

Wardrop Equilibrium Can Be Boundedly Rational: A New Behavioral Theory of Route Choice

Jiayang Li,^a Zhaoran Wang,^b Yu (Marco) Nie^{a,*}

^aDepartment of Civil and Environmental Engineering, Northwestern University, Evanston, Illinois 60208; ^bDepartment of Industrial Engineering and Management Science, Northwestern University, Evanston, Illinois 60208

*Corresponding author

Contact: jiayangli2024@u.northwestern.edu (JL); zhaoran.wang@northwestern.edu (ZW); y-nie@northwestern.edu,

<https://orcid.org/0000-0003-2083-470X> (Y(M)N)

Received: April 19, 2023

Revised: February 8, 2024; May 23, 2024

Accepted: June 3, 2024

Published Online in Articles in Advance:
July 19, 2024

<https://doi.org/10.1287/trsc.2023.0132>

Copyright: © 2024 INFORMS

Abstract. As one of the most fundamental concepts in transportation science, Wardrop equilibrium (WE) has always had a relatively weak behavioral underpinning. To strengthen this foundation, one must reckon with bounded rationality in human decision-making processes, such as the lack of accurate information, limited computing power, and suboptimal choices. This retreat from behavioral perfectionism in the literature, however, was typically accompanied by a conceptual modification of WE. Here, we show that giving up perfect rationality need not force a departure from WE. On the contrary, WE can be reached with global stability in a routing game played by boundedly rational travelers. We achieve this result by developing a day-to-day (DTD) dynamical model that mimics how travelers gradually adjust their route valuations, hence choice probabilities, based on past experiences. Our model, called cumulative logit (CumLog), resembles the classical DTD models but makes a crucial change; whereas the classical models assume that routes are valued based on the cost averaged over historical data, our model values the routes based on the cost accumulated. To describe route choice behaviors, the CumLog model only uses two parameters, one accounting for the rate at which the future route cost is discounted in the valuation relative to the past ones and the other describing the sensitivity of route choice probabilities to valuation differences. We prove that the CumLog model always converges to WE, regardless of the initial point, as long as the behavioral parameters satisfy certain mild conditions. Our theory thus upholds WE's role as a benchmark in transportation systems analysis. It also explains why equally good routes at equilibrium may be selected with different probabilities, which solves the instability problem posed by Harsanyi.

Funding: This research is funded by the National Science Foundation [Grants CMMI #2225087 and ECCS #2048075].

Supplemental Material: The online appendix is available at <https://doi.org/10.1287/trsc.2023.0132>.

Keywords: Wardrop equilibrium • bounded rationality • global stability • day-to-day dynamics

1. Introduction

Equilibrium is a fundamental instrument for understanding and analyzing social systems involving interactions between self-interested agents. A premise for transportation planning, for example, is that the individual decisions of travelers tend to bring a transportation system to equilibrium. In the simplest form, the interactions between the travelers can be framed as a noncooperative routing game in which each pursues, in the words of Simon (1955, p. 99), “the highest attainable point on his preference scale.” Wardrop (1952, p. 345) characterized the equilibrium of such a game as “the journey times on all the routes actually used are equal, and less than those which would be experienced by a single vehicle on any unused route.” He added that such equilibrium is appealing in practice because it

has a natural behavioral interpretation; “no driver can reduce his journey time by choosing a new route” (Wardrop 1952). Although Wardrop (1952) did not mention other behavioral assumptions, it is clear that his namesake equilibrium, like general equilibrium concepts in game theory (von Neumann 1928, Nash 1951), implicitly assumes *perfect rationality*, which means full and perfect information, well-defined preferences, and the capacity to compute and compare the utility of each alternative route. As Sheffi (1985, section 1.3) asserted in his celebrated book, Wardrop equilibrium (WE) implies that “motorists have full information” (i.e., they know the travel time on every possible route), that “they consistently make the correct decisions regarding route choice,” and that “they are identical in their behaviors.”

1.1. Critiques of Wardrop Equilibrium

Given the foundational role that WE plays in transportation science, its validity has been subjected to intense scrutiny. Most questions center on perfect rationality and stability.

1.1.1. Perfect Rationality. The questions about the validity of perfect rationality have been raised in economics since the 1950s (Simon 1955, Arrow 1966, Tversky and Kahneman 1985). By the early 1970s, transportation professionals began to express similar doubts. Dial (1971) cited transportation planners' desire to capture "the non-optimal behavior of trip makers" when choosing alternative routes of similar length. Daganzo and Sheffi (1977) observed that assuming that travelers always choose the shortest route can produce results unreasonably sensitive to inputs, especially in lightly congested networks. Clearly, the concern here is that because real travelers are no perfect *homo economicus*, sticking to an equilibrium that assumes that they are does not make much sense.

1.1.2. Global Stability. Beckmann, McGuire, and Winsten (1956) pointed out that to be useful, WE must be stable, or it "would be just an extreme state of rare occurrence." They mentioned both *local stability*, which ensures that the equilibrium can be restored after small perturbations, and *global stability*, which guarantees that it is reachable from any initial position (see Beckmann, McGuire, and Winsten 1956, section 3.3), although their focus was on the latter. Beckmann, McGuire, and Winsten (1956) suggested that WE may be achieved via an iterative adjustment process, in which travelers who actively search for better routes in one period base their decision on "the traffic conditions that prevailed in the preceding period." They speculated that WE is globally stable if the fraction of these "active" travelers decreases as time proceeds. Using a dynamical modeling framework that in some sense "operationalizes" this idea, many have investigated the global stability of WE since the 1980s (see, e.g., Smith 1984; Friesz et al. 1994; Zhang and Nagurney 1996; Yang and Zhang 2009; He, Guo, and Liu 2010; Guo et al. 2015).

1.1.3. Harsanyi's Instability Problem. Harsanyi's instability problem was extensively explored by game theoreticians but less known in transportation. To quote Harsanyi (1973), an equilibrium point, like WE, is inherently unstable "because any player can deviate without penalty from his equilibrium strategy even if all other players stick to theirs." To understand what he exactly meant, consider that at WE, travelers split between two routes of equal journey time at a ratio, say one to two. From the perspective of game theory, each traveler, in effect, adopts a mixed strategy that assigns a choice probability of 1/3 to one route and 2/3 to the other. However, no rational traveler should have the

incentive to stick to that mixed strategy other than a desire to keep the system at WE because the traveler can do equally well by shifting to a pure strategy that uses either route or any probabilistic mixtures of the two pure strategies.

The objection to perfect rationality was typically addressed by injecting into the model *bounded rationality*, which "takes into account the cognitive limitations of the decision maker—limitations of both knowledge and computational capacity" (Simon 1990). In transportation, bounded rationality is often linked specifically to Simon's satisficing theory (Simon 1955), which defines boundedly rational user equilibrium (BRUE) as a state where all travelers are content with, per their level of aspiration for perfection, the current (nonoptimal) travel choices (Mahmassani and Chang 1987; Mahmassani and Jou 2000; Lou, Yin, and Lawphongpanich 2010; Di et al. 2013). With the more liberal use of the term in Simon (1990), bounded rationality may also be interpreted as accepting perception errors and other sources of randomness in the system, leading to the concept of *stochastic user equilibrium* (SUE) in routing games (Daganzo and Sheffi 1977, Fisk 1980). Both SUE and BRUE are meant to be a distinct, if not better, alternative to WE. Importantly, the fact that these boundedly rational equilibria presumably converge to WE when random errors vanish or aspiration reaches the highest level does not offer a boundedly rational explanation for WE. This is because rationality is no longer bounded at the limit where SUE or BRUE becomes WE.

Using a day-to-day (DTD) dynamical model of a two-link network, Horowitz (1984) showed that the global stability of SUE depends on how travelers form their perception of current traffic conditions from past experiences. He also found that global stability is lost once random errors are set to zero. More specifically, because of discontinuity in the choice function, WE cannot be reached by his adjustment process. In a similar vein, Watling and Hazelton (2003) noted that the convergence to SUE through a dynamical process depends on a condition that becomes increasingly more stringent as perception errors become smaller. Indeed, as SUE converges to WE, it becomes impossible to meet the condition.

Harsanyi (1973) argued that bounded rationality could also solve the instability problem he identified. By assuming that each player's perception of other players' payoffs is subject to random errors, Harsanyi (1973) created a "disturbed game" that is always stable because its equilibrium only admits pure strategies. An example of Harsanyi's disturbed game is our routing game based on SUE, in which every traveler chooses the route believed to be the best (i.e., a pure strategy). Here, we note that there is only one such route for a traveler because the probability of having two or more routes that are deemed the best by the traveler is zero when

errors are continuous variables (see Daganzo and Sheffi 1977). As the random errors approach zero, the pure-strategy equilibrium of the disturbed game approaches the mixed-strategy equilibrium of the original game, and around the limit, the players would use their pure strategies approximately with the probabilities prescribed by the mixed-strategy equilibrium. However, this remedy, often known as Harsanyi's purification theorem, implies that the mixed-strategy equilibrium may only exist as an approximation to pure strategy equilibrium tied to exogenous random errors. Moreover, when errors are near zero, the difference in the payoffs between pure strategies diminishes, but the difference in the relative preferences for them, as manifested in the choice probabilities, may not. It is reasonable to expect a rational player to become increasingly indifferent to payoffs that become increasingly similar (Rosenthal 1989). However, this behavior is not captured at or near the limit of the perturbed game.

To recapitulate, the critiques on the perfect rationality assumption have generally led to an intellect exodus from WE. With bounded rationality, the definition of equilibrium is relaxed so that it can shift with the parameters chosen by the modeler to "bound" rationality. If the goal is to match observations—in terms of both individual route choices and aggregate traffic conditions—such flexibility is no doubt a blessing. On the flip side, it can also be a curse to allow the equilibrium to depend on latent behavioral variables that could vary in space, time, and population. In the case of BRUE, the equilibrium is a set rather than a singleton, even at the aggregate level and with fixed behavioral variables. Mahmassani and Chang (1987) described this lack of uniqueness as "the most disturbing question" because it "poses a dilemma for flow prediction in networks." Thus, it is hardly surprising that WE remains widely used in practice as a reliable fallback option for practitioners (Boyce and Williams 2015). WE also provides a benchmark against which alternative equilibrium concepts based on bounded rationality can be evaluated as it can often be viewed as their limit. This observation, however, lends no legitimacy to WE if one insists that rationality must always be bounded—recalling that the above limit is precisely where the bound on rationality is gone. Nor does it guarantee the global stability of WE through a behaviorally sound dynamical process compatible with bounded rationality.

1.2. Our Contribution

Motivated by the above theoretical gaps, here we set out to show that a suitable behavioral theory of route choice can resolve the seemingly innate conflict between WE and bounded rationality. Under mild conditions, the proposed theory guarantees global stability. That is, boundedly rational travelers can reach WE through a DTD dynamical process, *regardless of* initial conditions. Moreover, travelers' route choices at WE, paradoxically, are

compatible with both bounded and perfect rationality. In the parlance of game theory, this means the mixed-strategy equilibrium resulting from perfect rationality coincides with the probabilities of choosing pure strategies under bounded rationality. Therefore, the theory also solves Harsanyi's instability problem.

Our theory is built on a simple intuition; if two routes used at WE are assigned different choice probabilities, travelers must value these routes differently, even though their costs at WE are identical. To reconcile the ostensible contradiction in this statement, we conjecture that the route costs realized at WE are not the basis for deciding choice probabilities. Instead, travelers gradually build their valuation of each route through a DTD dynamical process. Consequently, the choice probabilities at WE reflect the *preferences accumulated through the entire history* of that process rather than just the experience at WE, which is achieved on the "last day" (i.e., the limit) of the process.

In our theory, the cost experienced on a route each day *accrues* to its valuation, whereas the DTD models in the literature typically view the valuation on a particular day as some average of the costs experienced up to that day (e.g., Horowitz 1984). As we shall see, this is a subtle but vital difference in a setting with an infinite horizon. On each day, travelers act with bounded rationality (i.e., they assign a choice probability to each route based on the valuations accumulated hitherto); the better the valuation, the larger the probability is. We shall prove that this dynamical process converges to WE for a rather broad class of cost accrual rules, including a naive addition rule (i.e., the valuation on day k equals the sum of the costs experienced on day $t = 1, \dots, k - 1$). Our proof requires the choice probabilities to be determined by the logit model (McFadden 1973), although other valuation-to-probability mappings may be considered as well.

When our dynamical process reaches WE, travelers would still choose the used routes with probabilities mapped by their "hidden" valuation of those routes. The benchmark route, which has the best valuation, receives the highest choice probability. Other routes are appraised against the benchmark. Travelers may be less inclined to use a route if it has a worse valuation than other routes in accordance with the notion of bounded rationality. They may also leave many routes unused. In our theory, these routes are interpreted as "unacceptable," which mathematically, means that their valuation is unboundedly worse than the benchmark's valuation.

Although our contribution is largely theoretical, the proposed dynamical process does provide a prototype algorithm for solving a broad class of routing games. This practical value of global stability analysis has been recognized early in Beckmann, McGuire, and Winsten (1956). An algorithm based on a DTD dynamical

process is simple because it requires no more information than route travel costs to operate. Therefore, it can handle routing games with more general features, such as user heterogeneity and spatial interactions in travel costs.

1.2.1. Organization. The rest of the paper is organized as follows. We discuss related works in Section 2. Section 3 sets up the routing game and describes the WE and the basic DTD dynamical model. In Section 4, we present and interpret the cumulative logit (CumLog) model. In Section 5, we prove the global stability of the CumLog model. Section 6 reports numerical experiments, and Section 7 concludes the study.

