My fundamental goal is to understand the role of timing, network structure, and uncertainty in making better decisions for complex networked systems.

My work characterizes **resource-allocation policies** for, and within, such heterogeneous **social, biological, and technological systems**. I integrate the fields of optimal control and dynamic programming, (nonlinear) optimization, games, and graph spectra to specify optimal decisions and to characterize the resulting evolution of networked systems. This both allows the design of **better policies** (e.g., for curbing the spread of epidemics under resource constraints) and the design of **better incentives** for autonomous decision-makers (e.g., for the persistence of ethical behavior in an organization). The unifying theme across my work has been stressing that in dynamic processes, *when* one takes an action is as important as how it is targeted and ignoring the temporal element of decision-making can lead to sub-optimal outcomes

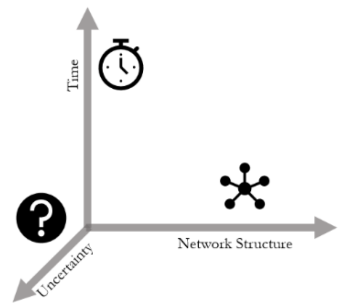

Fig. 1. Decision-making in complex systems should exploit timing, uncertainty, and network structure.

The effects of network structure on the spread of disease and opinions, the functioning of the power-grid, and security of technological networks have become clearer through the past two decades, leading to the foundation of the fields of complexity and network science. The effect of these network structures on decision-making and control is less clear, as more granular information leads to an explosion in the number of decision variables. Decision-making becomes harder when there are complex evolving processes acting on networks (e.g., epidemics), and when the network itself evolves. Deciding when, as well as how, to act further expands the decision-space.

As more data becomes available to decision-makers, models that simplify heterogeneities and over-provision for worst-case scenarios become less palatable. However, in an increasingly connected world where the decisions of a few can be amplified to shape global events, e.g., through the spread of “fake news” and the emergence of global pandemics due to antibiotic misuse, it is essential that we use the full measure of available information, including network effects, in policy-making.

My interdisciplinary research program can be broken down into the following 3 thrusts:

1. Developing tractable dynamic decision-making policies for complex systems;
   - *E.g.,* creating network-based metrics to evaluate dynamic treatment allocations in an epidemic.
2. Creating analysis frameworks for incentive-driven spreading of social phenomena;
   - *E.g.,* modeling the spread of unethical behavior in organizations.
3. Understanding the trade-off between identification, learning, and effectiveness in social systems;
   - *E.g.,* creating adaptive methods for network-based targeting policies in the field under uncertainty.

The metrics I will develop in the first thrust will provide dynamic decision-making tools for real-world policy-making. The frameworks I will create in the second thrust will help anticipate decisions made by strategic actors, allowing a second-order incentive-design approach to shaping networked outcomes. In parallel to the first two thrusts, in the third I will create a link with practical network-based policy-making by studying the effect of real-world information limitations on the performance of these policies, as well as adapting their approaches to be suitable to the incorporation of new information (i.e., learning).

Summary of Contributions

I have sought to illuminate decision-making structures that simplify complex dynamic decisions in biological, social, and technical contexts while stressing that in such processes, *when* one takes an action is as important as how it is targeted and ignoring the temporal element can lead to sub-optimal outcomes.

My work has primarily encompassed three subject-areas: **epidemics, social influence,** and **energy-aware design of technological networks**. In these areas, networks and non-linear dynamics play significant roles in amplifying local decisions that depend on time-varying states. These issues lead to complicated policy questions, where simple dynamic decision-making structures have the most use.

Epidemics: My work on epidemics started during my PhD at Penn through work on epidemic-like message spreading protocols in wireless networks. I then studied more general heterogeneous networked epidemic models and showed that optimal coordinated curative actions involve offering vaccinations and treatments as early as possible and not waiting to take action, regardless of the topology of the network [1, 2]. I also proved this structure persists for curbing malware spread in computer networks, where the remedy (patches) can, unlike vaccines, spread immunity. I delved further into studying stealthy complex malware, like Stuxnet, that aimed to spread in order to reach a particular target while avoiding detection. I showed that the most damaging attack would be one that spreads the most virulent variant of the malware first, at the risk of discovery [3, 4]. Surprisingly, this remains the most damaging attack even if the network implements various popular quarantine policies upon malware detection. In my thesis, I created an overarching taxonomy of heterogeneity for epidemics and characterized the optimal control of epidemics in the presence of heterogeneity (Figure 2, [5]). As a postdoctoral researcher at the Yale School of Public Health, I am designing efficient vaccine schedules and contact-tracing policies under real-world constraints. I plan to continue developing networked models and optimal dynamic allocation mechanisms to mitigate the spread of epidemics, focusing on understanding the direct relationship between network structure and allocations (Thrust 1) as well as creating valid counterfactuals for policy evaluation and adaptation in the field (Thrust 2). I will also continue applying epidemic modeling techniques to similar problems in networked technological systems, such as the spread of data and malware in wireless systems.

