Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

When Frictions Disappear: Computational Accessibility and Decentralized Adjustment

Stephen Lewarne

Slewarne@franciscan.edu

Franciscan University of Steubenville

Anton Kamenov

Research Article

Keywords: Equilibrium Accessibility, Equilibrium Selection, Decentralized Adjustment, Computational Economics, Path Dependence, Feasibility Constraints, Economic Dynamics

Posted Date: July 3rd, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10101567/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

When Frictions Disappear: Computational Accessibility and Decentralized Adjustment

Stephen Lewarne

(Corresponding Author)

Franciscan University of Steubenville Ohio

slewarne@franciscan.edu

Country Affiliation: USA

and

Anton Kamenov

NEXI Technologies Inc.

Antonkamenov@yahoo.com

Country Affiliation: USA

Abstract

Much of economic theory analyzes outcomes through equilibrium existence and efficiency, often abstracting from how such outcomes are identified in decentralized systems. This paper studies equilibrium selection in environments with discrete actions and shared feasibility constraints. In such settings, welfare-relevant allocations may be easy to verify once proposed yet difficult to identify through decentralized adjustment.

We develop a combinatorial allocation economy in which agents choose among discrete actions subject to resource constraints. A sequential best-response process is operational and converges to stable outcomes, but the resulting equilibria need not be welfare-maximizing. The underlying reason is that local improvement and global equilibrium selection constitute distinct problems: decentralized agents can optimize over currently feasible actions without identifying welfare-relevant equilibria in the full feasible set.

To formalize this distinction, I show that welfare-relevant equilibrium selection can be computationally inaccessible, providing a sharp separation between equilibrium existence and operational accessibility. As a result, realized outcomes depend on adjustment processes rather than on existence conditions alone. This generates path dependence, dispersion across

equilibria, and persistent effects of policy interventions even in the absence of traditional frictions.

Simulations indicate that decentralized adjustment converges rapidly while attaining only 75.6 percent of benchmark welfare on average, demonstrating that accessibility gaps can arise even when convergence itself is nearly immediate.

The contribution is to identify computational accessibility as a distinct determinant of equilibrium outcomes and to show how equilibrium existence, equilibrium accessibility, and equilibrium selection can diverge in decentralized economic systems.

Keywords: Equilibrium Accessibility; Equilibrium Selection; Decentralized Adjustment; Computational Economics; Path Dependence; Feasibility Constraints; Economic Dynamics

JEL Codes: C61, C63, C72, D50, D61, L14

Funding

No funding was received for this research.

1. Why Policy Remains Relevant After Frictions Disappear

Much of modern economic policy analysis begins from the presence of frictions. Markets may fail to reach efficient outcomes because information is asymmetric, search is costly, transactions are costly, prices are sticky, contracts are incomplete, or agents cannot coordinate on superior equilibria. The canonical examples are familiar: asymmetric information in Akerlof's lemons model, transaction costs in Coase and Williamson, search frictions in Diamond, and nominal adjustment costs in Mankiw's menu-cost model (Akerlof 1970; Coase 1937; Williamson 1975; Diamond 1982; Mankiw 1985).

This framework has been extraordinarily productive. It explains why disclosure rules, contracting institutions, monetary stabilization, antitrust policy, procurement design, and market regulation may matter. The underlying intuition is that policy becomes relevant when some obstacle prevents decentralized actors from reaching otherwise feasible gains from trade or efficient allocations.

Technological change complicates this picture. Search costs have fallen sharply in many markets. Digital platforms reduce matching costs. Artificial intelligence may further reduce information asymmetries, automate contracting tasks, improve forecasting, and lower the cost of coordination. To the extent that these developments weaken traditional frictions, one might expect the policy case to weaken as well.

This paper argues that such a conclusion is premature. Even if information becomes abundant, contracting becomes easier, and adjustment costs fall, decentralized economies may still face a different kind of constraint: combinatorial feasibility. In many production and allocation environments, the difficulty is not that agents lack information or that prices adjust slowly. The difficulty is that the number of feasible configurations grows rapidly with the size of the system. Welfare-improving allocations may exist and may be easy to verify once proposed, yet remain difficult to identify, coordinate upon, or attain through decentralized adjustment.

Economic analysis has traditionally focused on equilibrium existence and efficiency, with less emphasis on the accessibility of such outcomes. Yet for many economic questions, reachability may matter as much as existence. A welfare-improving allocation that cannot be reliably discovered through decentralized adjustment may have little practical relevance, regardless of its theoretical desirability.

This distinction shifts the policy question. Traditional frictions often impede the identification, coordination, or implementation of desirable allocations. Combinatorial constraints introduce a different difficulty: even when agents possess substantial information and act rationally,

decentralized adjustment may fail to assemble welfare-relevant feasible configurations.. A market may therefore converge to an allocation that is feasible and locally stable but not welfare-relevant, not because agents are irrational or uninformed, but because the globally superior configuration is not operationally accessible.

The argument connects Hayek's insight about dispersed knowledge to a different limitation. Hayek emphasized that economic knowledge is distributed and that prices help coordinate local information (Hayek 1945). The issue here is not primarily that information is dispersed. Even when the relevant information is known, the feasible allocation problem may be operationally difficult. This is the setting in which results from combinatorial optimization and computational economics become economically relevant. The issue is not merely whether desirable allocations exist, but whether decentralized systems can reliably identify and attain them. (Garey and Johnson 1979; Papadimitriou 1994; Sandholm 2002).

NP-hardness provides one formal way to separate equilibrium existence from operational accessibility. It establishes that the existence of a welfare-relevant allocation does not imply the existence of a tractable procedure capable of identifying it. More importantly, it shows why reliable equilibrium selection cannot generally be assumed. The existence of a welfare-relevant equilibrium and the ability of decision-makers to identify it are separate questions.

This distinction is not unique to decentralized systems. Even in a centralized decision-making setting or under perfect collusion, finding a welfare-relevant equilibrium remains an NP-hard combinatorial problem. Governments face the same computational limits. The contribution of this paper is therefore not to show that centralized actors can solve such problems, but to identify a distinct economic implication of computational accessibility. Traditional policy analysis asks how outcomes change when information, incentives, or adjustment costs change. This paper asks a different question: what happens when welfare-improving allocations exist, but decentralized systems cannot reliably discover them? The answer is that equilibrium existence, equilibrium accessibility, and equilibrium selection can diverge. In such environments, adjustment dynamics rather than equilibrium conditions become an important determinant of realized outcomes.

The central claim of the paper is that economic outcomes are determined not only by which allocations are feasible or which equilibria exist, but by which allocations are operationally accessible through decentralized adjustment processes.

A second claim is that the continuing relevance of policy need not rest only on asymmetric information, menu costs, or other conventional frictions. Decentralized adjustment may fail to identify welfare-improving allocations even when such allocations exist. Under these conditions, policy may remain relevant not because governments can reliably identify welfare-maximizing

outcomes, but because interventions can influence the paths through which decentralized systems explore feasible alternatives.

2. Combinatorial Constraints in Modern Economic Systems

The argument developed in this paper does not depend on any particular industry or institutional setting. Rather, it concerns a structural feature that appears across a wide range of contemporary economic systems: the allocation of scarce resources across large numbers of feasible combinations. In such environments, economic outcomes depend not merely on prices or information, but on the selection of feasible configurations from an expanding combinatorial set. This structure appears across a variety of applied settings, including supply chains, logistics networks, infrastructure systems, and procurement environments.

2.1 Supply Chains and Logistics Networks

Supply-chain design problems routinely involve supplier selection, facility location, inventory placement, transportation routing, and capacity allocation. These decisions interact through shared constraints and are typically formulated as mixed-integer optimization problems (Simchi-Levi, Kaminsky, and Simchi-Levi 2008; Snyder 2006). Likewise, logistics networks give rise to classical vehicle-routing, scheduling, and allocation problems whose computational complexity grows rapidly with network size (Dantzig and Ramser 1959; Toth and Vigo 2014).

The key economic feature is that feasible allocations are interdependent. Selecting one supplier, route, or capacity commitment may eliminate a large number of alternative configurations. As a result, decentralized adaptation occurs within a constrained and evolving feasible set rather than over a continuous choice space.

2.2 Energy, Infrastructure, and Capacity Allocation

Modern infrastructure systems exhibit similar characteristics. Electric power systems require the allocation of generation, transmission capacity, storage, and interconnection access while satisfying engineering and reliability constraints. Unit commitment and transmission-constrained dispatch are typically formulated as large-scale mixed-integer optimization problems (Wood, Wollenberg, and Sheblé 2013; Bienstock 2015).