1.2.2. Notation. We use $\mathbb{R}$ and $\mathbb{R}_+$ to denote the set of real numbers and nonnegative real numbers, respectively, and we use $\bar{\mathbb{R}} = \mathbb{R} \cup \{\infty, -\infty\}$ to denote the set of extended real numbers. For a vector $\mathbf{a} \in \mathbb{R}^n$, we denote $\|\mathbf{a}\|_p$ as its ℓ_p norm and denote $\text{supp}(\mathbf{a}) = \{i : a_i > 0\}$ as its support. For a matrix $A \in \mathbb{R}^{n \times m}$, we denote $\|A\|_p$ as its matrix norm induced by the vector ℓ_p norm. For two vectors $\mathbf{a}, \mathbf{b} \in \mathbb{R}^n$, their inner product is denoted as $\langle \mathbf{a}, \mathbf{b} \rangle$. For a finite set $\mathcal{A}$, we write $|\mathcal{A}|$ as the number of elements in $\mathcal{A}$ and $2^{\mathcal{A}}$ as the set of all subsets of $\mathcal{A}$.

2. Related Studies

Our work focuses on the stability analysis of Wardrop equilibrium under the assumption of bounded rationality. In this section, we review the works that consider bounded rationality (Section 2.1) and equilibrium stability (Section 2.2) in game theory and transportation. Given the immensity of the literature that touches upon these topics, we limit our attention to those that are directly related to our work.

2.1. Bounded Rationality

Bounded rationality is the idea that human decisions are affected by “the knowledge that decision-makers do and don’t have of the world, their ability or inability to evoke that knowledge when it is relevant, to work out the consequences of their actions, to conjure up possible courses of action, to cope with uncertainty, and to adjudicate among their many competing wants” (Simon 2000). This concept has been researched extensively by game theoreticians (see Section 2.1.1) and transportation researchers (see Section 2.1.2).

2.1.1. Application in Games. The auction game of Vickrey (1961), the Bayesian game of Harsanyi (1968), and the disturbed game of Harsanyi (1973) are earlier examples of games in which players are only boundedly rational in the sense that they must deal with incomplete or imperfect information. The ε -perfect equilibrium of Selten (1975) and the ε -proper

equilibrium of Myerson (1978) also assume bounded rationality because they allow for the possibility that players choose suboptimal strategies. van Damme (1987) examined why a player may mistakenly choose suboptimal strategies. He hypothesized that making a mental effort could help avoid such mistakes at the expense of a so-called “control cost” (van Damme 1987). The trade-off between finding the optimal strategy and minimizing this effort leads to a new game with bounded rationality.

In an attempt to resolve the instability problem identified by Harsanyi (1973) (see Section 2.2.3), Rosenthal (1989) suggested another boundedly rational alternative to the standard game script. Rather than assuming that players choose the best strategy with a probability of one (which implies perfect rationality), he argued that it is sufficient if equally good strategies are played with equal probabilities and better strategies are “played with probabilities not lower than worse strategies” (Rosenthal 1989). This idea was further developed by McKelvey and Palfrey (1995) into the quantal response equilibrium (QRE; also known as boundedly rational Nash equilibrium) model, which essentially assigns choice probabilities to strategies based on the random utility theory (McFadden 1973). Since the 1990s, the QRE game has been extended to deal with—among other things—auction games (Goeree, Holt, and Palfrey 2002) and capacity allocation games (Chen, Su, and Zhao 2012) as well as Markov games (Chen et al. 2022).

The idea that players adopt inferior strategies with positive probability may also be viewed as a trade-off between exploration (gathering new information from uncharted territory) and exploitation (making the best use of information available). Björnerstedt and Weibull (1994) argued that players may need to try inferior strategies in order to ensure that they are indeed suboptimal. They proposed an imitative dynamical process (more on this in Section 2.2.3) that allows players to use demonstrably suboptimal strategies throughout the process, not because of mistakes or imperfection but because of the need for exploration. We note in passing that the exploration-exploitation trade-off is central to many machine learning (ML) algorithms, particularly bandit algorithms (Lattimore and Szepesvári 2020) and reinforcement learning (RL) algorithms (Sutton and Barto 2018). Classical exploration strategies in ML include the random selection strategy—selecting every strategy with at least a small probability, commonly known as “ ε -greedy”—and the Boltzmann exploration strategy (Kocsis and Szepesvári 2006), which assumes that the probability of each pure strategy to be selected is proportional to its exponential cost (mathematically, it is equivalent to the logit choice model). In noncooperative games, if the players gradually weigh less toward exploration than exploitation, the learning process may

be guided toward Nash equilibrium (Cohen, Héliou, and Mertikopoulos 2017).

2.1.2. Application in Transportation. The STOCH algorithm of Dial (1971) was probably the first attempt to replace traffic assignment based on WE with something that recognizes “the non-optimal behavior of trip-makers” (i.e., bounded rationality). For what it was designed to do (i.e., performing a logit-based loading in an acyclic network), the algorithm was a remarkable success. However, Dial (1971) did not conceive an alternative equilibrium concept. More importantly, when applied in traffic assignment, the logit model’s reliance on the independence of irrelevant alternative (IIA) assumption can produce nonsensical results (Florian and Fox 1976). Daganzo and Sheffi (1977) proposed eliminating the IIA dependence by subjecting travelers to a normally distributed perception error on each link. This leads to the so-called probit model that is free of most problematic predictions of its logit counterpart but is much more computationally demanding (it usually requires Monte Carlo simulation) (see, e.g., Sheffi and Powell 1981). The high computational cost of the probit model has motivated many to seek remedies within the logit framework. Most efforts aim to account for crossroute correlations, such as C-logit (Cascetta et al. 1996), path-size logit (Ben-Akiva and Bierlaire 1999), and generalized nested logit (Wen and Koppelman 2001).

Daganzo and Sheffi (1977) also introduced a boundedly rational version of WE, called stochastic user equilibrium, at which “no traveler *believes* he can improve his travel time by unilaterally changing routes.” Clearly, bounded rationality here refers to travelers’ inability to receive (or perceive) accurate information. Fisk (1980) and Sheffi and Powell (1981) established equivalent mathematical formulations for SUE problems, respectively, based on the logit model and the probit model. Conceptually, the logit-based SUE model is quite similar to the QRE model discussed in the previous section (although the QRE model was originally developed for n -person games); see Di and Liu (2016, section 4.2.1) for a comparison.

Bounded rationality may also be incorporated through Simon’s satisficing theory (Simon 1955). In the context of the morning commute, Mahmassani and Chang (1987) introduced boundedly rational user equilibrium, which is attained when all travelers are satisfied with their choices (i.e., the gap between their current cost and their optimal cost is within an *indifference band* that reflects their aspiration level). They established the conditions for the existence of a BRUE and highlighted the nonuniqueness of such equilibrium. Hu and Mahmassani (1997) incorporated indifference bands of tolerable “schedule delay” into a simulation-assignment model to study the day-to-day evolution of

network flows under real-time information and reactive signal control. Using data collected from a virtual laboratory experiment, Mahmassani and Liu (1999) confirmed the existence of the indifference band; that is, travelers would not switch routes unless the improvement in trip time exceeds a certain threshold. Mahmassani and Jou (2000) took the above virtual experiment approach one step further. They compared the findings from such experiments with those obtained from field surveys to determine the transferability of the insights. The fact that BRUE is not unique has inspired studies that attempt to characterize the BRUE set (Di et al. 2013) or to build an equilibrium selection model (Lou, Yin, and Lawphongpanich 2010). Han, Szeto, and Friesz (2015) formulated a BRUE problem that considers within-day dynamics (i.e., including both departure time and route choices) as a variational inequality problem (VIP) and proposed several solution algorithms.

2.2. Stability

The concept of stability is front and center in equilibrium analysis because equilibrium may be short lived and difficult to reach without stability, thereby rendering it a useless construct. The stability of an equilibrium can be tested by the following questions. (i) Can the equilibrium be restored after small perturbations (local stability)? (ii) Can the equilibrium be reached from any initial position (global stability)? (iii) Can agents deviate from the equilibrium without penalty (Harsanyi’s instability)? We shall focus on questions (ii) and (iii) above in this section (note that global stability implies local stability). In Sections 2.2.1 and 2.2.2, we review, respectively, classical dynamical models for games and in transportation, which were developed largely to answer the question of stability. Section 2.2.3 deals with the instability of Harsanyi (1973).

2.2.1. Learning and Evolution in Games. Brown (1951) proposed an iterative process, called fictitious play, for solving certain finite games. His method assumes that a player in each round simply responds to what they have “learned” about the other player’s strategy, represented as the empirical frequency of plays in the previous rounds. This is often viewed as the origin of the learning-based dynamical methods in games (Fudenberg and Tirole 1991). The convergence of a fictitious play to mixed equilibrium was established by Robinson (1951) for two-person zero-sum finite games and by Miyasawa (1961) for two-person general-sum finite games with two pure strategies. However, Shapley (1964) showed that fictitious play could not ensure convergence in general two-person games, thereby casting doubts on the global stability of mixed equilibrium of finite games. For finite games with bounded rationality, the stability of equilibrium is easier to establish. For

example, Fudenberg and Kreps (1993) proved that the equilibrium of the disturbed game studied in Harsanyi (1973) can be reached through a learning-based dynamical process. Similarly, Chen, Friedman, and Thisse (1997) established conditions under which players can reach the quantal response equilibrium of McKelvey and Palfrey (1995) through fictitious play.

Bush and Mosteller (1955) suggested the strategies that have worked well in the past tend to be used more as the positive experience is “reinforced” through learning (Cross 1973). In reinforcement learning, players need not form beliefs about others’ strategies; instead, they simply update their strategies in response to realized rewards. RL algorithms may be linked to human behaviors in two ways (see Camerer 2011 for more details). The first assumes that the reward received by each player directly affects their future probability of choosing the same action; the higher the current reward, the greater the increases are in the choice probability in the future (e.g., Cross 1973, Arthur 1991). The other interpretation posits that a player’s probability of selecting each action is determined by a “score” associated with the action; the higher the current reward, the larger the increment is in the score (Erev and Roth 1998, Camerer and Hua Ho 1999).

Another line of thinking in the stability analysis for games originated from emulating biological evolution. Using game theory, Smith (1982) argued that seemingly counterintuitive behaviors (e.g., cooperation and altruism) can evolve and persist in a population because they are evolutionarily stable strategies. Because the theory applies to nonhuman species (Smith and Price 1973), its validity does not rely on any form of human rationality (it is difficult to imagine ants as utility-maximizing creatures). The imitative dynamical process proposed by Björnerstedt and Weibull (1994) was an early application of evolutionary mechanisms—selection, mutation, and replication—in human competitions. It builds on a simple assumption; players tend to imitate the successful behavior of others. Specifically, players in each round switch from their current strategy a to a pure strategy b with a certain transmission probability, which increases with the utility of b as well as the number of players selecting b in the last round. Hence, a pure strategy is more attractive if it is not only more profitable but also, more popular. Björnerstedt and Weibull (1994) proved that their imitative dynamical process is locally stable. More recently, Li et al. (2022) proved the global convergence. On the one hand, the imitative process differs from fictitious play in that it implies bounded rationality. On the other hand, unlike RL, it allows players to actively learn about and act on others’ strategies. The reader is referred to Weibull (1997) and Sandholm (2010) for details on the evolutionary game theory.

2.2.2. Dynamical Models in Routing Games. The study of the route choice adjustment process, referred to as dynamical models in transportation, can be traced back to the stability analysis of WE by Beckmann, McGuire, and Winsten (1956). Most dynamical models operate on one of the following two mechanisms: (i) a discrete-time mechanism that maps travelers’ valuation of available routes to route choice in discrete decision epochs, each representing one round of the routing game (Horowitz 1984) (as the epoch is naturally a day in transportation, these models are often referred to as day-to-day models), and (ii) a continuous-time mechanism in which the decision epoch is reduced to zero so that the relation between the change rate of route flows and the current route costs may be represented as an ordinary differential equation (Smith 1984). Our model falls into the first category. The reader may consult Watling and Hazelton (2003) and Cantarella et al. (2019) for comprehensive reviews of dynamical models.

At the core of the discrete-time mechanism is modeling how travelers re-evaluate and switch routes according to past experiences. As such learning processes do not involve anticipating other players’ strategies, it is largely driven by reinforcement. Horowitz (1984) proposed that on a given day, travelers may value a route based on a weighted average of either all experienced costs on that route before that day or the cost and the valuation on the previous day. He showed that the global stability of the DTD process, even when equilibrium exists and is unique, depends on how travelers incorporate past experience into the present route valuation. Instability ensues when the rate of adjustment is not properly selected (e.g., too much weight is given to either the recent past or the distant past). The schemes of Horowitz (1984) have since been extensively studied in transportation, with most efforts centering on tuning the weights in route evaluation, enriching behavioral contents, or establishing the existence and stability of equilibrium (e.g., Cascetta 1989, Cascetta and Cantarella 1993, Cantarella and Cascetta 1995, Watling 1999, Watling and Hazelton 2003, Cantarella and Watling 2016).

Because traffic conditions are subject to perception errors in the model of Horowitz (1984), his DTD process presumably converges to SUE. He did suggest that the model may be employed to study WE when perception errors vanish but concluded that his stability results could not be readily extended to the deterministic case. Indeed, although it is well known that the fixed point of the DTD dynamical process of Horowitz (1984) is SUE (see, e.g., Watling and Hazelton 2003, section 3), establishing its convergence to WE remains elusive even at the limit (i.e., when errors become zero). The primary difficulty, as noted in Watling and Hazelton (2003), is that a discrete choice model without

additional modeling devices cannot properly distribute travelers to a set of minimal and equal-cost routes according to the proportions prescribed by WE.