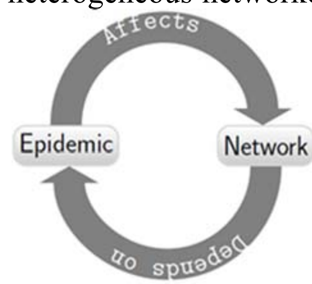

Fig. 2. Heterogeneity types:
1- Network heterogeneity,
2- Resource heterogeneity, &
3- Epidemic heterogeneity.

Social Influence: In my postdoctoral work at Cornell, I showed how political and marketing campaigns should optimally divide their budget between channels at different times to maximize purchase decisions and votes [6-8]. This is more complex than targeting to curb epidemics, as the effect of advertising over a channel is limited by reach and fit with the target audience, as well as possible reinforcement/disbelief from the social networks of targeted individuals. I showed that an optimal campaign should initially prioritize reach, and only focus on targeting likely voters/purchasers late on in the decision cycle, while using each channel in waves (i.e., cycling the use of channels) [8]. As a postdoc at Yale Institute for Network Science, I investigated social influence and targeting with complex behavioral models and limited network information. Many social targeting approaches require a complete mapping of a social network, which is expensive. I showed how limited network information can be used in scalable algorithms to choose optimal seed sets to create cascades [9-12], as well as how the stability of a social group can be quantified, even without knowledge of group norms, using a targeting heuristic [13, 14]. I proved that possible realizations of uncertain networks should be sampled in direct proportion to our priors to find influence maximizing sets [12] and that fairness norms can make a group less stable [14].

I also investigated differing organizational whistleblowing policies using game theory for a corporate partner, showing that networked policies lead to better organizational outcomes and that surprisingly, whistleblowing only thrives when punishments for unethical behavior are relatively light [15, 16]. I will continue my work on the effect of incentives on network-based decision-making in social systems (Thrust 2), with a focus on validating models through online experimentation platforms and integrating possible network adaptations into the policy design (Thrust 3).

Energy-Aware Design: As a PhD student at Penn working on delay-tolerant wireless networks, where end-to-end connectivity is rare and messages must be relayed to reach their destination I showed how the remaining energy in nodes should be optimally utilized in the message forwarding decision at each instant to guarantee message delivery while maximizing the lifetime of the network, leading to a simple and deployable DTN message-forwarding algorithm [17, 18]. As an intern at NEC Labs, I proposed a patent-pending framework for the use of grid-scale batteries, showing how their optimal use in a microgrid can be mapped into a problem of pricing battery power across time, and creating a novel, principled way for obtaining the prices from price (using shadow prices), demand, and generation data [19, 20]. As a postdoctoral associate at the Yale Institute for Network Science, I have developed distributed methods to create reliable wireless network backbones in hostile environments, such as battlefields, where reliability is measured through the energy endowment of nodes chosen as part of the backbone [21-24]. In the future, I will apply more of the distributed decision-making and adaptation methods developed for such technological systems to social systems (Thrust 2) while considering the relationship between interdependent social and technological networks, such as the electrical grid and transportation systems, and the effect of relevant resource constraints.

Future Research Directions

In the next 5 years, I plan to investigate 3 research thrusts that will propel my long-term research goal: understanding the role of time, autonomy, and model uncertainty in designing effective network policies.

Thrust 1 - Developing tractable dynamic decision-making policies for complex systems

I will investigate dynamic network-based decision-making metrics derived directly from the resource allocation framework, *a la* [8]. An important insight of complex systems models has been that treatments can be amplified by a system's internal dynamics to cause an outsize effect on behavior, e.g., network-based vaccination and quarantine in epidemics and social media targeting. The bulk of prior work on this topic substitutes notions of centrality to the structure of a network for importance in terms of influencing decision-making, and artificially assumes that one is constrained to interact with the network at a specific point in time, whereas one can execute dynamic policies relying on coordinated contingent actions. The dynamic resource allocation framework avoids these pitfalls while limiting decision-space explosion.

I will then develop approximations to these metrics that only require local network information and can incorporate a multiplicity of relevant networks simultaneously [21-24]. This is especially important in social systems where mapping the global network is prohibitively costly or there is significant measurement error.