Recent interconnection queues in electricity markets provide a particularly clear example. The challenge is often not a lack of information regarding potential projects, but the difficulty of allocating scarce transmission and grid-access capacity among a large number of competing feasible projects. Similar allocation issues arise in transportation infrastructure, airport slot allocation, port congestion, and telecommunications networks.

2.3 Procurement, Allocation, and Resource Assignment

Procurement systems provide another example. Governments and firms must frequently select feasible combinations of projects, contracts, vendors, and schedules subject to budgetary and operational constraints. In combinatorial auctions and procurement environments, determining welfare-maximizing allocations is itself a computationally difficult problem (Rothkopf, Pekec, and Harstad 1998; Sandholm 2002).

More generally, many allocation problems in economics share a common structure: resources are scarce, actions are discrete, and the value of a particular decision depends on the set of other decisions simultaneously implemented. The resulting problem is not merely one of optimization, but of selecting among large numbers of mutually interacting feasible configurations.

2.4 Common Structure

These examples differ substantially in institutional detail. Some involve private firms, others public agencies. Some rely on markets, others on administrative allocation. Yet they share a common structural characteristic. Economic outcomes depend on selecting feasible combinations of discrete actions under shared constraints.

This observation does not by itself establish computational difficulty. Nor does it imply that decentralized systems cannot adapt successfully. It does suggest, however, that many economically important environments possess a combinatorial structure that differs fundamentally from the continuous settings traditionally emphasized in equilibrium analysis.

The next section turns to the computational implications of that structure. Existing results from combinatorial optimization, allocation theory, and computational economics imply that welfare-relevant allocations may be difficult to identify even when they are easy to verify once proposed.

3. Why Accessibility Matters: Evidence from Contemporary Allocation Systems

The preceding section argued that many economically important environments possess a combinatorial structure. That observation alone does not establish that computational accessibility is economically relevant. A reasonable objection is that firms, markets, and institutions routinely solve complex allocation problems. Why should economists care whether some allocations are computationally difficult in principle?

The answer is that a growing number of contemporary economic systems exhibit persistent allocation problems that are difficult to explain solely through traditional frictions. In many cases, participants possess substantial information about the opportunities available to them.

The difficulty lies in selecting among large numbers of feasible but mutually interacting alternatives.

3.1 Electric Interconnection Queues

Electric power systems provide perhaps the clearest example. Across the United States, thousands of proposed generation and storage projects remain in interconnection queues awaiting access to transmission infrastructure (Rand et al. 2024). Developers often know the technology, expected output, financing requirements, and market opportunities associated with their projects. The difficulty is not ignorance. It is determining which combination of projects can be accommodated given transmission constraints, upgrade requirements, and network interactions.

The resulting problem is highly combinatorial. The feasibility of any individual project depends on the collection of other projects simultaneously granted access to the network. Queue reforms repeatedly focus on sequencing, prioritization, and coordination rather than on information provision alone. The challenge is fundamentally one of allocation.

3.2 Semiconductor Capacity

The semiconductor shortages of the early 2020s illustrate a similar pattern (U.S. Department of Commerce 2022; Varas et al. 2021). Market participants broadly understood the nature of the shortage. The problem was not a lack of awareness that semiconductor fabrication capacity had become scarce.

The difficulty was determining how limited fabrication, testing, transportation, and assembly resources should be allocated across a vast number of potential production combinations (U.S. Department of Commerce 2022). Automobile manufacturers, consumer electronics firms, industrial producers, and governments all competed for overlapping capacity. The relevant question was not whether demand existed. It was how scarce capacity should be distributed across competing feasible uses.

3.3 Defense Procurement and Industrial Capacity

Defense procurement provides another example (U.S. Government Accountability Office 2018; U.S. Department of Defense 2024). Governments often face decisions involving large numbers of mutually dependent projects competing for engineering capacity, manufacturing capability, skilled labor, testing facilities, certification resources, and supplier networks.

In many cases, the challenge is not identifying worthwhile projects. Rather, it is assembling a feasible collection of projects that can be completed simultaneously given underlying resource constraints. Recent assessments of the U.S. defense industrial base identify production bottlenecks, workforce constraints, supplier visibility and concentration, and limited

manufacturing capacity as persistent challenges. Small changes in priorities can alter which combinations of programs remain feasible and which capabilities remain available for future use.

The resulting problem extends beyond budgeting or project selection. Decisions concerning one program affect the resources available to others, creating interdependencies that resemble the accessibility challenges observed in infrastructure systems and production networks. The relevant issue is therefore not simply which projects are desirable, but which combinations of projects can realistically be pursued simultaneously.

4. Computational Accessibility and Welfare-Relevant Allocation

The environments discussed above share a common characteristic: they require the selection of feasible combinations of discrete actions subject to shared constraints. The relevant question is therefore not simply whether desirable allocations exist, but whether such allocations can be identified through decentralized decision-making.

The argument differs from the classic socialist calculation debate. Hayek (1945) emphasized that economically relevant information is dispersed and therefore difficult to aggregate. The present paper asks a different question. Suppose the relevant information is known. Suppose agents understand the available technologies, constraints, and opportunities. Does it follow that welfare-relevant allocations are operationally accessible? The answer may still be no. Accessibility depends not only on information but also on the computational structure of the allocation problem itself.

A large literature in computer science, operations research, and computational economics demonstrates that many allocation problems possess precisely this feature. Combinatorial auctions, packing problems, scheduling problems, routing problems, and resource-allocation problems are often computationally difficult even when feasibility can be verified quickly once a candidate allocation is proposed (Garey and Johnson 1979; Papadimitriou 1994; Sandholm 2002).

The distinction between finding and verifying is fundamental. Given a candidate allocation, it is often straightforward to determine whether constraints are satisfied and to calculate its welfare. The difficulty lies in identifying the allocation itself from among a large number of feasible possibilities.

In computational economics, complexity results for equilibrium computation show that equilibrium objects may be difficult to compute even when their defining properties are mathematically well understood (Deng, Papadimitriou, and Safra 2003; Chen, Dai, Du, and Teng

2009). These results differ in important ways from the environments considered here, but they reinforce a common point: economic existence theorems do not automatically imply computational accessibility.

The present paper focuses on a more specific question. Suppose a welfare-improving allocation exists. Suppose further that agents are fully rational, possess complete information, and face no adjustment costs. Does it follow that decentralized behavior will reliably identify and attain that allocation?

The answer is not obvious. In combinatorial environments, the set of feasible allocations may grow rapidly with system size. The existence of a welfare-superior allocation therefore does not imply that decentralized agents can discover it, coordinate upon it, or move toward it through local adjustments.

This distinction motivates the concept of operational accessibility. An allocation is operationally accessible if decentralized adjustment can plausibly discover and attain it using information and computational resources available to economic actors.[1] An allocation may therefore be:

1. feasible,
2. welfare-superior,
3. easy to verify once proposed,

yet still fail to be operationally accessible. Section 10 provides simulation examples illustrating precisely this possibility.

The remainder of the paper explores the implications of this possibility. Rather than assuming that welfare-relevant allocations are automatically attainable, the analysis asks what decentralized systems do when such allocations are computationally difficult to identify. The answer, I argue, is that local adjustment processes become the economically relevant object. In such environments, observed outcomes may be governed less by the existence of globally desirable allocations than by the dynamics through which feasible allocations are constructed.

5. What Decentralized Economies Do Instead

The preceding section argued that welfare-relevant allocations may be computationally inaccessible even when they are easy to verify once proposed. This observation raises a natural

[1] The distinction between desirable outcomes and operationally attainable outcomes has a long history in economics. In spirit, the present argument resembles the broader insight that the existence of an optimal plan does not imply its implementation (Kydland and Prescott 1977). Unlike their analysis, however, the mechanism here does not arise from commitment problems or dynamic inconsistency. It arises from computational accessibility and combinatorial feasibility.

question. If globally desirable allocations cannot be reliably identified, what process governs decentralized economic outcomes?

Traditional equilibrium analysis often leaves this question implicit. The standard approach is to characterize an equilibrium and then study its properties. Adjustment is treated as secondary. This methodology is often justified because, in many environments, equilibrium serves as a reasonable approximation of observed outcomes. If agents can discover and coordinate upon efficient allocations, then the adjustment process itself may have limited independent importance.

A long tradition in economics has nevertheless emphasized that outcomes depend not only on equilibrium conditions but also on the procedures by which decisions are made. Simon (1955, 1957) argued that economic agents frequently rely on tractable decision rules rather than exhaustive optimization. Hayek (1945) emphasized that decentralized systems operate through local information rather than global knowledge. More recently, Rubinstein (1998) examined the implications of computational limits for economic decision making, while Young (1998) and related work in adaptive dynamics studied how decentralized adjustment can shape long-run outcomes.

The present paper approaches the issue from a different direction. Agents may be fully informed and fully rational, yet welfare-relevant allocations may remain operationally inaccessible through decentralized adjustment.