By shrinking the decision epoch to zero, the continuous-time models center on moving flows between routes based on current costs. Behaviorally, this aggregate change is explained as travelers switching to routes that promise to lower their current costs. For example, the scheme proposed by Smith (1984) moves flow between every pair of routes at a rate proportional to the product of the flow on the higher-cost route and the cost difference. Using the Lyapunov theory, he proved that this simple scheme leads to a globally stable dynamical system (i.e., it produces a solution trajectory converging to WE regardless of the initial solution) if the route cost function is monotone. Similarly, the Brown–von Neumann–Nash (BNN) scheme employed by Yang (2005) shifts flow away from routes with travel costs above the weighted average of all routes. He also offered a behavioral explanation by interpreting the parameters in the BNN scheme as what Cantarella and Cascetta (1995) would call switching choice probability and route choice probability. Another widely used scheme is the so-called projected dynamical system, which may be viewed as a continuous-time version of the projection algorithm for variational inequality problems (Dupuis and Nagurney 1993). The idea is to change the route flows at a rate equal to the projection of the negative route cost vector onto the feasible set. Nagurney and Zhang (1997) noted that this movement can be driven by “travelers’ incentive to avoid more costly routes ... so that the sum of the flows equal the travel demand,” although the direct linkage between the projection operation and actual route switching behaviors is somewhat abstract. Projected dynamical processes were also employed to establish the stability of routing games with elastic demands (Friesz et al. 1994, Zhang and Nagurney 1996) based on similar conditions used in Smith (1984). The common requirement for global stability is the monotonicity of the route cost function. It is worth noting that the terminology used in these papers is asymptotic global stability, which equals stability defined in Smith (1984) and global stability concerned herein. Yang and Zhang (2009) showed that each of the above continuous-time processes is a *rational behavior adjustment process* (RBAP), which means that their flow-shifting scheme always leads to a strict reduction in the total cost with a sufficiently small step size. This observation gave rise to a class of continuous-time DTD models operating at the link level (He, Guo, and Liu 2010; Guo and Liu 2011; Di et al. 2015; Guo et al. 2015). These models have fewer behavioral contents than their route-based counterparts as their flow-shifting schemes usually rely on a target link flow pattern obtained from solving an optimization problem to

meet the RBAP requirement. More recently, Smith and Watling (2016) and Xiao et al. (2019) incorporated logit dynamics into continuous-time models and established their convergence to SUE.

2.2.3. Harsanyi’s Instability Problem. Harsanyi (1973) noted that a mixed-strategy Nash equilibrium of a finite game is inherently unstable because at the equilibrium, players can switch among equally good strategies (any of the pure strategies contained in the mixed strategy or their combinations) without penalty. If players cannot be compelled by their self-interest to *always* follow the prescription of the mixed strategy, it is difficult to sustain the equilibrium. In their celebrated book on equilibrium selection, Harsanyi and Selten (1988, section 1.6) named this *instability problem* one of the main difficulties with the concept of equilibrium in game theory. We note that WE is affected by the instability problem of Harsanyi (1973) as it is also a mixed-strategy equilibrium of a finite game.

A common remedy to the instability problem is bounded rationality, which typically means introducing random errors into payoffs (travel costs). Examples include the disturbed game of Harsanyi (1973), the SUE model of Daganzo and Sheffi (1977), and the QRE model of McKelvey and Palfrey (1995). Random errors suppress Harsanyi’s instability problem because they reduce the probability of having two pure strategies with identical costs to zero. Strictly speaking, this approach does not fix the instability problem in mixed-strategy equilibrium. It only posits that the existence of such equilibrium may be justified as an approximation to pure strategy equilibrium of the perturbed models.

Another remedy is to assume that players would never switch to an equally good strategy. Björnerstedt and Weibull (1994) argued that this assumption is implicit in Nash’s prescription. In the transportation literature, the assumption has been widely used to develop dynamical models; see, for instance, the rule of Smith (1984) that allows flow shifting to occur between two routes only when their costs are strictly different. Behaviorally, this may be explained as inertia or the tendency to settle with one’s current choices, especially when further search promises no additional benefits. However, inertia implies that travelers would never explore inferior routes. If travelers do, as assumed in most models, base their choices on what they have learned from past experience, ruling out exploration altogether seems a strong assumption that is necessary only because otherwise, Nash equilibrium (or WE) would be cursed with Harsanyi’s instability. Is the inertia assumption necessary? According to Björnerstedt and Weibull (1994), the answer is no. In their imitative dynamical process, an equally good strategy adopted by more players provides more “successful samples” for other players to imitate, thus attracting more

players in the next round. Under this mechanism, players would continue to switch between equally good strategies even at equilibrium, but the imitative dynamical process ensures that these movements do not push the system away from equilibrium. Thus, it resolves Harsanyi's instability problem without resorting to the inertia assumption or perturbation-based approximation.

2.3. Summary

To summarize what was reviewed earlier, bounded rationality may come from a decision maker's (i) inability to access accurate information (e.g., perception error), (ii) content with a suboptimal choice compatible with their level of aspiration, (iii) erroneous choices and effort to avoid them, or (iv) desire to explore seemingly suboptimal choices. In this paper, we interpret bounded rationality as imperfect choices—in the sense that decision makers allow themselves to use suboptimal strategies based on their valuation—which may be explained by any of the above four behavioral sources.

Our reading of the literature did not uncover a boundedly rational, behavior-driven, and globally stable dynamical process that can converge to WE, although such processes do exist for SUE. The convergence of continuous-time dynamical models to WE is well known. However, these models depend on perfect rationality and highly simplified learning and choice behaviors. They also need the assumption of inertia to overcome Harsanyi's instability problem.

The dynamical process proposed herein precisely fills this gap. On the one hand, it explicitly incorporates bounded rationality, learning behaviors, and individual choices. On the other hand, it always converges, for any given initial point, to a WE. Like the imitative dynamical process (Björnerstedt and Weibull 1994), our process also achieves immunity to Harsanyi's instability problem without assuming inertia. Unlike imitative dynamics, however, we do not assume that travelers know the flows on each route (i.e., the basis for imitation), which is not public information in the context of routing games.

3. Problem Setting

A routing game takes place on a transportation network modeled as a directed graph $\mathcal{G}(\mathcal{N}, \mathcal{E})$, where $\mathcal{N}$ and $\mathcal{E}$ are the set of nodes and links, respectively. Let $\mathcal{W} \subseteq \mathcal{N} \times \mathcal{N}$ be the set of origin-destination (OD) pairs and $\mathcal{K} \subseteq 2^{\mathcal{E}}$ be the set of available routes connecting all OD pairs. We use $\mathcal{K}_w \subseteq \mathcal{K}$ to denote the set of routes connecting $w \in \mathcal{W}$ and $\mathcal{E}_k \subseteq \mathcal{E}$ to denote the set of all links on route $k \in \mathcal{K}$. Also, let $\Sigma_{w,k}$ be the OD-route incidence with $\Sigma_{w,k} = 1$ if the route $k \in \mathcal{K}_w$ and zero otherwise, and let $\Lambda_{e,k}$ be the link-route incidence, with $\Lambda_{e,k} = 1$ if $e \in \mathcal{E}_k$ and zero otherwise. We write

$\Lambda = (\Lambda_{e,k})_{e \in \mathcal{E}, k \in \mathcal{K}}$ and $\Sigma = (\Sigma_{w,k})_{w \in \mathcal{W}, k \in \mathcal{K}}$. Let $\mathbf{d} = (d_w)_{w \in \mathcal{W}}$ be a vector with d_w denoting the number of travelers between $w \in \mathcal{W}$. All travelers are identical, and their route choice strategies are represented by a vector $\mathbf{p} = (p_k)_{k \in \mathcal{K}}$, where p_k equals the probability that they select $k \in \mathcal{K}_w$. The feasible region for $\mathbf{p}$ can then be written as $\mathcal{P} = \{\mathbf{p} \in \mathbb{R}_+^{|\mathcal{K}|} : \Sigma \mathbf{p} = \mathbf{1}\}$. The equilibrium of the routing game is characterized as travelers adopting a mixed strategy $\mathbf{p}$ that minimizes their own travel costs. To simplify the discussion, we assume that travelers between the OD pair adopt the same mixed strategy. According to the law of large numbers, $\mathbf{p}$ can hence be equivalently viewed as the proportion of travelers selecting each route.

Let $\mathbf{f} = (f_k)_{k \in \mathcal{K}}$ and $\mathbf{x} = (x_e)_{e \in \mathcal{E}}$, with f_k and x_e being the flow (i.e., number of travelers) on route k and link e , respectively. It follows $\mathbf{f} = \text{diag}(\Sigma^T \mathbf{d}) \mathbf{p}$ and $\Lambda \mathbf{f} = \mathbf{x}$. Further define $\mathbf{u} = (u_e)_{e \in \mathcal{E}}$ as a vector of link cost determined by a continuously differentiable function $u(\mathbf{x}) = (u_e(\mathbf{x}))_{e \in \mathcal{E}}$ (our analysis in the following sections does not require $\nabla u(\mathbf{x})$ be a diagonal or symmetric matrix). Then, the vector of route cost $\mathbf{c} = \Lambda^T \mathbf{u}$. To summarize, the route cost function $c : \mathcal{P} \rightarrow \mathbb{R}^{|\mathcal{K}|}$ can be defined as $c(\mathbf{p}) = \Lambda^T \mathbf{u} = \Lambda^T u(\Lambda \mathbf{f}) = \Lambda^T u(\bar{\Lambda} \mathbf{p})$, where $\bar{\Lambda} = \Lambda \text{diag}(\Sigma^T \mathbf{d})$.

3.1. Wardrop Equilibrium

A Wardrop equilibrium (Wardrop 1952) of the routing game can be defined as follows.

Definition 1 (Wardrop Equilibrium). A route choice strategy $\mathbf{p}^* \in \mathcal{P}$ is a WE strategy if $c_k(\mathbf{p}^*) > \min_{k' \in \mathcal{K}_w} c_{k'}(\mathbf{p}^*)$ implies that $p_k^* = 0$ for all $w \in \mathcal{W}$ and $k \in \mathcal{K}_w$.

In other words, a route included in a WE strategy must have the minimum cost. It is widely accepted that travelers must be perfectly rational to reach and keep a WE strategy, which in our context, means they always know the precise values of all route costs and consistently make correct choices accordingly (Sheffi 1985).

The WE routing game has an equivalent variational inequality problem (Dafermos 1980).

Proposition 1 (VIP Formulation of WE). A route choice strategy $\mathbf{p}^* \in \mathcal{P}$ is a WE strategy if and only if it solves the following VIP. Find $\mathbf{p}^* \in \mathcal{P}$ such that

$$\langle \mathbf{c}(\mathbf{p}^*), \mathbf{p} - \mathbf{p}^* \rangle \geq 0, \quad \forall \mathbf{p} \in \mathcal{P}. \quad (1)$$

We shall denote the solution set to the VIP (1) as $\mathcal{P}^*$, referred to as the WE strategy set. It is well known that $\mathcal{P}^*$ is a singleton if $c(\mathbf{p})$ is strongly monotone and a polyhedron if only $u(\mathbf{x})$ is (Dafermos 1980). In the latter case, although many WE strategies may exist, they must correspond to the same link flow $\mathbf{x}^*$. The questions that concern us here are *whether and how a WE strategy can always be achieved under reasonable assumptions of route choice behaviors*.

Since Beckmann, McGuire, and Winsten (1956), many have asked these questions and have largely settled the “whether” part. Specifically, it has been established that a WE strategy can be reached through a continuous-time dynamical process starting at any initial point, provided that the route cost function $c(p)$ is monotone (see, e.g., Smith 1984, Dupuis and Nagurney 1993, Friesz et al. 1994, Nagurney and Zhang 1997). Yet, the answer to the question of “how” is complicated by two issues. The first is behavioral. Note that travelers in the continuous-time models are supposed to be highly rational; they have accurate knowledge of and act on the most recent route costs to perpetually switch from higher-cost routes to lower-cost ones until a WE is reached. Therefore, continuous-time models have a limited capacity to accommodate such behaviors as learning from past experiences, exploring suboptimal routes, and indifference to equal-cost routes. Second and perhaps more importantly, global stability is secured by implicitly applying an arbitrarily small rate of adjustment to the prescribed “direction” of route flow changes (Watling 1999). However, whether the analysis is employed for the purpose of explaining real-world route-switching behaviors or of developing an equilibrium-finding algorithm, that rate cannot always be arbitrarily small. The questions are as follows. What is a suitable magnitude of the adjustment at a given time? How fast should this magnitude decrease as time proceeds? How is this pattern of time-varying adjustments related to route choice behaviors? The discrete-time dynamical models are better equipped to address these questions.

3.2. Discrete-Time Dynamical Model

At its core, a discrete-time dynamical or day-to-day model keeps track of travelers’ route valuation vector on day t , denoted as $s^t \in \mathbb{R}^{|\mathcal{K}|}$ ($t = 0, 1, \dots$), which is mapped to their route choice strategy $p^t \in \mathcal{P}$ by a function. A commonly used route choice function is based on the logit model derived from random utility theory (McFadden 1973, Ben-Akiva and Lerman 1985). Given a scalar $r > 0$, a logit-based route choice function $q_r : \mathbb{R}^{|\mathcal{K}|} \rightarrow \mathcal{P}$ gives $p^t = q_r(s^t)$, where

$$p_k^t = \frac{\exp(-r \cdot s_k^t)}{\sum_{k' \in \mathcal{K}_w} \exp(-r \cdot s_{k'}^t)}, \quad \forall k \in \mathcal{K}. \quad (2)$$

By manipulating how s^t is constructed and updated, the DTD models can represent a wide range of learning and choice behaviors. Below, we briefly review the two most popular models.

