

Approximation and Convergence of Large Atomic Congestion Games

Citation for published version (APA):

Cominetti, R., Scarsini, M., Schröder, M., & Stier-Moses, N. (2023). Approximation and Convergence of Large Atomic Congestion Games. *Mathematics of Operations Research*, 48(2), 784-811. <https://doi.org/10.1287/moor.2022.1281>

Document status and date:

Published: 01/05/2023

DOI:

[10.1287/moor.2022.1281](https://doi.org/10.1287/moor.2022.1281)

Document Version:

Publisher's PDF, also known as Version of record

Document license:

Taverne

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.



Mathematics of Operations Research

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Approximation and Convergence of Large Atomic Congestion Games

Roberto Cominetti, Marco Scarsini, Marc Schröder, Nicolás Stier-Moses

To cite this article:

Roberto Cominetti, Marco Scarsini, Marc Schröder, Nicolás Stier-Moses (2022) Approximation and Convergence of Large Atomic Congestion Games. Mathematics of Operations Research

Published online in Articles in Advance 31 Aug 2022

. <https://doi.org/10.1287/moor.2022.1281>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2022, INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>



Approximation and Convergence of Large Atomic Congestion Games

Roberto Cominetti,^a Marco Scarsini,^b Marc Schröder,^c Nicolás Stier-Moses^{d,*}

^aFacultad de Ingeniería y Ciencias, Universidad Adolfo Ibáñez, Santiago, Peñalolén 7941169, Región Metropolitana, Chile; ^bDipartimento di Economia e Finanza, Luiss University, 00197 Roma RM, Italy; ^cSchool of Business and Economics, Maastricht University, 6211 LK Maastricht, Netherlands; ^dCore Data Science, Meta, Menlo Park, California 94025

*Corresponding author

Contact: roberto.cominetti@uai.cl,  <https://orcid.org/0000-0001-9442-344X> (RC); marco.scarsini@luiss.it,

 <https://orcid.org/0000-0001-6473-794X> (MarcoS); m.schröder@maastrichtuniversity.nl,  <https://orcid.org/0000-0002-0048-2826> (MarcS);

nstier@utdt.edu,  <https://orcid.org/0000-0001-6380-1516> (NS-M)

Received: June 16, 2020

Revised: May 16, 2021; December 11, 2021

Accepted: April 7, 2022

Published Online in Articles in Advance:
August 31, 2022

MSC2020 Subject Classification: Primary:
91A14; secondary: 91A07, 91A43, 60B10

<https://doi.org/10.1287/moor.2022.1281>

Copyright: © 2022 INFORMS

Abstract. We consider the question of whether and in what sense, Wardrop equilibria provide a good approximation for Nash equilibria in atomic unsplittable congestion games with a large number of small players. We examine two different definitions of small players. In the first setting, we consider games in which each player’s weight is small. We prove that when the number of players goes to infinity and their weights to zero, the random flows in all (mixed) Nash equilibria for the finite games converge in distribution to the set of Wardrop equilibria of the corresponding nonatomic limit game. In the second setting, we consider an increasing number of players with a unit weight that participate in the game with a decreasingly small probability. In this case, the Nash equilibrium flows converge in total variation toward Poisson random variables whose expected values are Wardrop equilibria of a different nonatomic game with suitably defined costs. The latter can be viewed as symmetric equilibria in a Poisson game in the sense of Myerson, establishing a plausible connection between the Wardrop model for routing games and the stochastic fluctuations observed in real traffic. In both settings, we provide explicit approximation bounds, and we study the convergence of the price of anarchy. Beyond the case of congestion games, we prove a general result on the convergence of large games with random players toward Poisson games.

Funding: R. Cominetti gratefully acknowledges the support of Proyecto Anillo [Grant ANID/PIA/ACT192094]. M. Scarsini’s work was partially supported by the Gruppo Nazionale per l’Analisi Matematica, la Probabilità e le loro Applicazioni, Istituto Nazionale di Alta Matematica [Grant 2020 “Random walks on random games”] and the Progetti di Rilevante Interesse Nazionale 2017 [Grant “Algorithms, Games, and Digital Markets”]. This research also received partial support from the European Cooperation in Science and Technology [Action European Network for Game Theory].

Keywords: unsplittable atomic congestion games • nonatomic congestion games • Wardrop equilibrium • Poisson games • symmetric equilibrium • price of anarchy • price of stability • total variation

1. Introduction

Nonatomic congestion games were introduced by Wardrop [63] as a model for traffic networks with many drivers, in which each single agent has a negligible impact on congestion. The model is stated in terms of continuous flows, which are easier to analyze compared with a discrete model with a finite but large number of players. The heuristic justification is that a continuous flow model is a natural approximation to a game with many players, each one having a negligible weight. Although this argument is intuitive and plausible, the question of whether nonatomic games are the limit of atomic games is only formally addressed in special cases, mainly for atomic splittable games with homogeneous players as in Haurie and Marcotte [29].

Motivated by applications such as road traffic and telecommunications, it is also important to consider the issue of approximating *unsplittable* routing games. In them, players must route a given load over a single path, which can be chosen either deterministically or at random using a mixed strategy. We consider the more general class of *atomic unsplittable congestion games* (not necessarily routing games) and allow for heterogeneous players. The main question we address is whether Nash equilibria (NE) for these games are well-approximated by a

Wardrop equilibrium of a limiting nonatomic congestion game. A convergence result provides a stronger support for Wardrop's model as an approximation for large games with many small players, especially if we can estimate the distance between the corresponding equilibrium solutions.

In a *weighted congestion game*, there is a finite number of players who are characterized by a type and a weight. The type determines the set of feasible strategies for the player, and the weight determines the player's impact on the costs. Moreover, there is a finite set of resources, and each strategy corresponds to a subset of these resources. Players of a given type have the same set of available strategies. A strategy profile for all players induces a flow on each strategy equal to the total weight of players choosing it as well as a load on each resource equal to the total weight of players using that resource as part of their strategy. The cost of using a resource is a weakly increasing function of its load, and the cost of a strategy is additive over its resources. This defines a finite cost-minimization game. As shown by Rosenthal [53], every congestion game in which all the players have the same weight admits a potential and has equilibria in pure strategies. With heterogeneous weights, only mixed equilibria are guaranteed to exist (Nash [46]).

As a special case of congestion games, a routing game features a finite directed network whose edges represent the resources. The origin–destination (OD) pairs encode the types, and the corresponding origin–destination paths provide the strategies. To illustrate, consider a routing game over a simple network composed of two parallel edges with the same strictly increasing cost function $c(\cdot)$ as shown in Figure 1, and suppose that there are n players who need to choose an edge to route an identical weight of $w \equiv d/n$. Here, d denotes the total weight or demand. In a symmetric equilibrium, every player randomizes by choosing each route with probability $1/2$. Consequently, the number of players on each route is distributed as a $\text{Binomial}(n, 1/2)$ random variable. The total load on each edge is, therefore, d/n times a $\text{Binomial}(n, 1/2)$ and converges in distribution to $d/2$, which is precisely the Wardrop equilibrium for a total demand of d units of flow. Note that, already in this simple example, there is a multitude of other equilibria for which n_1 and n_2 players, respectively, choose the upper and lower edges for sure (with $n_1, n_2 \leq n/2$), and the remaining players (if any) randomize appropriately so as to equalize the expected cost of both routes. This includes the special case of a pure equilibrium in which half of the players take each route, up to one unit if n is odd. As the number of players n tends to infinity, all these different equilibria converge to the unique Wardrop equilibrium with a $(d/2, d/2)$ split of the flow.

In real networks, players are confronted to make decisions while facing multiple sources of uncertainty. In particular, even if the population of potential drivers might be known, the subset of drivers that are actually on the road at any given time is random. In fact, the contribution of an additional car to congestion is small but not negligible, and the congestion experienced by an agent depends basically on how many drivers are on the road at the same time. To that point, Angelidakis et al. [3] and Cominetti et al. [17] study *Bernoulli congestion games* in which players participate with an independent probability. If we focus on a small time interval, the probability that any given player participates in the game during that interval emerges naturally as a small parameter. Motivated by this model, we ask the question whether a congestion game with a large but random number of players and small participation probabilities can be approximated by a nonatomic congestion game. To address this question, we study the convergence of the Nash equilibria for a sequence of Bernoulli games toward the Wardrop equilibria of some nonatomic game. Taking the limit in this setting yields a different limit game in which the random loads on the resources converge to a family of Poisson random variables, whose expected values can be again characterized as Wardrop equilibria of a suitably defined nonatomic game or, alternatively, as an equilibrium of a Poisson game in the sense of Myerson [43]. This establishes a novel and precise connection between the Wardrop model and Poisson games. To provide some insight, consider again the parallel-edge example of Figure 1 except that now each of the n players has a unit weight but is present in the game with a small probability d/n . In this case, the effective demand is a random variable $D^n \sim \text{Binomial}(n, d/n)$, which converges as $n \rightarrow \infty$ to a random variable $D \sim \text{Poisson}(d)$. Also, in a symmetric equilibrium in which each player chooses an edge

Figure 1. (Color online) Parallel edge network.


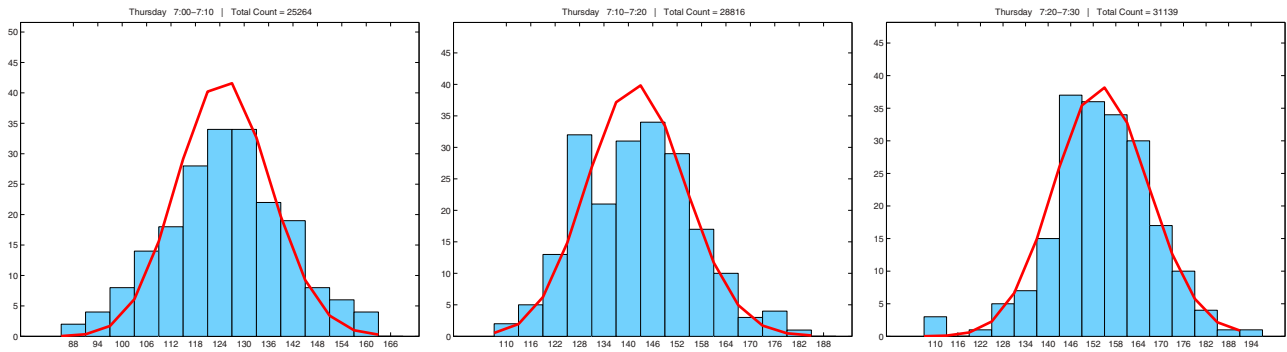


Figure 2. (Color online) Traffic counts. Dublin 2017–2018, Thursday 7:00–7:30 a.m. Data from Transport Infrastructure Ireland, <https://www.tii.ie/roads-tolling/operations-and-maintenance/traffic-count-data/>.



uniformly at random, the load on each route is distributed as a $\text{Binomial}(n, d/(2n))$ and converges to a $\text{Poisson}(d/2)$, whose expected value is again $d/2$. We show that this convergence holds for all Nash equilibria and for general congestion games with an increasing number of heterogeneous players who are active with different vanishing probabilities.

Note that, in both cases—vanishing weights and vanishing probabilities—the equilibria of the nonatomic limit games characterize the *expected values* of the random loads. However, in the Poisson regime, the resource loads remain random in the limit, whereas in the Wardrop limit these random loads converge to a constant. The fact that the Poisson limit retains some variability makes it more suitable to model the traffic flows observed in real networks. As a matter of fact, real traffic flows exhibit stochastic fluctuations that are empirically confirmed to be close to Poisson distributions, at least under moderate congestion conditions. As an illustration, Figure 2 shows the histograms of traffic counts over three consecutive 10-minute intervals, observed every Thursday during two years on a specific road segment in Dublin. The red curves give the expected counts for Poisson distributions with the same mean. These histograms reveal a persistent day-to-day variability of traffic flows so that a model predicting a random distribution is one step closer to reality compared with the point estimates provided by the Wardrop equilibrium. Our results establish a theoretically sound connection between the Wardrop and Poisson equilibria, which, combined, seem to provide a more sensible model for the traffic flows observed in real networks.

1.1. Our Contribution

After introducing the relevant classes of congestion games in Section 2, Section 3 deals with nonatomic approximations of weighted atomic congestion games. Under mild and natural conditions, we show that Nash equilibria in a sequence of weighted congestion games converge to a Wardrop equilibrium of a limiting nonatomic game. More precisely, if the number of players grows to infinity and their weights tend to zero in a way that the aggregate demands converge, then, for any sequence of mixed equilibria in the finite games, the random variables that represent the resource loads converge in distribution to a (deterministic) Wardrop equilibrium for a nonatomic limit game. We stress that players are not assumed to be symmetric and that they may have different weights and strategy sets. As long as their weights converge to zero, the random resource loads converge in distribution to some constants, which are precisely a Wardrop equilibrium for the limit game. This provides a strong support to Wardrop’s model as a sensible approximation for large unsplitable congestion games with many small players.

Section 4 focuses on approximations of Bernoulli congestion games. It considers sequences of Bayesian Nash equilibria for Bernoulli congestion games when the participation probability of players tends to zero and establishes the convergence of the resource loads toward a family of Poisson variables whose expected values are Wardrop equilibria of another nonatomic game with suitably defined costs. In Section 5, we unveil the connection between this nonatomic limit game and a Poisson game in the sense of Myerson [43], namely, the Wardrop equilibria are in one-to-one correspondence with the equilibria of an associated Poisson game with countably many players and demands distributed as Poisson random variables. We stress that Poisson games were originally introduced axiomatically and not as a limit of a sequence of finite games. We close this gap by proving a general result on convergence of sequences of games with Bernoulli players (not necessarily congestion games) toward Poisson games. This constitutes a novel alternative justification for Poisson games.

In Section 6, we derive nonasymptotic bounds for the distance between the distributions of the equilibrium loads in the finite games and the Wardrop and Poisson counterparts. These bounds provide explicit estimates for using these simpler models to approximate the equilibria in the finite games. Finally, in Section 7, we turn to investigate the convergence of the price of anarchy (PoA) and the price of stability (PoS)—as measures of the inefficiency of equilibria—for sequences of weighted congestion games and Bernoulli congestion games.

In summary, the main contributions of this paper are the following:

- a. A formal statement and proof of the fact that Nash equilibria of weighted atomic unsplittable congestion games converge to Wardrop equilibria of nonatomic congestion games as the weights of all players vanish. This is achieved under very weak conditions and holds even when weights are different (in that case, pure Nash equilibria (PNE) may not exist).
- b. A formal statement and proof of the fact that, as the participation probability of each player vanishes, the random loads in Bayesian Nash equilibria of Bernoulli congestion games converge to Poisson random variables whose expectations are Wardrop equilibria of a suitably defined nonatomic congestion game. Again, this is achieved under weak conditions; for instance, the participation probability of different players may differ.
- c. A connection between Poisson and Wardrop equilibria for congestion games and a formal asymptotic justification of Poisson games as limits of Bernoulli games.
- d. Nonasymptotic bounds for the distance between the equilibria in weighted and Bernoulli congestion games to the corresponding Wardrop and Poisson equilibria.
- e. A proof of the convergence of the PoA and PoS for sequences of weighted and Bernoulli congestion games.

1.2. Related Literature

Convergence of weighted congestion games toward nonatomic games has been considered before, mainly for the case in which players can split their weight over the available strategies. Haurie and Marcotte [29] are the first to study such convergence issues. In a setting of atomic routing games with splittable flows and elastic demands, they prove that, when players on each OD pair are identically replicated n times, in the limit when n grows, the splittable equilibria converge toward Wardrop equilibria. More recently, Jacquot and Wan [30] consider splittable routing on parallel networks with heterogeneous players, and Jacquot and Wan [31] study the approximation of nonatomic aggregative games by a sequence of finite splittable games. Also, the relation between Nash equilibria and a Wardrop-like notion of equilibrium in aggregative games with finitely many players is studied by Pacagnan et al. [50].

We are not aware of any results on convergence of equilibria for unsplittable weighted congestion games. The closest is Milchtaich [38] who studies limits of *finite crowding games* with an increasing number of players n with identical weights $1/n$, that is, unweighted congestion games with singleton strategies and type-dependent costs. Considering pure equilibria only, he establishes the convergence of the per-type aggregate strategy loads toward an equilibrium of a large crowding game. In another related result, Sandholm [57] proves that infinite potential games can be obtained as limits of finite potential games. His results are related to ours because every unweighted congestion game is a potential game (Rosenthal [53]), and conversely, every finite potential game is isomorphic to an unweighted congestion game (Monderer and Shapley [41]). The difference is that, here, we consider the more general class of weighted congestion games, which, in general, do not have a potential structure. Moreover, Sandholm [57] studies only the convergence of the potential functions and does not address the convergence of equilibria. We also mention Feldman et al. [22], who establish asymptotic upper bounds for the PoA in unsplittable games when the number n of players increases and their weight is $1/n$. Their results are based on the so-called (λ, μ) -smoothness in the large, which provides upper bounds for the PoA but, again, without addressing the convergence of equilibria. In a similar vein, Chen et al. [15] consider nonatomic congestion games as limits of atomic unsplittable games in order to define altruism and to investigate its impact on the price of anarchy.

In the present paper, we consider unsplittable weighted congestion games and present approximation results for the flows themselves, which is a stronger statement and, therefore, technically harder to accomplish. In fact, whereas equilibria for splittable congestion games (as well as pure equilibria in the unsplittable case) are conveniently described in terms of flows and loads that live in the same finite dimensional space as Wardrop equilibria, in the unsplittable case with equilibria in mixed strategies, these flows and loads become random so that limits must be properly understood in terms of convergence of random variables. This can be used later to derive the convergence of aggregate metrics, such as the PoA.

Besides routing games, auctions are another class of games in which there is significant previous work on asymptotic properties when the number of players becomes large. Satterthwaite and Williams [58] and Rustichini et al. [56] study the efficiency loss in markets as the number of bidders grows. Caserta and de Vries [14], Lambert and Shoham [34], Blumrosen and Holenstein [9], and Fibich and Gaviot [23] use asymptotic analysis

to understand the convergence of revenue in various auction settings. As in our work, several of these references compute equilibria or approximate equilibria as the auctions become larger and use those characterizations to derive limiting results. Although convergence might appear natural at first sight, some asymptotic results are fragile and only hold under carefully stated assumptions. For instance, Roberts and Postlewaite [52] show that the gain from deviating from competitive behavior in an exchange economy does not need to diminish as the number of agents in the economy goes to infinity. Also, Swinkels [61] and Feldman et al. [22] show that the inefficiency in auctions does not need to disappear in the limit with a large number of objects and players.

Several older papers consider the stochastic aspects of traffic, both theoretically and empirically, including the role of the Poisson distribution for modeling it (see, e.g., Adams [1], Breiman [11, 12], Buckley [13], Mayne [36], Miller [40], Oliver [49]). More recently, various authors study congestion games with stochastic features, focusing on the efficiency of equilibria under incomplete information. For instance, Gairing et al. [25] study the inefficiency of equilibria for congestion games in which the weight of a player is private information. Looking at the cost uncertainty, Nikolova and Stier-Moses [47] and Piliouras et al. [51] consider players' risk attitudes in nonatomic and atomic cases, respectively. Roughgarden [55] shows that whenever player types are independent, the inefficiency bounds for complete information games extend to Bayesian Nash equilibria of the incomplete information game. Wang et al. [62] and Correa et al. [18] look at similar questions for nonatomic routing games. This trend does not only include congestion games: Stidham [60] studies the efficiency of some classical queueing models on various networks, whereas Hassin et al. [28] examine a queueing model with heterogeneous agents and study how the inefficiency of equilibria varies with the intensity function. Closer to our results on Bernoulli games, Angelidakis et al. [3] consider congestion games with stochastic players who are risk-averse, restricting their attention to the case of parallel edges. In the same spirit, Cominetti et al. [17] study Bernoulli congestion games in which each player i takes part in the game independently with probability p_i and find sharp bounds for the PoA as a function of the maximum p_i .