This distinction matters because decentralized actors rarely solve global allocation problems directly. Firms do not evaluate every possible supply-chain configuration before selecting suppliers. Grid operators do not compare every feasible allocation of transmission capacity. Procurement systems do not examine every possible combination of contracts and projects. In practice, decisions are made through a sequence of local revisions, substitutions, reallocations, and adjustments.

The relevant economic question therefore becomes procedural rather than existential. Instead of asking whether a desirable allocation exists, one asks how decentralized systems move through the space of feasible allocations. The question is not what an omniscient planner could compute, but what decentralized agents can implement through feasible adjustment processes.

This observation suggests a simple principle:

When globally desirable allocations are computationally inaccessible, decentralized systems rely on locally feasible improvements.

The principle is closely related to themes that appear throughout the literature on adaptive behavior, learning, and evolutionary dynamics. Agents typically evaluate alternatives that are

currently available rather than exhaustively searching all feasible possibilities. Decisions are made incrementally. Existing commitments constrain future options. As a result, realized outcomes depend not only on preferences and technologies but also on the path through which feasible allocations are constructed (Young 1998).

Importantly, local adjustment need not imply irrationality. Agents may optimize over currently feasible opportunities while the system as a whole fails to attain a welfare-superior allocation. Local optimization and global inefficiency are therefore compatible.

The distinction between local and global reasoning is familiar in operations research and combinatorial optimization. Many computationally difficult problems admit efficient local-search procedures that improve candidate solutions without guaranteeing globally optimal outcomes (Papadimitriou and Steiglitz 1998; Aarts and Lenstra 2003). The interpretation here is different. The objective is not to design an algorithm. Rather, the local-search analogy provides a useful way to think about decentralized economic adjustment. Firms revise actions using information available in their immediate neighborhood of feasible alternatives rather than solving the global allocation problem.

This perspective has an important implication. The economically relevant object may not be the welfare-maximizing allocation itself. Instead, it may be the dynamic process through which feasible allocations are generated, revised, and stabilized. If decentralized adjustment systematically explores only a small subset of the feasible allocation space, then equilibrium existence alone provides an incomplete description of economic outcomes.

The contribution is not the existence of local optima, which are familiar throughout optimization and game theory. The contribution is that decentralized equilibrium dynamics operating under feasibility constraints may systematically select such outcomes even when welfare-superior equilibria exist elsewhere in the feasible set.

The remainder of the paper studies perhaps the simplest formal representation of this idea: sequential best-response under feasibility constraints. Firms revise actions one at a time, taking the current actions of others as given and selecting the highest-valued action that remains feasible. The resulting process is computationally trivial at the level of the individual firm. Yet, as the following sections demonstrate, the outcomes reached by this dynamic can differ systematically from welfare-relevant allocations that exist elsewhere in the feasible set.

The central implication is that computational inaccessibility shifts attention away from equilibrium existence and toward adjustment dynamics. Once welfare-relevant allocations are no longer assumed to be operationally accessible, the process of decentralized adaptation becomes a primary object of economic analysis rather than a secondary detail.

6. A Simple Model of Decentralized Adjustment

The previous sections argued that welfare-relevant allocations may be computationally inaccessible even when they exist. If so, equilibrium existence alone is insufficient to predict outcomes. The relevant question becomes how decentralized systems move through the space of feasible allocations.

The model below is a minimal combinatorial allocation economy. Agents are fully informed, fully rational, and face no adjustment costs. The objective is to isolate the role of feasibility constraints. The approach is related to the literature on adaptive adjustment and decentralized learning (Young 1998) but differs in emphasis. The issue is not why agents fail to optimize. The issue is what decentralized optimization implies when globally desirable allocations may be difficult to identify.

6.1 Economic Environment

Let

$$N = \{1, \dots, n\}$$

denote a finite set of firms and

$$R = \{1, \dots, r\}$$

a finite set of resources.

Each resource $r \in R$ has capacity

$$q_r \in \mathbb{Z}_+.$$

Resources are exogenous capacity units. They may represent supplier slots, logistics channels, transmission capacity, certification opportunities, procurement commitments, infrastructure access, compute capacity, or other scarce inputs.

Each firm $i \in N$chooses one action from a finite action set

$$A_i.$$

An action represents a complete sourcing or allocation plan.

Associated with each action $a \in A_i$ is a resource-requirement vector

$$d(a) = (d_1(a), \dots, d_r(a)),$$

where

$$d_r(a) \in \{0,1\}.$$

A profile of actions is denoted

$$a = (a_1, \dots, a_n).$$

Feasibility

A profile is feasible if aggregate resource use does not exceed available capacity:

$$\sum_{i \in N} d_r\,(a_i) \leq q_r \forall r \in R.$$

Feasibility is the only source of interaction among firms. The absence of direct payoff interaction is intentional. The objective is to isolate the role of shared feasibility constraints as a source of economic interdependence. A firm's action affects the opportunity sets available to other firms through the allocation of scarce resources rather than through preferences or beliefs. Many economically important allocation environments—including supplier selection, transmission access, procurement systems, and capacity allocation—possess precisely this structure.

Payoffs

Each firm has payoff function

$$u_i \colon A_i \to \mathbb{Z}.$$

Payoffs depend only on the firm's chosen action. A sourcing plan may therefore have greater value because it provides access to superior suppliers, lower transportation costs, higher output, or greater reliability. The source of value is not important for the analysis. What matters is that firms rank feasible actions according to their payoffs.

Welfare

Aggregate welfare is defined as

$$W(a) = \sum_{i \in N} u_i\,(a_i).$$

The welfare measure is intentionally simple. The objective is not to construct a comprehensive welfare criterion but to distinguish welfare-relevant allocations from allocations generated by decentralized adjustment.

6.2 Equilibrium

The benchmark concept is a pure feasible equilibrium.

Definition 1 (Pure Feasible Equilibrium)

A feasible profile

$$a^*$$

is a pure feasible equilibrium if no firm can profitably switch to another action that remains feasible given the actions of all other firms.

Formally,

$$u_i(a_i^*) \geq u_i(b)$$

for every firm $i \in N$and every alternative action $b \in A_i$such that

$$(b, a_{-i}^*)$$

is feasible.

Equilibrium therefore combines two requirements.

First, all capacity constraints must be satisfied.

Second, no feasible profitable unilateral deviation may exist.

Notice that equilibrium does not require firms to identify the welfare-maximizing allocation. It requires only that firms optimize over their currently feasible opportunities.

6.3 Sequential Best Response Under Feasibility

The central object of the paper is not equilibrium itself but the process by which decentralized systems move through the feasible allocation space.

Definition 2 (Sequential Best Response)

Let

$$\{i_t\}_{t\geq 0}$$

be a sequence selecting one firm in each period.

Starting from a feasible profile

$$a^0,$$

firm i_t chooses

$$a_{i_t}^{t+1} \in \arg\max_{b\in A_{i_t}} \{u_{i_t}(b): (b, a_{-i_t}^{t}) \text{ is feasible}\}.$$

All other firms retain their previous actions:

$$a_j^{t+1} = a_j^t \forall j \neq i_t.$$

The rule is deliberately local. Firms do not solve a global allocation problem. They evaluate only actions currently available to them.

Definition 3 (Absorbing State)

A feasible profile a is absorbing if no firm possesses a profitable feasible deviation.

By construction, every absorbing state is a pure feasible equilibrium.

6.4 Properties of Local Adjustment

The first question is whether the adjustment process is well defined. Specifically, does decentralized adaptation preserve feasibility and converge to a stable outcome?

Proposition 1 (Feasibility Preservation)

If a^0 is feasible, then every subsequent profile generated by sequential best response remains feasible.

Proof

Suppose $a^t$is feasible. Firm $i_t$chooses only among actions satisfying

$$\left(b, a_{-i_t}^{t}\right)$$

feasible.

Therefore the updated profile a^{t+1}also satisfies all resource constraints. By induction, feasibility is preserved for every t.

□

The second question is whether decentralized adjustment eventually stabilizes.

Proposition 2 (Finite Convergence)

Assume payoffs are integer-valued and firms move only when a strict improvement is available. Then sequential best response converges in finite time to an absorbing state.

Proof

Recall that aggregate welfare is defined as

$$W(a) = \sum_{i \in N} u_i\,(a_i).$$

Whenever a firm makes a strictly profitable move under sequential best response,

$$u_{i_t}\left(a_{i_t}^{t+1}\right) > u_{i_t}\left(a_{i_t}^{t}\right).$$

Since the actions of all other firms remain unchanged,

$$W(a^{t+1}) - W(a^{t}) = u_{i_t}\left(a_{i_t}^{t+1}\right) - u_{i_t}\left(a_{i_t}^{t}\right) > 0.$$

Because the action sets are finite, the number of feasible profiles is finite. Consequently, the set of attainable welfare values is finite and $W$is bounded above. Since payoffs are integer-valued, each profitable deviation increases welfare by at least one unit. Therefore, only finitely many welfare-improving moves are possible.