The weighted-average model assumes s^t be a weighted average of the costs received in the past by setting

$$s^t = \sum_{i=0}^{t-1} \eta^i \cdot c(p^i), \quad \text{with} \quad \sum_{i=0}^{t-1} \eta^i = 1, \quad (3)$$

where $\eta^i \geq 0$ weighs how the cost received on day i ($0 \leq i \leq t-1$) affects the valuation on day t . Thus, the entire history of past experiences is allowed to affect the present-day decision.

The successive-average (SA) model, as a simplification of the weighted-average model, sets

$$s^t = (1 - \eta^t) \cdot s^{t-1} + \eta^t \cdot c(p^{t-1}), \quad (4)$$

where $\{\eta^t \in (0, 1); t = 1, 2, \dots\}$ is a sequence of constants. In this model, the past experience is condensed into yesterday’s valuation. This decision mode imposes a much lower information burden on travelers as it claims no direct memory of the experience prior to yesterday.

In the literature, both models were initially discussed by Horowitz (1984). The latter can be viewed as a special case of the former, noting that recursively applying Equation (4) yields

$$s^t = \eta^t \cdot c(p^{t-1}) + (1 - \eta^t) \cdot \eta^{t-1} \cdot c(p^{t-2}) + \dots + \prod_{i=2}^t (1 - \eta^i) \cdot \eta^1 \cdot c(p^0). \quad (5)$$

Horowitz (1984) assumed the travelers’ perception of s^t is subject to a random error $\epsilon^t \in \mathbb{R}^{|\mathcal{K}|}$. He considered two possibilities for the distribution of ϵ^t ; the first (Model 1) assumes that the distribution of ϵ^t is independent of t , whereas the second (Model 2) treats ϵ^t as the sum of the perception errors in the past (hence, a function of t). Model 1 is much easier to analyze because it allows us to treat the parameter r in the logit Model (2) as a time-invariant constant. In this case, if (p^t, s^t) converges to a fixed point $(\hat{p}, \hat{s})$, then we have $\hat{s} = c(\hat{p})$ and $\hat{p} = q_r(\hat{s})$, and hence, $\hat{p} = q_r(c(\hat{p}))$. Following Daganzo and Sheffi (1977), a route choice p^* that satisfies this equation is a stochastic user equilibrium. The global stability of SUE under the successive-average model has been extensively studied. For a two-link network, Horowitz (1984) analyzed the global stability of the dynamical Model (4) under the assumptions that the perception error is nonzero and the cost function is both monotone and Lipschitz continuous. Here, we note that the original Lipschitz continuous assumption of Horowitz (1984) is imposed on a composite function that combines $c(p)$ and the distribution function of ϵ^t . This assumption can always be satisfied when $c(p)$ is Lipschitz continuous, and the perception errors are not reduced to zero (that is, the stochastic model is not degraded to a deterministic one). He proved that the model is globally stable if (i) $\sum_{i=1}^{\infty} \eta_i = \infty$ and (ii) η_t becomes sufficiently small for a sufficiently large t . Cantarella and Cascetta (1995) extended Horowitz’s analysis to general networks but limited the stability analysis to the case where η^t is time invariant. Like Cascetta and Cantarella (1993), they introduced the

switching probability α to describe the likelihood that a traveler would even consider route choice on a given day. Thus, the model in Horowitz (1984) can be viewed as a special case of their model when $\alpha = 1$ on every day. The global stability of their model can be guaranteed when either α or η is sufficiently small; a small α may be interpreted as strong habitual inertia, whereas a sufficiently small η is the same requirement as in Horowitz (1984). Watling (1999) analyzed the local stability of the model with a constant η and investigated the possibility of applying Lyapunov's theory to determine its domain of attractions.

It is well known that SUE can be made arbitrarily close to WE by letting $r \rightarrow \infty$ (Fisk 1980; Erlander 1998; Mamun, Xu, and Yin 2011). However, it remains an open question whether the stability result of SUE is applicable to WE at that limit. As noted by Watling and Hazelton (2003, example 4), given a fixed $r > 0$, the stability of the successive-average model requires η to be within $(0, \bar{\eta}_r]$, where $\bar{\eta}_r$ decreases to zero when $r \rightarrow \infty$. Thus, attempting to approximate WE with SUE by choosing an arbitrarily large r is problematic as the feasible range of η needed for convergence vanishes at the limit.

Is it possible to design a scheme that coordinates the increase of r and the decrease of η so that the final stationary point is steered toward WE? To the best of our knowledge, this question has not received much attention in the literature. Even if such a coordinated scheme can be identified, a more important question is how to make sense of it *behaviorally*. Specifically, why would the travelers couple the changes in r and η in such a way as prescribed by the stability analysis?

In the next section, we shall propose an alternative to the classical DTD dynamical models that is deceptively simple at first glance but holds promise to answer the above questions.

4. Cumulative Logit Model

We are now ready to propose a new DTD dynamical system that is dubbed, for the reasons that will soon become clear, the cumulative logit model. In developing the CumLog model, we were inspired by the conjecture in Beckmann, McGuire, and Winsten (1956) about the study of stability (the emphasis is our own).

Through a simple and plausible model, one can get a rough picture of the *minimum of conditions* that must be met in order that the adjustment process should converge to equilibrium. (Beckmann, McGuire, and Winsten 1956)

Indeed, the overarching goal of this study is to develop that simple and plausible model envisioned by Beckmann, McGuire, and Winsten (1956) and to identify “the minimum of conditions” that ensures the convergence of a dynamical adjustment process to WE.

The CumLog model adopts the basic framework of the DTD Model (3). That is, on day t , travelers update

their route valuation vector s^t based on the route costs on day $t - 1$ and select a route choice strategy p^t according to s^t . Unlike (3), s^t in the CumLog model is updated, starting from some $s^0 \in \mathbb{R}^{|\mathcal{K}|}$, as follows:

$$s^t = s^{t-1} + \eta^t \cdot c(p^{t-1}), \quad (6)$$

where $\{\eta^t, t = 1, 2, \dots\}$ is a sequence of positive constants. Moreover, travelers determine their strategy on day t according to the logit model by setting $p^t = q_r(s^t)$ according to Equation (2). The simplest interpretation of CumLog is that travelers value routes based on the cost received and *accumulated* over the entire history up to $t - 1$. Indeed, in the special case of $\eta = 1$, travelers literally add up all received costs without ever discounting the experiences in the distant past.


Upon noticing the suspicious similarity between Schemes (6) and (4), some readers may understandably question the plausibility of our central claim: that Scheme (6) can somehow ensure convergence to WE under mild conditions, whereas as widely asserted in the literature, Scheme (4) cannot. Therefore, in what follows, we shall first explain why averaging and accumulating route costs are fundamentally different in the dynamical process (Section 4.1). Section 4.2 provides a behavioral interpretation of the CumLog model. An illustrative example is given in Section 4.3.

4.1. Difference Between Average and Accumulation

Let us revisit the thought experiment used to demonstrate Harsanyi's instability problem in Section 1. Suppose a routing game in a two-route network converges to a WE strategy that assigns route 1 and route 2 the choice probabilities of $1/3$ and $2/3$, respectively. Panels (a) and (b) in Figure 1 depict, respectively, how the flows and costs on the two routes gradually reach the WE through a dynamic adjustment process. The details of the process need not concern us here. Suffice it to say that at the end of the process, the costs on both routes are identical, and the probability of choosing route 1 becomes $1/3$, which implies that a third of the travelers end up using that route.

Figure 1(c) compares the evolution of the difference in the two route valuation schemes (average versus accumulation) accompanying the dynamical process. In the average scheme, the difference in the route evaluations is bound to vanish regardless of the weights if the process ends at the WE. Given the same route valuations, it is difficult to see why at WE, travelers seem to prefer route 2 over route 1 with a 2:1 margin as suggested by the mixed strategy. Indeed, as Harsanyi (1973) pointed out, there is no way to explain that preference other than insisting that travelers prefer WE itself. This seemingly illogical preference is at the heart of Harsanyi's instability. Nor can this problem be explained away by bounded rationality. Note that

Figure 1. (Color online) Illustration of a Convergence Process Toward WE in a Two-Route Routing Game



Notes. (a) Evolution of flows. (b) Evolution of costs. (c) Valuation differences.

the errors in the valuation would affect choice probability only when there is a nonzero difference in the “deterministic” part of the valuations. At WE, the difference becomes zero. Hence, no errors could, on their own, swing travelers one way or the other. Unlike in the average scheme, the discrepancy in route valuations converges to a nonzero constant in the accumulation scheme (see the solid line with x markers in Figure 1(c)). The constant equals the shaded area in Figure 1(b), which visualizes the valuation difference accrued through the dynamical process. This cumulative difference then explains the mixed WE strategy; the travelers prefer route 2 at equilibrium because they value it substantially (but not infinitely) more than route 1, although the equilibrium route costs are the same.

Remark 1. In behavioral economics, the reinforcement learning model proposed by Erev and Roth (1998) assumes that players in a finite game have the highest propensity to choose the pure strategy that gives the greatest total reward in the past. This assumption is similar to our assumption that the route with a smaller cumulative route cost is selected with a greater probability.

4.2. Behavioral Interpretation

The route choice behaviors implied by the CumLog model can be summarized as follows.

- On each day, travelers choose *afresh* a mixed-route choice strategy based on current route valuations through a logit model. No additional assumptions are needed regarding the behavioral inertia—the reluctance to make changes once a choice becomes habitual—that is often explicitly modeled by a switching probability in the DTD literature (e.g., Cantarella and Cascetta 1995).
- Route choices are driven by relative valuation rather than absolute valuation as dictated by the logit model. The benchmark is the “best” route that receives the highest choice probability. Other routes are appraised against the benchmark. They shall also be selected with a strictly positive probability as long as the valuation difference between these routes and the

benchmark is finite; they may be selected with a zero probability if their valuations are deemed unacceptable (i.e., *infinitely* worse than the benchmark).

Here is the rationale behind our claim that *the CumLog model is boundedly rational*; rather than committing to never moving to a route with a worse valuation than they currently enjoy, the travelers consistently assign a nonzero probability to suboptimal routes with an *acceptable* valuation.

The two parameters of the CumLog model can be linked to route choice behaviors as follows.

- The parameter r characterizes how travelers’ route choice strategy p^t is determined from the valuation vector s^t . It measures the trade-off between exploration and exploitation; the larger the parameter r , the more exploitative the travelers (meaning that they are less likely to explore suboptimal routes). Thus, r will be referred to as the *exploitation parameter*. In this study, the parameter r is fixed at a constant value. One may interpret this setting in one of two ways. (i) The perception errors are independent of t , which is the assumption used to justify in Horowitz (1984, model 1), or (ii) travelers’ propensity for accepting suboptimal routes or their desired balance between exploration and exploitation is time invariant (Fudenberg and Kreps 1993, Kocsis and Szepesvári 2006). Our stability analysis is agnostic on the interpretation of the exploitation parameter.

- The parameter η^t regulates how the valuation vector s^t is updated. Before the routing game is played, travelers have an initial route valuation vector s^0 . If no prior preference exists, then they simply set $s_k^0 = 0$ for all $k \in \mathcal{K}$. On day $t \geq 1$, the travelers update the valuation s^t by adding to it a cost vector $\eta^t \cdot c(p^{t-1})$, where $\eta^t \geq 0$ is the weight on day t . The weight η^t measures the impact of the cost received on day $t - 1$ on the travelers’ valuation on day t . Behaviorally, it captures how quickly travelers become disposed to ignore the latest information and “settle down.” Thus, η^t will be referred to as the *proactivity measure*. The larger the η^t , the less passive the travelers. As we shall see, the stability of the CumLog model depends on the asymptotic

behavior of the proactivity measure η^t . Not all feasible sequences of η^t guarantee convergence to WE. For example, if travelers stop incorporating new information into route valuation too soon (indicating a rapid descent to extreme passivity), CumLog may stabilize quickly but at a place far away from WE. We shall consider two asymptotic rules for the proactivity measure in this study. In the first, $\eta^t = \eta > 0$ for all $t \geq 0$ (i.e., the level of proactivity remains at a constant level through the entire process). The second rule dictates that η^t monotonically decreases to zero as $t \rightarrow \infty$. Thus, the costs received by the travelers will have a progressively diminishing impact on their route choice. Another rule, in which η^t converges to some constant $\eta > 0$ as $t \rightarrow \infty$, may be inferred from the above two.

4.3. Illustrative Example

We close by illustrating CumLog with a simple routing game played on a network with three parallel links connecting an OD pair. The cost functions on the links are $u_1(x_1) = x_1$, $u_2(x_2) = x_2 + 1$, and $u_3(x_3) = x_3 + 2.25$. The total demand is $d = 3$. It can be easily verified the WE conditions dictate that the three links be selected with probabilities $p_1^* = 2/3$, $p_2^* = 1/3$, and $p_3^* = 0$. The WE strategy is unique in this case because all route cost functions are strictly increasing.

Setting $r = 0.25$, fixing $\eta^t = 1$, and starting from $s^0 = 0$, we run the CumLog model from day 0 to day 12 and report the convergence process in Figure 2. A WE is reached after day 12, with the proportions of travelers selecting each route converging rather precisely to $2/3$, $1/3$, and 0 , respectively, and the costs on the two routes included in the mixed strategy, routes 1 and 2, become identical.