Games with a random number of players are introduced by Myerson [42–45] with the main goal of analyzing elections with a large number of voters. In his seminal paper, Myerson [43] shows that the case in which the number of players has a Poisson distribution is of particular relevance, being the only case in which an environmental equivalence holds, that is, the belief of a player of any type about the type profile of the other active players coincides with the belief of an external game theorist. Myerson [44] deals with large Poisson games in which the parameter of the Poisson distribution diverges. His approach differs from ours in the sense that he starts with a Poisson distribution, axiomatically justified, and lets the expectation of this distribution go to infinity. In our case, we start with a finite number of players and let their number diverge in such a way that, in the limit, we get a Poisson distribution but not necessarily with a large parameter. Our derivation of Poisson games as limits of finite Bernoulli games seems to be new. The asymptotic approach we take to Poisson games is based on results from Poisson approximation theory. In probability, results about this topic abound. For an overview of this literature, we refer the reader to the books and surveys by Barbour et al. [6], Barbour and Chen [4], and Novak [48].

Several other papers study the properties of games with a random number of players and Poisson games in particular. Among them, Milchtaich [39] provides a sophisticated analysis of general games with population uncertainty. De Sinopoli and Pimienta [19], De Sinopoli et al. [20], and Meroni and Pimienta [37] deal with various properties of equilibria in Poisson games, such as stability and existence of equilibria in undominated strategies. Other papers apply Poisson games in different settings, not necessarily related to elections. Lim and Matros [35] study contests with finitely many players in which each player takes part in the game independently with the same probability, whereas Du and Gong [21] use a Poisson game to model parking congestion and propose a decentralized and coordinated online parking mechanism to reduce congestion. Let us finally mention Kordonis and Papavassilopoulos [32] and Bernhard and Deschamps [8], who study dynamic games in which players arrive at random over time.

1.3. Organization of the Paper

The paper is organized as follows. Section 2 describes the different versions of the congestion games that we study and introduces the basic notation. Sections 3 and 4 study the convergence of sequences of weighted and Bernoulli congestion games, respectively. Section 5 examines Poisson games in greater generality and focuses on convergence to them. Section 6 provides the nonasymptotic bounds on the rate of convergence of the finite congestion games to their limits, followed by Section 7 in which we present convergence results for the PoA and PoS. Section 8 presents some conclusions and possible directions for future work. Appendices A and B include auxiliary results and proofs used in Sections 6 and 7, respectively, whereas Appendix C provides a short summary

of known results on Poisson approximations for sums of Bernoulli random variables that are used to derive our results. Finally, Appendix D contains a glossary of notations for ease of reference.

2. Congestion Games: Definitions and Notations

This section introduces the basic models of congestion games and the variations that are studied in this paper. Informally, a congestion game is played over a set of resources whose costs depend on the mass of players using the resource. Each player chooses a subset of resources among the allowed subsets, seeking to achieve the minimum possible cost. Throughout this paper, we consider a finite set of resources \mathcal{E} , where each $e \in \mathcal{E}$ is associated with a weakly increasing continuous cost function $c_e : \mathbb{R}_+ \rightarrow \mathbb{R}_+$. We also fix a finite set of types \mathcal{T} , where each type $t \in \mathcal{T}$ is associated with a set of feasible actions $\mathcal{S}_t \subseteq 2^{\mathcal{E}}$, which describes the pure strategies.

To name a standard example of a congestion game, routing games capture the topology of a network structure. These games are defined over the set of edges of a finite graph, encoded by \mathcal{E} , and the set of OD pairs, encoded by \mathcal{T} . The set of actions \mathcal{S}_t contains the feasible paths for the OD pair representing type t , and the costs $c_e(\cdot)$ represent the delays when traversing an edge e .

The structural objects

$$\mathcal{G} = (\mathcal{E}, (c_e)_{e \in \mathcal{E}}, \mathcal{T}, (\mathcal{S}_t)_{t \in \mathcal{T}}) \quad (1)$$

are the same in all the congestion games considered hereafter, and the only differences are in how we describe the set of players and their behavior. In the nonatomic framework, players are considered to be infinitesimally small, and the model is stated in terms of the aggregate mass of players that use each strategy and resource. In contrast, for weighted congestion games as well as for Bernoulli congestion games, we have a finite set of players who behave strategically. The rest of this section describes these three different models precisely.

As a guide for the notation used in the sequel, we write $\Delta(\mathcal{S}_t)$ for the simplex of all probability distributions over the strategy set \mathcal{S}_t . We use capital letters for random variables and lowercase for their expected values. For instance, X_e represents a random load on a resource $e \in \mathcal{E}$ with expected value $x_e = \mathbb{E}[X_e]$, and we add a hat \hat{x}_e when referring to an equilibrium.

2.1. Nonatomic Congestion Games

A nonatomic congestion game Γ^∞ is given by a pair $(\mathcal{G}, \mathbf{d})$, where \mathcal{G} stands for the structural objects of the game as in (1), and $\mathbf{d} = (d_t)_{t \in \mathcal{T}}$ is a vector of demands with $d_t \geq 0$ representing the aggregate demand of type t . The *total demand* is given by the sum over all types $d_{\text{tot}} = \sum_{t \in \mathcal{T}} d_t$.

Each demand d_t is split over the corresponding strategies \mathcal{S}_t and induces loads on the resources. Specifically, a *strategy flow* vector $\mathbf{y} := (y_{t,s})_{t \in \mathcal{T}, s \in \mathcal{S}_t}$ and a *resource load* vector $\mathbf{x} := (x_e)_{e \in \mathcal{E}}$ are called *feasible* if they satisfy the following constraints:

$$\forall t \in \mathcal{T}, \quad d_t = \sum_{s \in \mathcal{S}_t} y_{t,s} \text{ with } y_{t,s} \geq 0, \quad \text{and} \quad \forall e \in \mathcal{E}, \quad x_e = \sum_{t \in \mathcal{T}} \sum_{s \in \mathcal{S}_t} y_{t,s} \mathbb{1}_{\{e \in s\}}. \quad (2)$$

The set of such feasible *flow-load pairs* (\mathbf{y}, \mathbf{x}) is denoted by $\mathcal{F}(\mathbf{d})$. Note that the resource loads \mathbf{x} are uniquely defined by the strategy flows \mathbf{y} but not vice versa. Nevertheless, instead of only considering strategy flows, we refer to flow-load pairs because some concepts are easier to express in terms of flows, whereas others are defined in terms of loads. The notation and nomenclature are inspired by routing games, and most of our examples are of this type because they are intuitive and well-studied. However, we use these terms in the more general setting of congestion games even when there is no network structure of resources and strategies.

A *Wardrop equilibrium* is defined as a feasible flow-load pair $(\hat{\mathbf{y}}, \hat{\mathbf{x}}) \in \mathcal{F}(\mathbf{d})$ for which the prevailing cost of all used strategies is minimal or, mathematically,

$$\forall t \in \mathcal{T}, \quad \forall s, s' \in \mathcal{S}_t \quad \hat{y}_{t,s} > 0 \Rightarrow \sum_{e \in s} c_e(\hat{x}_e) \leq \sum_{e \in s'} c_e(\hat{x}_e). \quad (3)$$

Because only strategies with minimum cost are used, all the strategies used by any given type have the same cost, and as a consequence, any strategy flow decomposition of $\hat{\mathbf{x}}$ yields an equilibrium. The set of Wardrop equilibria of the nonatomic game Γ^∞ is denoted by $\text{WE}(\Gamma^\infty)$.

2.2. Weighted Congestion Games

A weighted congestion game is defined as a tuple $\Gamma_W = (\mathcal{G}, (w_i, t_i)_{i \in \mathcal{N}})$, where \mathcal{N} is a finite set of players and each player $i \in \mathcal{N}$ has a weight $w_i \in \mathbb{R}_+$ and a type $t_i \in \mathcal{T}$ that determines the player's strategy set \mathcal{S}_{t_i} . Our use of the term "type" slightly differs from what is common in game theory, in which a type is a random variable

associated to each player, whose distribution is common knowledge but whose realization is private information. Here, the players' types are deterministic, and the type of each player is common knowledge.

The aggregate demand for each type t and the total demand are given by

$$d_t = \sum_{i:t_i=t} w_i, \quad \text{and} \quad d_{\text{tot}} = \sum_{t \in \mathcal{T}} d_t = \sum_{i \in \mathcal{N}} w_i. \quad (4)$$

Let $\sigma = (\sigma_i)_{i \in \mathcal{N}}$ be a *mixed strategy profile*, in which $\sigma_i \in \Delta(\mathcal{S}_i)$ represents the mixed strategy used by player $i \in \mathcal{N}$, and let $\Sigma := \times_{i \in \mathcal{N}} \Delta(\mathcal{S}_i)$ be the set of mixed strategy profiles. Call \mathbf{E}_σ the expectation with respect to the product probability measure $\mathbf{P}_\sigma := \times_{i \in \mathcal{N}} \sigma_i$ over the set of pure strategy profiles $\mathcal{S} = \times_{i \in \mathcal{N}} \mathcal{S}_i$. If S_i is the random strategy of player i , whose distribution is σ_i , then the probability that player i uses a given resource e is

$$\sigma_{i,e} := \mathbf{P}_\sigma(e \in S_i) = \sum_{s \in \mathcal{S}_i} \sigma_i(s) \mathbb{1}_{\{e \in s\}}. \quad (5)$$

Accordingly, the strategy flows $Y_{t,s}$ and the resource loads X_e become random variables

$$Y_{t,s} = \sum_{i:t_i=t} w_i \mathbb{1}_{\{S_i=s\}} \quad \text{and} \quad X_e = \sum_{i \in \mathcal{N}} w_i \mathbb{1}_{\{e \in S_i\}}, \quad (6)$$

given by the random realizations of S_i . Their expected values are

$$y_{t,s} := \mathbf{E}_\sigma[Y_{t,s}] = \sum_{i:t_i=t} w_i \sigma_i(s) \quad \text{and} \quad x_e := \mathbf{E}_\sigma[X_e] = \sum_{i \in \mathcal{N}} w_i \sigma_{i,e}.$$

A straightforward calculation shows that the pair (\mathbf{y}, \mathbf{x}) satisfies (2) so that $(\mathbf{y}, \mathbf{x}) \in \mathcal{F}(\mathbf{d})$.

To take the perspective of a fixed player i , we assume the player already selected resource e and define the conditional load

$$X_{i,e} := w_i + \sum_{j \neq i} w_j \mathbb{1}_{\{e \in S_j\}}. \quad (7)$$

Using this, the expected cost of player i conditional on the player using the resource e is

$$\mathbf{E}_\sigma[c_e(X_e) \mid e \in S_i] = \mathbf{E}_\sigma[c_e(X_{i,e})].$$

A strategy profile $\hat{\sigma} \in \Sigma$ is a *mixed Nash equilibrium* if

$$\forall i \in \mathcal{N}, \quad \forall s, s' \in \mathcal{S}_i, \quad \hat{\sigma}_i(s) > 0 \Rightarrow \sum_{e \in s} \mathbf{E}_{\hat{\sigma}}[c_e(X_{i,e})] \leq \sum_{e \in s'} \mathbf{E}_{\hat{\sigma}}[c_e(X_{i,e})].$$

The set of mixed Nash equilibria of Γ_W is denoted by $\text{MNE}(\Gamma_W)$.

When all the players have the same weight $w_i \equiv w$, Rosenthal [53] proves that Γ_W is a potential game, and as a consequence, pure Nash equilibria are guaranteed to exist (see also Monderer and Shapley [41]). Also, Fotakis et al. [24] show that every weighted congestion game with affine costs admits an exact potential. Conversely, Harks et al. [27] prove that, if \mathcal{C} is a class of cost functions such that every weighted congestion game with costs in \mathcal{C} admits a potential, then \mathcal{C} only contains affine functions. Existence of pure equilibria in weighted congestion games is further studied by Harks and Klimm [26]. Beyond these cases, one can only guarantee the existence of equilibria in mixed strategies (Nash [46]).

2.3. Bernoulli Congestion Games

In a weighted congestion game, the randomness arises only from the players' mixed strategies. In this section, we add another stochastic element: players may not be present in the game. A Bernoulli congestion game is a congestion game in which each player $i \in \mathcal{N}$ has a unit weight $w_i \equiv 1$ but takes part in the game only with some probability $u_i \in (0, 1]$ and otherwise remains inactive and incurs no cost. The participation events, assumed to be independent, are encoded in random variables $U_i \sim \text{Bernoulli}(u_i)$, which indicate whether player $i \in \mathcal{N}$ is active or not. A Bernoulli congestion game is denoted by $\Gamma_B = (\mathcal{G}, (u_i, t_i)_{i \in \mathcal{N}})$.

The framework is similar to a weighted congestion game in which the w_i 's are replaced by random weights $U_i \in \{0, 1\}$ with expected value u_i so that the per-type demands become the random variables $D_t = \sum_{i:t_i=t} U_i$ with expected values $d_t = \mathbf{E}[D_t] = \sum_{i:t_i=t} u_i$. The formulas are, therefore, very similar with w_i replaced by U_i or by u_i when taking expectations. Nevertheless, later we see that the two classes of games behave quite differently in some respects.

Let $\sigma = (\sigma_i)_{i \in \mathcal{N}} \in \Sigma$ be a profile of mixed strategies. We assume that each player chooses a mixed strategy before the actual realization of these random variables so that no player knows for sure who is present in the

game. Now, randomness is induced by both the random participation and these mixed strategies. This is described by a discrete probability space $(\Omega, 2^\Omega, \mathbf{P}_\sigma)$, where $\Omega = \{0, 1\}^{\mathcal{N}} \times \mathcal{S}$ with $\mathcal{S} = \times_{i \in \mathcal{N}} \mathcal{S}_i$ as before, and \mathbf{P}_σ is now the probability measure induced by σ and by the random participation of players; that is, for $\omega \in \{0, 1\}^{\mathcal{N}}$ and $s \in \mathcal{S}$, we have

$$\mathbf{P}_\sigma(\omega, s) = \prod_{i \in \mathcal{N}} \mathbf{P}_i(\omega_i) \sigma_i(s_i),$$

with $\mathbf{P}_i(1) = u_i$ and $\mathbf{P}_i(0) = 1 - u_i$. The corresponding expectation operator is denoted by \mathbf{E}_σ .

As before, $\mathbb{1}_{\{e \in S_i\}}$ is a Bernoulli random variable indicating whether the random strategy S_i includes resource e with (5) still in place. Additionally, let $U_{i,e} = U_i \mathbb{1}_{\{e \in S_i\}}$ indicate whether player i is active and chooses resource e , for which we have $\mathbf{E}_\sigma[U_{i,e}] = u_i \sigma_{i,e}$. Then, the total number of active players of type $t \in \mathcal{T}$ using strategy $s \in \mathcal{S}_t$ and the total load on resource $e \in \mathcal{E}$ are now the random variables

$$Y_{t,s} = \sum_{i:t_i=t} U_i \mathbb{1}_{\{S_i=s\}}, \quad \text{and} \quad X_e = \sum_{i \in \mathcal{N}} U_{i,e}. \quad (8)$$

The expected strategy flows and resource loads are

$$y_{t,s} := \mathbf{E}_\sigma[Y_{t,s}] = \sum_{i:t_i=t} u_i \sigma_i(s), \quad \text{and} \quad x_e := \mathbf{E}_\sigma[X_e] = \sum_{i \in \mathcal{N}} u_i \sigma_{i,e}. \quad (9)$$

Once again, the pair (\mathbf{y}, \mathbf{x}) satisfies (2) so that $(\mathbf{y}, \mathbf{x}) \in \mathcal{F}(\mathbf{d})$.

When $U_{i,e} = 1$, conditional on player i selecting resource $e \in \mathcal{E}$, its load is $X_{i,e} = 1 + Z_{i,e}$, where $Z_{i,e}$ represents the number of other players using that resource, that is,

$$Z_{i,e} = \sum_{j \neq i} U_{j,e}. \quad (10)$$

Then, the conditional expected cost for player $i \in \mathcal{N}$ when using this resource is

$$\mathbf{E}_\sigma[c_e(X_e) \mid U_{i,e} = 1] = \mathbf{E}_\sigma[c_e(X_{i,e})] = \mathbf{E}_\sigma[c_e(1 + Z_{i,e})]. \quad (11)$$

Note that, in this setting, all players have unit weight so that the resource loads are integer-valued, and therefore, the costs $c_e : \mathbb{N}^* \rightarrow \mathbb{R}_+$ need only be defined over the positive integers $\mathbb{N}^* = \mathbb{N} \setminus \{0\}$.

A strategy profile $\hat{\sigma} \in \Sigma$ is a Bayesian Nash equilibrium if

$$\forall i \in \mathcal{N}, \forall s, s' \in \mathcal{S}_i \quad \hat{\sigma}_i(s) > 0 \Rightarrow \sum_{e \in \mathcal{S}} \mathbf{E}_{\hat{\sigma}}[c_e(1 + Z_{i,e})] \leq \sum_{e \in \mathcal{S}'} \mathbf{E}_{\hat{\sigma}}[c_e(1 + Z_{i,e})].$$

The set of all Bayesian Nash equilibria of Γ_B is denoted by $\text{BNE}(\Gamma_B)$.

Remark 1. Cominetti et al. [17, proposition 3.3] show that Γ_B is a potential game, which drives the class of Bernoulli congestion games apart from the class of weighted congestion games, which admit a potential only in special cases. In particular, a Bernoulli congestion game always has equilibria in pure strategies. However, here we consider both pure and mixed equilibria.

3. Convergence of Weighted Congestion Games

Now that we laid out the games we are considering, we proceed to study the convergence of weighted to nonatomic congestion games when the number of players increases and their weights decrease. The only stochastic element appearing in this section is the fact that players randomize by considering mixed strategies. Under mild and natural conditions, we prove that the equilibrium resource loads converge in distribution to constants. These constants are equal to the resource loads prevailing under a Wardrop equilibrium of the corresponding nonatomic game.

We consider a sequence of weighted congestion games

$$\Gamma_W^n = (\mathcal{G}, (w_i^n, t_i^n)_{i \in \mathcal{N}^n})$$

in which the elements that vary over the sequence are the set \mathcal{N}^n of players, their weights w_i^n , and their types t_i^n . We study if and how the equilibria for this sequence converge. All the notations in Section 2.2 remain in place by simply adding the superscript n . We assume that the number of players goes to infinity and the sequence of weights goes to zero in such a way that the aggregate demand for each type converges. In other words, as there are more players, no player becomes dominant, and the demands remain bounded. This is captured by the

following asymptotic behavior as $n \rightarrow \infty$:

$$|\mathcal{N}^n| \rightarrow \infty, \tag{12a}$$

$$w^n := \max_{i \in \mathcal{N}^n} w_i^n \rightarrow 0, \tag{12b}$$

$$d_t^n := \sum_{i: t_i^n=t} w_i^n \rightarrow d_t. \tag{12c}$$

Under these conditions, the random loads at equilibrium for the sequence of games Γ_W^n converge in distribution to the loads of a Wardrop equilibrium of the corresponding nonatomic limit game.