The adjustment process therefore terminates after finitely many steps at a profile with no profitable feasible deviation. By Definition 3, that profile is an absorbing state.

□

The convergence result is related to finite-improvement arguments in adaptive dynamics and potential games (Monderer and Shapley 1996). However, the significance here is not the

existence of a potential function itself. Rather, it is that decentralized adjustment remains operational even when welfare-relevant allocations may not be.

This observation creates the central tension explored in the remainder of the paper. Local adjustment can be computationally trivial, while welfare-relevant allocation may remain computationally inaccessible. The next section develops this distinction formally and examines the consequences of the resulting computational–dynamic gap.

7. Computational Accessibility and Welfare-Relevant Equilibrium Selection

The previous section established that decentralized adjustment is operational. Starting from any feasible allocation, firms can evaluate locally feasible alternatives and converge to a stable outcome through sequential best response. The existence of such a process does not imply, however, that the resulting allocation is welfare-relevant.

The key distinction is between local accessibility and global accessibility. Local accessibility concerns whether agents can identify improving actions within their currently feasible opportunity set. Global accessibility concerns whether a welfare-relevant equilibrium can be identified among all feasible equilibria.

The distinction can now be stated formally. Local adjustment is a problem of improvement from a given feasible allocation. Welfare-relevant equilibrium selection is a problem of choosing among the set of feasible equilibria. These are different objects. A firm may be able to identify a profitable feasible deviation using only local information, while the economy as a whole may remain unable to identify which equilibrium satisfies a given welfare criterion. The first problem concerns operational adjustment. The second concerns global allocation. The possibility that these two problems differ is the central computational issue addressed in this section.

To formalize this distinction, consider the following decision problem.

Definition 4 (Welfare-Relevant Equilibrium Selection)

Given:

- a combinatorial allocation economy E,
- a welfare threshold B,

determine whether there exists a pure feasible equilibrium a^*satisfying

$$W(a^*) \geq B.$$

The problem differs fundamentally from the local adjustment problem studied in Section 5. Local adjustment asks whether agents can improve upon currently feasible actions. Welfare-

relevant equilibrium selection asks whether there exists a stable allocation satisfying a specified welfare criterion.

The distinction matters because the existence of an equilibrium does not imply the existence of a procedure capable of identifying welfare-relevant equilibria efficiently.

Theorem 1 (Computational Inaccessibility)

There exist combinatorial allocation economies in which welfare-relevant equilibrium selection is NP-hard.

Proof

The result follows by reduction from the classical 3-SET PACKING problem (Garey and Johnson 1979).

Consider an instance consisting of a finite universe U, a collection of three-element subsets

$$\mathcal{C} = \{S_1, \dots, S_m\},$$

and an integer K.

Construct a combinatorial allocation economy in which:

- each element of U corresponds to a resource,
- each set S_j corresponds to a sourcing plan,
- resource capacities equal one,
- productive actions yield payoff one,
- outside options yield payoff zero.

Under this construction, a welfare level of at least K is attainable if and only if there exist K pairwise disjoint sets in the original packing problem.

Determining whether a welfare-relevant equilibrium exists therefore solves the underlying packing problem.

Since 3-SET PACKING is NP-complete, welfare-relevant equilibrium selection is NP-hard.

A complete proof is provided in Appendix A.

□

Verification versus Identification

The significance of Theorem 1 lies not in the difficulty of verifying equilibrium once proposed. Verification remains straightforward.

Given a candidate profile a, one can directly check feasibility:

$$\sum_{i \in N} d_r\,(a_i) \leq q_r \forall r \in R,$$

and verify that no feasible profitable unilateral deviation exists.

The difficulty lies in identifying welfare-relevant equilibria among the set of feasible allocations.

This distinction mirrors a broader theme in computational economics. Equilibrium existence and computational accessibility are separate questions (Deng, Papadimitriou, and Safra 2003; Chen, Dai, Du, and Teng 2009). The present setting differs because the source of difficulty is not continuous fixed-point computation, but the combinatorial structure generated by discrete actions and shared feasibility constraints.

Economic Interpretation

The proposition has an immediate economic interpretation.

A welfare-superior equilibrium may exist.

It may be easy to verify once discovered.

It may even be stable once reached.

Yet decentralized actors may be unable to identify it through local adjustment.

The existence of welfare-relevant equilibria therefore provides only limited information concerning realized outcomes. If such equilibria are computationally inaccessible, actual economic behavior will be governed less by the set of possible equilibria and more by the adjustment processes through which allocations are constructed.

This observation motivates the next section. If local adjustment is computationally easy while welfare-relevant equilibrium selection is computationally difficult, the two processes need not produce the same outcome. The resulting divergence is the central economic mechanism of the paper.

Proposition 3 (Limits of Decentralized Equilibrium Selection)

Suppose there exists a decentralized procedure that, for every economy in the class defined above, identifies a welfare-relevant equilibrium whenever one exists and does so in time polynomial in the size of the economy.

Then that procedure solves the decision problem ALLOCATION–EQUILIBRIUM–WELFARE in polynomial time.

Consequently, unless (P = NP), no decentralized polynomial-time procedure can be guaranteed to identify welfare-relevant equilibria across all instances of the economy.

Proof

Any procedure capable of identifying a welfare-relevant equilibrium whenever one exists can be used to answer the decision problem defined in Definition 4. The result therefore follows from Theorem 1.

The significance of Theorem 1 is therefore not merely that a computationally difficult problem exists. The result places limits on what decentralized equilibrium-selection procedures can be expected to accomplish. The issue is not whether a welfare-relevant equilibrium exists, but whether decentralized agents can reliably identify it through computationally feasible adjustment processes.

Why Sequential Best Response

The preceding results do not imply that agents are restricted to sequential best response or that richer search procedures are unavailable. Rather, they establish a structural limitation: no decentralized polynomial-time procedure can be guaranteed to identify welfare-relevant equilibria across all instances of the environment. Sequential best response is therefore studied as a minimal operational benchmark for decentralized adjustment. The purpose is not to restrict behavior, but to characterize the outcomes generated by tractable, locally optimizing procedures when welfare-relevant equilibrium selection is computationally inaccessible. Richer adjustment rules may alter the set of reachable allocations, but they do not eliminate the underlying distinction between local adjustment and welfare-relevant equilibrium selection established in Theorem 1.

8. The Computational–Dynamic Gap

The previous sections established two distinct results.

First, decentralized adjustment is operational. Firms can evaluate locally feasible alternatives and converge to a stable allocation through sequential best response.

Second, welfare-relevant equilibrium selection may be computationally inaccessible. Identifying a stable allocation satisfying a welfare criterion can require solving an NP-hard problem.

Taken separately, neither result is particularly surprising. The first follows from standard finite-improvement reasoning. The second follows from the combinatorial structure of discrete allocation problems.

The important observation is that the two results concern different objects.

The local dynamic studied in Section 5 asks:

What action should I choose next?

The computational problem studied in Section 6 asks:

Which equilibrium is welfare-relevant?

There is no reason these two questions should have the same answer.

This section establishes that they generally do not.

Theorem 2 (Computational–Dynamic Gap)

There exist supplier-choice economies in which:

1. a welfare-superior equilibrium exists;
2. sequential best response converges in finite time;
3. the equilibrium reached by sequential best response is not welfare-maximizing.

Proof

Consider an economy with three firms

$$N = \{A, B, C\}$$

and two resources

$$R = \{r_1, r_2\}.$$

Capacities satisfy

$$q_{r_1} = q_{r_2} = 1.$$

Each firm possesses two actions:

$$H_i$$

and

$$L_i.$$

The high action requires both resources:

$$d(H_i) = (1,1),$$

while the low action requires neither:

$$d(L_i) = (0,0).$$

Payoffs are

$$u_A(H_A) = 3,$$
$$u_B(H_B) = 4,$$
$$u_C(H_C) = 4,$$

and

$$u_i(L_i) = 1 \forall i.$$

Because every high action requires both resources, at most one firm can choose a high action in any feasible allocation.

Consider the initial profile

$$(L_A, L_B, L_C).$$

Suppose firm A moves first.

Since

$$u_A(H_A) = 3 > 1 = u_A(L_A),$$

firm Aswitches to H_A.

The resulting allocation is

$$(H_A, L_B, L_C).$$

This allocation is feasible.

Neither B nor C can switch to its high action because both resources are already exhausted.

Therefore, no profitable feasible deviation exists.

The allocation is absorbing.