In Figure 2(b), the shadowed areas with and without hatches highlight the difference in valuation between routes 1 and 2 and that between routes 2 and 3, respectively. Route 1 is always the lowest-cost route throughout the process, and thus, it is always selected by most travelers. The shadowed area without hatches in Figure 2(b) approaches a constant value as t increases.

Consequently, the relative preference for route 1 over route 2 became stabilized, indicating that route 2 is an inferior but acceptable option. The red area in Figure 2(b), however, grew to infinity as $t \rightarrow \infty$, which means that route 3 became infinitely worse than route 2 and eventually was abandoned.

This example exhibits another distinction between the CumLog model and the classical DTD models. Even with a finite exploitation parameter r , the CumLog model is capable of identifying and eliminating the routes that no WE strategy should use. The classical models, however, are obliged by their averaging scheme to keep a positive flow on every route unless $r \rightarrow \infty$. This result can be expected from the fact that the limiting point of these models is SUE rather than WE.

5. Global Stability


In this section, we present and prove the main stability result concerning the CumLog model proposed in Section 4. Simply put, the objective is to show that under mild conditions, the DTD dynamical Model (6) always converges to a WE regardless of the initial point. The following assumptions describe some of the conditions on the link cost function $u(x)$, whose domain (the set of feasible link flows) is written as $\mathcal{X} = \{x : \mathbb{R}^{|\mathcal{A}|} : x = \bar{\Lambda}p, p \in \mathcal{P}\}$.

Assumption 1. The link cost function $u(x)$ is twice continuously differentiable on $\mathcal{X}$.

Assumption 2. For all $x \in \mathcal{X}$, the symmetric parts of both $\nabla u(x)$ and $(\nabla u(x))^2$ are positive semidefinite.

Assumption 2 is satisfied as long as $u(x)$ is monotone and $\nabla u(x)$ is symmetric. If $\nabla u(x)$ is asymmetric, Assumption 2 still holds if $\nabla u(x)$ is not “too asymmetric” (i.e., the antisymmetric part does not exceed the symmetric part) (see Hammond and Magnanti 1987 for a more rigorous description). As the assumptions require neither $u(x)$ to be strongly monotone nor $\nabla u(x)$ to be symmetric, our analysis can be applied to a broad class of routing games, including those with nonseparable and

Figure 2. (Color online) Illustration of a Convergence Process Toward WE in a Three-Route Routing Game



Notes. (a) Route choices. (b) Travel costs.

nonstrictly increasing cost functions. The next proposition characterizes the route cost function $c(\mathbf{p})$ and the WE strategy set $\mathcal{P}^*$.

Proposition 2. *Under Assumptions 1 and 2, (i) there exists $L \geq 0$ such that $c(\mathbf{p})$ is $1/4L$ -cocoercive on $\mathcal{P}$ (i.e., $\langle c(\mathbf{p}') - c(\mathbf{p}), \mathbf{p}' - \mathbf{p} \rangle \geq 1/4L \cdot \|c(\mathbf{p}') - c(\mathbf{p})\|_2^2$ for all $\mathbf{p}, \mathbf{p}' \in \mathcal{P}$); (ii) $\langle c(\mathbf{p}), \mathbf{p} - \mathbf{p}^* \rangle \geq 0$ for all $\mathbf{p} \in \mathcal{P}$ and $\mathbf{p}^* \in \mathcal{P}^*$, with equality for any given $\mathbf{p}^* \in \mathcal{P}^*$ if and only if $\mathbf{p} \in \mathcal{P}^*$; and (iii) $\mathcal{P}^*$ is convex and compact.*

Proof. Following the standard convex analysis techniques, we can show Assumptions 1 and 2 $\Rightarrow$ (1) $\Rightarrow$ (2) $\Rightarrow$ (3). The reader is referred to Online Appendix A.1 for details. $\square$

We note that when $\nabla c(\mathbf{p})$ is symmetric, L is the Lipschitz constant of $c(\mathbf{p})$ on $\mathcal{P}$ (i.e., any $L \geq \max_{\mathbf{p} \in \mathcal{P}} \|\nabla c(\mathbf{p})\|_2$ can be used to fulfill the requirement). The case with asymmetric $\nabla c(\mathbf{p})$ requires $L = \max_{w \in \mathcal{W}} \{d_w\} \cdot H \cdot \|\mathbf{A}\|_2^2$, where H is the Lipschitz constant of $u(\mathbf{x})$ on $\mathcal{X} = \{\mathbf{x} \in \mathbb{R}_+^{|\mathcal{A}|} : \mathbf{x} = \bar{\mathbf{A}}\mathbf{p}, \mathbf{p} \in \mathcal{P}\}$.

We are now ready to give the main stability result.

Theorem 1. *Under Assumptions 1 and 2, suppose that $s^0 < \infty$; then, $\mathbf{p}^t$ in the CumLog Model (6) converges to a fixed point $\mathbf{p}^* \in \mathcal{P}^*$, the solution set to the VIP (1) if either of the following two conditions is satisfied: (i) $\lim_{t \rightarrow \infty} \eta^t = 0$ and $\lim_{t \rightarrow \infty} \sum_{i=0}^t \eta^i = \infty$ or (ii) $\eta^t = \eta < 1/2rL$ for all $t \geq 0$.*

Below, Section 5.1 interprets the convergence result, whereas Section 5.2 sketches the proof of Theorem 1 (the complete proof is given in Online Appendices A.2–A.5).

5.1. Interpretation of the Convergence Result

5.1.1. Convergence Conditions

The two convergence conditions given by Theorem 1 can be interpreted as follows.

- The first part of condition (i) in Theorem 1 implies that the attention paid to the latest route cost decreases with t and eventually disappears altogether. In other words, travelers tend to settle down in the long run. The second part states that convergence may be at risk if travelers settle down too soon. Specifically, the decreasing rate of η^t cannot be faster than $\mathcal{O}(t^{-1})$. The following are a few examples that meet this “not-too-soon” requirement: $\eta^t = \mathcal{O}(1/t)$, $\eta^t = \mathcal{O}(1/\sqrt{t})$, or $\eta^t = \mathcal{O}(1/\log(t))$. The second part is introduced to ensure that the routes not included in any WE strategy become infinitely worse than the best route when $t \rightarrow \infty$. If the condition is violated, say, for example, $\eta^t = \mathcal{O}(t^{-(1+\delta)})$ for some $\delta > 0$, the valuation s^t would remain bounded when $t \rightarrow \infty$. The monotone convergence theorem then guarantees the convergence of s^t . Denoting its limit as $\bar{s}$, we can show that the difference

between any elements of $\bar{s}$ is bounded because $\bar{s}$ itself is bounded. As a result, per the logit model, all routes are bound to receive positive flow as $t \rightarrow \infty$, which in general, violates the WE conditions.

- Condition (ii) in Theorem 1 means that travelers would never stop incorporating new information into route valuations. Rather, their propensity for proactivity is maintained at a constant level below a certain threshold. Compared with condition (i) in Theorem 1, this is a weaker requirement, and thus, the convergence under it is more difficult to establish. Condition (ii) in Theorem 1 may also be interpreted as given a fixed level of proactivity η , the exploitation parameter r must not exceed $1/2\eta L$. At first glance, this upper bound on r is puzzling because one would expect that a larger r makes it easier to reach WE because in theory, WE is the limiting case of SUE when $r \rightarrow \infty$. However, a moment of reflection reveals that in a dynamical process, a large and constant r means that travelers tend to underexplore the route space, which might prevent them from reaching WE. In learning theory, it is well known that a sufficient level of exploration in the early stage is critical to effective exploitation in the later stage (Lattimore and Szepesvári 2020).

5.1.2. Significance. The above interpretation of the convergence conditions indicates that CumLog can reach Wardrop equilibrium under relatively mild assumptions about boundedly rational route choice behaviors. This finding challenges the long-standing belief in transportation research that WE has a shaky behavioral foundation (see, e.g., Watling and Hazelton 2003).

Our theory also resolves Harsanyi’s instability problem, which manifests as follows. When WE is reached, all used routes are equally good. Hence, breaking the tie arbitrarily is not against anyone’s interest. However, if travelers do break the tie arbitrarily, the system cannot stay at WE. The CumLog model solves this dilemma by allowing travelers to assign a different valuation (hence, a different choice probability) to routes that have identical costs at WE. Specifically, when CumLog reaches a WE strategy $\mathbf{p}^*$, given any OD pair $w \in \mathcal{W}$ and two WE routes $k, k' \in \mathcal{K}_w$ between that pair, travelers’ valuation difference between routes k and k' would be $-(\log(p_k^*) - \log(p_{k'}^*))/r$. This valuation difference explains why two identical-cost routes are selected with different probabilities.

5.2. Sketch of the Proof

To sketch the proof of Theorem 1, let us define $\mathbf{p}_w = \{p_k\}_{k \in \mathcal{K}_w}$ and the negative entropy function as $\phi_w(\mathbf{p}_w) = \langle \mathbf{p}_w, \log(\mathbf{p}_w) \rangle$. Also, define $\tilde{\eta}^t \equiv r \cdot \eta^t$ and $\mathcal{Q}_w^t = \{\mathbf{p}_w \in \mathcal{P}_w : \text{supp}(\mathbf{p}_w) \subseteq \text{supp}(\mathbf{p}_w^t)\}$. We shall prove Theorem 1 by showing that there exists $\mathbf{p}^* \in \mathcal{P}^*$ such that the distance between $\mathbf{p}^*$ and $\mathbf{p}^t$ converges to zero when $t \rightarrow \infty$. Because the decision variable is defined on

$\mathcal{P} = \prod_{w \in \mathcal{W}} \mathcal{P}_w$, the Cartesian product of a set of probability simplex $\mathcal{P}_w = \{\mathbf{p}_w \in \mathbb{R}_+^{|\mathcal{K}_w|} : \mathbf{1}^T \mathbf{p}_w = 1\}$, the Kullback–Leibler divergence (also known as the statistical distance) is a natural measure given by

$$\begin{aligned} D(\mathbf{p}^*, \mathbf{p}^t) &= \sum_{w \in \mathcal{W}} D_w(\mathbf{p}_w^*, \mathbf{p}_w^t) \\ &= \sum_{w \in \mathcal{W}} \phi_w(\mathbf{p}_w) - \phi_w(\mathbf{p}_w^t) - \langle \nabla \phi_w(\mathbf{p}_w^t), \mathbf{p}_w - \mathbf{p}_w^t \rangle. \end{aligned} \quad (7)$$

The proof is divided into three steps.

Step 1. We first establish the relation between $\mathbf{p}^{t+1}$ and $\mathbf{p}^t$ in the CumLog dynamical process.

Lemma 1. *Given any $r > 0$, $\eta^t > 0$, and $\mathbf{s}^t \in \bar{\mathbb{R}}^{|\mathcal{K}|}$, let $\mathbf{p}^t = q_r(\mathbf{s}^t)$. Then, for a vector $\tilde{\mathbf{p}} \in \mathcal{P}$, we have $\tilde{\mathbf{p}} = q_r(\mathbf{s}^t + \eta^t \cdot c(\mathbf{p}^t))$ if and only if*

$$\langle \nabla \phi_w(\mathbf{p}_w^t) - \nabla \phi_w(\tilde{\mathbf{p}}_w), \mathbf{p}_w - \tilde{\mathbf{p}}_w \rangle \leq \tilde{\eta}^t \cdot \langle c_w(\mathbf{p}^t), \mathbf{p}_w - \tilde{\mathbf{p}}_w \rangle, \quad \forall \mathbf{p}_w \in \mathcal{Q}_w^t, \quad \forall w \in \mathcal{W}. \quad (8)$$

Proof. The proof follows from two well-known results: (1) the equivalence between the logit model and a convex program (Proposition A.14 in the Online Appendix) and (2) the equivalence between the convex program and a VIP (Lemma A.3 in the Online Appendix). The reader is referred to Online Appendix A.2 for the complete proof. $\square$

Step 2. We next link $D(\mathbf{p}^*, \mathbf{p}^{t+1})$ to $D(\mathbf{p}^*, \mathbf{p}^t)$.

Lemma 2. *If $s^0 < \infty$, then given any $\mathbf{p}^* \in \mathcal{P}^*$, $\mathbf{p}^t$ and $\mathbf{p}^{t+1}$ in the CumLog dynamical process satisfy*

$$D(\mathbf{p}^*, \mathbf{p}^{t+1}) \leq D(\mathbf{p}^*, \mathbf{p}^t) - \frac{1}{2} \cdot \|\mathbf{p}^t - \mathbf{p}^{t+1}\|_2^2 + \tilde{\eta}^t \cdot \langle c(\mathbf{p}^t), \mathbf{p}^* - \mathbf{p}^{t+1} \rangle. \quad (9)$$

Proof. As $s^0 < \infty$, we have $\mathbf{s}^t < \infty$, and thus, $\mathcal{Q}_w^t = \mathcal{P}_w$ for all $t \geq 0$. Invoking Lemma 1 then completes the proof. See Online Appendix A.3 for details. $\square$

Proposition 3. *Under Assumptions 1 and 2, if $s^0 < \infty$, then given any $\mathbf{p}^* \in \mathcal{P}^*$, $\mathbf{p}^t$ and $\mathbf{p}^{t+1}$ in the CumLog dynamical process satisfy*

$$D(\mathbf{p}^*, \mathbf{p}^{t+1}) \leq D(\mathbf{p}^*, \mathbf{p}^t) - \frac{1 - 2\tilde{\eta}^t L}{2} \cdot \|\mathbf{p}^t - \mathbf{p}^{t+1}\|_2^2. \quad (10)$$

Proof. As per Proposition 2, $c(\mathbf{p})$ is $1/4L$ -cocoercive. The result then follows from Lemma 2 and the cocoercivity of $c(\mathbf{p})$. See Online Appendix A.4 for the complete proof. $\square$

Step 3. We finally prove the convergence of $\mathbf{p}^t$ toward a point in $\mathcal{P}^*$ under conditions (i) and (ii) in Theorem 1.

