-Sided Markets. *Journal of the European Economic Association*, *1*(4), 990-1029.

Rochet, J.-C., & Tirole, J. (2004a). *Defining two-sided markets* (Tech. Rep.). Citeseer.

Rochet, J.-C., & Tirole, J. (2004b). Two-sided markets: an overview. *Institut d'Economie Industrielle working paper*.

Rochet, J.-C., & Tirole, J. (2006). Two-sided markets: a progress report. *The RAND Journal of Economics*, *37*(3), 645–667.

Scheid, A., Boursier, E., Durmus, A., Moulines, E., & Jordan, M. (2025). Learning contracts in hierarchical multi-agent systems. *arXiv preprint arXiv:2501.19388*.

Simonetto, A., Monteil, J., & Gambella, C. (2019). Real-time city-scale ridesharing via linear assignment problems. *Transportation Research Part C: Emerging Technologies*, *101*, 208 - 232.

Weyl, E. G. (2010, September). A price theory of multi-sided platforms. *American Economic Review*, *100*(4), 1642-72. Retrieved from `https://www.aeaweb.org/articles?id=10.1257/aer.100.4.1642`

Yan, C., Zhu, H., Korolko, N., & Woodard, D. (2020). Dynamic pricing and matching in ride-hailing platforms. *Naval Research Logistics (NRL)*, *67*(8), 705–724.