Its welfare is

$$W(H_A, L_B, L_C) = 3 + 1 + 1 = 5.$$

Now consider

$$(L_A, H_B, L_C).$$

This allocation is also feasible and absorbing.

Its welfare is

$$W(L_A, H_B, L_C) = 1 + 4 + 1 = 6.$$

Likewise

$$(L_A, L_B, H_C)$$

is feasible, absorbing, and yields welfare six.

Thus, a welfare-superior equilibrium exists, yet decentralized adjustment can converge to a lower-welfare equilibrium.

The significance of the result is not that local optima exist. Such outcomes are familiar throughout optimization theory. The significance is that decentralized equilibrium dynamics can systematically select them as equilibrium outcomes even when welfare-superior equilibria exist elsewhere in the feasible set.

□

Corollary 1 (Path Dependence)

Different update orders can generate different absorbing equilibria.

Proof

From the same initial allocation

$$(L_A, L_B, L_C),$$

if firm A moves first, welfare converges to five.

If firm B or C moves first, welfare converges to six.

Therefore, realized outcomes depend upon adjustment history.

□

Interpretation

The significance of Theorem 2 does not lie in the numerical example itself. The example is intentionally small and serves only to illustrate a more general mechanism. A decentralized economy may possess a welfare-superior equilibrium, a simple adjustment process, and guaranteed convergence, yet still fail to attain the welfare-superior outcome.

The source of this failure is not irrationality, imperfect information, adjustment costs, or strategic deception. Rather, it arises because local improvement and global welfare selection are fundamentally different computational tasks. Firms can optimize over currently feasible opportunities and converge to a stable allocation without identifying welfare-superior equilibria elsewhere in the feasible set.

The result demonstrates that convergence and welfare attainment are distinct concepts. A decentralized system may converge reliably while still failing to reach the highest-welfare equilibrium available within the economy. The economy converges because firms can evaluate nearby alternatives.

The economy fails to attain the welfare-superior allocation because no decentralized agent solves the global allocation problem. This distinction constitutes the computational–dynamic gap. The existence of such a gap implies that equilibrium existence alone is insufficient for prediction. Understanding realized outcomes requires attention not only to the equilibria that exist, but also to the adjustment processes through which feasible allocations are explored. The

next section examines an important consequence of this mechanism: temporary policy interventions can generate persistent effects even when traditional frictions are absent.

9. Policy Persistence Without Traditional Frictions

The previous section established that decentralized adjustment can converge to absorbing equilibria that are not welfare-maximizing. The existence of such allocations has an important implication. If outcomes depend upon the path through which feasible allocations are reached, then temporary interventions can have lasting consequences even after the intervention itself is removed.

This observation differs from the traditional rationale for policy persistence. In standard accounts, persistent effects arise because agents face adjustment costs, information frictions, contractual rigidities, or other barriers that slow movement toward equilibrium. Here, persistence emerges from a different source. The allocation reached by decentralized adjustment may itself alter the future set of feasible actions.

The result is a form of history dependence generated by feasibility rather than by friction.

9.1 Benchmark Equilibrium Reasoning

Consider the standard equilibrium benchmark.

Suppose a temporary policy changes payoffs for a finite period. Examples include:

- a temporary subsidy,
- a tariff,
- preferential procurement,
- temporary transmission priority,
- a short-lived regulatory exemption.

Once the policy is removed, the underlying economic environment returns to its original state.

Under conventional equilibrium reasoning, one would expect the allocation to return as well. If the original primitives are restored, the original equilibrium should again characterize outcomes.

This logic implicitly assumes that the relevant equilibrium remains operationally accessible after the policy ends.

The computational-dynamic gap suggests that this assumption need not hold.

9.2 A Temporary Policy Shock

Return to the three-firm economy introduced in Section 7.

Baseline payoffs satisfy

$$u_A(H_A) = 3,$$
$$u_B(H_B) = 4,$$
$$u_C(H_C) = 4,$$

and

$$u_i(L_i) = 1 \forall i.$$

The welfare-maximizing equilibria allocate the scarce resources to either firm Bor firm C, generating welfare

$$W = 6.$$

Now suppose a temporary policy increases the payoff to firm A's high action during periods

$$t \in [0, T].$$

Specifically,

$$u_A^P(H_A) = 3 + s,$$

where

$$s > 1.$$

During the policy period,

$$u_A^P(H_A) > u_B(H_B)$$

and

$$u_A^P(H_A) > u_C(H_C).$$

Firm Atherefore becomes the highest-valued claimant to the scarce resources.

Suppose the economy begins from

$$(L_A, L_B, L_C).$$

If firm A updates during the policy period, it chooses H_A.

The resulting allocation is

$$(H_A, L_B, L_C).$$

When the policy expires, payoffs return to baseline.

However,

$$u_A(H_A) = 3 > 1 = u_A(L_A).$$

Firm Atherefore remains at H_A.

Firms $B$and Cwould prefer their high actions, but those actions remain infeasible because the scarce resources are already occupied.

The allocation therefore remains

$$(H_A, L_B, L_C),$$

even after the policy has disappeared.

Proposition 4 (Policy Persistence)

There exist supplier-choice economies and temporary policy interventions such that:

1. the policy changes the allocation reached by decentralized adjustment;
2. the allocation remains absorbing after the policy is removed;
3. the post-policy allocation differs from the welfare-maximizing equilibrium under the restored baseline environment.

Proof

The construction above provides an example.

During the policy period, firm A acquires the scarce resources.

After the policy expires, the resulting allocation remains feasible and absorbing.

Yet welfare remains

$$W = 5$$

rather than

$$W = 6.$$

Thus a temporary intervention generates a persistent allocation effect.

□

9.3 Interpretation

The proposition does not imply that governments can solve computationally difficult allocation problems, nor does it imply that intervention necessarily improves welfare. Rather, it identifies a different channel through which policy may matter. When welfare-relevant allocations are computationally inaccessible, policy can influence outcomes by altering the path through which feasible allocations are constructed. In this setting, policy affects the trajectory of adjustment rather than directly solving the underlying allocation problem.

This distinction becomes increasingly important in environments characterized by combinatorial constraints. As information frictions decline and computational tools improve, inefficiencies may increasingly arise from the structure of the allocation problem itself rather than from traditional market failures. Policy effectiveness and policy persistence therefore need not disappear with the elimination of conventional frictions. They may instead emerge from the computational structure of decentralized coordination.

10. Simulation Illustration

The examples and formal results establish the existence of a computational–dynamic gap. The purpose of the simulation below is to illustrate that the mechanism is not confined to small constructed examples. Even in richer environments with many firms, resources, and actions, decentralized adjustment can converge rapidly while selecting widely different absorbing outcomes.

10.1 Simulation design

The simulated economy has

$$n = 50$$

firms and

$$r = 100$$

resources in the benchmark specification. Each resource has unit capacity:

$$q_r = 1$$

for every r.

Each firm has five randomly generated sourcing plans. Each sourcing plan requires three distinct resources. Payoffs are drawn independently from the discrete uniform distribution on

$$\{1, \dots, 10\}.$$

The simulation proceeds in two stages.

First, for a fixed economy, feasible initial profiles are generated using a greedy randomized assignment procedure. Starting from each feasible initial profile, the economy evolves according to sequential best-response dynamics. At each sweep, firms update in random order and switch only when a strictly higher-payoff feasible sourcing plan is available. The process terminates when no firm possesses a profitable feasible deviation.

Second, the same economy is evaluated against a computed benchmark allocation. Because exact enumeration is infeasible at this scale, the benchmark is obtained through repeated randomized global search combined with greedy welfare-improving revisions. The resulting benchmark is not interpreted as a proof of the global optimum. Rather, it provides a high-welfare reference allocation against which decentralized outcomes can be compared. Because the benchmark is generated through randomized global search rather than exact optimization, it should not be interpreted as the true welfare maximum. Consequently, the reported efficiency ratios represent upper bounds on decentralized performance and lower bounds on the true welfare loss relative to the global optimum.

The principal performance measure is the efficiency ratio

$$\frac{W^{BR}}{W^{B}},$$

where

$$W^{BR}$$

denotes terminal welfare generated by sequential best response and

$$W^B$$

denotes the welfare associated with the computed benchmark allocation.

In addition to welfare levels, the simulation reports convergence speed, welfare dispersion across runs, the number of distinct absorbing equilibria, and the efficiency ratio relative to the benchmark allocation.

10.2 Simulation Results

Across 100 runs using identical primitives but different feasible initial conditions and update orders, all runs converged. The process generated 100 distinct absorbing equilibria, implying that every run converged to a different stable outcome.

Terminal welfare ranged from

$$93$$

to

$$188,$$

with mean welfare equal to

$$142.18.$$

The mean number of update sweeps required for convergence was

$$2.05,$$

and no run required more than

$$4$$

sweeps.