Finally, I will incorporate evaluation, validation, and statistical tests for accuracy directly into the framework from the outset. Recently, the practical value of network targeting has been questioned due to poor performance in trials arising from over-simplifications in theoretical studies and a mismatch between promises and outcomes. The resource allocation framework allows policymakers the flexibility to adapt policies to varying goals. Validating the framework in field studies and online experiments is the long-term goal of this thrust.

Thrust 2- Creating analysis frameworks for incentive-driven spreading of social phenomena

I will study the spread of behaviors among complex decision-makers, incorporating biases relevant to the context (from literature in sociology and psychology), and will investigate the effects of network structure on the spread of behaviors. While mechanistic epidemic models are adequate to explain the spread of some pathogens and behaviors, they perform less well in describing spreading processes that are due to complex human decisions. Network and evolutionary games offer promising approaches to creating identifiable models for the spread of behavior, including through a focus on networked externalities

Furthermore, I will investigate the design of incentives and interventions to help spread beneficial actions. Of particular interest is the case where both the network and the incentives can be designed or altered simultaneously to optimize outcomes. In my previous work, I have investigated social comparison and social norms [13, 14, 25], and emotional reasoning [26] and their effect on social choices and behaviors. Aggregating these effects that manifest at differing population scales (i.e., at the psychological, interpersonal, and networked levels) requires the creation of more complex network games.

Finally, I intend to validate models of decision-making, as well as the usefulness of interventions, through simple experiments on real subjects, as enabled through online experimentation platforms such as Breadboard. The ultimate goal of this thrust is to understand the effect of network structure on population-level behavior change and the learning of successful strategies, and to help these insights guide policy-design for public health interventions.

Thrust 3 – Understanding the trade-off between identification, learning, and effectiveness in social systems

I will develop adaptive resource allocation tools that focus on optimizing the *trajectory* of a spreading process while at the same time minimizing the uncertainty cone of possible realizations resulting from the interventions, using approaches from dynamical systems and stochastic processes. Work on real-world complex system decision-making with information limitations has focused primarily on the robust maximization paradigm, which optimizes against a worst-case realization of uncertainty. This, however, can lead to overprovisioning and a suboptimal allocation of resources.

Furthermore, I will investigate the effect of information limitations and systematic biases in the way information is gathered and represented on decision-making in the first two thrusts. Most insights on the effect of targeted interventions in networked systems depend on a complete characterization of network structure, assuming what is observed is either the whole existing network or is a representative sample. Interactions between the observed network and unobserved existing contacts can complicate decision-making, rendering many efficient algorithms useless in practical settings.

Finally, I will incorporate the effect of information accumulation in the decision-making framework using approaches from reinforcement learning. It is necessary to offer decision-making algorithms and tools that can adapt and learn from incoming data with more diverse sets of guarantees, as well as the ability to consider counterfactuals. The ultimate goal of this thrust is to incorporate incoming data into a simultaneous modeling and decision-making ecosystem, allowing decision-makers to tune the trade-off between robustness and optimality according to their domain knowledge and policy goals, with the certainty that the model will itself adapt its model and decisions to incoming data.