Proof. Proposition 3 indicates that $\lim_{t \rightarrow \infty} D(\mathbf{p}^*, \mathbf{p}^t)$ exists for all $\mathbf{p}^* \in \mathcal{P}^*$ under both conditions (i) and (ii) in Theorem 1. All that is left to prove is to find a convergent subsequence $\{\mathbf{p}^{t_j}\} \subseteq \{\mathbf{p}^t\}$ with $\mathbf{p}^{t_j} \rightarrow \mathbf{p}^*$ when $j \rightarrow \infty$. Denote $\hat{\mathbf{p}}$ as the limit of $\mathbf{p}^{t_j}$. Then, $\mathbf{p}^{t_j} \rightarrow \hat{\mathbf{p}}$ implies that $D(\hat{\mathbf{p}}, \mathbf{p}^{t_j}) \rightarrow 0$. As $\lim_{t \rightarrow \infty} D(\hat{\mathbf{p}}, \mathbf{p}^t)$ exists, it follows $D(\hat{\mathbf{p}}, \mathbf{p}^t) \rightarrow 0$, which implies that $\mathbf{p}^t \rightarrow \hat{\mathbf{p}}$. Under condition (i) given in Theorem 1, we prove by contradiction; if there does not exist a subsequence of $\{\mathbf{p}^t\}$ that converges to $\mathcal{P}^*$, properties (ii) and (iii) given in Proposition 2 cannot both hold. Under condition (ii) in Theorem 1, we first invoke the Bolzano–Weierstrass theorem to exact a convergent subsequence and then, prove that its limit must belong to $\mathcal{P}^*$. The reader is referred to Online Appendix A.5 for details. $\square$

Remark 2 (On Infinite Valuations). In our model, when either condition (i) or condition (ii) in Theorem 1 is met, $\mathbf{s}^t$ converges to infinity when $t \rightarrow \infty$. Although an infinite value might seem unrealistic, a key feature in the logit model is that the choice probability depends on the relative, rather than absolute, valuations of alternatives. In our context, this may be interpreted as travelers monitoring the differences between the elements in the valuation vector $\mathbf{s}^t$. Thus, as long as these differences are finite, the model will give correct results. One can also modify the original model to avoid such a nuanced interpretation. To simplify the discussion, assume $|\mathcal{W}| = 1$ (i.e., there is only one OD pair). Consider the following two behavioral variants.

- Variant 1. On each day t , after updating route valuations, travelers further adjust them such that the best-valued route is normalized to zero (i.e., replacing $\mathbf{s}^t$ by $\mathbf{s}^t - \min(\mathbf{s}^t)$).
- Variant 2. One each day t , the travelers update route valuations by $\mathbf{s}^t = \mathbf{s}^{t-1} + \eta^t \cdot (c(\mathbf{p}^{t-1}) - \min(c(\mathbf{p}^{t-1})))$ so that the valuation of the route with the lowest cost on day t remains unchanged.

Mathematically, both variants are identical to the original CumLog model. However, they would ensure that travelers' valuation of routes used at WE would be bounded. However, the valuation on the non-WE routes would still grow to infinity, signifying that they are unacceptable.

Remark 3 (Related Convergence Results). The two conditions given in Theorem 1 are similar to the conditions detailed in Horowitz (1984, theorem 1). The difference is twofold. First, his proof is established only for the two-link network. Second and most important, his dynamical system, based on an average rather than a cumulative scheme, converges to SUE rather than a WE. Horowitz (1984) pondered the possibility of extending his stability result to the WE case (what he called “the deterministic model”). His conclusion was negative because letting $r \rightarrow \infty$ (the equivalent of zero perception errors) not only violates Lipschitz continuity

but also voids the argument that the convergence of the average cost difference to the equilibrium value must imply the convergence to the correct equilibrium (see Horowitz 1984, example 4). The CumLog model also shares a similar mathematical structure with a classic online convex programming method known as the dual-averaging (DA) algorithm (Xiao 2009). Recently, Mertikopoulos and Zhou (2019) studied the convergence of the DA algorithm in continuous games. The convergence conditions given by them are similar to condition (i) in Theorem 1, although they were proved using a different technique. Condition (ii) in Theorem 1, to the best of our knowledge, has never been rigorously shown to ensure the convergence of the DA algorithm in games or VIPs.

Remark 4 (On the Convergence of the Successive-Average Model to WE). At the end of Section 3.2, we ask whether properly coupling the increase of r and the decrease of η would steer the successive-average Model (4) toward WE. With Theorem 1, this question can now be partially answered. Consider the following two DTD models.

- Model I. $\mathbf{s}^t = \mathbf{s}^{t-1} + c(\mathbf{p}^{t-1})$ and $\mathbf{p}^t = q_r(\mathbf{s}^t)$.
- Model II. $\mathbf{s}^t = (1 - \eta^t) \cdot \mathbf{s}^{t-1} + \eta^t \cdot c(\mathbf{p}^{t-1})$ and $\mathbf{p}^t = q_{r^t}(\mathbf{s}^t)$ with $\eta^t = 1/(t+1)$ and $r^t = r \cdot (t+1)$.

Model I is a CumLog model, and Model II is a successive-average model. Mathematically, the two models are identical in the sequence of $\{\mathbf{p}^t\}_{t=0}$ that they generate (the proof is omitted for brevity). Because Model I is a CumLog model with $\eta = 1$, its convergence is guaranteed for sufficiently small r as per condition (ii) in Theorem 1. This means that the successive-average model converges to WE as well if η^t decreases at a rate of $\mathcal{O}(1/t)$ and r^t increases at a rate of $\mathcal{O}(t)$. However, this peculiar coupling between η^t and r^t in Model II does not seem to have a reasonable behavioral explanation. In Section 6.2, we shall show that very slight modifications of the coupling mechanism could lead to vastly different convergence patterns.

6. Numerical Examples

The proposed CumLog dynamical process is tested on two small networks: a three-node, four-link (3N4L) network (Friesz et al. 1990) and the Sioux-Falls network (Leblanc 1975).

The 3N4L network, as shown in Figure 3, has three nodes, four links, and one OD pair. It has four routes connecting the origin (node 1) and the destination (node 3). For ease of reference, let us say that route 1 uses links 2 and 4, route 2 uses links 1 and 4, route 3 uses links 2 and 3, and route 4 uses links 1 and 3. The number of travelers from node 1 to node 3 is 10. Given the flow x_a on link a , we model its costs as $u_a = h_a + w_a \cdot x_a^4$, where $[h_1, h_2, h_3, h_4]^T = [4, 20, 1, 30]^T$ and $[w_1, w_2, w_3, w_4]^T = [1, 5, 30, 1]^T$.

The Sioux-Falls network has 24 nodes, 76 links, and 528 OD pairs. We refer the readers to Leblanc (1975) for the topology, travel demand, and cost function of the Sioux-Falls network.

On the two networks, we shall perform four sets of experiments; the first three run on the 3N4L network, whereas the fourth runs on both the 3N4L network and the Sioux-Falls network. Section 6.1 tests the CumLog model under various behavioral parameters. In Section 6.2, we explore how the classical successive-average model, like those studied by Horowitz (1984), can be redirected toward a WE under the guidance of our theory. In Section 6.3, we test the CumLog model with heterogeneous travelers who exhibit different sensitivity to route valuations. Finally, Section 6.4 investigates the difference between the WE strategies reached by the CumLog model starting from different initial points.

For a route choice strategy $\mathbf{p}$, we use the relative gap of its corresponding link flow $\mathbf{x} \in \mathcal{X} = \{\mathbf{x} : \mathbf{x} = \Lambda \mathbf{p}, \mathbf{p} \in \mathcal{P}\}$, denoted as $\delta(\mathbf{x})$, to assess its distance from WE. The relative gap is computed as

$$\delta(\mathbf{x}) = -\frac{\langle u(\mathbf{x}), \mathbf{x}' - \mathbf{x} \rangle}{\langle u(\mathbf{x}), \mathbf{x} \rangle}, \quad \mathbf{x}' \in \arg \min_{\mathbf{x}'' \in \mathcal{X}} \langle u(\mathbf{x}), \mathbf{x}'' \rangle. \quad (11)$$

6.1. Test of Convergence Conditions

We first examine conditions (i) and (ii) given in Theorem 1. To test condition (i) in Theorem 1, we set $\eta^t = 1/(t+1)$ for $r = 10, 20$, and 40 ; to test condition (ii) in Theorem 1, we fix $\eta^t = 1$ for $r = 0.25, 0.5, 1, 2.5$. For each setting, we run the CumLog model starting from $\mathbf{s}^0 = [0, 0, 0, 0]^T$ until one of the following criteria is met. (i) The number of iterations reaches 120, (ii) the relative gap drops below 10^{-9} , or (iii) the algorithm begins to diverge.

The convergence patterns reported in Figures 4 and 5 generally agree with the prediction of Theorem 1. Under condition (i) in Theorem 1, convergence is ensured regardless of the value of r . Interestingly, the larger the value of r , the slower the convergence is at the beginning, and the faster the convergence is at the end (see Figure 4). For condition (ii) in Theorem 1, the convergence can only be guaranteed when r is sufficiently small. In this case, we can observe from Figure 5 that the convergence rate increases with r when $r \leq 1$. Indeed,

Figure 3. (Color online) The 3N4L Network







Figure 4. (Color online) Convergence Pattern of the CumLog Model Under Conditions (i) in Theorem 1 $\eta^t = 1/(t+1)$ 

setting $r = 1$ and $\eta = 1$ delivers the best performance among all scenarios; it reaches the target with less than 30 iterations. However, when $r = 2.5$, the CumLog model failed to converge.

We next fix $r = 1$ and examine how the decreasing rate of η^t affects the convergence performance. We set $\eta^t = (t+1)^\alpha$ and report the convergence patterns corresponding to $\alpha = -0.5, -0.25, 0, 0.25$ in Figure 6. The convergence of the model under $\alpha = -0.5, -0.25, 0$ is guaranteed by Theorem 1. The faster the η^t decreases with t , the more quickly travelers tend to settle down, and the slower the convergence is. When the decreasing rate is zero (i.e., η^t becomes a constant), the convergence is the fastest. However, when the trend is reversed and η^t begins to increase with t ($\alpha = 0.25$), the process quickly diverges (see the line with dot markers in Figure 6). This is expected as neither condition (i) in Theorem 1 nor condition (ii) in Theorem 1 would be satisfied with $r = 1$ and $\eta^t = (t+1)^{0.25}$.

6.2. Convergence of Revised Successive-Average Models to WE

In the second experiment, we attempt to manipulate the classical successive-average model using our theory so that it converges to WE rather than SUE. The basic

Figure 5. (Color online) Convergence Pattern of the CumLog Model Under Condition (ii) in Theorem 1 with $\eta^t = 1$ **Figure 6.** (Color online) Convergence Pattern of the CumLog Model with $r = 1$ and $\eta^t = (t+1)^\alpha$ 

idea is to gradually raise the value of r (the exploitation parameter) while reducing η (the proactivity measure), as discussed earlier in Remark 4. We test five models, all based on the SA process $s^t = (1 - \eta^t) \cdot s^{t-1} + \eta^t \cdot c(p^{t-1})$ and $p^t = q_{r^t}(s^t)$. However, η^t and r^t are set differently, as detailed below.

Model A. $\eta^t = 1/(t+1)$ and $r^t = t+1$.

Model B. $\eta^t = 1/(t+1)^{0.99}$ and $r^t = t+1$.

Model C. $\eta^t = 1/(t+1)^{1.01}$ and $r^t = t+1$.


Model D. $\eta^t = 1/(t+1)^{0.99}$ and $r^t = (t+1)^{0.99}$.

Model E. $\eta^t = 1/(t+1)^{1.01}$ and $r^t = (t+1)^{1.01}$.

Model A, as pointed out in Remark 4, is equivalent to the CumLog model with $r = 1$ and $\eta = 1$. Hence, its convergence to WE is guaranteed by Theorem 1 and already confirmed in Section 6.1. What we try to examine is the robustness of the “perturbed” successive-average model. Specifically, what happens if we slightly perturb the changing rates of the two parameters? If the manipulated model is robust, then such perturbations should not have a significant impact on the convergence pattern. Compared with Model A, Models B and C keep the same increasing rate for r^t but slightly modify the decreasing rate of η^t ; Models D and E change both r^t and η^t but keep $r^t \cdot \eta^t = 1$ as in Model A. We start all models from $s^0 = [0, 0, 0, 0]^T$ and report the convergence pattern in Figure 7.