To evaluate the welfare significance of these outcomes, the simulation compares decentralized adjustment against a computed benchmark allocation. The benchmark achieved welfare.

$$W^B = 188.$$

Because the benchmark may fall short of the true optimum, the reported efficiency ratios should be interpreted conservatively. They represent upper bounds on decentralized performance, and true welfare losses may exceed those documented here.

The mean efficiency ratio

$$\frac{W^{BR}}{W^B}$$

was

$$0.756,$$

implying that decentralized adjustment attained approximately 75.6 percent of benchmark welfare on average. The lowest observed efficiency ratio was

$$0.495,$$

while the highest was

$$1.000.$$

Table 1 summarizes the benchmark simulation results.

Table 1. Benchmark Simulation Results

Statistic	Value
Firms	50
Resources	100
Actions per Firm	5
Simulation Runs	100
Distinct Absorbing States	100
Lowest Welfare	93
Highest Welfare	188
Mean Welfare	142.18
Welfare Spread	95
Mean Convergence Sweeps	2.05
Maximum Convergence Sweeps	4
Computed Benchmark Welfare	188
Mean Efficiency Ratio	0.756
Minimum Efficiency Ratio	0.495
Maximum Efficiency Ratio	1.000

Figure 1. Distribution of Terminal Welfare Across Decentralized Adjustment Runs

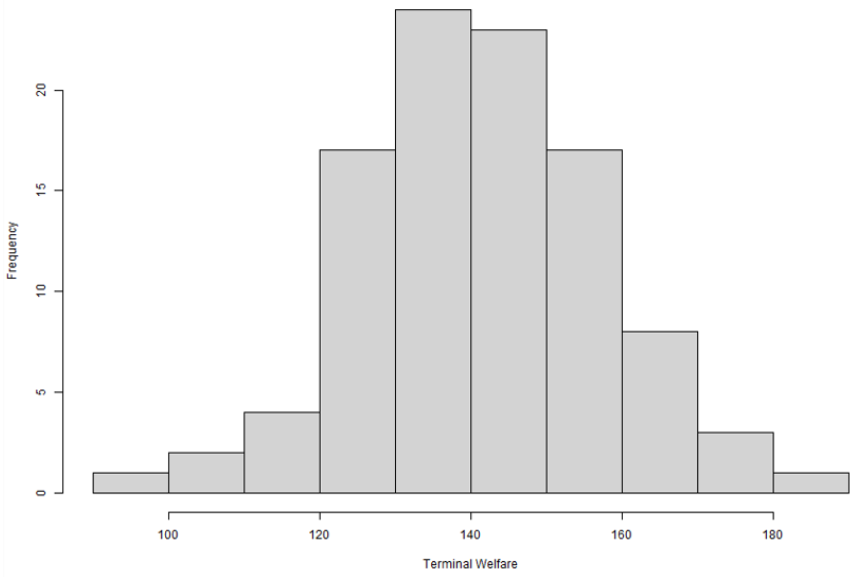


Figure 1 reports terminal welfare across 100 decentralized adjustment runs using identical primitives but different feasible initial conditions and update orders. All runs converge rapidly to absorbing equilibria, yet the resulting welfare levels vary substantially. The distribution illustrates the computational–dynamic gap: decentralized adjustment reliably converges, but convergence does not imply attainment of benchmark welfare. The figure therefore demonstrates that accessibility and convergence are distinct concepts.

10.3 Interpretation

The simulation reveals a striking contrast between convergence and efficiency. In the benchmark case, decentralized adjustment converges in slightly more than two update sweeps on average yet still leaves approximately one-quarter of benchmark welfare unrealized. Decentralized adjustment is extremely fast, requiring slightly more than two update sweeps on average. Yet the resulting allocations achieve only

$$75.6\%$$

of benchmark welfare on average and, in some runs, less than

$$50\%$$

of benchmark welfare.

The source of this gap is not slow adjustment. Nor is it imperfect information, bounded rationality, or adjustment costs. All firms are fully informed and optimize over currently feasible opportunities. The welfare loss arises because local improvement and global allocation selection are distinct computational tasks.

The simulation therefore illustrates the computational–dynamic gap identified in the formal model. Decentralized systems may converge quickly and reliably while still failing to identify welfare-relevant allocations.

10.4 Resource Density and Accessibility

The preceding simulation considered a single benchmark environment. To examine how accessibility changes with the structure of the allocation problem, the simulation was repeated under varying levels of resource availability. Specifically, the number of resources was varied while holding constant the number of firms, actions per firm, and payoff distributions. The objective is to determine whether the severity of the computational–dynamic gap depends on the density of feasibility constraints.

Table 2. Resource Density and Accessibility

Resources	Mean Welfare	Mean Efficiency Ratio	Welfare Spread	Mean Convergence Sweeps
75	105.63	0.669	66	1.40
100	138.23	0.681	82	1.97
125	164.14	0.717	82	2.18
150	190.98	0.746	77	2.56

As Table 2 shows resource availability increases from

$$r = 75$$

to

$$r = 150,$$

mean welfare rises from 105.63 to 190.98. More importantly, the mean efficiency ratio rises from 0.669 to 0.746. Thus decentralized adjustment captures a larger share of benchmark welfare when resource constraints become less restrictive. Although welfare dispersion and efficiency do not vary monotonically across every parameter value, the overall pattern is consistent with the theory: greater resource availability is associated with higher average efficiency and a smaller accessibility gap.

Welfare dispersion does not vary monotonically across parameter values. The welfare spread rises from 66 to 82 as resource availability increases from $r = 75$ to $r = 100$, remains at 82 for $r = 125$, and then falls slightly to 77 at $r = 150$. However, the economically relevant measure in the present context is accessibility rather than dispersion alone. Mean efficiency rises systematically across all four cases.

The result is consistent with the paper's central argument. When resources are scarce, actions compete more intensely for feasibility. Local decisions therefore eliminate larger portions of the feasible allocation space, increasing the likelihood that decentralized adjustment converges to welfare-inferior absorbing equilibria. As resource availability expands, feasibility constraints become less binding and the accessibility gap narrows.

At the same time, convergence remains extremely rapid. Mean convergence sweeps increase only modestly, from 1.40 to 2.56. The results therefore suggest that accessibility problems arise not because decentralized adjustment is slow, but because local adjustment and global allocation selection remain distinct tasks even when convergence is nearly immediate.

11. Economic Implications

The analysis developed above suggests a different perspective on economic policy. Much of modern policy analysis is organized around the presence of frictions. Markets fail because information is incomplete, prices adjust slowly, contracts are costly, or incentives are distorted. Remove these frictions and decentralized adjustment is generally expected to move the economy toward efficient outcomes.

The environments studied here point to a different possibility. Welfare-relevant allocations may remain difficult to attain even when traditional frictions are weak or absent. The obstacle is not necessarily informational. It is combinatorial.

11.1 Existence Is Not Enough

A recurring theme in economics is that the existence of an equilibrium provides a benchmark for prediction and welfare analysis. The results above suggest that existence alone may be insufficient.

A welfare-superior equilibrium may:

exist,
be stable,

and

be easy to verify once identified,

yet remain operationally inaccessible through decentralized adjustment.

The distinction is important because economic outcomes are generated by adjustment processes rather than existence theorems.

In this sense, the relevant question is not merely:

What allocations are possible?

but

Which allocations can decentralized systems actually reach?

11.2 A Different View of Market Failure

Traditional market failures arise from information problems, externalities, public goods, monopoly power, or contractual frictions.

The mechanism identified here is different.

The economy may fail to attain welfare-relevant outcomes even when:

- information is complete,
- agents are rational,
- contracts are enforceable,
- adjustment costs are absent.

The source of inefficiency is the interaction between discrete choices and shared feasibility constraints.

This does not replace traditional market-failure theory. It complements it.

11.3 Policy as Path Selection

The analysis also suggests a different interpretation of policy.

Policy is often viewed as correcting distortions. Taxes, subsidies, regulations, and procurement rules alter incentives in order to move the economy toward a preferred equilibrium.

In combinatorial environments, policy may instead function as a path-selection mechanism.

Temporary interventions can:

- change the order in which scarce resources are claimed,
- alter which allocations become feasible,
- influence which absorbing state is reached.

The effect is not necessarily to solve the global allocation problem. Rather, policy influences the trajectory through which decentralized adjustment proceeds.

11.4 Implications for Technological Change

The argument has particular relevance for discussions of artificial intelligence and digital coordination.

Many observers expect technological progress to weaken traditional justifications for policy intervention by reducing search costs, information asymmetries, forecasting errors, and

transaction costs. To the extent that these frictions diminish, one might expect decentralized systems to move more reliably toward efficient outcomes.