Theorem 1. Let Γ_W^n be a sequence of weighted congestion games satisfying the assumptions in (12), and let $\hat{\sigma}^n \in \text{MNE}(\Gamma_W^n)$ be an arbitrary sequence of mixed Nash equilibria. Then, the corresponding sequence of expected flow-load pairs $(\hat{\mathbf{y}}^n, \hat{\mathbf{x}}^n)$ is bounded and every accumulation point $(\hat{\mathbf{y}}, \hat{\mathbf{x}})$ is a Wardrop equilibrium of the nonatomic congestion game $\Gamma^\infty = (\mathcal{G}, (d_t)_{t \in \mathcal{T}})$. Furthermore, along every convergent subsequence, the random flow-load pairs $(\mathbf{Y}^n, \mathbf{X}^n)$ converge in distribution toward $(\hat{\mathbf{y}}, \hat{\mathbf{x}})$.

Proof. As for every mixed strategy profile, the expected flow-load pair $(\hat{\mathbf{y}}^n, \hat{\mathbf{x}}^n)$ belongs to $\mathcal{F}(\mathbf{d}^n)$ so that we have (2), adding the superscript n to all the terms involved. Because $d_t^n \rightarrow d_t$, it follows from (2) that the sequence $(\hat{\mathbf{y}}^n, \hat{\mathbf{x}}^n)$ is bounded, and then, passing to the limit in these equations, we conclude that every accumulation point $(\hat{\mathbf{y}}, \hat{\mathbf{x}})$ belongs to $\mathcal{F}(\mathbf{d})$.

Now, take a convergent subsequence and, for simplicity, rename it so that the full sequence converges $(\hat{\mathbf{y}}^n, \hat{\mathbf{x}}^n) \rightarrow (\hat{\mathbf{y}}, \hat{\mathbf{x}})$. From (6), we get

$$\text{Var}_{\hat{\sigma}^n}[X_e^n] = \text{Var}_{\hat{\sigma}^n} \left[\sum_{i \in \mathcal{N}^n} w_i^n \mathbb{1}_{\{e \in S_i^n\}} \right] = \sum_{i \in \mathcal{N}^n} (w_i^n)^2 \hat{\sigma}_{i,e}^n (1 - \hat{\sigma}_{i,e}^n) \leq \frac{1}{4} \sum_{i \in \mathcal{N}^n} (w_i^n)^2.$$

It can be shown similarly that

$$\text{Var}_{\hat{\sigma}^n}[Y_{t,s}^n] \leq \frac{1}{4} \sum_{i \in \mathcal{N}^n} (w_i^n)^2.$$

Conditions (12), in turn, imply

$$\sum_{i \in \mathcal{N}^n} (w_i^n)^2 \leq w^n \sum_{i \in \mathcal{N}^n} w_i^n = w^n \sum_{t \in \mathcal{T}} d_t^n \rightarrow 0,$$

so that $\text{Var}_{\hat{\sigma}^n}[Y_{t,s}^n] \rightarrow 0$ and $\text{Var}_{\hat{\sigma}^n}[X_e^n] \rightarrow 0$, from which convergence in distribution follows.

It remains to show that $(\hat{\mathbf{y}}, \hat{\mathbf{x}})$ is a Wardrop equilibrium; that is, we need to establish (3). If $\hat{y}_{t,s} > 0$, then for n large enough, we have $\hat{y}_{t,s}^n > 0$, and there exists some player $i \in \mathcal{N}^n$ of type $t_i^n = t$ with $\hat{\sigma}_i^n(s) > 0$. Note that i actually depends on n , so we should write i^n . For the sake of simplicity, we omit the superscript. The equilibrium condition in Γ_W^n implies that, for each alternative strategy $s' \in \mathcal{S}_t$ and for this player i ,

$$\sum_{e \in \mathcal{S}} \mathbb{E}_{\hat{\sigma}^n}[c_e(X_{i,e}^n)] \leq \sum_{e \in \mathcal{S}'} \mathbb{E}_{\hat{\sigma}^n}[c_e(X_{i,e}^n)]. \tag{13}$$


Because $|X_{i,e}^n - X_e^n| \leq w_i^n \leq w^n \rightarrow 0$, it follows that $X_{i,e}^n \rightarrow \hat{x}_e$ in distribution. Moreover, the loads $X_e^n \geq 0$ are bounded above by the total demands $d_{\text{tot}}^n = \sum_{t \in \mathcal{T}} d_t^n$, which converge to $d_{\text{tot}} = \sum_{t \in \mathcal{T}} d_t$ so that both X_e^n and $X_{i,e}^n$ are uniformly bounded. It follows that $\mathbb{E}_{\hat{\sigma}^n}[c_e(X_{i,e}^n)] \rightarrow c_e(\hat{x}_e)$, and letting $n \rightarrow \infty$ in (13), we obtain (3) as required. \square

Example 1. We illustrate the previous result on the Wheatstone network shown in Figure 3. There is a single OD pair and $n \geq 2$ identical players, each one with weight $w_i \equiv 1/n$ so that the total demand is $d_{\text{tot}} = 1$. For each n , there is a unique symmetric pure Nash equilibrium in which all players take the zigzag path e_1, e_3, e_5 and pay a cost equal to two. For $n = 2$, we also have a pure equilibrium in which one player takes the upper path, whereas the other takes the lower path, and they both pay $3/2$. For $n \geq 3$, the only mixed equilibrium (modulo permutation of the players) is when $n - 1$ players take the zigzag path and the last player mixes in any possible way over the three paths. In the limit when $n \rightarrow \infty$, all these equilibria converge to the Wardrop equilibrium of the corresponding nonatomic game with a unit flow over the zigzag path.

For games with strictly increasing costs $c_e(\cdot)$, the Wardrop equilibrium loads \hat{x}_e are unique, which yields the following direct consequence.

Corollary 1. Suppose that the resource costs $c_e(\cdot)$ are strictly increasing. Then, for every sequence Γ_W^n of weighted congestion games satisfying (2) and each $\hat{\sigma}^n \in \text{MNE}(\Gamma_W^n)$, the random loads X_e^n converge in distribution to the unique Wardrop equilibrium loads \hat{x}_e of the nonatomic limit game $\Gamma^\infty = (\mathcal{G}, (d_t)_{t \in \mathcal{T}})$.

Figure 3. (Color online) Wheatstone network.



As a complement to these results, in Section 6.1, we derive explicit bounds on the approximation and rate of convergence under additional conditions on the costs. Also, in Section 7 we prove that the PoA for the sequence of weighted congestion games Γ_W^n converges to the PoA of the limit game Γ^∞ .

4. Convergence of Bernoulli Congestion Games

In this section, we study the convergence of Bernoulli congestion games to nonatomic games when the number of players increases and their participation probabilities decrease. We show that the equilibrium flows and loads converge in total variation toward Poisson random variables whose expected values are a Wardrop equilibrium of a corresponding nonatomic game with suitably defined cost functions. Nevertheless, we highlight that our results prove also that the distributions of the random loads converge.

We proceed to study the convergence of Bayesian Nash equilibria $\hat{\sigma}^n \in \text{BNE}(\Gamma_B^n)$ for a sequence of Bernoulli congestion games

$$\Gamma_B^n = (\mathcal{G}, (u_i^n, t_i^n)_{i \in \mathcal{N}^n}).$$

The elements that vary over this sequence are the set \mathcal{N}^n of players, their participation probabilities u_i^n , and their types t_i^n as well as the underlying probability spaces encoding the random participation of players and their mixed strategies. We assume that the number of players goes to infinity while the sequence of participation probabilities goes to zero in such a way that the aggregate expected demand for each type converges. This is captured by the following asymptotic behavior as $n \rightarrow \infty$:

$$|\mathcal{N}^n| \rightarrow \infty, \tag{14a}$$

$$u^n := \max_{i \in \mathcal{N}^n} u_i^n \rightarrow 0, \tag{14b}$$

$$d_t^n := \sum_{i: t_i^n = t} u_i^n \rightarrow d_t. \tag{14c}$$

Our analysis is based on results on Poisson approximations for sums of Bernoulli random variables. In particular, under (14), the per-type random demands $D_t^n = \sum_{i: t_i^n = t} U_i^n$ are known to converge in total variation to a Poisson variable $D_t \sim \text{Poisson}(d_t)$ (see, e.g., Adell and Lekuona [2, corollary 3.1]). We recall that a sequence of probability measures \mathbf{Q}^n on \mathbb{N} converges in total variation to the probability measure \mathbf{Q} whenever $\rho_{\text{TV}}(\mathbf{Q}^n, \mathbf{Q}) \rightarrow 0$, where

$$\rho_{\text{TV}}(\mathbf{Q}^n, \mathbf{Q}) := \sup_{A \subset \mathbb{N}} |\mathbf{Q}^n(A) - \mathbf{Q}(A)| = \frac{1}{2} \sum_{k \in \mathbb{N}} |\mathbf{Q}^n(k) - \mathbf{Q}(k)|. \tag{15}$$

Similarly, a sequence of integer-valued random variables T^n converges in total variation to T if their distributions satisfy $\rho_{\text{TV}}(\mathcal{L}(T^n), \mathcal{L}(T)) \rightarrow 0$. Convergence in total variation is stronger than convergence in distribution—which was the concept used in Section 3—and is suitable for situations in which the limit distribution is discrete. Appendix C collects further results on Poisson approximations that are used in our subsequent analysis.

Proposition 1. *Let Γ_B^n be a sequence of Bernoulli congestion games satisfying the conditions in (14) and σ^n an arbitrary sequence of mixed strategy profiles. Let $Y_{i,s}^n$ and X_e^n be the random flows and loads given by (8) with $U_i = U_i^n$, $t_i = t_i^n$, and $S_i = S_i^n$ drawn independently at random according to σ_i^n , and let $(\mathbf{y}^n, \mathbf{x}^n)$ be the vector of expected flows and loads. Then,*

- a. *The sequence $(\mathbf{y}^n, \mathbf{x}^n)_{n \in \mathbb{N}}$ is bounded, and each accumulation point (\mathbf{y}, \mathbf{x}) belongs to $\mathcal{F}(\mathbf{d})$.*

- b. Along any convergent subsequence of $(\mathbf{y}^n, \mathbf{x}^n)$, the random flows $Y_{t,s}^n$ and loads X_e^n converge in total variation to $Y_{t,s} \sim \text{Poisson}(y_{t,s})$ and $X_e \sim \text{Poisson}(x_e)$, respectively.
 c. The limit Poisson variables $(Y_{t,s})_{t \in \mathcal{T}, s \in \mathcal{S}_t}$ are independent.

Proof.

- a. Because the expected demands d_t^n are convergent, this follows directly from (9), which in this case, reads

$$d_t^n = \sum_{s \in \mathcal{S}_t} y_{t,s}^n \quad \text{and} \quad x_e^n = \sum_{t \in \mathcal{T}} \sum_{s \in \mathcal{S}_t} y_{t,s}^n \mathbb{1}_{\{e \in \mathcal{S}\}}.$$

- b. Take a convergent subsequence and for simplicity rename it to be the full sequence $(\mathbf{y}^n, \mathbf{x}^n) \rightarrow (\mathbf{y}, \mathbf{x})$. Considering a Poisson variable $V_e^n \sim \text{Poisson}(x_e^n)$ and using Theorem C.1(a), we have

$$\rho_{\text{TV}}(\mathcal{L}(X_e^n), \mathcal{L}(V_e^n)) \leq u^n,$$

so that, using the triangle inequality and (C.1), we conclude

$$\rho_{\text{TV}}(\mathcal{L}(X_e^n), \mathcal{L}(X_e)) \leq u^n + |x_e^n - x_e| \rightarrow 0. \tag{16}$$

A similar argument shows that $Y_{t,s}^n$ converges in total variation to $Y_{t,s}$.

- c. Consider the joint moment-generating function M^n of the random variables $Y_{t,s}^n$,

$$M^n(\boldsymbol{\lambda}) = \mathbb{E}_{\boldsymbol{\sigma}^n} \left[\exp \left(\sum_{t \in \mathcal{T}} \sum_{s \in \mathcal{S}_t} \lambda_{t,s} Y_{t,s}^n \right) \right],$$

where $\boldsymbol{\lambda} = (\lambda_{t,s})_{t \in \mathcal{T}, s \in \mathcal{S}_t}$. Using the fact that

$$Y_{t,s}^n = \sum_{i: t_i^n = t} U_i^n \mathbb{1}_{\{S_i^n = s\}},$$

with $U_i^n \mathbb{1}_{\{S_i^n = s\}}$ independent across players and types (although not across strategies), we have

$$\begin{aligned} M^n(\boldsymbol{\lambda}) &= \mathbb{E}_{\boldsymbol{\sigma}^n} \left[\exp \left(\sum_{t \in \mathcal{T}} \sum_{s \in \mathcal{S}_t} \sum_{i: t_i^n = t} \lambda_{t,s} U_i^n \mathbb{1}_{\{S_i^n = s\}} \right) \right] \\ &= \prod_{t \in \mathcal{T}} \prod_{i: t_i^n = t} \mathbb{E}_{\boldsymbol{\sigma}^n} \left[\exp \left(U_i^n \sum_{s \in \mathcal{S}_t} \lambda_{t,s} \mathbb{1}_{\{S_i^n = s\}} \right) \right] \\ &= \prod_{t \in \mathcal{T}} \prod_{i: t_i^n = t} \left(1 - u_i^n + u_i^n \mathbb{E}_{\boldsymbol{\sigma}^n} \left[\exp \left(\sum_{s \in \mathcal{S}_t} \lambda_{t,s} \mathbb{1}_{\{S_i^n = s\}} \right) \right] \right) \\ &= \prod_{t \in \mathcal{T}} \prod_{i: t_i^n = t} \left(1 - u_i^n + u_i^n \sum_{s \in \mathcal{S}_t} \sigma_i^n(s) \exp(\lambda_{t,s}) \right). \end{aligned}$$

Taking logarithms and using the fact that $\ln(1 + y) = y + O(y^2)$, it follows that

$$\begin{aligned} \ln M^n(\boldsymbol{\lambda}) &= \sum_{t \in \mathcal{T}} \sum_{i: t_i^n = t} \ln \left(1 + u_i^n \sum_{s \in \mathcal{S}_t} \sigma_i^n(s) [\exp(\lambda_{t,s}) - 1] \right) \\ &= \sum_{t \in \mathcal{T}} \sum_{i: t_i^n = t} \left(\sum_{s \in \mathcal{S}_t} u_i^n \sigma_i^n(s) [\exp(\lambda_{t,s}) - 1] + O((u_i^n)^2) \right) \\ &= \sum_{t \in \mathcal{T}} \sum_{s \in \mathcal{S}_t} y_{t,s}^n [\exp(\lambda_{t,s}) - 1] + \sum_{i \in \mathcal{N}^n} O((u_i^n)^2) \end{aligned}$$

In view of the conditions in (14), the sum $\sum_{i \in \mathcal{N}^n} O((u_i^n)^2)$ converges to zero, and therefore,

$$\lim_{n \rightarrow \infty} M^n(\boldsymbol{\lambda}) = \prod_{t \in \mathcal{T}} \prod_{s \in \mathcal{S}_t} \exp(y_{t,s} [\exp(\lambda_{t,s}) - 1]),$$

which is the moment-generating function of a family of independent Poisson random variables with parameters $y_{t,s}$. \square

Remark 2. Even if the strategy flows $Y_{t,s}$ are independent, in general, the resource loads X_e are not. For instance, in a trivial network with a single OD pair connected by a unique path, all the edges carry exactly the same load.

With these preliminary facts, we proceed to introduce the nonatomic congestion game $\check{\Gamma}^\infty$ that characterizes the limit of the sequence Γ_B^n . Recall that the costs $c_e: \mathbb{N}^* \rightarrow \mathbb{R}_+$ in Bernoulli games are only defined over the positive integers. The game $\check{\Gamma}^\infty = (\check{\mathcal{G}}, (d_t)_{t \in \mathcal{T}})$ is given by the limiting demands d_t and the auxiliary costs $\check{c}_e: \mathbb{R}_+ \rightarrow \mathbb{R}_+ \cup \{+\infty\}$, where $\check{c}_e(x)$ is defined for $x \geq 0$ by taking a random variable $X \sim \text{Poisson}(x)$ and setting

$$\check{c}_e(x) := \mathbb{E}[c_e(1 + X)]. \quad (17)$$

Notice that, similar to (11), the +1 is because a player must account for the player's own presence on the resource. To ensure that these expected costs are well-defined and smooth, we impose the following mild condition:

$$\exists \nu \in \mathbb{N} \text{ and } d_{\max} > d_{\text{tot}} \text{ with } \mathbb{E}[\Delta^2 c_e(1 + V)] \leq \nu \text{ for all } e \in \mathcal{E} \text{ and } V \sim \text{Poisson}(d_{\max}), \quad (18)$$

where $\Delta^2 c_e(k) = c_e(k+2) - 2c_e(k+1) + c_e(k)$. This condition holds in particular for costs with subexponential growth $c_e(k) \leq b \exp(ak)$ for some $a, b \geq 0$, which include all polynomials. The condition fails for fast growing functions such as $c_e(k) = k!$ or $c_e(k) = \exp(\exp(k))$. The following result summarizes the main consequences of (18).

Lemma 1. Assume (18) and let

$$\zeta := (e^{d_{\max}} - 1)\nu + \max_{e \in \mathcal{E}} (c_e(2) - c_e(1)), \quad (19)$$

$$\Lambda(u) := \frac{\nu d_{\max}}{2} \frac{ue^u}{(1-u)^2} + u\zeta. \quad (20)$$

Then, the costs $\check{c}_e(\cdot)$ are finite and of class C^2 on $[0, d_{\max}]$ with $0 \leq \check{c}'_e(x) \leq \zeta$ for all $x \in [0, d_{\max}]$. Moreover, let Γ_B be a Bernoulli congestion game with $d_{\text{tot}} \leq d_{\max}$ and let $u = \max_{i \in \mathcal{N}} u_i$. Let X_e be the random loads in a mixed strategy profile σ , and $Z_{i,e} = X_e - U_{i,e}$ be the loads excluding player i . Then, the expected values $x_e = \mathbb{E}_\sigma[X_e]$ and $z_{i,e} = \mathbb{E}_\sigma[Z_{i,e}]$ satisfy $|z_{i,e} - x_e| = \mathbb{E}_\sigma[U_{i,e}] \leq u$, and we have

$$|\mathbb{E}_{\hat{\sigma}}[c_e(1 + Z_{i,e})] - \check{c}_e(x_e)| \leq \Lambda(u). \quad (21)$$

Proof. The smoothness and the bound $0 \leq \check{c}'_e(x) \leq \zeta$ for the auxiliary costs follow directly from Proposition C.2(b) and Corollary C.1(a). In particular, $\check{c}_e(\cdot)$ is ζ -Lipschitz, and then (21) follows by using a triangle inequality and Theorem C.1(b). \square

With all these preliminaries, we are ready to prove the convergence of equilibria for Bernoulli congestion games.