Surprisingly, although Model A converges quickly as expected, none of its four slightly perturbed versions were able to converge at a similar speed—not even close. In fact, based on the trend, it is unclear whether they would ever converge to a point sufficiently close to WE. Figure 7 indicates that after 100,000 days (or 274 years), they are still far away from reaching the target precision (relative gap of 10^{-9}). We do not know what caused this dramatic slowdown when the perturbation moves the parameters so slightly away from the trajectory charted by Theorem 1. Indeed, if we compare Model A with Models B and C, the only difference is that η^t is changed from $(t+1)^{-1}$ in Model A to $(t+1)^{-0.99}$ in Model B—which slightly slows down the

Figure 7. (Color online) Convergence Pattern of the SA Model with r^t and η^t of Different Changing Rates



decreasing rate of η^t —and $(1+t)^{-1.01}$ in Model C—which slightly speeds up how η^t is decreased. Yet, Model A converges within less than a month, at least four to five orders of magnitude faster than both Models B and C. Regardless of the cause, the phenomenon draws a sharp contrast with the robustness of the CumLog model against the changing rate in η^t . In Figures 4–6, the CumLog model’s convergence speed varies within a much narrower range, despite much greater variations applied to the parameters.

6.3. User Heterogeneity at WE

We next construct an experiment in which travelers differ from each other in terms of route choice behaviors. Specifically, travelers are divided into four classes (classes 1–4) with the exploitation parameter r set to 0.01, 0.1, 1, and 10, respectively. As noted before, a larger r suggests a smaller perception error, a greater sensitivity to route evaluations, or a stronger propensity for exploitation depending on the preferred interpretation of the modeler. The total demand remains the same as in the first two experiments but is equally allocated to the four classes. The travelers from different classes are identical in every aspect except for the value of r . Their initial valuations of the routes are $[0, 0, 0, 0]^T$, and their proactivity measure η^t is set to a constant of one. We


ran the CumLog model for 1,000 iterations and reached a relative gap below 10^{-14} . The WE strategy obtained corresponds to exactly the same WE link flow as in the homogeneous case. Figure 8 reports the WE route choice strategies of each class. We can see that all but class 4 rank the four routes in the same order: route 2 > route 1 > route 4 > route 3. Class 4 is different only because it does not use the two lower-ranked routes (routes 4 and 3). This indicates that, in this setting, different travelers value the routes the same way but react to the valuations distinctively. Moreover, the class with a larger r is more concentrated on the higher-ranked routes. Class 1 was almost indifferent among the four routes, whereas class 4 completely abandoned routes 4 and 3.

What is remarkable about the above result is that it illustrates that WE is compatible with not only bounded rationality but also user heterogeneity. At a WE, some travelers may stick to one or very few routes because they are too rational (or cost sensitive) to tolerate inferior routes. On the other end of the spectrum are those who are open to exploring all acceptable options with similar probabilities, even the routes with much worse valuations. Still more travelers would fall between the two extremes. The choices of the travelers, as diverse as they are, still result in the same network traffic conditions as if everyone behaves identically and rationally. Therefore, the CumLog model allows us to simultaneously accept that WE approximately exists in the real world at the aggregate level and reject the implausible implication that every traveler must be same and perfectly rational.

6.4. Nonuniqueness of WE Strategies

Our last experiment, performed on both the 3N4L network and the Sioux-Falls network is devised to demonstrate that the CumLog model may reach WE strategies of different properties when initialized from different points. Note that in both networks, the link cost function $u(x)$ is strictly monotone on $\mathcal{X}$. According to Sheffi (1985), under such conditions, the link flow at WE is

Figure 8. (Color online) WE Route Choice Strategies of Travelers with Different Exploitation Parameters ($r = 0.01, 0.1, 1, 10$)



Notes. (a) Class 1: $r = 0.01$. (b) Class 2: $r = 0.1$. (c) Class 3: $r = 1$. (d) Class 4: $r = 10$.

unique. Denoting the unique link flow at WE as s^* , the set of route choice strategies at WE can be represented as a polyhedron:

$$\mathcal{P}^* = \{p^* \in \mathcal{P} : \bar{\Lambda} p^* = x^*\}. \quad (12)$$

A useful property for differentiating $p^* \in \mathcal{P}^*$ is their entropy, which may be interpreted as its likelihood of realization given the information known to the modeler (e.g., satisfying the WE conditions) (Wilson 2011). Mathematically, entropy may be defined as $\Phi(p^*) = -\langle \text{diag}(q)p^*, \log(p^*) \rangle$. Different $p^* \in \mathcal{P}^*$ may use different set of routes. For example, an interior point of the polyhedron (12) uses every route that *may* be used at WE, observing a “no-route-left-behind” policy (Bar-Gera and Boyce 1999). A standard WE algorithm, however, usually finds only a subset of all possible WE routes (i.e., it admits a solution on the boundary of the polyhedron (12)). In what follows, we explore how the choice of the initial strategy p^0 affects the location and entropy of the equilibrium route choice strategy p^* as well as the set of used routes.

In the 3N4L network, the set of route choice strategies at WE can be written as

$$\mathcal{P}^* = \{p^* : p^* = [0.3 - \lambda, 0.4 - \lambda, 0.3 + \lambda, \lambda]^T, \lambda \in [0, 0.3]\}.$$

To visualize the difference between these solutions, we plot the relation between the entropy of all $p^* \in \mathcal{P}^*$ and their corresponding value of λ in Figure 9. It can be seen that the entropy of p^* first increases and then decreases with λ ; the entropy peaks at $\lambda = 0.12$, which corresponds to $\bar{p}^* = [0.18, 0.28, 0.42, 0.12]^T$ or the “maximum-entropy” solution.

We then generate a random sample of 2,000 s^0 from a normal distribution and run CumLog starting from an initial point corresponding to each s^0 in the sample. In each run, we set $r = 1$ and $\eta^t = 1$, and we terminate it if the relative equilibrium gap is smaller than 10^{-6} . Figure 10 plots the histogram of the entropy of p^* reached by our model. If all of the 2,000 initial solutions end up at

Figure 9. (Color online) Relation Between $\Phi(p^*)$ for All $p^* \in \mathcal{P}$ and Their Corresponding Value of λ





Figure 10. (Color online) Distribution of $\Phi(p^*)$ for p^* Reached by 2,000 s^0




Note. The red line highlights $\Phi(p^*)$ for p^* reached by $s^0 = 0$.

the same WE strategy, then the entropy values would be concentrated at a single point in the histogram. Instead, we find that they spread out between a minimum of 10.77 and a maximum of 12.84. Moreover, the equilibrium solution with the highest entropy among the 2,000 points corresponds to the initial solution $s^0 = 0$ (i.e., all valuations are initially set to zero, representing zero information on all routes). A closer look reveals that this solution is indeed the maximum-entropy solution (i.e., $\bar{p}^* = [0.18, 0.28, 0.42, 0.12]^T$).

We then investigate how s^0 affects the set of routes used at the WE reached by the CumLog model on the Sioux-Falls network. Using the methods by Tobin and Friesz (1988) and Xie and Nie (2019), we find that the maximum and minimum numbers of routes to be used at WE are 770 and 557, respectively. This suggests that a WE algorithm may locate a solution whose number of used routes is anywhere between 557 and 770. To perform the test, we use a set of 1,238 routes that contains all 770 routes that may be used by a WE strategy. The initial valuation s^0 is randomly sampled from a normal distribution, and the sample size is set to 2,000. The all-zero initialization $s^0 = 0$ is employed as a benchmark. In all runs, we set $r = 2.5$, $\eta^t = 1$, and the maximum number of days to 1,000.

Figure 11 shows how the relative gap and the total number of routes actively used by travelers (a route is actively used if its probability of being selected is no less than 10^{-6}) change with the number of days. We can see that the convergence pattern is affected by the initial point, but the impact on the convergence rate is insignificant. The number of actively used paths also descends quickly to the lower bound. Interestingly, the CumLog model never “accidentally” eliminates a potential WE route, nor does it ever fail to exclude routes that are not supposed to be there; the number of used routes at the solution reached by CumLog is always 770. If this property can be established analytically, it will make the CumLog model a suitable algorithm for finding all WE routes.

Figure 11. (Color online) Convergence Pattern of the CumLog Model with Different s^0



Notes. The shadowed area represents the collection of the convergence curves corresponding to 2,000 random s^0 ; the solid line represents the convergence curve corresponding to $s^0 = 0$. The black dashed line in panel (b) represents the maximum number of routes that may be used by a WE strategy. (a) Relative gap. (b) Number of used routes.

7. Conclusions

As one of the most fundamental concepts in transportation science, Wardrop equilibrium was the cornerstone of countless large mathematical models that were built in the past six decades to plan, design, and operate transportation systems around the world. However, like Nash equilibrium, its more famous cousin, WE has always had a somewhat flimsy behavioral foundation. The efforts to strengthen this foundation have largely centered on reckoning with the imperfections in human decision-making processes, such as the lack of accurate information, limited computing power, and suboptimal choices. This retreat from behavioral perfectionism was typically accompanied by a conceptual expansion of equilibrium. In place of WE, for example, transportation researchers had defined such generalized equilibrium concepts as stochastic user equilibrium and boundedly rational user equilibrium. Invaluable as these alternatives are to enrich our understanding of equilibrium and to advance modeling and computational tools, they advocate for the abandonment of WE, predicated on its incompatibility with real behaviors. Our study aims to demonstrate that giving up perfect rationality need not force a departure from WE. To this end, we construct a day-to-day dynamical model that mimics how travelers gradually adjust their valuations of routes, hence the choice probabilities, based on past experiences.

Our model, called cumulative logit, resembles the classical DTD models but makes a crucial change; whereas the classical models assume that routes are valued based on the cost averaged over historical data, our model values the routes based on the cost accumulated. To describe route choice behaviors, the CumLog model only uses two parameters, one accounting for the rate at which the future route cost is discounted in the valuation relative to the past ones (the proactivity measure) and the other describing the sensitivity of route choice probabilities to valuation differences (the

exploitation parameter). We prove that CumLog always converges to WE, regardless of the initial point, as long as the proactivity measure either shrinks to zero at a sufficiently slow pace as time proceeds or is held at a sufficiently small constant value.

By equipping WE with a route choice theory compatible with bounded rationality, we uphold its role as a benchmark in transportation systems analysis. Compared with the incumbents, our theory requires no modifications of WE as a result of behavioral accommodation. This simplicity helps to avoid the complications that come with a “moving benchmark,” be it caused by a multitude of equilibria or the dependence of equilibrium on certain behavioral traits. Moreover, by offering a plausible explanation for travelers’ preferences among equal-cost routes at WE, the theory resolves the theoretical challenge posed by Harsanyi’s instability problem. Note that we lay no claim on the behavioral truth about route choices. Real-world routing games take place in such complicated and ever-evolving environments that they may never reach a true stationary state, much less the prediction of a mathematical model riddled with a myriad of assumptions. Indeed, a relatively stable traffic pattern in a transportation network may be explained as a point in a BRUE set, an SUE tied to properly calibrated behavioral parameters, or simply, a crude WE reached by CumLog. Although more empirical research is needed to vet our theory and compare it with existing ones, we should no longer write off WE simply because it adheres to behavioral perfectionism.

Other than satisfying theoretical interests, the CumLog model may also be used as a prototype algorithm for solving routing games. On large networks, the convergence of the CumLog model may be relatively slow (see, for example, Figure 11(a)). This is hardly surprising given that no higher-order information (e.g., the derivative of route cost) is employed. However, if the goal is to find a good approximate solution quickly, then a CumLog-based algorithm can be quite competitive

thanks to its simplicity (only route costs are needed), flexibility (easy extension to more general settings), and stability (relatively weak requirements for convergence). To be sure, the current CumLog model is still far away from a practical algorithm for WE routing games. Of the missing components, the most important is an efficient route-generation scheme. We leave the development of such an algorithm to future investigations.

Numerical experiments in Section 6.4 revealed a few noteworthy phenomena. First, the CumLog model is capable of identifying all routes that may be used by any WE strategy. Second, the WE strategy resulting from the dynamical process is closely related to the initial route valuation. In particular, it seems that an all-zero initial valuation leads to the entropy-maximizing (or most likely) WE strategy. Does this mean the CumLog model can be used to guide the selection of a unique WE strategy, especially in locating the most likely one? We also leave this question to a future study.

Acknowledgments

The authors are grateful for the valuable comments offered by Mr. Boyi Liu and Prof. Hani Mahmassani at Northwestern University and Prof. Yafeng Yin at the University of Michigan, Ann Arbor. The remaining errors are the authors' own.