The analysis developed here suggests a more nuanced possibility. Improvements in information processing do not eliminate combinatorial complexity. A system may possess abundant information and powerful computational tools while still facing a vast space of feasible allocations. In such environments, the difficulty is not necessarily knowing what opportunities exist. The difficulty is selecting among large numbers of mutually interacting possibilities.

This distinction matters because technological progress affects different constraints differently. Information may become cheaper. Search may become easier. Communication may become faster. Yet the number of feasible configurations in complex allocation systems may continue to grow. In some settings, the combinatorial challenge may become more important precisely because other frictions have become less important.

The implication is not that technology is ineffective. On the contrary, improved computational tools may expand the set of allocations that are operationally accessible. The point is that accessibility itself becomes an increasingly important economic variable. As traditional frictions recede, the question shifts from whether desirable allocations exist to whether decentralized systems can reliably discover and attain them.

11.5 Beyond the Illustrative Cases

Although the paper focuses on supplier-choice environments, the underlying logic is broader.

The same structure appears in:

- infrastructure allocation,
- transmission access,
- procurement systems,
- logistics networks,
- certification bottlenecks,
- platform ecosystems,
- health-care capacity allocation.

What these environments share is not a common industry but a common computational structure.

The relevant economic problem is the selection of feasible combinations under shared constraints.

11.6 The Central Lesson

The central lesson of the paper is that computational accessibility deserves a place alongside existence and efficiency in economic analysis.

For much of economics, the progression has been:

$$\text{Existence} \rightarrow \text{Efficiency} \rightarrow \text{Policy}.$$

The argument developed here suggests a different sequence:

$$\text{Existence} \rightarrow \text{Accessibility} \rightarrow \text{Adjustment} \rightarrow \text{Outcomes}.$$

When welfare-relevant allocations are computationally inaccessible, adjustment dynamics become as important as equilibrium existence itself.

The central lesson is that equilibrium existence, equilibrium accessibility, and equilibrium selection are distinct concepts. Traditional economic analysis focuses primarily on the first. The environments studied here suggest that the second and third may often be equally important.

The analysis has focused on deterministic adjustment processes. A large literature on stochastic stability examines equilibrium selection under noisy dynamics (Kandori, Mailath, and Rob 1993; Young 1993). Those results typically characterize long-run selection among equilibria in environments with relatively simple strategic structure. Whether analogous selection results extend to combinatorial allocation environments with shared feasibility constraints remains an open question. The present analysis suggests that computational accessibility may play an important role in determining which equilibria remain reachable under such dynamics.

12. Conclusion

The traditional case for policy intervention is frequently grounded in frictions. Markets fail because information is imperfect, adjustment is costly, contracts are incomplete, or incentives are distorted. Remove these frictions and decentralized behavior is generally expected to move the economy toward efficient outcomes.

This paper has explored a different possibility.

Many economically important allocation environments possess a combinatorial structure. Supply chains, logistics networks, infrastructure systems, procurement environments, and capacity-allocation problems all require the selection of feasible combinations of discrete actions subject to shared constraints. In such settings, welfare-relevant allocations may exist without being operationally accessible.

The central distinction is between existence and accessibility. An allocation may be feasible, stable, and welfare-superior, yet remain difficult to identify through decentralized adjustment. The relevant economic question therefore shifts from whether desirable outcomes exist to whether decentralized systems can reliably reach them.

To study this issue, the paper introduced a simple model of sequential best response under feasibility constraints. The resulting adjustment process is operational and converges to stable outcomes. However, the equilibrium selected by local adjustment need not be welfare-maximizing. The paper therefore identifies a computational–dynamic gap: decentralized systems can converge reliably while simultaneously failing to attain welfare-relevant allocations.

This gap generates several implications. It provides a mechanism for inefficiency without relying on informational frictions or adjustment costs. It generates path dependence through evolving feasibility constraints. It implies that temporary policy interventions can have persistent effects by altering the trajectory of decentralized adjustment.

The broader implication is that policy relevance need not disappear as technology reduces traditional frictions. Improvements in information processing, contracting, and communication may weaken some sources of market failure while leaving combinatorial allocation problems largely intact. In such environments, the continuing importance of policy may derive less from correcting informational imperfections and more from influencing the paths through which feasible allocations are constructed. The traditional economic intuition is that more choice improves welfare. In combinatorial environments, however, larger opportunity sets can make welfare-relevant allocations harder to identify. In this limited sense, small may sometimes be beautiful.

The paper does not argue that governments can solve computationally difficult allocation problems better than markets. Governments face the same computational constraints. The claim is more limited. When welfare-relevant allocations are computationally inaccessible, realized outcomes depend on adjustment processes as much as on equilibrium existence. Understanding those processes therefore becomes a central task of economic analysis.

Appendix A. Computational Complexity of Welfare-Relevant Equilibrium Selection

This appendix establishes the computational inaccessibility result discussed in Section 6. The objective is not to introduce a new complexity class or computational framework, but to demonstrate that allocation-relevant equilibrium selection can inherit the hardness of well-known combinatorial allocation problems.

A.1 Welfare-Relevant Equilibrium Selection

Recall the supplier-choice environment introduced in Section 5.

A welfare-relevant equilibrium is a pure feasible equilibrium a^*satisfying

$$W(a^*) \geq B,$$

where $B$is a specified welfare threshold.

Define the decision problem:

ALLOCATION–EQUILIBRIUM–WELFARE

Input:

- a supplier-choice economy E,
- a welfare threshold B.

Question:

Does there exist a pure feasible equilibrium a^*such that

$$W(a^*) \geq B?$$

The purpose of the theorem below is to establish that this question may be computationally difficult even when equilibrium verification is straightforward.

A.2 Membership in NP

The problem belongs to NP.

Given a candidate profile

$$a^* = (a_1^*, \dots, a_n^*),$$

verification requires three checks.

Feasibility

For every resource r,

$$\sum_{i \in N} d_r\,(a_i^*) \leq q_r.$$

Equilibrium

For every firm $i \in N$and every alternative action $b \in A_i$, verify that either

$$(b, a_{-i}^*)$$

is infeasible or

$$u_i(a_i^*) \geq u_i(b).$$

Welfare

Compute

$$W(a^*) = \sum_{i \in N} u_i\,(a_i^*)$$

and verify

$$W(a^*) \geq B.$$

Each step requires time polynomial in the size of the explicit input representation.

Therefore

$$\text{ALLOCATION} - \text{EQUILIBRIUM} - \text{WELFARE} \in NP.$$

A.3 NP-Hardness

The hardness result follows by reduction from the classical 3-SET PACKING problem (Garey and Johnson 1979).

An instance of 3-SET PACKING consists of:

$$U = \{u_1, \dots, u_m\},$$

a collection

$$\mathcal{C} = \{S_1, \dots, S_k\},$$

where each set satisfies

$$| S_j | = 3,$$

and an integer K.

The question is whether there exist K pairwise disjoint sets.

A.4 Construction

Construct a supplier-choice economy as follows.

Firms

Create

$$N = \{1, \dots, K\}.$$

Resources

For every element $u \in U$, create a resource

$$r_u.$$

Set

$$q_{r_u} = 1.$$

Actions

Each firm possesses:

- one action corresponding to every set $S_j \in \mathcal{C}$,
- one outside option O_i.

Thus

$$A_i = \mathcal{C} \cup \{O_i\}.$$

Resource Requirements

For each set action S_j,

$$d_{r_u}(S_j) = \begin{cases} 1 & \text{if } u \in S_j, \\ 0 & \text{otherwise.} \end{cases}$$

Outside options use no resources:

$$d_{r_u}(O_i) = 0.$$

Payoffs

For every productive action,

$$u_i(S_j) = 1.$$

For outside options,

$$u_i(O_i) = 0.$$

Welfare Threshold

Set

$$B = K.$$

A.5 Correctness

Suppose there exist K pairwise disjoint sets.

Assign one such set to each firm.

Because the sets are disjoint,

$$\sum_{i \in N} d_{r_u}(a_i) \leq 1 \forall u.$$

The profile is feasible.

Every firm receives payoff one, implying

$$W(a) = K.$$

No profitable feasible deviation exists because all productive actions yield identical payoff.

Hence a welfare-relevant equilibrium exists.

Conversely, suppose a welfare-relevant equilibrium exists.

Since

$$W(a^*) \geq K$$

and there are exactly K firms with payoffs bounded above by one, every firm must receive payoff one.

Therefore every firm selects a productive action.

Feasibility implies that no two selected actions share a resource.

Consequently the selected actions correspond to K pairwise disjoint sets.

Thus the original 3-SET PACKING instance is a yes-instance.

Theorem A1

ALLOCATION–EQUILIBRIUM–WELFARE is NP-complete.

Proof

Membership in NP follows from Section A.2.