Theorem 2. Let Γ_B^n be a sequence of Bernoulli congestion games satisfying (14) and (18), and let $\hat{\sigma}^n \in \text{BNE}(\Gamma_B^n)$ be an arbitrary sequence of Bayesian Nash equilibria. Then, the corresponding sequence of expected flow-load pairs $(\hat{\mathbf{y}}^n, \hat{\mathbf{x}}^n)$ is bounded and every accumulation point $(\hat{\mathbf{y}}, \hat{\mathbf{x}})$ is a Wardrop equilibrium for the nonatomic congestion game $\check{\Gamma}^\infty$. Furthermore, along every such convergent subsequence, the random flows $Y_{t,s}^n$ and loads X_e^n converge in total variation to Poisson limits $Y_{t,s} \sim \text{Poisson}(\hat{y}_{t,s})$ and $X_e \sim \text{Poisson}(\hat{x}_e)$, where the variables $Y_{t,s}$ are independent.

Proof. From Proposition 1, we have that $(\hat{\mathbf{y}}^n, \hat{\mathbf{x}}^n)$ is bounded and also that, along any convergent subsequence, $Y_{t,s}^n$ and X_e^n converge in total variation to Poisson limits $Y_{t,s} \sim \text{Poisson}(\hat{y}_{t,s})$ and $X_e \sim \text{Poisson}(\hat{x}_e)$ with $Y_{t,s}$ independent. It remains to show that every accumulation point is a Wardrop equilibrium for $\check{\Gamma}^\infty$. Take a convergent subsequence and rename it to be the full sequence $(\hat{\mathbf{y}}^n, \hat{\mathbf{x}}^n) \rightarrow (\hat{\mathbf{y}}, \hat{\mathbf{x}})$. We must prove that, for each type t and each pair of strategies s, s' in \mathcal{S}_t , we have

$$\hat{y}_{t,s} > 0 \Rightarrow \sum_{e \in \mathcal{S}} \check{c}_e(\hat{x}_e) \leq \sum_{e \in \mathcal{S}'} \check{c}_e(\hat{x}_e). \quad (22)$$

A strict inequality $\hat{y}_{t,s} > 0$ implies that, for all n large enough, we have $\hat{y}_{t,s}^n > 0$, and there must be a player $i = i^n$ of type $t_i^n = t$ with $\hat{\sigma}_i^n(s) > 0$. The equilibrium condition in Γ_B^n implies that, for each alternative strategy $s' \in \mathcal{S}_t$ for player i , we have

$$\sum_{e \in \mathcal{S}} \mathbb{E}_{\hat{\sigma}^n} [c_e(1 + Z_{i,e}^n)] \leq \sum_{e \in \mathcal{S}'} \mathbb{E}_{\hat{\sigma}^n} [c_e(1 + Z_{i,e}^n)]. \quad (23)$$

Using Lemma 1, we get

$$|\mathbb{E}_{\hat{\sigma}^n}[c_e(1 + Z_{i,e}^n)] - \check{c}_e(x_e^n)| \leq \Lambda(u^n) \rightarrow 0,$$

and because $x_e^n \rightarrow \hat{x}_e$, it follows that $\mathbb{E}_{\hat{\sigma}^n}[c_e(1 + Z_{i,e}^n)] \rightarrow \check{c}_e(\hat{x}_e)$. Thus, letting $n \rightarrow \infty$ in (23), we obtain (22) as required. \square

Example 2. Consider a sequence $(\Gamma_{\mathbb{B}}^n)_{n \in \mathbb{N}}$ of games played on the Wheatstone network in Figure 3 with n identical players and $u_i \equiv u = 1/n$ so that the total expected demand is $d_{\text{tot}} = 1$. Notice that the conditional expected cost of a player $i \in \mathcal{N}$ when using e_1 or e_5 equals $1 + u(x - 1)$, where x is the number of players using the edge. It follows that the zigzag path (e_1, e_3, e_5) is strictly dominated by a linear combination of the upper and lower paths, and no player uses the zigzag path. This is in sharp contrast with the weighted congestion games described in Example 1, in which all but one player chose the zigzag path. This difference is explained by the fact that each player has a unit weight and a significant impact on the costs.

For an even number of players $n = 2k$, there is a unique PNE (modulo permutations of players), in which half of the players choose the upper path and the other half take the lower path. In this equilibrium, the expected cost for each player is $(2.5n - 1)/n^2$. When $n = 2k + 1$ is odd, there is an NE in which each edge gets k players for sure, and the last player randomizes arbitrarily between these two paths. For all n , there is also a symmetric mixed equilibrium in which every player mixes with $q = 1/2$ between the upper and lower paths, and the expected cost of each player is $(5n - 1)/(2n^2)$. In general, one can prove that every mixed Nash equilibrium is of the form where k_1 players take the upper path, k_2 players take the lower path, and the remaining $k_3 = n - k_1 - k_2$ players, with $k_3 - 1 > |k_2 - k_1|$, randomize between these two paths with exactly the same probabilities $(q, 1 - q)$ with $q = \frac{1}{2}(1 + \frac{k_2 - k_1}{k_3 - 1})$ so as to equalize their costs. As $n \rightarrow \infty$, all these equilibria converge to the Wardrop equilibrium of a nonatomic game where the costs of e_1 and e_5 become $\check{c}(x) = 1 + x$. In this Wardrop equilibrium, half of the demand goes on the upper path and the remaining half on the lower path.

Under a mild additional condition, the equilibrium loads of the nonatomic limit game are unique, and every sequence of equilibria converges. The following is a direct consequence of Theorem 2 and the strict monotonicity of $\check{c}_e(\cdot)$ established in Corollary C.1(b).

Corollary 2. Let $c_e : \mathbb{N}^* \rightarrow \mathbb{N}$ be weakly increasing and nonconstant and assume that conditions (18) hold. Then, the extended functions $\check{c}_e(\cdot)$ are strictly increasing, and the loads \hat{x}_e are the same in every Wardrop equilibrium for $\check{\Gamma}^\infty$. Moreover, for every sequence $\Gamma_{\mathbb{B}}^n$ of Bernoulli congestion games satisfying the conditions in (14) and every sequence $\hat{\sigma}^n \in \text{BNE}(\Gamma_{\mathbb{B}}^n)$, the random loads X_e^n converge in total variation to a random variable $X_e \sim \text{Poisson}(\hat{x}_e)$.

Theorem 2 and its Corollary 2 provide one of our main results. They show that the equilibrium flows in Bernoulli congestion games converge in total variation toward Poisson variables whose expected values coincide with the Wardrop equilibria of an auxiliary nonatomic game $\check{\Gamma}^\infty$. In Section 5, we supplement this result by showing that the limit game $\check{\Gamma}^\infty$ can be interpreted as a Poisson game in the sense of Myerson [43] and the equilibria of such Poisson games can similarly be determined by means of a Wardrop equilibrium of $\check{\Gamma}^\infty$. This provides a novel justification for Poisson games as well as an alternative interpretation of Wardrop equilibria.

In addition to these results, Section 6.2 presents explicit bounds for the approximation and rate of convergence of the distributions of the resource loads, and Section 7 shows that the PoA for the sequence of Bernoulli congestion games $\Gamma_{\mathbb{B}}^n$ converges to the PoA in the nonatomic limit $\check{\Gamma}^\infty$.

5. Poisson Games

In this section, we expose the connection between the nonatomic game $\check{\Gamma}^\infty$ obtained in Section 4 as a limit of Bernoulli congestion games and Poisson games in the sense of Myerson [43]. We first formally introduce Myerson's framework and Poisson games and conclude that Poisson congestion games arise naturally as a limit of Bernoulli congestion games (which is essentially proved in Theorem 2), and the equilibria of the Poisson game coincide with the Wardrop equilibria of the nonatomic game. This is a novel justification for Wardrop's model and shows that equilibria in Poisson congestion games can be calculated from the solution of a related Wardrop equilibrium. Later, we consider more general games with bounded costs and establish the convergence of sequences of games with Bernoulli players toward Poisson games, providing a microfoundation for Poisson games.

5.1. Games with Population Uncertainty

Games with population uncertainty are introduced by Myerson [43] as a model for situations in which players do not know exactly which players participate in the game but share an awareness of the number and type of players that might be present. These games are described by a finite set of types \mathcal{T} and a joint probability distribution μ over $\mathbb{N}^{\mathcal{T}}$ that characterizes the random number of players of each type $N = (N_t)_{t \in \mathcal{T}}$ that take part in the game.¹ Each player of type $t \in \mathcal{T}$ has a corresponding strategy set \mathcal{S}_t and a cost $C(Y^{-t}; t, s)$ that depends on the chosen action $s \in \mathcal{S}_t$ and on the strategy flows Y^{-t} induced by the actions of the *other players* as described subsequently. Altogether, a game with population uncertainty is defined by the tuple $\Gamma_P = (\mathcal{T}, \mu, (\mathcal{S}_t)_{t \in \mathcal{T}}, C)$.

Because players cannot distinguish the specific identities of the other players and only know the joint probability distribution μ according to which $N = (N_t)_{t \in \mathcal{T}}$ is drawn, they are treated symmetrically by assuming that all players of type t adopt the same mixed strategy $\sigma_t \in \Delta(\mathcal{S}_t)$, which only depends on their type. The measure μ and the strategy profile $\sigma = (\sigma_t)_{t \in \mathcal{T}}$ determine the distribution of the strategy flows $Y = (Y_{t,s})_{t \in \mathcal{T}, s \in \mathcal{S}_t}$, where $Y_{t,s}$ is the number of players of type t that choose the action s . Indeed, conditionally on the event $N_t = \bar{n}_t$, the flows $Y_t = (Y_{t,s})_{s \in \mathcal{S}_t}$ for type $t \in \mathcal{T}$ are distributed across strategies as independent multinomials $Y_t \sim \text{Multinomial}(\bar{n}_t, \sigma_t)$. Hence, for each $n = (n_{t,s})_{t \in \mathcal{T}, s \in \mathcal{S}_t}$ the distribution of Y is given by

$$P_{\mu, \sigma}(Y = n) = \mu(\bar{n}) \prod_{t \in \mathcal{T}} \left(\bar{n}_t! \prod_{s \in \mathcal{S}_t} \frac{\sigma_t(s)^{n_{t,s}}}{n_{t,s}!} \right), \quad (24)$$

where $\bar{n} := (\bar{n}_t)_{t \in \mathcal{T}}$ with $\bar{n}_t := \sum_{s \in \mathcal{S}_t} n_{t,s}$.

Now, in order to determine the expected cost for a generic player of type t who is active in the game, we need the conditional distribution assessed by such a player for the strategy flows Y^{-t} induced by the other players. To this end, let $N^{-t} = (N_{t'}^{-t})_{t' \in \mathcal{T}}$ be the random vector giving the number of players for each type, excluding this active player of type t . As argued by Myerson [42, 43] (see also Milchtaich [39]), the *posterior distribution* of N^{-t} can be identified with the Palm distribution of μ viewed as a finite point process, that is,

$$\mu(\bar{n} | t) := P_{\mu, \sigma}(N^{-t} = \bar{n}) = \frac{(\bar{n}_t + 1) \mu(\bar{n} + \delta_t)}{E_{\mu}[N_t]},$$

where $\bar{n} + \delta_t$ denotes the vector \bar{n} with \bar{n}_t replaced by $\bar{n}_t + 1$. Accordingly, the posterior distribution of the strategy flows $Y^{-t} = (Y_{t',s}^{-t})_{t' \in \mathcal{T}, s \in \mathcal{S}_{t'}}$ induced by the remaining players is given by (24) with $\mu(\bar{n})$ replaced by $\mu(\bar{n} | t)$. The expected cost for a player of type t is computed according to this posterior distribution, and an equilibrium is then defined as a family of type-dependent mixed strategies $\hat{\sigma} = (\hat{\sigma}_t)_{t \in \mathcal{T}}$ with $\hat{\sigma}_t \in \Delta(\mathcal{S}_t)$ such that

$$\forall t \in \mathcal{T}, \forall s, s' \in \mathcal{S}_t \quad \hat{\sigma}_t(s) > 0 \Rightarrow E_{\mu, \hat{\sigma}}[C(Y^{-t}; t, s)] \leq E_{\mu, \hat{\sigma}}[C(Y^{-t}; t, s')]. \quad (25)$$

In order for the expectations to be well-defined, Myerson [43] assumes the functions $C(\cdot; t, s)$ to be bounded. For congestion games, this requires the costs $c_e(\cdot)$ to be bounded. Fortunately, this can be relaxed, and we only require the much weaker condition in (18). Later, when considering more general Poisson games, we go back to Myerson's assumptions.

Remark 3. The Bernoulli congestion games in Section 2.3 fall in the framework of games with population uncertainty in which $N_t = \sum_{i:t_i=t} U_i$ is the sum of independent nonhomogeneous Bernoulli random variables. However, we consider not only strategies defined by the player's type, but also asymmetric equilibria in which players choose their strategies individually.

Example 3. Consider the game Γ_B^n on the Wheatstone network of Figure 3 in Example 2. Recall that there are n players—all of the same type t —and each one is present with probability $1/n$. Assume that all players play a mixed strategy in which, with probability $1/2$, they choose the upper path and with probability $1/2$ they choose the lower path. Let $Y = (Y_1, Y_2, Y_3)$ denote the random vector that gives the number of players on the paths $s_1 = (e_1, e_4)$, $s_2 = (e_1, e_3, e_5)$, and $s_3 = (e_2, e_5)$, respectively. Then, $Y_1, Y_3 \sim \text{Binomial}(n, (2n)^{-1})$ and $Y_2 \equiv 0$, whereas the corresponding posterior distributions are $Y_1^{-t}, Y_3^{-t} \sim \text{Binomial}(n-1, (2n-2)^{-1})$ and $Y_2^{-t} \equiv 0$.

5.2. Poisson Games

An important subclass of games with population uncertainty is obtained when the N_t s are independent Poisson variables $N_t \sim \text{Poisson}(d_t)$ with $d_t > 0$, that is,

$$\mu(\bar{n}) = P(N_t = \bar{n}_t, \forall t \in \mathcal{T}) = \prod_{t \in \mathcal{T}} e^{-d_t} \frac{d_t^{\bar{n}_t}}{\bar{n}_t!}. \quad (26)$$

It is not difficult to see that such Poisson games are characterized by the fact that the posteriors $\mu(\cdot | t)$ coincide with μ for every $t \in \mathcal{T}$. Moreover, in this case, the loads $Y_{t,s}$ are also independent with $Y_{t,s} \sim \text{Poisson}(d_t \sigma_t(s))$. In fact, as shown in Myerson [43, theorem 1], in a game with population uncertainty, the variables $Y_{t,s}$ are independent if and only if the game is Poisson.

It turns out that the nonatomic game $\check{\Gamma}^\infty$ in Section 4, obtained as a limit of a sequence of Bernoulli congestion games, can be interpreted as a Poisson game defined by the costs

$$C(Y^{-t}; t, s) = \sum_{e \in \mathcal{S}} c_e(1 + X_e^{-t}), \quad X_e^{-t} = \sum_{t' \in \mathcal{T}} \sum_{s \in \mathcal{S}_{t'}} Y_{t',s}^{-t} \mathbb{1}_{\{e \in \mathcal{S}\}}. \quad (27)$$

We state this observation in the following result.

Theorem 3. Let σ be a strategy profile in the Poisson game defined by the costs in (27) and the demands given in (26) with $d_t > 0$ satisfying (18). Define (\mathbf{y}, \mathbf{x}) as $y_{t,s} = d_t \sigma_t(s)$ and $x_e = \sum_{t \in \mathcal{T}} \sum_{s \in \mathcal{S}_t} y_{t,s} \mathbb{1}_{\{e \in \mathcal{S}\}}$. Then, σ is an equilibrium in the Poisson game if and only if (\mathbf{y}, \mathbf{x}) is a Wardrop equilibrium for the nonatomic game $\check{\Gamma}^\infty$ in Section 4.

Proof. It suffices to note that the posterior distribution of Y^{-t} is Poisson with independent components and expected values $E_{\mu, \sigma}[Y_{t',s}^{-t}] = y_{t',s}$ so that $X_e^{-t} \sim \text{Poisson}(x_e)$. Then, taking expectation, we get precisely

$$E_{\mu, \sigma}[C(Y^{-t}; t, s)] = \sum_{e \in \mathcal{S}} E[c_e(1 + X_e^{-t})] = \sum_{e \in \mathcal{S}} \check{c}_e(x_e), \quad (28)$$

where $\check{c}_e(\cdot)$ is defined as in (17). \square

Note that the costs in (27) depend on the strategy flows Y^{-t} only through the aggregate resource loads X_e^{-t} , whereas Myerson [43] considers more general cost functions $C(Y^{-t}; t, s)$. However, as already mentioned, Myerson requires costs to be bounded to ensure that their expected values are well-defined, whereas for congestion games, this boundedness can be relaxed and replaced by the conditions in (18).

5.3. Convergence of Bernoulli Games with Bounded Costs

As mentioned in the introduction, Myerson [43] introduces Poisson games axiomatically. The following variation of Theorem 2 provides a justification for Poisson games as limits of a sequence of finite games with population uncertainty of Bernoulli type. The result goes beyond the separable cost structure (27) of congestion games, allowing the cost of a player to depend on the player's type and action and the full vector of strategy flows. To compensate, the costs are required to be bounded.

Consider a finite set of types \mathcal{T} with strategy sets \mathcal{S}_t for $t \in \mathcal{T}$ and a sequence of games Γ^n with finitely many players $i \in \mathcal{N}^n$ with types t_i^n and probabilities u_i^n of being active. Let U_i^n be Bernoulli random variables with $P(U_i^n = 1) = u_i^n$, and let player i choose $S_i^n \in \mathcal{S}_{t_i^n}$ at random using a mixed strategy σ_i^n . Let Y^n be the random vector of strategy flows $Y_{t,s}^n = \sum_{j: t_j^n = t} U_j^n \mathbb{1}_{\{S_j^n = s\}}$, and define $Y^{-i,n}$ similarly by excluding player i . The expected cost of an action $s \in \mathcal{S}_{t_i^n}$ for player i is given by $E_{\hat{\sigma}^n}[C(Y^{-i,n}; t_i^n, s)]$. A Nash equilibrium $\hat{\sigma}^n$ is defined as usual by the condition

$$\forall i \in \mathcal{N}^n, \forall s, s' \in \mathcal{S}_{t_i^n} \quad \hat{\sigma}_i^n(s) > 0 \Rightarrow E_{\hat{\sigma}^n}[C(Y^{-i,n}; t_i^n, s)] \leq E_{\hat{\sigma}^n}[C(Y^{-i,n}; t_i^n, s')].$$

Theorem 4. Consider a sequence of games Γ^n as before with bounded cost functions $C(\cdot; t, s)$ and assume the conditions in (14) with $d_t > 0$. Then, for every sequence $\hat{\sigma}^n$ of Nash equilibria, the expected loads $y_{t,s}^n = E_{\hat{\sigma}^n}[Y_{t,s}^n]$ are bounded, and each accumulation point $\hat{\mathbf{y}} = (\hat{y}_{t,s})_{t \in \mathcal{T}, s \in \mathcal{S}_t}$ corresponds to an equilibrium $\hat{\sigma}$ in the Poisson game by setting $\hat{\sigma}_t(s) = \hat{y}_{t,s}/d_t$ for all $s \in \mathcal{S}_t$.