References

- Arrow KJ (1966) Exposition of the theory of choice under uncertainty. *Synthese* 16(3/4):253–269.
- Arthur WB (1991) Designing economic agents that act like human agents: A behavioral approach to bounded rationality. *Amer. Econom. Rev.* 81(2):353–359.
- Bar-Gera H, Boyce D (1999) Route flow entropy maximization in origin-based traffic assignment. Ceder A, ed. *Proc. 14th Internat. Sympos. Transportation Traffic Theory* (Emerald Group Publishing Limited, Leeds, UK), 397–415.
- Beckmann M, McGuire C, Winsten C (1956) *Studies in the Economics of Transportation* (Yale University Press, New Haven, CT).
- Ben-Akiva M, Bierlaire M (1999) Discrete choice methods and their applications to short term travel decisions. Hall RW, ed. *Handbook of Transportation Science* (Springer, Cham, Switzerland), 5–33.
- Ben-Akiva ME, Lerman SR (1985) *Discrete Choice Analysis: Theory and Application to Travel Demand* (MIT Press, Cambridge, MA).
- Björnerstedt J, Weibull JW (1994) Nash equilibrium and evolution by imitation. Technical report, Stockholm University, Stockholm.
- Boyce DE, Williams HC (2015) *Forecasting Urban Travel: Past, Present and Future* (Edward Elgar Publishing, Northampton, MA).
- Brown GW (1951) Iterative solution of games by fictitious play. Koopmans TC, ed. *Activity Analysis of Production and Allocation* (Wiley, New York), 374–376.
- Bush R, Mosteller F (1955) *Stochastic Models for Learning* (Wiley, New York).
- Camerer CF (2011) *Behavioral Game Theory: Experiments in Strategic Interaction* (Princeton University Press, Princeton, NJ).
- Camerer C, Hua Ho T (1999) Experience-weighted attraction learning in normal form games. *Econometrica* 67(4):827–874.
- Cantarella GE, Cascetta E (1995) Dynamic processes and equilibrium in transportation networks: Toward a unifying theory. *Transportation Sci.* 29(4):305–329.
- Cantarella GE, Watling DP (2016) Modelling road traffic assignment as a day-to-day dynamic, deterministic process: A unified approach to discrete-and continuous-time models. *EURO J. Transportation Logist.* 5(1):69–98.
- Cantarella G, Watling D, De Luca S, Di Pace R (2019) *Dynamics and Stochasticity in Transportation Systems: Tools for Transportation Network Modelling* (Elsevier, Amsterdam).
- Cascetta E (1989) A stochastic process approach to the analysis of temporal dynamics in transportation networks. *Transportation Res. Part B Methodological* 23(1):1–17.
- Cascetta E, Cantarella GE (1993) Modelling dynamics in transportation networks: State of the art and future developments. *Simulation Practice Theory* 1(2):65–91.
- Cascetta E, Nuzzolo A, Russo F, Vitetta A (1996) A modified logit route choice model overcoming path overlapping problems: Specification and some calibration results for interurban networks. Lesort JB, ed. *Proc. 13th Internat. Sympos. Transportation Traffic Theory* (Pergamon Press, Oxford, UK), 697–711.
- Chen HC, Friedman JW, Thisse JF (1997) Boundedly rational Nash equilibrium: A probabilistic choice approach. *Games Econom. Behav.* 18(1):32–54.
- Chen Y, Su X, Zhao X (2012) Modeling bounded rationality in capacity allocation games with the quantal response equilibrium. *Management Sci.* 58(10):1952–1962.
- Chen S, Yang D, Li J, Wang S, Yang Z, Wang Z (2022) Adaptive model design for Markov decision process. Chaudhuri K, Jegelka S, Song L, Szepesvari C, Niu G, Sabato S, eds. *Proc. 39th Internat. Conf. Machine Learn.* (PMLR, New York), 3679–3700.
- Cohen J, Héliou A, Mertikopoulos P (2017) Learning with bandit feedback in potential games. *Adv. Neural Inform. Processing Systems* 30:6372–6381.
- Cross JG (1973) A stochastic learning model of economic behavior. *Quart. J. Econom.* 87(2):239–266.
- Dafermos S (1980) Traffic equilibrium and variational inequalities. *Transportation Sci.* 14(1):42–54.
- Daganzo CF, Sheffi Y (1977) On stochastic models of traffic assignment. *Transportation Sci.* 11(3):253–274.
- Di X, Liu HX (2016) Boundedly rational route choice behavior: A review of models and methodologies. *Transportation Res. Part B Methodological* 85:142–179.
- Di X, Liu HX, Ban XJ, Yu JW (2015) Submission to the DTA 2012 special issue: On the stability of a boundedly rational day-to-day dynamic. *Networks Spatial Econom.* 15(3):537–557.
- Di X, Liu HX, Pang JS, Ban XJ (2013) Boundedly rational user equilibria (BRUE): Mathematical formulation and solution sets. *Transportation Res. Part B Methodological* 57:300–313.
- Dial RB (1971) A probabilistic multipath traffic assignment model which obviates path enumeration. *Transportation Res.* 5(2):83–111.
- Dupuis P, Nagurny A (1993) Dynamical systems and variational inequalities. *Ann. Oper. Res.* 44:7–42.
- Erev I, Roth AE (1998) Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *Amer. Econom. Rev.* 88:848–881.
- Erlander S (1998) Efficiency and the logit model. *Ann. Oper. Res.* 82:203–218.
- Fisk C (1980) Some developments in equilibrium traffic assignment. *Transportation Res. Part B Methodological* 14(3):243–255.
- Florian M, Fox B (1976) On the probabilistic origin of Dial's multipath traffic assignment model. *Transportation Res.* 10(5):339–341.
- Friesz TL, Tobin RL, Cho HJ, Mehta NJ (1990) Sensitivity analysis based heuristic algorithms for mathematical programs with variational inequality constraints. *Math. Programming* 48:265–284.
- Friesz TL, Bernstein D, Mehta NJ, Tobin RL, Ganjalizadeh S (1994) Day-to-day dynamic network disequilibrium and idealized traveler information systems. *Oper. Res.* 42(6):1120–1136.

- Fudenberg D, Kreps DM (1993) Learning mixed equilibria. *Games Econom. Behav.* 5(3):320–367.
- Fudenberg D, Tirole J (1991) *Game Theory* (MIT Press, Cambridge, MA).
- Goeree JK, Holt CA, Palfrey TR (2002) Quantal response equilibrium and overbidding in private-value auctions. *J. Econom. Theory* 104(1):247–272.
- Guo X, Liu HX (2011) Bounded rationality and irreversible network change. *Transportation Res. Part B Methodological* 45(10):1606–1618.
- Guo RY, Yang H, Huang HJ, Tan Z (2015) Link-based day-to-day network traffic dynamics and equilibria. *Transportation Res. Part B Methodological* 71:248–260.
- Hammond JH, Magnanti TL (1987) Generalized descent methods for asymmetric systems of equations. *Math. Oper. Res.* 12(4):678–699.
- Han K, Szeto WY, Friesz TL (2015) Formulation, existence, and computation of boundedly rational dynamic user equilibrium with fixed or endogenous user tolerance. *Transportation Res. Part B Methodological* 79:16–49.
- Harsanyi JC (1968) Games with incomplete information played by “Bayesian” players part II. Bayesian equilibrium points. *Management Sci.* 14(5):320–334.
- Harsanyi JC (1973) Games with randomly disturbed payoffs: A new rationale for mixed-strategy equilibrium points. *Internat. J. Game Theory* 2(1):1–23.
- Harsanyi JC, Selten R (1988) *A General Theory of Equilibrium Selection in Games* (MIT Press, Cambridge, MA).
- He X, Guo X, Liu HX (2010) A link-based day-to-day traffic assignment model. *Transportation Res. Part B Methodological* 44(4):597–608.
- Horowitz JL (1984) The stability of stochastic equilibrium in a two-link transportation network. *Transportation Res. Part B Methodological* 18(1):13–28.
- Hu TY, Mahmassani HS (1997) Day-to-day evolution of network flows under real-time information and reactive signal control. *Transportation Res. Part C Emerging Tech.* 5(1):51–69.
- Kocsis L, Szepesvári C (2006) Bandit based Monte-Carlo planning. Fürnkranz J, Scheffer T, eds. *Proc. 17th Eur. Conf. Machine Learn.* (Springer-Verlag, Berlin, Germany), 282–293.
- Lattimore T, Szepesvári C (2020) *Bandit Algorithms* (Cambridge University Press, Cambridge, UK).
- Leblanc LJ (1975) An algorithm for the discrete network design problem. *Transportation Sci.* 9(3):183–199.
- Li J, Yu J, Wang Q, Liu B, Wang Z, Nie YM (2022) Differentiable bilevel programming for Stackelberg congestion games. Preprint, submitted. September 15, <https://arxiv.org/abs/2209.07618>.
- Lou Y, Yin Y, Lawphongpanich S (2010) Robust congestion pricing under boundedly rational user equilibrium. *Transportation Res. Part B Methodological* 44(1):15–28.
- Mahmassani HS, Chang GL (1987) On boundedly rational user equilibrium in transportation systems. *Transportation Sci.* 21(2):89–99.
- Mahmassani HS, Jou RC (2000) Transferring insights into commuter behavior dynamics from laboratory experiments to field surveys. *Transportation Res. Part A Policy Practice* 34(4):243–260.
- Mahmassani HS, Liu YH (1999) Dynamics of commuting decision behaviour under advanced traveller information systems. *Transportation Res. Part C Emerging Tech.* 7(2–3):91–107.
- Mamun MS, Xu H, Yin Y (2011) Select zone analysis for traffic impact studies. *Transportation Res. Rec.* 2263(1):123–130.
- McFadden D (1973) Conditional logit analysis of qualitative choice behaviour. Zarembka P, ed. *Frontiers in Econometrics* (Academic Press, Cambridge, MA), 105–142.
- McKelvey RD, Palfrey TR (1995) Quantal response equilibria for normal form games. *Games Econom. Behav.* 10(1):6–38.
- Mertikopoulos P, Zhou Z (2019) Learning in games with continuous action sets and unknown payoff functions. *Math. Programming* 173:465–507.
- Miyasawa K (1961) On the convergence of the learning process in a 2×2 non-zero-sum two-person game. Technical report, Princeton University Press, Princeton, NJ.
- Myerson RB (1978) Refinements of the Nash equilibrium concept. *Internat. J. Game Theory* 7(2):73–80.
- Nagurney A, Zhang D (1997) Projected dynamical systems in the formulation, stability analysis, and computation of fixed-demand traffic network equilibria. *Transportation Sci.* 31(2):147–158.
- Nash J (1951) Non-cooperative games. *Ann. Math.* 54(2):286–295.
- Robinson J (1951) An iterative method of solving a game. *Ann. Math.* 54(2):296–301.
- Rosenthal RW (1989) A bounded-rationality approach to the study of noncooperative games. *Internat. J. Game Theory* 18(3):273–292.
- Sandholm WH (2010) *Population Games and Evolutionary Dynamics* (MIT Press, Cambridge, MA).
- Selten R (1975) Reexamination of the perfectness concept for equilibrium points in extensive games. *Internat. J. Game Theory* 4(1):25–55.
- Shapley LS (1964) Some topics in two-person games. Dresher M, Shapley LS, Tucker AW, eds. *Advances in Game Theory* (Princeton University Press, Princeton, NJ), 1–29.
- Sheffi Y (1985) *Urban Transportation Networks* (Prentice-Hall, Upper Saddle River, NJ).
- Sheffi Y, Powell W (1981) A comparison of stochastic and deterministic traffic assignment over congested networks. *Transportation Res. Part B Methodological* 15(1):53–64.
- Simon HA (1955) A behavioral model of rational choice. *Quart. J. Econom.* 69(1):99–118.
- Simon HA (1990) Bounded rationality. Eatwell J, Milgate M, Newman P, eds. *Utility and Probability* (Palgrave Macmillan, London), 15–18.
- Simon HA (2000) Bounded rationality in social science: Today and tomorrow. *Mind Soc.* 1:25–39.
- Smith JM (1982) *Evolution and the Theory of Games* (Cambridge University Press, Cambridge, UK).
- Smith MJ (1984) The stability of a dynamic model of traffic assignment: An application of a method of Lyapunov. *Transportation Sci.* 18(3):245–252.
- Smith JM, Price GR (1973) The logic of animal conflict. *Nature* 246(5427):15–18.
- Smith MJ, Watling DP (2016) A route-swapping dynamical system and Lyapunov function for stochastic user equilibrium. *Transportation Res. Part B Methodological* 85:132–141.
- Sutton RS, Barto AG (2018) *Reinforcement Learning: An Introduction* (MIT Press, Cambridge, MA).
- Tobin RL, Friesz TL (1988) Sensitivity analysis for equilibrium network flow. *Transportation Sci.* 22(4):242–250.
- Tversky A, Kahneman D (1985) *The Framing of Decisions and the Psychology of Choice* (Springer, Cham, Switzerland).
- van Damme E (1987) *Stability and Perfection of Nash Equilibria* (Springer, Cham, Switzerland).
- Vickrey W (1961) Counterspeculation, auctions, and competitive sealed tenders. *J. Finance* 16(1):8–37.
- von Neumann J (1928) Zur theorie der gesellschaftsspiele. *Math. Ann.* 100(1):295–320.
- Wardrop JG (1952) Some theoretical aspects of road traffic research. *Proc. Institution Civil Engineers* 1(3):325–362.
- Watling D (1999) Stability of the stochastic equilibrium assignment problem: A dynamical systems approach. *Transportation Res. Part B Methodological* 33(4):281–312.
- Watling D, Hazelton ML (2003) The dynamics and equilibria of day-to-day assignment models. *Networks Spatial Econom.* 3(3):349–370.
- Weibull JW (1997) *Evolutionary Game Theory* (MIT Press, Cambridge, MA).
- Wen CH, Koppelman FS (2001) The generalized nested logit model. *Transportation Res. Part B Methodological* 35(7):627–641.
- Wilson A (2011) *Entropy in Urban and Regional Modelling*, vol. 1 (Routledge, Abingdon-on-Thames, UK).

- Xiao L (2009) Dual averaging method for regularized stochastic learning and online optimization. *Adv. Neural Inform. Processing Systems* 22:2116–2124.
- Xiao F, Shen M, Xu Z, Li R, Yang H, Yin Y (2019) Day-to-day flow dynamics for stochastic user equilibrium and a general Lyapunov function. *Transportation Sci.* 53(3):683–694.
- Xie J, Nie Y (2019) A new algorithm for achieving proportionality in user equilibrium traffic assignment. *Transportation Sci.* 53(2): 566–584.
- Yang F (2005) An evolutionary game theory approach to the day-to-day traffic dynamics. PhD thesis, University of Wisconsin–Madison, Madison.
- Yang F, Zhang D (2009) Day-to-day stationary link flow pattern. *Transportation Res. Part B Methodological* 43(1):119–126.
- Zhang D, Nagurney A (1996) On the local and global stability of a travel route choice adjustment process. *Transportation Res. Part B Methodological* 30(4):245–262.