NP-hardness follows from the reduction in Sections A.3–A.5.

Therefore

$$\text{ALLOCATION} - \text{EQUILIBRIUM} - \text{WELFARE}$$

is NP-complete.

□

Appendix B. Interpreting the Simulation Results

The simulation results reported in Section 10 are intended to illustrate the economic significance of the computational–dynamic gap. Because the benchmark economy contains many firms, resources, and feasible allocations, it is useful to provide an intuitive interpretation of what the simulation is actually measuring.

The simulated economy contains 50 firms competing for access to 100 scarce resources. Each firm possesses five alternative sourcing plans, and each sourcing plan requires access to a subset of the available resources. Some plans generate higher payoffs than others, but not all plans can be implemented simultaneously because firms compete for overlapping resources.

The simulation begins from a feasible allocation and allows firms to revise their sourcing plans through sequential best response. At each step, a firm examines the opportunities currently available to it and switches to a higher-payoff plan whenever such a plan remains feasible. Firms therefore behave rationally and optimize over the opportunities immediately available to them. No firm possesses imperfect information, and no adjustment costs are imposed.

The resulting adjustment process converges rapidly. In the benchmark simulations, convergence occurs in slightly more than two update sweeps on average. From the perspective of decentralized adaptation, the economy functions smoothly. Firms identify better opportunities, revise their actions, and eventually reach a stable allocation.

The surprising result is that rapid convergence does not imply high welfare. Although the adjustment process reliably reaches an absorbing equilibrium, decentralized adjustment attains only 75.6 percent of benchmark welfare on average. This gap emerges despite the absence of informational frictions, adjustment costs, or irrational behavior.

The reason is that firms solve optimization problems using information and opportunities available within their current feasible set. Each firm asks a practical question: given the current allocation of resources, what is the best feasible action available to me? The resulting adjustment process is fully rational and locally optimizing. The difficulty arises not because firms are uninformed or irrational, but because identifying welfare-relevant allocations may require solving a computational problem that no decentralized agent can tractably solve. Consequently, decentralized adjustment can converge to stable allocations that are locally optimal yet inferior to higher-welfare allocations elsewhere in the feasible set.

A simple analogy is a crowded parking lot. Each driver rationally chooses the best available parking space at the moment of arrival. The process converges quickly because every driver eventually parks. Yet the resulting arrangement need not minimize total walking distance across all drivers. The system reaches a stable outcome without necessarily reaching the highest-welfare allocation.

The simulation demonstrates the same principle in a more general allocation environment. Convergence and welfare attainment are distinct concepts. A decentralized economy may adapt successfully, stabilize quickly, and satisfy all feasibility constraints while still failing to reach welfare-relevant allocations that exist elsewhere in the feasible set.

The simulation therefore provides a larger-scale illustration of the computational–dynamic gap established analytically in Theorem 2. Equilibrium existence identifies what outcomes are possible. Decentralized adjustment determines which outcomes are actually reached. The simulations illustrate that these two objects need not coincide, even in environments where adaptation itself is nearly immediate.

Generative AI Statement

Generative AI tools were used during the preparation of this manuscript for limited research and editorial assistance. These uses included literature discovery, bibliographic searches, exploration of simulation design and robustness checks, and review of draft text for clarity, organization, and completeness of coverage.

The authors designed the research, developed the theoretical framework, specified the model, conducted and evaluated the simulations, verified all mathematical arguments and references, interpreted the results, and are solely responsible for the content of the manuscript. All conclusions, interpretations, errors, and opinions are the authors' alone.

References

Aarts, Emile H. L., and Jan Karel Lenstra, eds. 2003. *Local Search in Combinatorial Optimization*. Princeton: Princeton University Press.

Akerlof, George A. 1970. "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism." *Quarterly Journal of Economics* 84 (3): 488–500.

Bienstock, Daniel. 2015. *Electrical Transmission System Cascades and Vulnerability: An Operations Research Viewpoint*. Philadelphia: Society for Industrial and Applied Mathematics.

Chen, Xi, Decheng Dai, Ye Du, and Shang-Hua Teng. 2009. "Settling the Complexity of Arrow–Debreu Equilibria in Markets with Additively Separable Utilities." In *Proceedings of the 50th Annual IEEE Symposium on Foundations of Computer Science*, 273–282.

Coase, Ronald H. 1937. "The Nature of the Firm." *Economica* 4 (16): 386–405.

Dantzig, George B., and John H. Ramser. 1959. “The Truck Dispatching Problem.” *Management Science* 6 (1): 80–91.

Deng, Xiaotie, Christos H. Papadimitriou, and Shmuel Safra. 2003. “On the Complexity of Price Equilibria.” *Journal of Computer and System Sciences* 67 (2): 311–324.

Diamond, Peter A. 1982. “Aggregate Demand Management in Search Equilibrium.” *Journal of Political Economy* 90 (5): 881–894.

Garey, Michael R., and David S. Johnson. 1979. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. New York: W. H. Freeman.

Hayek, F. A. 1945. “The Use of Knowledge in Society.” *American Economic Review* 35 (4): 519–530.

Kydland, Finn E., and Edward C. Prescott. 1977. “Rules Rather than Discretion: The Inconsistency of Optimal Plans.” *Journal of Political Economy* 85 (3): 473–491.

Mankiw, N. Gregory. 1985. “Small Menu Costs and Large Business Cycles: A Macroeconomic Model of Monopoly.” *Quarterly Journal of Economics* 100 (2): 529–537.

Monderer, Dov, and Lloyd S. Shapley. 1996. “Potential Games.” *Games and Economic Behavior* 14 (1): 124–143.

Papadimitriou, Christos H. 1994. *Computational Complexity*. Reading, MA: Addison-Wesley.

Papadimitriou, Christos H., and Kenneth Steiglitz. 1998. *Combinatorial Optimization: Algorithms and Complexity*. Mineola, NY: Dover. Originally published 1982.

Rand, Joseph, Nick Manderlink, Will Gorman, Ryan H. Wiser, Joachim Seel, Julie Mulvaney Kemp, Seongeun Jeong, and Fritz Kahrl. 2024. *Queued Up: 2024 Edition—Characteristics of Power Plants Seeking Transmission Interconnection as of the End of 2023*. Berkeley, CA: Lawrence Berkeley National Laboratory.

Rothkopf, Michael H., Aleksandar Pekeč, and Ronald M. Harstad. 1998. “Computationally Manageable Combinational Auctions.” *Management Science* 44 (8): 1131–1147.

Rubinstein, Ariel. 1998. *Modeling Bounded Rationality*. Cambridge, MA: MIT Press.

Sandholm, Tuomas. 2002. “Algorithm for Optimal Winner Determination in Combinatorial Auctions.” *Artificial Intelligence* 135 (1–2): 1–54.

Simon, Herbert A. 1955. “A Behavioral Model of Rational Choice.” *Quarterly Journal of Economics* 69 (1): 99–118.

Simon, Herbert A. 1957. *Models of Man: Social and Rational*. New York: Wiley.

Simchi-Levi, David, Philip Kaminsky, and Edith Simchi-Levi. 2008. *Designing and Managing the Supply Chain: Concepts, Strategies, and Case Studies*. 3rd ed. Boston: McGraw-Hill/Irwin.

Snyder, Lawrence V. 2006. "Facility Location under Uncertainty: A Review." *IIE Transactions* 38 (7): 547–564.

Toth, Paolo, and Daniele Vigo, eds. 2014. *Vehicle Routing: Problems, Methods, and Applications*. 2nd ed. Philadelphia: Society for Industrial and Applied Mathematics.

U.S. Department of Commerce. 2022. *Results from Semiconductor Supply Chain Request for Information*. Washington, DC: U.S. Department of Commerce.

U.S. Department of Defense 2024. *2023 National Defense Industrial Strategy*. Washington, DC: U.S. Department of Defense.

U.S. Government Accountability Office. 2018. *Defense Industrial Base: Integrating Existing Supplier Data and Addressing Workforce Challenges Could Improve Risk Analysis*. GAO-18-435. Washington, DC: GAO.

Varas, Antonio, Raj Varadarajan, Ramiro Palma, Jimmy Goodrich, and Falan Yinug. 2021. *Strengthening the Global Semiconductor Supply Chain in an Uncertain Era*. Boston: Boston Consulting Group and Semiconductor Industry Association.

Williamson, Oliver E. 1975. *Markets and Hierarchies: Analysis and Antitrust Implications*. New York: Free Press.

Wood, Allen J., Bruce F. Wollenberg, and Gerald B. Sheblé. 2013. *Power Generation, Operation, and Control*. 3rd ed. Hoboken, NJ: Wiley.

Young, H. Peyton. 1998. *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions*. Princeton: Princeton University Press.