Proof. The boundedness of the expected strategy loads $y_{t,s}^n$ follows from $\sum_{s \in \mathcal{S}_t} y_{t,s}^n = d_t^n \rightarrow d_t$. Take a convergent subsequence and rename it so that $\mathbf{y}^n \rightarrow \hat{\mathbf{y}}$, and define $\hat{\sigma}$ as in the statement. If $\hat{\sigma}_t(s) > 0$ for some $s \in \mathcal{S}_t$, then $\hat{y}_{t,s} > 0$ so that, for all n large, we have $y_{t,s}^n > 0$ and there must be a player $i^n \in \mathcal{N}^n$ with type t and $\hat{\sigma}_{i^n}^n(s) > 0$. The equilibrium condition for i^n implies that

$$\forall s' \in \mathcal{S}_t \quad E_{\hat{\sigma}^n}[C(Y^{-i^n,n}; t, s)] \leq E_{\hat{\sigma}^n}[C(Y^{-i^n,n}; t, s')]. \quad (29)$$

As in Proposition 1, it follows that the variables $Y^{-i^n,n}$ converge in total variation and, hence, in distribution, to a random vector \mathbf{Y} with independent Poisson components $Y_{t,s} \sim \text{Poisson}(\hat{y}_{t,s})$. Letting $n \rightarrow \infty$ in (29), we obtain (25), from which the result follows. \square

Remark 4. Theorem 4 remains valid when some of the demands converge to zero $d_t^n \rightarrow d_t = 0$ by considering the Poisson game restricted to the nontrivial types with $d_t > 0$.

Remark 5. Despite their similarity, Theorems 2 and 4 are independent, and neither one follows from the other. In fact, Theorem 4 allows for more general forms of the costs as long as they are bounded, whereas Theorem 2 can handle unbounded costs provided that they have the specific additive structure in (27) and under the conditions in (18).

6. Approximation Bounds and Rates of Convergence

In this section, we establish nonasymptotic bounds for the distance between the distribution of the edge loads at equilibrium in the finite games and the loads in the corresponding Wardrop and Poisson games. These bounds provide quantitative estimates on how well a Wardrop or Poisson equilibrium approximates the equilibrium of the finite game. We highlight that these results require the costs to have derivatives bounded away from zero.

As in previous sections, we analyze separately the case of weighted and Bernoulli congestion games. In both settings, we exploit estimates for the distance between Wardrop equilibria in nonatomic games with different demands as well as between exact and approximate equilibria. Recall that an ε -approximate Wardrop equilibrium $(\mathbf{y}, \mathbf{x}) \in \mathcal{F}(\mathbf{d})$, or ε -Wardrop equilibrium, is defined exactly like a Wardrop equilibrium up to an additive error:

$$\forall t \in \mathcal{T}, \forall s, s' \in \mathcal{S}_t \quad y_{t,s} > 0 \Rightarrow \sum_{e \in \mathcal{S}} c_e(x_e) \leq \sum_{e \in \mathcal{S}'} c_e(x_e) + \varepsilon. \quad (30)$$

Proposition 2. Let $(\hat{\mathbf{y}}, \hat{\mathbf{x}})$ be a Wardrop equilibrium for a nonatomic congestion game Γ^∞ with $d_{\text{tot}} \leq d_{\text{max}}$. Suppose that there exists some constant $c'_{\min} > 0$ such that

$$\forall e \in \mathcal{E}, \forall x \in [0, d_{\text{max}}] \quad c'_{\min} \leq c'_e(x),$$

and let $\Xi = \sqrt{2C/c'_{\min}}$ with $C \geq \sum_{e \in \mathcal{S}} c_e(d_{\text{max}})$ for all strategies $s \in \bigcup_{t \in \mathcal{T}} \mathcal{S}_t$.

a. If $(\mathbf{y}, \mathbf{x}) \in \mathcal{F}(\mathbf{d})$ is an ε -approximate Wardrop equilibrium, then $\|\mathbf{x} - \hat{\mathbf{x}}\|_2 \leq \sqrt{\varepsilon d_{\text{max}}/c'_{\min}}$.

b. If $(\hat{\mathbf{y}}', \hat{\mathbf{x}}')$ is a Wardrop equilibrium for perturbed demands \mathbf{d}' with $d'_{\text{tot}} \leq d_{\text{max}}$, then $\|\hat{\mathbf{x}}' - \hat{\mathbf{x}}\|_2 \leq \Xi \cdot \sqrt{\|\mathbf{d}' - \mathbf{d}\|_1}$.

Proof. See Appendix A. \square

6.1. Weighted Congestion Games

Theorem 1 shows that the random loads X_e^n in the finite weighted congestion games converge in distribution to the Wardrop equilibrium loads \hat{x}_e . Under additional assumptions that call for lower and upper bounds on the derivatives of the cost functions, we can find nonasymptotic estimates for the rate of convergence.

We first show that any mixed Nash equilibrium of a weighted congestion game Γ_W is close to a Wardrop equilibrium in a nonatomic game Γ^∞ with the same aggregate demands. The distance depends—in a way that the theorem makes precise—on the topology of the instance, on the cost functions, and on the magnitude of the weights.

Theorem 5. Let Γ_W be a weighted congestion game with aggregate demands $(d_t)_{t \in \mathcal{T}}$ given by (4) with $d_{\text{tot}} \leq d_{\text{max}}$. Assume that there exist constants $c'_{\max} \geq c'_{\min} > 0$ such that

$$\forall e \in \mathcal{E}, \forall x \in [0, d_{\text{max}}], \quad c'_{\min} \leq c'_e(x) \leq c'_{\max}$$

and suppose in addition that the derivatives $c'_e(\cdot)$ are L -Lipschitz continuous for some $L \geq 0$. Define $\theta = \sqrt{d_{\text{max}}/4} + \sqrt{\kappa d_{\text{max}}(2c'_{\max} + Ld_{\text{max}}/4)/c'_{\min}}$ with $\kappa \geq |s|$ for all strategies $s \in \bigcup_{t \in \mathcal{T}} \mathcal{S}_t$. Let X_e be the random load in a mixed Nash equilibrium $\hat{\boldsymbol{\sigma}} \in \text{MNE}(\Gamma_W)$ and \hat{x}_e the unique resource load in the Wardrop equilibrium for the nonatomic game $\Gamma^\infty = (\mathcal{G}, (d_t)_{t \in \mathcal{T}})$ with the same aggregate demands as Γ_W . Letting $w = \max_{i \in \mathcal{N}} w_i$, we have

$$\|X_e - \hat{x}_e\|_{L^2} \leq \theta \cdot \sqrt{w}. \quad (31)$$

As a reminder, the L^2 norm is defined as $\|X\|_{L^2} = \sqrt{\mathbf{E}[X^2]}$.

Proof. See Appendix A. \square

An unfortunate feature of the estimate (31) is the presence of the lower bound c'_{\min} in the coefficient θ . In particular, for fixed demands, the left-hand side of (31) remains bounded, whereas the bound on the right diverges as $c'_{\min} \rightarrow 0$ so that this estimate is not always tight. However, for arbitrary demands, the following example

shows that the discrepancy between the equilibrium loads in a weighted atomic game and in the Wardrop equilibrium can be arbitrarily large when c'_{\min} is small, and both the right- and left-hand sides of (31) tend to infinity although at different rates. Although this example was constructed to illustrate the dependence on c'_{\min} , the upper bound parameter θ depends on various other parameters. A full study of the rate of convergence in relation to each of them is a topic for future research.

Example 4. Consider the network in Figure 4 with demands d_1, d_2 from two different origins and a common destination. Let v and z denote the fractions of the demands sent over the central path.

Equalizing the costs of the central and outer paths, we get the linear system

$$\begin{aligned} (1 + 2\beta)\widehat{v} + \widehat{z} &= \beta d_1 \\ \widehat{v} + (1 + 2\beta)\widehat{z} &= \beta d_2, \end{aligned} \tag{32}$$

whose solution is a Wardrop equilibrium provided that $\widehat{v} \geq 0$ and $\widehat{z} \geq 0$. Similarly, for an atomic unsplittable game with weights $w_i \equiv w$ and demands $d_1 = n_1 w, d_2 = n_2 w$, integer multiples of w , we have that $v = k_1 w$ and $z = k_2 w$ with $k_1, k_2 \in \mathbb{N} \setminus \{0\}$ represent a pure Nash equilibrium if and only if

$$\begin{aligned} (1 + 2\beta)v + z &= \beta d_1 + \varepsilon_1 \\ x + (1 + 2\beta)z &= \beta d_2 + \varepsilon_2 \end{aligned} \quad \text{with } \varepsilon_1, \varepsilon_2 \in [-(1 + \beta)w, \beta w]. \tag{33}$$

In particular, for $\varepsilon_1 = 0$ and $\varepsilon_2 = -(1 + \beta)w$ with $1/\beta \in \mathbb{N}$, it follows from (33) that the demands $d_1 = w[2k_1 + (k_1 + k_2)/\beta]$ and $d_2 = w[2k_2 + 1 + (k_1 + k_2 + 1)/\beta]$ are integer multiples of w . Now, solving (32) and (33), we get

$$\begin{pmatrix} v \\ z \end{pmatrix} - \begin{pmatrix} \widehat{v} \\ \widehat{z} \end{pmatrix} = \frac{w}{4\beta} \begin{pmatrix} 1 \\ -(1 + 2\beta) \end{pmatrix},$$

so that the nonnegativity conditions $\widehat{v} \geq 0$ and $\widehat{z} \geq 0$ are equivalent to $k_1 \geq 1/(4\beta)$ and $k_2 \geq -(1 + 2\beta)/(4\beta)$. Hence, taking $k_1 = \lceil 1/(4\beta) \rceil$ and $k_2 = 1$, we have that (v, z) and $(\widehat{v}, \widehat{z})$ are, respectively, Nash and Wardrop equilibria. Their load difference is of the order of $O(w/\beta)$, which diverges as $\beta \rightarrow 0$. Note that here $c'_{\min} = \beta, c'_{\max} = 1, L = 0, \kappa = 2$, and $d_{\max} = d_1 + d_2 \sim w/(2\beta^2)$, so that the bound (31) becomes $\theta\sqrt{w} \sim \sqrt{2}w/\beta^{3/2}$, which also diverges as $\beta \rightarrow 0$ although at a faster rate compared with the actual error $O(w/\beta)$. We do not know an instance when (31) is tight.

Combining Proposition 2 and Theorem 5, we can derive an explicit estimate for the L^2 distance between the random resource loads X_e^n in a sequence of mixed Nash equilibria for Γ_W^n , and the loads in the Wardrop equilibrium of the nonatomic limit game Γ^∞ .


The resulting bound in (34) has an additional term that involves the distance between the demands along the sequence Γ_W^n and those in Γ^∞ . Thus, having an equilibrium close to the limit requires small weights and aggregate demands that are close to the limiting demands.

Corollary 3. Let Γ_W^n be a sequence of weighted congestion games satisfying (12) with $d_{\text{tot}}^n \leq d_{\max}$ for all n . Let the costs $c_e(\cdot)$ and θ be as in Theorem 5 and Ξ as in Proposition 2. Let X_e^n be the random loads in a sequence of mixed Nash equilibria $\widehat{\sigma}^n \in \text{MNE}(\Gamma_W^n)$ and \widehat{x}_e the unique resource loads in the Wardrop equilibrium for the nonatomic limit game $\Gamma^\infty = (\mathcal{G}, (d_t)_{t \in \mathcal{T}})$. Then, with w^n defined by (12b), we have

$$\|X_e^n - \widehat{x}_e\|_{L^2} \leq \theta \cdot \sqrt{w^n} + \Xi \cdot \sqrt{\|d^n - d\|_1}. \tag{34}$$

Proof. See Appendix A. \square

Figure 4. (Color online) A network with two OD pairs and unbounded errors for $\beta \rightarrow 0$. Labels over the arcs denote cost functions, whereas labels close to origins denote path flows.



6.2. Bernoulli Congestion Games

Analogously to Section 6.1, we now find nonasymptotic estimates for the rate of convergence in Bernoulli congestion games. Notice that (16) already provides a simpler—although difficult to quantify—bound on the distance between the resource loads in the finite games Γ_B^n and the Poisson limits. Indeed, although u^n is a primitive parameter of the model and its proximity to zero is readily available, the second term $|\widehat{x}_e^n - \widehat{x}_e|$ is only known to converge to zero asymptotically. The following estimates provide more explicit bounds expressed directly in terms of the primitives of the model.

We first show that the distance between a Bayesian Nash equilibrium for a Bernoulli congestion game Γ_B and a Wardrop equilibrium for the nonatomic game $\check{\Gamma}^\infty$ with the same aggregate demands is bounded by an expression that depends on the topology and the function $\Lambda(u)$ defined in (20) evaluated at the maximum participation probability $u = \max_{i \in \mathcal{N}} u_i$.

Theorem 6. *Let Γ_B be a Bernoulli congestion game with expected demands $d_t = \sum_{i:t_i=t} u_i$ and $d_{\text{tot}} \leq d_{\text{max}}$. Suppose that (18) holds and that $\check{c}'_e(x) \geq \check{c}'_{\min} > 0$ for all $x \in [0, d_{\text{max}}]$. Let $\check{\theta} = \sqrt{2\kappa d_{\text{max}}/\check{c}'_{\min}}$ with $\kappa \geq |s|$ for all strategies $s \in \cup_{t \in \mathcal{T}} \mathcal{S}_t$. Let X_e be the random loads in a Bayesian Nash equilibrium $\widehat{\sigma} \in \text{BNE}(\Gamma_B)$ with expected flow loads (\mathbf{y}, \mathbf{x}) , and $\widehat{X}_e \sim \text{Poisson}(\widehat{x}_e)$ with \widehat{x} the unique Wardrop equilibrium loads in the nonatomic game $\check{\Gamma}^\infty$ with costs $\check{c}_e(\cdot)$ and with the same aggregate demands d_t as Γ_B . Then, $\|\mathbf{x} - \widehat{\mathbf{x}}\|_2 \leq \check{\theta} \cdot \sqrt{\Lambda(u)}$ and*

$$\rho_{\text{TV}}(\mathcal{L}(X_e), \mathcal{L}(\widehat{X}_e)) \leq u + \check{\theta} \cdot \sqrt{\Lambda(u)}. \quad (35)$$

Proof. See Appendix A. \square

Remark 6. The condition $\check{c}'_e(x) \geq \check{c}'_{\min}$ holds when $c_e(k+1) \geq c_e(k) + \check{c}'_{\min}$ for all $e \in \mathcal{E}$ and $k \geq 1$. A milder assumption is $0 < \delta := \min_{e \in \mathcal{E}} c_e(2) - c_e(1)$ in which case one can take $\check{c}'_{\min} = \delta e^{-d_{\text{max}}}$.

Example 5. Similarly to the case of weighted congestion games, for $\check{c}'_{\min} = \beta \approx 0$, the distance $\|\mathbf{x} - \widehat{\mathbf{x}}\|_2$ between the loads in a Bernoulli congestion game and the corresponding Wardrop equilibrium can be arbitrarily large. Indeed, consider the graph in Figure 4 and a Bernoulli game with identical participation probabilities $u_i \equiv u$ and expected demands $d_1 = n_1 u$, $d_2 = n_2 u$, integer multiples of u . In this setting, $v = k_1 u$ and $z = k_2 u$ with $k_1, k_2 \in \mathbb{N} \setminus \{0\}$, is a pure Bayes–Nash equilibrium if and only if

$$\begin{aligned} (1 + 2\beta)v + z &= \beta d_1 + \varepsilon_1 \\ v + (1 + 2\beta)z &= \beta d_2 + \varepsilon_2 \end{aligned} \quad \text{with } \varepsilon_1, \varepsilon_2 \in [-1 - \beta u, -1 + (1 + \beta)u]. \quad (36)$$

Taking $\varepsilon_1 = -1 + (1 + \beta)u$ and $\varepsilon_2 = -1 - \beta u$ with $1/\beta \in \mathbb{N}$ and $1/(\beta u) \in \mathbb{N}$, the demands are indeed integer multiples of u , whereas solving (32) and (36), we get

$$\begin{pmatrix} v \\ z \end{pmatrix} - \begin{pmatrix} \widehat{v} \\ \widehat{z} \end{pmatrix} = \frac{u}{4\beta(1 + \beta)} \begin{pmatrix} 1 + 2\beta(2 + \beta - 1/u) \\ -1 - 2\beta(1 + \beta + 1/u) \end{pmatrix}.$$

Hence, by taking $k_1 \geq (1 + 2\beta(2 + \beta - 1/u))/(4\beta(1 + \beta))$ and $k_2 = 1$, we have that $(\widehat{v}, \widehat{z})$ is nonnegative, and therefore, it is a Wardrop equilibrium with $\|\mathbf{x} - \widehat{\mathbf{x}}\|_2 \sim O(u/\beta)$ for u fixed and $\beta \approx 0$. The bound in Theorem 6 is again not tight because $\check{\theta} \sqrt{\Lambda(u)} \sim O(u/\beta^{3/2})$.

Using Theorem 6, we may estimate the distance between the random loads in a sequence Γ_B^n of Bernoulli congestion games and the corresponding Wardrop equilibrium in the limit game $\check{\Gamma}^\infty$. The bound in (37) has an additional term that involves the distance between the demands along the sequence Γ_B^n and the demands in $\check{\Gamma}^\infty$. Hence, having an equilibrium close to the limit requires small participation probabilities together with the aggregate demands being close to the limiting demands.

Corollary 4. *Let Γ_B^n be a sequence of Bernoulli congestion games satisfying (14) and (18) with $d_{\text{tot}}^n \leq d_{\text{max}}$. Suppose that $\check{c}'_e(x) \geq \check{c}'_{\min} > 0$ for all $x \in [0, d_{\text{max}}]$ and let $\check{\theta} = \sqrt{2\kappa d_{\text{max}}/\check{c}'_{\min}}$ and $\check{\Xi} = \sqrt{2C/\check{c}'_{\min}}$, where $\kappa \geq |s|$ and $C \geq \sum_{e \in \mathcal{E}} \check{c}_e(d_{\text{max}})$ for all $s \in \cup_{t \in \mathcal{T}} \mathcal{S}_t$. Let X_e^n be the random load in a sequence of Bayes–Nash equilibria $\widehat{\sigma}^n \in \text{BNE}(\Gamma_B^n)$, and $X_e \sim \text{Poisson}(\widehat{x}_e)$*

with \widehat{x}_e the unique resource load in the Wardrop equilibrium for the nonatomic limit game $\check{\Gamma}^\infty$. Then,

$$\rho_{\text{TV}}(\mathcal{L}(X_e^n), \mathcal{L}(X_e)) \leq u^n + \check{\theta} \cdot \sqrt{\Lambda(u^n)} + \check{\xi} \cdot \sqrt{\|\check{d}^n - \check{d}\|_1}. \quad (37)$$

Proof. See Appendix A. \square

7. Convergence of the Inefficiency of Equilibria

In this section, we study the convergence of the inefficiency of equilibria as captured by the PoA and the PoS. Let us begin by recalling these notions. We measure the social cost of a strategy profile as the sum of all players' costs. This provides us with a yardstick with which we can quantify the efficiency of equilibria as first proposed by Koutsoupias and Papadimitriou [33]. The *price of anarchy* is the worst case ratio of the social cost of the equilibrium to the optimum. Here, the worst case is taken with respect to all possible equilibria. The *price of stability* is defined accordingly with respect to the best equilibria.