References

1. M. H. R. Khouzani, S. ESHGHI, S. Sarkar, S. S. Venkatesh, *"Optimal patching in clustered epidemics of malware," in* IEEE ITA Workshop*, 2012.*
2. S. ESHGHI, M. H. R. Khouzani, S. Sarkar, S. S. Venkatesh, *"Optimal patching in clustered malware epidemics," in* **IEEE/ACM Trans. on Networking***, 24(1), 283–298, 2016.*
3. S. ESHGHI, S. Sarkar, S. S. Venkatesh, *"Visibility-aware optimal contagion of malware epidemics," in* IEEE ITA Workshop*, 2015.*
4. S. ESHGHI, S. Sarkar, S. S. Venkatesh, *"Visibility-aware optimal contagion of malware epidemics," in* **IEEE Trans. on Automatic Control***, 62(10), 5205-5212, 2017.*
5. S. ESHGHI, *"Optimal control of epidemics in the presence of heterogeneity,"* ***PhD thesis****, University of Pennsylvania, 2015.*
6. S. ESHGHI, V. M. Preciado, S. Sarkar, S. S. Venkatesh, Q. Zhao, R. D'Souza, A. Swami, *"Spread, then target, and advertise in waves: optimal capital allocation across advertising channels,"* IEEE ITA Workshop*, 2017.*
7. S. ESHGHI, L. Tassiulas, *"Efficient dynamic centrality metrics for election advertising - a case study"*, in Yale Day of Data*, 2017.*
8. S. ESHGHI, V. M. Preciado, S. Sarkar, S. S. Venkatesh, Q. Zhao, R. D'Souza, A. Swami, *"Spread, then target, and advertise in waves: optimal budget allocation across advertising channels," in* **IEEE/ACM Trans. on Network Science and Engineering***, 2018.*
9. S. Stein, S. ESHGHI, S. Maghsudi, L. Tassiulas, R. K. E. Bellamy, N. R. Jennings, *"Heuristic algorithms for influence maximization in partially observable social networks," in* Proceedings of SocInf*, 2017.*
10. S. ESHGHI, S. Maghsudi, V. Restocchi, L. Tassiulas, R. K. E. Bellamy, N. R. Jennings, S. Stein, *"Influence maximisation in partially observable social networks," in* Proceedings of Conference on Complex Networks Workshop*, 2017, 98-100.*
11. S. Stein, S. ESHGHI, S. Maghsudi, L. Tassiulas, R. K. E. Bellamy, N. R. Jennings, *"Influence maximisation beyond organisational boundaries," in* Proceedings of IEEE DAIS Workshop*, 2017.*
12. S. ESHGHI, S. Maghsudi, V. Restocchi, S. Stein, L. Tassiulas*, "Efficient influence maximization under partial network visibility," in preparation.*
13. S. ESHGHI, G. R. Williams, G. B. Colombo, L. D. Turner, D. G. Rand, R. M. Whitaker, L. Tassiulas, *"Mathematical models for social group behavior," in IEEE* DAIS Workshop*, 2017.*
14. S. ESHGHI, G. R. Williams, G. B. Colombo, L. D. Turner, D. G. Rand, R. M. Whitaker, L. Tassiulas, *"Social group stability and fracture," in* Proceedings of the 2017 IEEE Allerton Conference*, 478-485.*
15. S. ESHGHI, L. Tassiulas, *"Innovation, cheating, and whistleblowing - a game theoretic perspective,"*, in Proceedings of the 2018 IEEE CISS*, 2018.*
16. S. ESHGHI, L. Tassiulas, *"Whistleblowing games in networks"*, in Proceedings of the 2018 IEEE CDC*, 2018.*
17. M. H. R. Khouzani, S. ESHGHI, S. Sarkar, N. B. Shroff, S. S. Venkatesh, *"Optimal energy-aware epidemic routing in DTNs," in* Proceedings of Mobihoc '12*, 175–182, 2012.*
18. S. ESHGHI, M. H. R. Khouzani, S. Sarkar, N. B. Shroff, S. S. Venkatesh *"Optimal energy-aware epidemic routing in DTNs," in* **IEEE Trans. on Automatic Control***, 60(6), 1554–1569, 2015.*
19. S. ESHGHI, R. M. Patil, *"Optimal battery pricing and energy management for microgrids," in* Proceedings of the 2015 IEEE ACC*, 4994–5001, 2015.*
20. S. ESHGHI, R. M. Patil, R. Sharma, *"Optimal battery pricing and energy management for localized energy resources,"* **US Patent 20,160,093,002 (Application)**. *Mar-2016.*
21. D. Papakostas, S. ESHGHI, D. Katsaros, L. Tassiulas, *"Distributed algorithms for multi-layer connected edge-dominating sets"*, in Proceedings of the 2018 IEEE CDC*, 2018.*
22. D. Papakostas, S. ESHGHI, D. Katsaros, L. Tassiulas, *"Energy-aware distributed edge domination of multilayer networks"*, submitted to 2019 IEEE ACC*, 2018.*
23. D. Papakostas, S. ESHGHI, D. Katsaros, L. Tassiulas, *"Distributed algorithms for multi-layer connected edge-dominating sets"*, **IEEE Control Systems Letters (L-CSS)**, 3(1), 31-36, 2019.
24. D. Papakostas, S. ESHGHI, D. Katsaros, L. Tassiulas, *"Energy-efficient backbone formation in military multi-layer ad-hoc networks"*, **Ad Hoc Networks**, 81, 17-44, 2018.
25. R. K. E. Bellamy, G. B. Colombo, S. ESHGHI, G. de Mel, C. Giammanco, R. Morris, D. G. Rand, L. D. Turner, R. M. Whitaker, G. R. Williams, *"A computational framework for modelling inter-group behaviour using psychological theory," in* Proceedings of SPIE Defense + Security*, 2018.*
26. D. Mott, T. Kelley, C. Giammanco, S. ESHGHI, Y. Zhang, *"A framework for modelling the effect of emotion on uncritical reasoning," in* KSCO Workshop*, 2017.*