Starting with nonatomic congestion games, their *social cost* is given by

$$\forall (\mathbf{y}, \mathbf{x}) \in \mathcal{F}(\mathbf{d}) \quad \text{SC}(\mathbf{y}, \mathbf{x}) := \sum_{e \in \mathcal{E}} x_e c_e(x_e), \quad (38)$$

from where the social optimum is $\text{Opt}(\Gamma^\infty) := \min_{(\mathbf{y}, \mathbf{x}) \in \mathcal{F}(\mathbf{d})} \text{SC}(\mathbf{y}, \mathbf{x})$. It is well-known (see Beckmann et al. [7]) that, whenever the cost functions c_e are weakly increasing, all Wardrop equilibria have the same social cost so that defining $\text{Eq}(\Gamma^\infty) := \text{SC}(\widehat{\mathbf{y}}, \widehat{\mathbf{x}})$ for any $(\widehat{\mathbf{y}}, \widehat{\mathbf{x}}) \in \text{WE}(\Gamma^\infty)$, it follows that the PoA and PoS for a nonatomic game coincide and are given by

$$\text{PoA}(\Gamma^\infty) = \text{PoS}(\Gamma^\infty) := \frac{\text{Eq}(\Gamma^\infty)}{\text{Opt}(\Gamma^\infty)}.$$

The corresponding definitions for weighted and Bernoulli congestion games are similar, adjusted for the fact that now these games include stochastic realizations. The *expected social cost* is

$$\forall \boldsymbol{\sigma} \in \Sigma \quad \text{ESC}(\boldsymbol{\sigma}) := \mathbb{E}_{\boldsymbol{\sigma}} \left[\sum_{e \in \mathcal{E}} X_e c_e(X_e) \right], \quad (39)$$

where $(X_e)_{e \in \mathcal{E}}$ are the random resource loads induced by the mixed strategy profile $\boldsymbol{\sigma}$. The optimum cost is $\text{Opt}(\Gamma) := \min_{\boldsymbol{\sigma} \in \Sigma} \text{ESC}(\boldsymbol{\sigma})$, and the social optimum is $\widehat{\boldsymbol{\sigma}} \in \arg \min_{\boldsymbol{\sigma} \in \Sigma} \text{ESC}(\boldsymbol{\sigma})$. Considering the worst and best social cost at equilibrium, the price of anarchy and stability are

$$\text{PoA}(\Gamma) := \max_{\widehat{\boldsymbol{\sigma}} \in \text{MNE}(\Gamma)} \frac{\text{ESC}(\widehat{\boldsymbol{\sigma}})}{\text{Opt}(\Gamma)}, \quad \text{and} \quad \text{PoS}(\Gamma) := \min_{\widehat{\boldsymbol{\sigma}} \in \text{MNE}(\Gamma)} \frac{\text{ESC}(\widehat{\boldsymbol{\sigma}})}{\text{Opt}(\Gamma)}.$$

With these definitions, we may now establish the convergence of the PoA for a sequence of weighted congestion games Γ_W^n . The proof uses some auxiliary results on the convergence of social costs presented in Appendix B.

Theorem 7. Let Γ_W^n be a sequence of weighted congestion games satisfying the conditions in (12) and let $\widehat{\boldsymbol{\sigma}}^n \in \text{MNE}(\Gamma_W^n)$. Then, $\text{ESC}(\widehat{\boldsymbol{\sigma}}^n) \rightarrow \text{Eq}(\Gamma^\infty)$, and therefore, both the price of anarchy $\text{PoA}(\Gamma_W^n)$ and the price of stability $\text{PoS}(\Gamma_W^n)$ converge toward $\text{PoA}(\Gamma^\infty) = \text{PoS}(\Gamma^\infty)$.

Proof. From Theorem 1, we know that every accumulation point of $(\widehat{\mathbf{y}}^n, \widehat{\mathbf{x}}^n)$ is a Wardrop equilibrium for Γ^∞ ; hence, by Lemma B.1, the full sequence $\text{ESC}(\widehat{\boldsymbol{\sigma}}^n)$ converges to $\text{Eq}(\Gamma^\infty)$ as $n \rightarrow \infty$. On the other hand, Proposition B.1 shows that $\text{Opt}(\Gamma_W^n) \rightarrow \text{Opt}(\Gamma^\infty)$ from which the convergence of $\text{PoA}(\Gamma_W^n)$ and $\text{PoS}(\Gamma_W^n)$ follows at once. \square


Example 6. Consider the sequence of games in Example 1 and the different equilibria described there. The social cost is minimized by splitting half of the players between the upper and lower paths (up to one player when n is odd). For $n = 2$, we have $\text{PoA}(\Gamma_n) = 4/3$ and $\text{PoS}(\Gamma_n) = 1$, whereas for $n \geq 3$, by setting $\delta_n = 1$ if n is odd and $\delta_n = 0$ otherwise, we get $\text{PoA}(\Gamma^n) = 4n^2/(3n^2 + \delta_n)$ and $\text{PoS}(\Gamma^n) = (4n^2 - 2n + 2)/(3n^2 + \delta_n)$, both converging to $\text{PoA}(\Gamma^\infty) = \text{PoS}(\Gamma^\infty) = 4/3$ as $n \rightarrow \infty$.

The following analogous result holds for sequences of Bernoulli congestion games.

Theorem 8. Let Γ_B^n be a sequence of Bernoulli congestion games satisfying the conditions in (14) and (18). Then, for every sequence $\widehat{\boldsymbol{\sigma}}^n \in \text{BNE}(\Gamma_B^n)$, the expected social cost $\text{ESC}(\widehat{\boldsymbol{\sigma}}^n)$ converges to $\text{Eq}(\check{\Gamma}^\infty)$. As a consequence, both $\text{PoA}(\Gamma_B^n)$ and $\text{PoS}(\Gamma_B^n)$ converge to $\text{PoA}(\check{\Gamma}^\infty) = \text{PoS}(\check{\Gamma}^\infty)$.

Proof. From Lemma 1, we have that $\check{c}'_e(x) \geq 0$ for $x \in [0, d_{\max}]$ so that the extended costs $\check{c}'_e(\cdot)$ are weakly increasing, and therefore, the social cost $\text{SC}(\widehat{\mathbf{y}}, \widehat{\mathbf{x}}) \equiv \text{Eq}(\check{\Gamma}^\infty)$ is the same in every Wardrop equilibrium. Now, Theorem 2

Figure 5. (Color online) Pigou network.



states that every accumulation point of $(\widehat{y}^n, \widehat{x}^n)$ is a Wardrop equilibrium for $\check{\Gamma}^\infty$; hence, by Lemma B.2, the full sequence $\text{ESC}(\widehat{\sigma}^n)$ converges to $\text{Eq}(\check{\Gamma}^\infty)$ as $n \rightarrow \infty$. On the other hand, Proposition B.2 gives $\text{Opt}(\Gamma_B^n) \rightarrow \text{Opt}(\check{\Gamma}^\infty)$ from which the convergence of $\text{PoA}(\Gamma_B^n)$ and $\text{PoS}(\Gamma_B^n)$ follows at once. \square

Example 7. Consider the sequence of games Γ_B^n on the Wheatstone network in Example 2. The social cost is minimized by splitting half of the players between the upper and lower routes (up to one player for n odd). This strategy profile is also a pure Nash equilibrium so that $\text{PoS}(\Gamma_B^n) = 1$ for all $n \in \mathbb{N}$. As far as the PoA is concerned, the worst equilibrium occurs when each player chooses the upper and lower routes with probability 1/2. Setting $\delta_n = 1/n$ when n is odd and $\delta_n = 0$ otherwise, we obtain $\text{PoA}(\Gamma_B^n) = (5n - 1)/(5n - 2 + \delta_n)$, which converges to $\text{PoA}(\check{\Gamma}^\infty) = 1$ as $n \rightarrow \infty$.

Remark 7. For polynomial costs $c_e(\cdot)$ of degree at most d , we have that $\check{c}_e(x)$ are again polynomials of the same degree (though with different coefficients) so that the results in Roughgarden [54] imply

$$\lim_{n \rightarrow \infty} \text{PoA}(\Gamma_B^n) = \text{PoA}(\check{\Gamma}^\infty) \leq B(d) := \frac{(d + 1)\sqrt[d]{d + 1}}{(d + 1)\sqrt[d]{d + 1} - d}.$$

In fact, for $d = 1$, the bound $\text{PoA}(\Gamma_B^n) \leq 4/3$ is valid as soon as $u^n \leq 1/4$ (see Cominetti et al. [17]). For higher degrees, we conjecture the existence of a threshold for u^n under which $\text{PoA}(\Gamma_B^n)$ already falls below the nonatomic bound $B(d)$. The current result only implies that we have this as an asymptotic bound when $u^n \rightarrow 0$.

Remark 8. Examples 6 and 7 may suggest that the PoA in the nonatomic game obtained as a limit of weighted congestion games is larger than the PoA in the nonatomic limit game for Bernoulli congestion games. This is not true in general. Consider, for instance, the Pigou network in Figure 5 with a demand of one. The Wardrop equilibrium of the standard nonatomic game is optimal, so $\text{PoA}(\Gamma^\infty) = 1$. The Wardrop equilibrium of the nonatomic limit game of the Bernoulli game, in which the auxiliary cost function on the upper edge is now $\check{c}_1(x) = 1 + x$, sends all demand on the upper path, whereas in the social optimum, the demand is split over the upper and lower path. So, we have that $\text{PoA}(\check{\Gamma}^\infty) = 8/7$.

8. Conclusions and Perspectives

In this paper, we study the convergence of equilibria of atomic unsplitable congestion games with an increasing number of players toward a Wardrop equilibrium for a limiting nonatomic game. For the case in which players have vanishing weights, the random flow-load pairs (Y^n, X^n) at a mixed equilibrium in the finite games are shown to converge in distribution toward constant flow-load pairs $(\widehat{y}, \widehat{x})$, which are Wardrop equilibria for the nonatomic limit game. In contrast, if players have a fixed unit weight but are present in the game with vanishing probabilities, then (Y^n, X^n) converge in total variation to Poisson variables (Y, X) , whose expected values $(\widehat{y}, \widehat{x})$ are again characterized as a Wardrop equilibrium for a nonatomic congestion game with auxiliary cost functions. In this case, the limit variables (Y, X) can also be interpreted as an equilibrium for a Poisson game in the sense of Myerson [43]. Under additional conditions, we also established explicit estimates for the distance between the random loads (Y^n, X^n) and their limits $(\widehat{y}, \widehat{x})$ and (Y, X) , respectively. These convergence results are completed by showing that, in both frameworks, the PoA and the PoS converge to the PoA of the limit game.

We do not address the combined case in which both the weights and the presence probabilities vary across players. Such situations may be relevant for routing games in which cars and trucks have a different impact on traffic or in the presence of heterogeneous drivers that may be slower or faster, inducing more or less congestion. Other settings in which players are naturally heterogeneous arise in telecommunications, in which packets come in different sizes, and in processor sharing, in which tasks arriving to a server have different workloads. In such

cases, one might still expect to obtain a limit game that is likely to yield a *weighted nonatomic game*, possibly involving weighted sums of Poisson distributions.

Another direction not explored in this paper concerns the case of oligopolistic competition in which some players, for example, TomTom, Waze, FedEx, UPS, etc., may control a nonnegligible fraction of the demand, whereas simultaneously, another fraction of the demand behaves as individual selfish players. A natural conjecture is that in the vanishing weight limit one may converge to a composite game as those studied by Cominetti et al. [16] and Sorin and Wan [59] with the coexistence of atomic splittable and nonatomic players. Similarly, in the case of fixed weights and vanishing probabilities, one may expect to converge to some form of composite game involving Poisson random variables, which remains to be discovered.

Acknowledgments

This collaboration started in Dagstuhl at the Seminar on Dynamic Traffic Models in Transportation Science in 2018. Marco Scarsini is a member of the Gruppo Nazionale per l'Analisi Matematica, la Probabilità e le loro Applicazioni, Istituto Nazionale di Alta Matematica.

Appendix A. Approximation Bounds and Rates of Convergence: Proofs

Proof of Proposition 2.

a. Defining $C_s := \sum_{e \in \mathcal{S}} c_e(x_e)$ and $\eta_t := \min_{s \in \mathcal{S}_t} C_s$, we bound the squared distance by

$$\begin{aligned} \sum_{e \in \mathcal{E}} c'_{\min} (x_e - \widehat{x}_e)^2 &\leq \sum_{e \in \mathcal{E}} (c_e(x_e) - c_e(\widehat{x}_e))(x_e - \widehat{x}_e) \\ &\leq \sum_{e \in \mathcal{E}} c_e(x_e)(x_e - \widehat{x}_e) \\ &= \sum_{t \in \mathcal{T}} \sum_{s \in \mathcal{S}_t} \left(\sum_{e \in \mathcal{S}} c_e(x_e) \right) (y_{t,s} - \widehat{y}_{t,s}) \\ &= \sum_{t \in \mathcal{T}} \sum_{s \in \mathcal{S}_t} (C_s - \eta_t)(y_{t,s} - \widehat{y}_{t,s}) \\ &\leq \sum_{t \in \mathcal{T}} \sum_{s \in \mathcal{S}_t} (C_s - \eta_t) y_{t,s}. \end{aligned}$$

Here, in the second inequality, we drop the sum $\sum_{e \in \mathcal{E}} c_e(\widehat{x}_e)(x_e - \widehat{x}_e)$, which is nonnegative because $(\widehat{\mathbf{y}}, \widehat{\mathbf{x}})$ is a Wardrop equilibrium for Γ^∞ and $(\mathbf{y}, \mathbf{x}) \in \mathcal{F}(\mathbf{d})$. The first equality follows from expressing the resource loads in terms of the strategy flows and exchanging the order of summation and the second equality from the fact that $\sum_{s \in \mathcal{S}_t} (y_{t,s} - \widehat{y}_{t,s}) = 0$ for all types t as both solutions $(\widehat{\mathbf{y}}, \widehat{\mathbf{x}})$ and (\mathbf{y}, \mathbf{x}) are feasible. Now, using (30), we conclude

$$c'_{\min} \|\mathbf{x} - \widehat{\mathbf{x}}\|_2^2 \leq \sum_{t \in \mathcal{T}} \sum_{s \in \mathcal{S}_t} \varepsilon y_{t,s} = \varepsilon \sum_{t \in \mathcal{T}} d_t = \varepsilon d_{\text{tot}} \leq \varepsilon d_{\max}.$$

b. Proceeding as in the previous part, we have

$$\begin{aligned} c'_{\min} \|\widehat{\mathbf{x}}' - \widehat{\mathbf{x}}\|_2^2 &\leq \sum_{e \in \mathcal{E}} (c_e(\widehat{x}'_e) - c_e(\widehat{x}_e))(\widehat{x}'_e - \widehat{x}_e) \\ &= \sum_{e \in \mathcal{E}} c_e(\widehat{x}'_e)(\widehat{x}'_e - \widehat{x}_e) + \sum_{e \in \mathcal{E}} c_e(\widehat{x}_e)(\widehat{x}_e - \widehat{x}'_e). \end{aligned}$$

Let the two sums in the right-hand side be denoted by $\Psi(\widehat{\mathbf{x}}', \widehat{\mathbf{x}})$ and $\Psi(\widehat{\mathbf{x}}, \widehat{\mathbf{x}}')$, respectively. To bound $\Psi(\widehat{\mathbf{x}}', \widehat{\mathbf{x}})$, we exploit the fact that $\widehat{\mathbf{x}}'$ is a Wardrop equilibrium for the demand \mathbf{d}' . Because $(\widehat{\mathbf{y}}, \widehat{\mathbf{x}})$ is feasible for \mathbf{d} but not for \mathbf{d}' , for each type t with $d_t > 0$, we consider the rescaled flows $y_{t,s} = \widehat{y}_{t,s} d'_t / d_t$, whereas when $d_t = 0$, we simply take $y_{t,s} = \widehat{y}'_{t,s}$. Letting $x_e = \sum_{t \in \mathcal{T}} \sum_{s \in \mathcal{S}_t} y_{t,s} \mathbb{1}_{\{e \in \mathcal{S}\}}$ denote the corresponding resource loads, we have that $(\mathbf{y}, \mathbf{x}) \in \mathcal{F}(\mathbf{d}')$, and therefore,

$$\begin{aligned} \Psi(\widehat{\mathbf{x}}', \widehat{\mathbf{x}}) &= \sum_{e \in \mathcal{E}} c_e(\widehat{x}'_e)(\widehat{x}'_e - x_e) + \sum_{e \in \mathcal{E}} c_e(\widehat{x}_e)(x_e - \widehat{x}_e) \\ &\leq \sum_{e \in \mathcal{E}} c_e(\widehat{x}'_e)(x_e - \widehat{x}_e) \\ &= \sum_{t \in \mathcal{T}} \sum_{s \in \mathcal{S}_t} \left(\sum_{e \in \mathcal{S}} c_e(\widehat{x}'_e) \right) (y_{t,s} - \widehat{y}_{t,s}), \end{aligned}$$

where, in the inequality, we drop the first sum that is nonpositive because $(\widehat{\mathbf{y}}, \widehat{\mathbf{x}}')$ is a Wardrop equilibrium and (\mathbf{y}, \mathbf{x}) is feasible for \mathbf{d}' , whereas the last equality follows by expressing the resource loads in terms of the strategy flows and exchanging the order of the sums.

We now analyze each term in the outer sum over $t \in \mathcal{T}$. When $d_t > 0$, the inner double sum can be bounded as

$$\begin{aligned} \sum_{s \in \mathcal{S}_t} \left(\sum_{e \in \mathcal{S}} c_e(\widehat{x}'_e) \right) (y_{t,s} - \widehat{y}_{t,s}) &= \sum_{s \in \mathcal{S}_t} \left(\sum_{e \in \mathcal{S}} c_e(\widehat{x}'_e) \right) \widehat{y}_{t,s} (d'_t/d_t - 1) \\ &\leq \sum_{s \in \mathcal{S}_t} C \widehat{y}_{t,s} |d'_t/d_t - 1| \\ &= C |d'_t - d_t|, \end{aligned}$$

whereas when $d_t = 0$, we have $y_{t,s} = \widehat{y}'_{t,s}$ and $\widehat{y}_{t,s} = 0$ so that

$$\sum_{s \in \mathcal{S}_t} \left(\sum_{e \in \mathcal{S}} c_e(\widehat{x}'_e) \right) (y_{t,s} - \widehat{y}_{t,s}) = \sum_{s \in \mathcal{S}_t} \left(\sum_{e \in \mathcal{S}} c_e(\widehat{x}'_e) \right) \widehat{y}'_{t,s} \leq C d'_t = C |d'_t - d_t|.$$

Summing these estimates over all $t \in \mathcal{T}$, we get $\Psi(\widehat{x}', \widehat{x}) \leq C \|d' - d\|_1$. Symmetrically, we have $\Psi(\widehat{x}, \widehat{x}') \leq C \|d - d'\|_1$, from which the result follows. \square

Proof of Theorem 5. (This proof is out of order to prevent a forward reference in the next proof) Let $x_e = E_{\widehat{\sigma}}[X_e]$ so that

$$\|X_e - \widehat{x}_e\|_{L^2} \leq \|X_e - x_e\|_{L^2} + |x_e - \widehat{x}_e|, \quad (\text{A.1})$$

with

$$\|X_e - x_e\|_{L^2}^2 = \text{Var}_{\widehat{\sigma}}(X_e) = \sum_{i \in \mathcal{N}} w_i^2 \sigma_{i,e} (1 - \sigma_{i,e}) \leq \frac{1}{4} \sum_{i \in \mathcal{N}} w_i^2 \leq \frac{w d_{\text{tot}}}{4} \leq \frac{w d_{\text{max}}}{4}. \quad (\text{A.2})$$

With this bound in place, we proceed to estimate $|x_e - \widehat{x}_e|$ by showing that (\mathbf{y}, \mathbf{x}) is an ε -Wardrop equilibrium with

$$\varepsilon = \kappa(2c'_{\text{max}} + L d_{\text{max}}/4)w.$$

To this end, we first observe that, because the derivatives $c'_e(\cdot)$ are L -Lipschitz, we have

$$-\frac{1}{2}L(X_e - x_e)^2 \leq c_e(X_e) - c_e(x_e) - c'_e(x_e)(X_e - x_e) \leq \frac{1}{2}L(X_e - x_e)^2,$$

so that, taking expectations and noting that the expected value of the linear part vanishes, we get

$$|E_{\widehat{\sigma}}[c_e(X_e)] - c_e(x_e)| \leq \frac{1}{2} L \text{Var}_{\widehat{\sigma}}(X_e) \leq L w d_{\text{max}}/8. \quad (\text{A.3})$$

Now, if $y_{t,s} > 0$, then there exists some player $i \in \mathcal{N}$ with $t_i = t$ and $\widehat{\sigma}_i(s) > 0$ so that the equilibrium condition in Γ_W implies that, for each alternative strategy $s' \in \mathcal{S}_t$, we have

$$\sum_{e \in \mathcal{S}} E_{\widehat{\sigma}}[c_e(X_{i,e})] \leq \sum_{e \in \mathcal{S}'} E_{\widehat{\sigma}}[c_e(X_{i,e})]. \quad (\text{A.4})$$

Now, because $|X_{i,e} - X_e| \leq w$ and $c_e(\cdot)$ is c'_{max} -Lipschitz, by using (A.3) we get

$$|E_{\widehat{\sigma}}[c_e(X_{i,e})] - c_e(x_e)| \leq c'_{\text{max}} w + |E_{\widehat{\sigma}}[c_e(X_e)] - c_e(x_e)| \leq (c'_{\text{max}} + L d_{\text{max}}/8)w,$$

and then, we can approximate (A.4) with respect to costs:

$$\sum_{e \in \mathcal{S}} c_e(x_e) \leq \sum_{e \in \mathcal{S}'} c_e(x_e) + 2\kappa(c'_{\text{max}} + L d_{\text{max}}/8)w.$$

This shows that (\mathbf{y}, \mathbf{x}) is an ε -Wardrop equilibrium for the nonatomic game with demands \mathbf{d} , and then, invoking Proposition 2(a), we get

$$|x_e - \widehat{x}_e| \leq \|\mathbf{x} - \widehat{\mathbf{x}}\|_2 \leq \sqrt{\kappa d_{\text{max}}(2c'_{\text{max}} + L d_{\text{max}}/4)w/c'_{\text{min}}}. \quad (\text{A.5})$$

Plugging (A.5) and (A.2) into (A.1), we obtain the final estimate in (31). \square

Proof of Theorem 3. Take $(\widehat{\mathbf{y}}^n, \widehat{\mathbf{x}}^n)$, a Wardrop equilibrium for the nonatomic game Γ^∞ with demands \mathbf{d}^n . A triangle inequality gives $\|X_e^n - \widehat{x}_e\|_{L^2} \leq \|X_e^n - \widehat{x}_e^n\|_{L^2} + |\widehat{x}_e^n - \widehat{x}_e|$, so the result follows from the estimate $\|X_e^n - \widehat{x}_e^n\|_{L^2} \leq \theta \sqrt{w^n}$ in Theorem 5 and the bound $|\widehat{x}_e^n - \widehat{x}_e| \leq \Xi \sqrt{\|\mathbf{d}^n - \mathbf{d}\|_1}$ from Proposition 2(b). \square

Proof of Theorem 6. The bound $\|\mathbf{x} - \widehat{\mathbf{x}}\|_2 \leq \check{\theta} \sqrt{\Lambda(u)}$ follows from Proposition 2(a) because \mathbf{x} is an ε -Wardrop equilibrium for $\check{\Gamma}^\infty$ with $\varepsilon = 2\kappa\Lambda(u)$. Indeed, if $y_{t,s} > 0$, there exists some player $i \in \mathcal{N}$ with $t_i = t$ and $\widehat{\sigma}_i(s) > 0$ so that the equilibrium condition yields for all $s' \in \mathcal{S}_t$

$$\sum_{e \in \mathcal{S}} E_{\widehat{\sigma}}[c_e(1 + Z_{i,e})] \leq \sum_{e \in \mathcal{S}'} E_{\widehat{\sigma}}[c_e(1 + Z_{i,e})],$$

and then (21) implies the conditions for ε -Wardrop equilibrium

$$\sum_{e \in \mathcal{S}} \check{c}_e(x_e) \leq \sum_{e \in \mathcal{S}'} \check{c}_e(x_e) + 2\kappa\Lambda(u).$$

Now, taking $V_e \sim \text{Poisson}(x_e)$, Theorem C.1(a) gives $\rho_{\text{TV}}(\mathcal{L}(X_e), \mathcal{L}(V_e)) \leq u$, and (C.1) implies $\rho_{\text{TV}}(\mathcal{L}(V_e), \mathcal{L}(\widehat{X}_e)) \leq |x_e - \widehat{x}_e|$ so that (35) follows from a triangle inequality. \square

Proof of Corollary 4. This follows by considering $\widehat{X}_e^n \sim \text{Poisson}(\widehat{x}_e^n)$ with $(\widehat{\mathbf{y}}^n, \widehat{\mathbf{x}}^n)$, a Wardrop equilibrium for the game $\widehat{\Gamma}^\infty$ with demands d^n and then using a triangle inequality and applying Theorem 6 and Proposition 2(b). \square

Appendix B. Convergence of Social Costs

This appendix includes the auxiliary results on the convergence of social costs required to establish the convergence of the PoA and the PoS presented in Section 7.

B.1. Weighted Congestion Games

From Section 3, we know that the equilibria in a sequence of weighted congestion games converge to the set of Wardrop equilibria for the limit game. We prove that the corresponding social costs at equilibrium as well as the optimal social costs also converge. To this end, we start by proving that, for any sequence of converging expected flow–load pairs, the social cost of the sequence converges to that of the limiting flow–load pair.

Lemma B.1. *Let Γ_W^n be a sequence of weighted congestion games satisfying the conditions in (12) and σ^n an arbitrary sequence of mixed strategy profiles (not necessarily equilibria). Let $(Y_{i,s}^n, X_e^n)$ be the corresponding random flow–load pairs with expected values $(y_{i,s}^n, x_e^n)$. Then, along any subsequence of $(\mathbf{y}^n, \mathbf{x}^n)$ converging to some (\mathbf{y}, \mathbf{x}) , the expected social cost $\text{ESC}(\sigma^n)$ converges to $\text{SC}(\mathbf{y}, \mathbf{x}) = \sum_{e \in \mathcal{E}} x_e c_e(x_e)$.*

Proof. Take a convergent subsequence and rename it so that $(\mathbf{y}^n, \mathbf{x}^n) \rightarrow (\mathbf{y}, \mathbf{x})$. By conditioning on $\mathbb{1}_{\{e \in S_i^n\}}$, which indicates whether player i selects a strategy including e , we get

$$\text{ESC}(\sigma^n) = \sum_{e \in \mathcal{E}} \sum_{i \in \mathcal{N}^n} \mathbb{E}_{\sigma^n} [w_i^n \mathbb{1}_{\{e \in S_i^n\}} c_e(X_e^n)] = \sum_{e \in \mathcal{E}} \sum_{i \in \mathcal{N}^n} w_i^n \sigma_{i,e}^n \mathbb{E}_{\sigma^n} [c_e(X_{i,e}^n)].$$

Note that $|X_{i,e}^n - X_e^n| \leq w^n \rightarrow 0$ and $0 \leq X_e^n \leq d_{\text{tot}}^n$ with $d_{\text{tot}}^n \rightarrow d_{\text{tot}}$. Because $c_e(\cdot)$ is continuous and, hence, uniformly continuous on compact intervals, it follows that $\mathbb{E}_{\sigma^n} [c_e(X_{i,e}^n)] - \mathbb{E}_{\sigma^n} [c_e(X_e^n)]$ converges to zero uniformly in i , that is,

$$\delta_e^n := \max_{i \in \mathcal{N}^n} |\mathbb{E}_{\sigma^n} [c_e(X_{i,e}^n)] - \mathbb{E}_{\sigma^n} [c_e(X_e^n)]| \rightarrow 0.$$

Hence, using the identity

$$x_e^n = \sum_{i \in \mathcal{N}^n} w_i^n \sigma_{i,e}^n,$$

we obtain

$$\begin{aligned} \left| \text{ESC}(\sigma^n) - \sum_{e \in \mathcal{E}} x_e^n c_e(x_e^n) \right| &\leq \sum_{e \in \mathcal{E}} \sum_{i \in \mathcal{N}^n} w_i^n \sigma_{i,e}^n |\mathbb{E}_{\sigma^n} [c_e(X_{i,e}^n)] - c_e(x_e^n)| \\ &\leq \sum_{e \in \mathcal{E}} \sum_{i \in \mathcal{N}^n} w_i^n \sigma_{i,e}^n (\delta_e^n + |\mathbb{E}_{\sigma^n} [c_e(X_e^n)] - c_e(x_e^n)|) \\ &= \sum_{e \in \mathcal{E}} x_e^n (\delta_e^n + |\mathbb{E}_{\sigma^n} [c_e(X_e^n)] - c_e(x_e^n)|). \end{aligned}$$

The conclusion follows because $x_e^n \rightarrow x_e$ and $X_e^n \xrightarrow{\mathcal{D}} x_e$ so that $\mathbb{E}_{\sigma^n} [c_e(X_e^n)] \rightarrow c_e(x_e)$. \square

Using the previous lemma, we may derive the convergence of the optimal social cost.

Proposition B.1. *Let Γ_W^n be a sequence of weighted congestion games satisfying (12). Then, $\text{Opt}(\Gamma_W^n) \rightarrow \text{Opt}(\Gamma^\infty)$.*

Proof. Let $(\widetilde{\mathbf{y}}, \widetilde{\mathbf{x}})$ be a social optimum flow–load pair in the limiting game Γ^∞ . We convert the strategy flow $\widetilde{\mathbf{y}}$ into mixed strategies $\widetilde{\sigma}_t \in \Delta(\mathcal{S}_t)$ by setting

$$\forall s \in \mathcal{S}_t \quad \widetilde{\sigma}_t(s) = \widetilde{y}_{t,s} / d_t$$

when $d_t > 0$ and otherwise taking an arbitrary $\widetilde{\sigma}_t \in \Delta(\mathcal{S}_t)$ for each type with $d_t = 0$.

Let $\widetilde{\sigma}^n$ be the strategy profile for Γ_W^n in which player i plays $\widetilde{\sigma}_i^n = \widetilde{\sigma}_i$. For each t such that $d_t = 0$, we have $\widetilde{y}_{t,s}^n \rightarrow 0 = \widetilde{y}_{t,s}$ for all $s \in \mathcal{S}_t$, whereas when $d_t > 0$, we have

$$\widetilde{y}_{t,s}^n = \frac{\widetilde{y}_{t,s} d_t^n}{d_t} \rightarrow \widetilde{y}_{t,s}.$$

Hence, $(\widetilde{\mathbf{y}}^n, \widetilde{\mathbf{x}}^n)$ converges to $(\widetilde{\mathbf{y}}, \widetilde{\mathbf{x}})$, and Lemma B.1 implies that $\text{ESC}(\widetilde{\sigma}^n) \rightarrow \text{SC}(\widetilde{\mathbf{y}}, \widetilde{\mathbf{x}})$.

Now, take a sequence $\tilde{\sigma}^n$ of optimal mixed strategies in Γ_W^n and let $(\tilde{y}^n, \tilde{x}^n)$ be the corresponding expected loads. From the optimality of $\tilde{\sigma}^n$, we have $\text{ESC}(\tilde{\sigma}^n) \leq \text{ESC}(\tilde{\sigma}^n)$ so that

$$\limsup_{n \rightarrow \infty} \text{ESC}(\tilde{\sigma}^n) \leq \limsup_{n \rightarrow \infty} \text{ESC}(\tilde{\sigma}^n) = \text{SC}(\tilde{y}, \tilde{x}) = \text{Opt}(\Gamma^\infty). \tag{B.1}$$

On the other hand, taking a subsequence along which we attain the $\liminf_{n \rightarrow \infty} \text{ESC}(\tilde{\sigma}^n)$ and extracting a further subsequence so that $(\tilde{y}^n, \tilde{x}^n)$ converges to a certain limit (y, x) , it follows that

$$\liminf_{n \rightarrow \infty} \text{ESC}(\tilde{\sigma}^n) = \text{SC}(y, x) \geq \text{Opt}(\Gamma^\infty),$$

which, combined with (B.1), yields the result. \square

B.2. Bernoulli Congestion Games

The following are the analogous results for a sequence of Bernoulli congestion games.

Lemma B.2. *Let Γ_B^n be a sequence of Bernoulli congestion games satisfying the conditions in (14) and (18), and let σ^n be an arbitrary sequence of mixed strategies. Let $Y_{t,s}^n$ and X_e^n be the corresponding random loads with expected values $y_{t,s}^n$ and x_e^n . Then, along any subsequence of (y^n, x^n) converging to some (y, x) , the expected social cost $\text{ESC}(\sigma^n)$ converges to $\check{\text{SC}}(y, x) := \sum_{e \in \mathcal{E}} x_e \check{c}_e(x_e)$.*

Proof. Take a convergent subsequence and rename it so that $(y^n, x^n) \rightarrow (y, x)$. By conditioning on the event $U_{i,e}^n = 1$, we have

$$\text{ESC}(\sigma^n) = \sum_{e \in \mathcal{E}} \sum_{i \in \mathcal{N}^n} \mathbb{E}_{\sigma^n} [U_{i,e}^n c_e(X_e^n)] = \sum_{e \in \mathcal{E}} \sum_{i \in \mathcal{N}^n} u_i^n \sigma_{i,e}^n \mathbb{E}_{\sigma^n} [c_e(1 + Z_{i,e}^n)].$$

Using the identity $x_e^n = \sum_{i \in \mathcal{N}^n} u_i^n \sigma_{i,e}^n$ and invoking Lemma 1, we obtain

$$\begin{aligned} \left| \text{ESC}(\sigma^n) - \sum_{e \in \mathcal{E}} x_e^n \check{c}_e(x_e^n) \right| &\leq \sum_{e \in \mathcal{E}} \sum_{i \in \mathcal{N}^n} u_i^n \sigma_{i,e}^n |\mathbb{E}_{\sigma^n} [c_e(1 + Z_{i,e}^n)] - \check{c}_e(x_e^n)| \\ &\leq \sum_{e \in \mathcal{E}} \sum_{i \in \mathcal{N}^n} u_i^n \sigma_{i,e}^n \Lambda(u^n) \\ &= \sum_{e \in \mathcal{E}} x_e^n \Lambda(u^n) \rightarrow 0, \end{aligned}$$

and then the conclusion follows from $x_e^n \rightarrow x_e$ and the continuity of $\check{c}_e(\cdot)$. \square

Proposition B.2. *Let Γ_B^n be a sequence of Bernoulli congestion games satisfying (14) and (18). Then, $\text{Opt}(\Gamma_B^n) \rightarrow \text{Opt}(\check{\Gamma}^\infty)$.*

Proof. It suffices to repeat the proof of Proposition B.1 step by step, replacing $\text{SC}(\cdot)$ with $\check{\text{SC}}(\cdot)$ and invoking Lemma B.2 instead of Lemma B.1. \square

Appendix C. Poisson Approximation for Sums of Bernoulli Random Variables

This section collects some known facts on the Poisson approximation for sums of Bernoulli random variables. The main results are taken from Adell and Lekuona [2], Barbour and Hall [5], and Borisov and Ruzankin [10], suitably adapted to our goals. We recall that $X \sim \text{Poisson}(x)$ with parameter $x \geq 0$ if and only if

$$\mathbb{P}(X = k) = e^{-x} \frac{x^k}{k!} \quad \forall k \in \mathbb{N}.$$

As usual, we denote the law of X by $\mathcal{L}(X)$.

Two Poisson variables $X \sim \text{Poisson}(x)$ and $Y \sim \text{Poisson}(y)$ are close when $x \approx y$. In fact, their total variation distance (see (15)) can be estimated as (see Adell and Lekuona [2])

$$\rho_{\text{TV}}(\mathcal{L}(X), \mathcal{L}(Y)) \leq 1 - e^{-|x-y|} \leq |x-y|. \tag{C.1}$$

Given a function $h : \mathbb{N} \rightarrow \mathbb{R}$, for each $X \sim \text{Poisson}(x)$ with $\mathbb{E}[|h(X)|] < \infty$, we define

$$\check{h}(x) = \mathbb{E}[h(X)] = \sum_{k=0}^{\infty} h(k) e^{-x} \frac{x^k}{k!}.$$

We also define $h \mapsto \Delta h$ as the operator that takes the function h into $\Delta h(k) = h(k+1) - h(k)$. We state three useful facts that are used in the sequel.

Proposition C.1. *Let $X \sim \text{Poisson}(x)$. Then,*

- a. $\mathbb{E}[|h(Y)|] \leq e^{x-y} \mathbb{E}[|h(X)|]$ for each $Y \sim \text{Poisson}(y)$ with $y \leq x$.
- b. For each $j = 1, 2, \dots$, we have²

$$\mathbb{E}[|\Delta^j h(X)|] < \infty \iff \mathbb{E}[|h(X+j)|] < \infty \iff \mathbb{E}[X^j |h(X)|] < \infty.$$

Proof. Property (a) is just the monotonicity of $x \mapsto e^x \cdot \mathbb{E}[|h(X)|] = \sum_{k=0}^{\infty} |h(k)| \frac{x^k}{k!}$, whereas (b) can be found in Borisov and Ruzankin [10, proposition 1]. \square

Proposition C.2. Let $h : \mathbb{N} \rightarrow \mathbb{R}$ and $V \sim \text{Poisson}(d_{\max})$.

a. If $\mathbb{E}[|h(V)|] < \infty$, then $\check{h}(x)$ is well-defined and continuous for $x \in [0, d_{\max}]$.

b. If $\mathbb{E}[|\Delta^j h(V)|] < \infty$ for some $j \in \mathbb{N}$, then \check{h} is of class C^j on $[0, d_{\max}]$ and $\check{h}^{(j)}(x) = \mathbb{E}[\Delta^j h(X)]$.

Proof.

a. It suffices to note that the series

$$f(x) = e^x \check{h}(x) = \sum_{k=0}^{\infty} h(k) \frac{x^k}{k!}$$

is well-defined and continuous. This follows because the partial sums $f_n(x) = \sum_{k=0}^n h(k)x^k/k!$ converge uniformly to $f(x)$. Indeed,

$$\sup_{x \in [0, d_{\max}]} |f(x) - f_n(x)| \leq \sup_{x \in [0, d_{\max}]} \sum_{k=n+1}^{\infty} |h(k)| \frac{x^k}{k!} \leq \sum_{k=n+1}^{\infty} |h(k)| \frac{d_{\max}^k}{k!},$$

where the latter tail of the series tends to zero as $n \rightarrow \infty$ because $\mathbb{E}[|h(V)|] < \infty$.

b. Consider first the case $j = 1$. We note that the derivatives

$$f'_n(x) = \sum_{k=0}^{n-1} h(k+1) \frac{x^k}{k!}$$

converge uniformly toward $g(x) = \sum_{k=0}^{\infty} h(k+1)x^k/k!$. This follows from part (a) because, by Proposition C.1(b) with $j = 1$, we have $\mathbb{E}[|h(V+1)|] < \infty$. Hence, f is C^1 with $f'(x) = g(x)$, and then, $\check{h}(x) = e^{-x}f(x)$ is also in C^1 with

$$\check{h}'(x) = e^{-x}f'(x) - e^{-x}f(x) = \sum_{k=0}^{\infty} (h(k+1) - h(k))e^{-x} \frac{x^k}{k!} = \mathbb{E}[\Delta h(X)].$$

This establishes the case $j = 1$. Applying this property to Δh , we obtain the result for $j = 2$, and then, the cases $j = 3, 4, \dots$ follow by induction. \square

Corollary C.1. Let $V \sim \text{Poisson}(d_{\max})$ and suppose that $\mathbb{E}[|\Delta^2 h(V)|] \leq \nu < \infty$.

a. For all $x \in [0, d_{\max}]$, we have $|\check{h}'(x)| \leq (e^{d_{\max}} - 1)\nu + |h(1) - h(0)|$.

b. If $h(\cdot)$ is weakly increasing, then for all $x \in [0, d_{\max}]$, we have $\check{h}'(x) \geq 0$ with strict inequality when $h(\cdot)$ is nonconstant.

Proof. (a) Combining Propositions C.2(b) and C.1(a), we get $|\check{h}''(x)| \leq e^{d_{\max}-x}\nu$ so that, by integration, it follows that $|\check{h}'(x) - \check{h}'(0)| \leq (e^{d_{\max}} - 1)\nu$, and we conclude because $\check{h}'(0) = h(1) - h(0)$.

(b) This follows from Proposition C.2(b). \square

Let $S = X_1 + \dots + X_n$ be a sum of independent Bernoulli random variables with $\mathbb{P}(X_i = 1) = p_i \in (0, 1)$, and denote $x = \mathbb{E}[S] = p_1 + \dots + p_n$. Consider a Poisson variable $X \sim \text{Poisson}(x)$ with the same expectation. The following result shows that S and X are close when the p_i 's are small.

Theorem C.1. Let $p = \max\{p_1, \dots, p_n\}$. Then,

a. The following double inequality holds:

$$\rho_{\text{TV}}(\mathcal{L}(S), \mathcal{L}(X)) \leq (1 - e^{-x}) x^{-1} \sum_{i=1}^n p_i^2 \leq p.$$

b. If $h : \mathbb{N} \rightarrow \mathbb{R}$ is such that $\mathbb{E}[|\Delta^2 h(X)|] \leq \nu < \infty$, then

$$|\mathbb{E}[h(S)] - \mathbb{E}[h(X)]| \leq \frac{\nu x}{2} \frac{pe^p}{(1-p)^2}.$$

Proof. These properties follow from Barbour and Hall [5, theorem 1] and Borisov and Ruzankin [10, corollary 4], respectively. \square

Appendix D. List of Symbols

The following table contains the symbols used in the paper.

$\text{BNE}(\Gamma_B^n)$	Set of Bayesian Nash equilibria of the game Γ_B^n
c_e	Cost function of edge e
\check{c}_e	Cost function of edge e in the Poisson limit game, defined in (17)
C	Upper bound for $\sum_{e \in S} c_e(d_{\max})$ in Proposition 2 and $\sum_{e \in S} \check{c}_e(d_{\max})$ in Corollary 4
C_s	$\sum_{e \in S} c_e(x_e)$
c'_{\min}	Lower bound on $c'_e(\cdot)$
c'_{\max}	Upper bound on $c'_e(\cdot)$
\check{c}'_{\min}	Lower bound on $\check{c}'_e(\cdot)$
\mathbf{d}	Demand vector
d_t	Demand of type t
d_t^n	Aggregate demand of type t in the games Γ_W^n and Γ_B^n
d_{\max}	Upper bound on the demand
d_{tot}	Total demand
D	Random demand
e	Edge
\mathcal{E}	Set of resources
E_σ	Expectation induced by σ
$\text{Eq}(\Gamma^\infty)$	Equilibrium cost of Γ^∞
$\text{ESC}(\sigma)$	Expected social cost of σ , defined in (39)
$\mathcal{F}(d)$	Feasible pairs for demand d
\mathcal{G}	$(\mathcal{E}, (c_e)_{e \in \mathcal{E}}, \mathcal{T}, (\mathcal{S}_t)_{t \in \mathcal{T}})$, defined in (1)
L	Lipschitz constant for $c'_e(\cdot)$
$\mathcal{L}(X)$	Law of the random variable X
$\text{MNE}(\Gamma)$	Set of mixed Nash equilibria of Γ
\mathbf{n}	$(n_{t,s})_{t \in \mathcal{T}, s \in \mathcal{S}_t}$
N_e	$\sum_{i \in \mathcal{S}_i} U_i$, that is, random number of players who use resource e
N_t	Random number of players of type t in a Poisson game
N	$(N_t)_{t \in \mathcal{T}}$
\mathcal{N}	Set of players
$\text{Opt}(\Gamma^\infty)$	Optimum social cost in Γ^∞
$\text{Opt}(\Gamma)$	Optimum expected social cost in Γ
P_σ	Probability measure induced by σ
$\text{PNE}(\Gamma)$	Set of pure Nash equilibria of Γ
PoA	Price of anarchy
PoS	Price of stability
\mathcal{S}_t	Set of strategies for type t
\mathbf{s}	Strategy profile
s_i	Strategy of player i
S_i	Random strategy of player i
$\text{SC}(\mathbf{y}, \mathbf{x})$	Social cost of (\mathbf{y}, \mathbf{x}) , defined in (38)
\mathcal{T}	Set of types
t_i	Type of player i
u_i	Probability of player i being active
u^n	$\max_{i \in \mathcal{N}^n} u_i^n$, defined in (14b)
U_i	Indicator of player i being active
$U_{i,e}$	$U_i \mathbb{1}_{\{e \in \mathcal{S}_i\}}$
w_i	Weight of player i
w^n	$\max_{i \in \mathcal{N}^n} w_i^n$, defined in (12b)
\mathbf{x}	Load vector
x_e	Load of edge e
\mathbf{X}	Random load vector
X_e	Random load of edge e
$X_{i,e}$	$w_i + \sum_{j \neq i} w_j \mathbb{1}_{\{e \in \mathcal{S}_j\}}$, defined in (7)
$\tilde{\mathbf{x}}$	Optimum load in Γ^∞
$(\hat{\mathbf{y}}, \hat{\mathbf{x}})$	Equilibrium flow–load pair
\mathbf{y}	Flow vector
$y_{t,s}$	Flow of type t on strategy s
\mathbf{Y}	Random flow vector
$Y_{t,s}$	Random flow of type t on strategy s
$Z_{i,e}$	$\sum_{j \neq i} U_{j,e}$, defined in (10)

Appendix D. (continued)

Γ	Game
Γ_B	Bernoulli congestion game
Γ_P	Game with population uncertainty
Γ_W	Weighted congestion game
Γ^∞	Nonatomic congestion game
Δ	Difference operator
ζ	$(e^{d_{\max}} - 1)V + \max_{e \in \mathcal{E}} (c_e(2) - c_e(1))$, defined in (19)
η_t	$\min_{s \in \mathcal{S}_t} C_s$
θ	$\sqrt{d_{\max}/4} + \sqrt{2\kappa d_{\max}(c'_{\max} + Ld_{\max}/4)/c'_{\min}}$, defined in Theorem 5
$\check{\theta}$	$\sqrt{2\kappa d_{\max}/\check{c}'_{\min}}$, defined in Theorem 6
κ	cardinality of the largest feasible strategy $s \in \cup_{t \in \mathcal{T}} \mathcal{S}_t$, defined in Theorem 5
$\Lambda(u)$	$\frac{d_{\max}V}{2} \frac{ue^{u}}{(1-u)^2} + \zeta u$, defined in (20)
μ	Joint probability distribution of the number of players in a Poisson game
ν	Upper bound of $E[\Delta^2 c_e(1 + V)]$ in (18)
Ξ	$\sqrt{2C/c'_{\min}}$, defined in Proposition 2
$\check{\Xi}$	$\sqrt{2C/\check{c}'_{\min}}$, defined in Corollary 4
ρ_{TV}	Total variation distance, defined in (15)
σ	Mixed strategy profile
σ_i	Mixed strategy of player i
$\sigma_{i,e}$	$P_{\sigma}(e \in S_i)$, defined in (5)
$\hat{\sigma}$	Mixed Nash equilibrium
$\bar{\sigma}$	Optimum mixed strategy in Γ
Σ	Set of mixed strategy profiles
$\Delta(\mathcal{S}_t)$	Simplex of probability measures on \mathcal{S}_t
$\ x\ _2$	$\sqrt{\sum_{e \in \mathcal{E}} x_e^2}$
$\ d\ _1$	$\sum_{t \in \mathcal{T}} d_t $
$\ X\ _{L^2}$	$\sqrt{E[X^2]}$

Endnotes

¹ To avoid trivialities, we suppose that, for each $t \in \mathcal{T}$, the random variable N_t is not identically zero.

² Note that, if $E[X^j|h(X)] < \infty$ holds for a certain j , it also holds for $j' = 1, \dots, j$.

References

[1] Adams DJ (1936) Road traffic considered as a random series. *J. Institution Civil Engineers* 4(1):121–130.
 [2] Adell JA, Lekuona A (2005) Sharp estimates in signed Poisson approximation of Poisson mixtures. *Bernoulli* 11(1):47–65.
 [3] Angelidakis H, Fotakis D, Lianess T (2013) Stochastic congestion games with risk-averse players. Vöcking B, ed. *Algorithmic Game Theory*, Lecture Notes in Computer Science, vol. 8146 (Springer, Heidelberg, Germany), 86–97.
 [4] Barbour AD, Chen LHY, eds. (2005) *An Introduction to Stein’s Method* (Singapore University Press, Singapore).
 [5] Barbour AD, Hall P (1984) On the rate of Poisson convergence. *Math. Proc. Cambridge Philos. Soc.* 95(3):473–480.
 [6] Barbour AD, Holst L, Janson S (1992) *Poisson Approximation* (The Clarendon Press, Oxford University Press, New York).
 [7] Beckmann MJ, McGuire C, Winsten CB (1956) *Studies in the Economics of Transportation* (Yale University Press, New Haven, CT).
 [8] Bernhard P, Deschamps M (2017) On dynamic games with randomly arriving players. *Dynamic Games Appl.* 7(3):360–385.
 [9] Blumrosen L, Holenstein T (2008) Posted prices vs. negotiations: An asymptotic analysis. *Proc. Ninth ACM Conf. Electronic Commerce* (Association for Computing Machinery, New York), 49.
 [10] Borisov IS, Ruzankin PS (2002) Poisson approximation for expectations of unbounded functions of independent random variables. *Ann. Probab.* 30(4):1657–1680.
 [11] Breiman L (1962) On some probability distributions occurring in traffic flow. *Bull. Inst. Internat. Statist.* 39(4):155–161.
 [12] Breiman L (1963) The Poisson tendency in traffic distribution. *Ann. Math. Statist.* 34(1):308–311.
 [13] Buckley D (1967) Road traffic counting distributions. *Transportation Res.* 1(2):105–116.
 [14] Caserta S, de Vries C (2005) Auctions with numerous bidders. Tinbergen Institute Discussion Paper, No. 05-031/2, Amsterdam.
 [15] Chen PA, de Keijzer B, Kempe D, Schäfer G (2014) Altruism and its impact on the price of anarchy. *ACM Trans. Econom. Comput.* 2(4):1–45.
 [16] Cominetti R, Correa JR, Stier-Moses NE (2009) The impact of oligopolistic competition in networks. *Oper. Res.* 57(6):1421–1437.
 [17] Cominetti R, Scarsini M, Schröder M, Stier-Moses N (2019) Price of anarchy in stochastic atomic congestion games with affine costs. *Proc. ACM Conf. Economics and Computation* (Association for Computing Machinery, New York), 579–580.
 [18] Correa J, Hoeksma R, Schröder M (2019) Network congestion games are robust to variable demand. *Transportation Res. Part B* 119:69–78.
 [19] De Sinopoli F, Pimienta C (2009) Undominated (and) perfect equilibria in Poisson games. *Games Econom. Behav.* 66(2):775–784.
 [20] De Sinopoli F, Meroni C, Pimienta C (2014) Strategic stability in Poisson games. *J. Econom. Theory* 153:46–63.

- [21] Du L, Gong S (2016) Stochastic Poisson game for an online decentralized and coordinated parking mechanism. *Transportation Res. Part B* 87:44–63.
- [22] Feldman M, Immorlica N, Lucier B, Roughgarden T, Syrgkanis V (2016) The price of anarchy in large games. *Proc. 48th Annual ACM SIGACT Sympos. Theory Comput.* (ACM, New York), 963–976.
- [23] Fibich G, Gaviols A (2010) Large auctions with risk-averse bidders. *Internat. J. Game Theory* 39(3):359–390.
- [24] Fotakis D, Kontogiannis S, Spirakis P (2005) Selfish unsplitable flows. *Theoretical Comput. Sci.* 348(2–3):226–239.
- [25] Gairing M, Monien B, Tiemann K (2008) Selfish routing with incomplete information. *Theory Comput. Systems* 42(1):91–130.
- [26] Harks T, Klimm M (2012) On the existence of pure Nash equilibria in weighted congestion games. *Math. Oper. Res.* 37(3):419–436.
- [27] Harks T, Klimm M, Möhring RH (2011) Characterizing the existence of potential functions in weighted congestion games. *Theory Comput. Systems* 49(1):46–70.
- [28] Hassin R, Nowik I, Shaki YY (2018) On the price of anarchy in a single-server queue with heterogeneous service valuations induced by travel costs. *Eur. J. Oper. Res.* 265(2):580–588.
- [29] Haurie A, Marcotte P (1985) On the relationship between Nash-Cournot and Wardrop equilibria. *Networks* 15(3):295–308.
- [30] Jacquot P, Wan C (2018) Routing game on parallel networks: the convergence of atomic to nonatomic. Preprint, submitted April 9, <https://arxiv.org/abs/1804.03081>.
- [31] Jacquot P, Wan C (2022) Nonatomic aggregative games with infinitely many types. *Eur. J. Oper. Res.* 301(3):1149–1165.
- [32] Kordonis I, Papavasilopoulos GP (2015) LQ Nash games with random entrance: An infinite horizon major player and minor players of finite horizons. *IEEE Trans. Automatic Control* 60(6):1486–1500.
- [33] Koutsoupias E, Papadimitriou C (1999) Worst-case equilibria. *Lecture Notes in Computer Science*, vol. 1563 (Springer, Berlin), 404–413.
- [34] Lambert NS, Shoham Y (2007) Asymptotically optimal repeated auctions for sponsored search. *Proc. Ninth Internat. Conf. Electronic Commerce* (Association for Computing Machinery, New York), 55–64.
- [35] Lim W, Matros A (2009) Contests with a stochastic number of players. *Games Econom. Behav.* 67(2):584–597.
- [36] Mayne AJ (1954) Some further results in the theory of pedestrians and road traffic. *Biometrika* 41(3/4):375–389.
- [37] Meroni C, Pimienta C (2017) The structure of Nash equilibria in Poisson games. *J. Econom. Theory* 169:128–144.
- [38] Milchtaich I (2000) Generic uniqueness of equilibrium in large crowding games. *Math. Oper. Res.* 25(3):349–364.
- [39] Milchtaich I (2004) Random-player games. *Games Econom. Behav.* 47(2):353–388.
- [40] Miller AJ (1970) An empirical model for multilane road traffic. *Transportation Sci.* 4(2):164–186.
- [41] Monderer D, Shapley LS (1996) Potential games. *Games Econom. Behav.* 14(1):124–143.
- [42] Myerson RB (1998) Extended Poisson games and the Condorcet jury theorem. *Games Econom. Behav.* 25(1):111–131.
- [43] Myerson RB (1998) Population uncertainty and Poisson games. *Internat. J. Game Theory* 27(3):375–392.
- [44] Myerson RB (2000) Large Poisson games. *J. Econom. Theory* 94(1):7–45.
- [45] Myerson RB (2002) Comparison of scoring rules in Poisson voting games. *J. Econom. Theory* 103(1):219–251.
- [46] Nash JF Jr (1950) Equilibrium points in n -person games. *Proc. Natl. Acad. Sci. USA* 36(1):48–49.
- [47] Nikolova E, Stier-Moses NE (2014) A mean-risk model for the traffic assignment problem with stochastic travel times. *Oper. Res.* 62(2):366–382.
- [48] Novak SY (2019) Poisson approximation. *Probab. Surveys* 16:228–276.
- [49] Oliver RM (1961) A traffic counting distribution. *Oper. Res.* 9:802–810.
- [50] Paccagnan D, Gentile B, Parise F, Kamgarpour M, Lygeros J (2019) Nash and Wardrop equilibria in aggregative games with coupling constraints. *IEEE Trans. Automatic Control* 64(4):1373–1388.
- [51] Piliouras G, Nikolova E, Shamma JS (2016) Risk sensitivity of price of anarchy under uncertainty. *ACM Trans. Econom. Comput.* 5(1):1–27.
- [52] Roberts DJ, Postlewaite A (1976) The incentives for price-taking behavior in large exchange economies. *Econometrica* 44(1):115–127.
- [53] Rosenthal RW (1973) A class of games possessing pure-strategy Nash equilibria. *Internat. J. Game Theory* 2:65–67.
- [54] Roughgarden T (2003) The price of anarchy is independent of the network topology. *J. Comput. System Sci.* 67(2):341–364.
- [55] Roughgarden T (2015) The price of anarchy in games of incomplete information. *ACM Trans. Econom. Comput.* 3(1):1–20.
- [56] Rustichini A, Satterthwaite MA, Williams SR (1994) Convergence to efficiency in a simple market with incomplete information. *Econometrica* 62(5):1041–1063.
- [57] Sandholm WH (2001) Potential games with continuous player sets. *J. Econom. Theory* 97(1):81–108.
- [58] Satterthwaite MA, Williams SR (1989) The rate of convergence to efficiency in the buyer's bid double auction as the market becomes large. *Rev. Econom. Stud.* 56(4):477–498.
- [59] Sorin S, Wan C (2016) Finite composite games: Equilibria and dynamics. *J. Dynamic Games* 3(1):101–120.
- [60] Stidham S (2014) The price of anarchy for a network of queues in heavy traffic. Pulat PS, Sarin SC, Uzsoy R, eds. *Essays in Production, Project Planning and Scheduling: A Festschrift in Honor of Salah Elmaghraby* (Springer, Boston), 91–121.
- [61] Swinkels JM (2001) Efficiency of large private value auctions. *Econometrica* 69(1):37–68.
- [62] Wang C, Doan XV, Chen B (2014) Price of anarchy for non-atomic congestion games with stochastic demands. *Transportation Res. Part B* 70:90–111.
- [63] Wardrop JG (1952) Some theoretical aspects of road traffic research. *Proc. Inst. Civil Engineers, Part II*, vol. 1, 325–378.