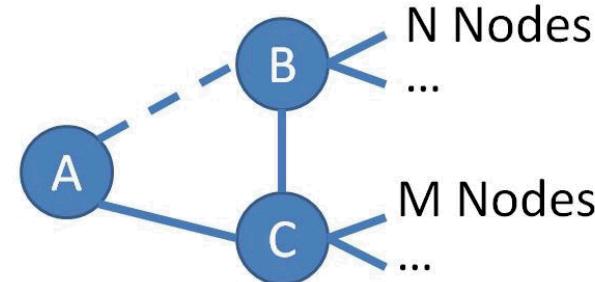
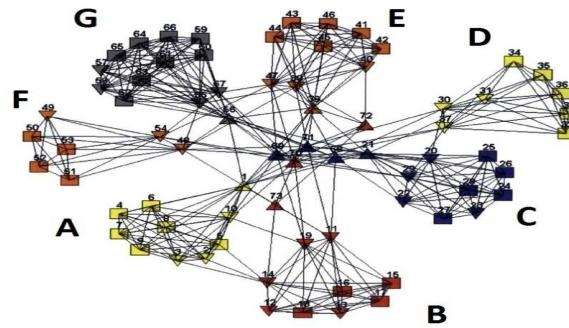


Exceptional service in the national interest



Bayesian Security Games for Controlling Contagion



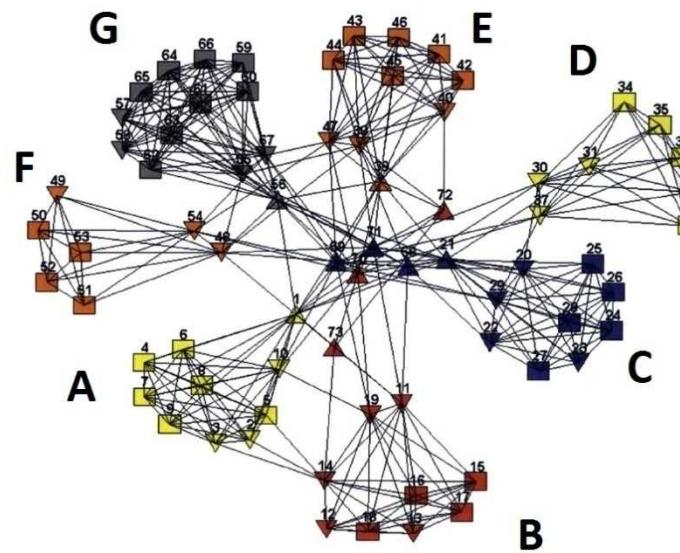
Jason Tsai¹, Yundi Qian¹, Yevgeniy Vorobeychik², Christopher Kiekintveld³, and Milind Tambe¹



Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

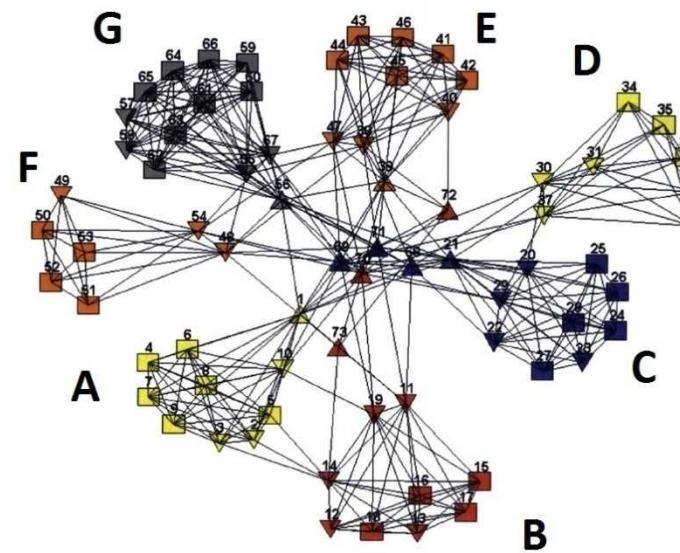
Setting

- Start with a network: edges represent possible influence transmission
- Two players, both aim to maximize influence, i.e., convert the majority of the nodes to their agenda



Motivation: Counterinsurgency

- U.S. vs Taliban in Afghanistan
- Taliban attempts to gain local support
- U.S. attempts to mitigate Taliban support

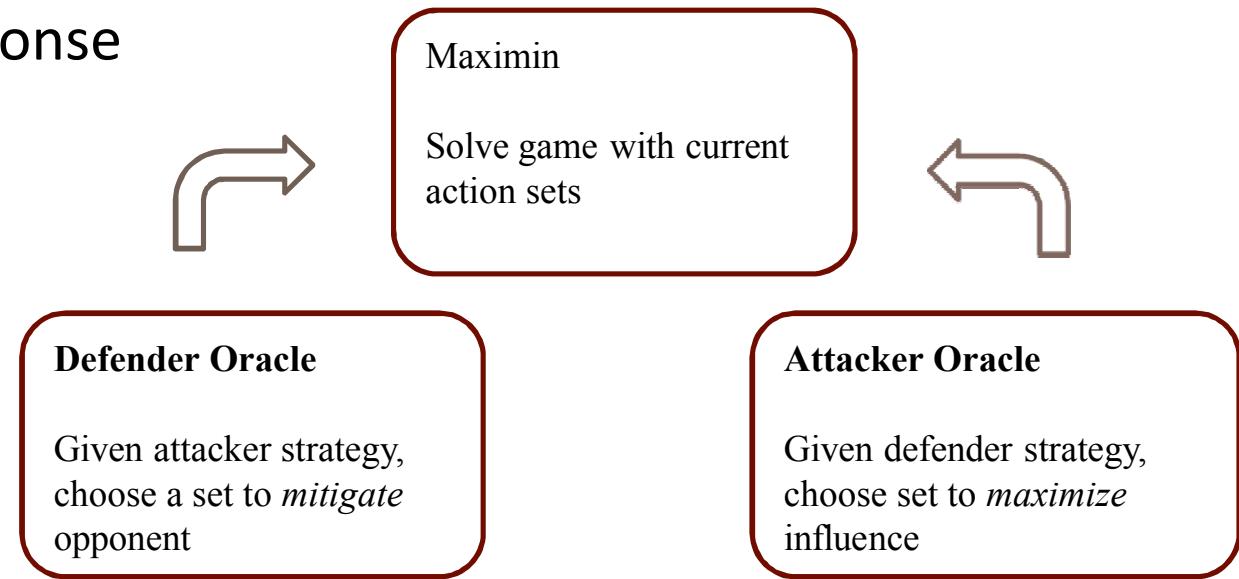


Point of Departure: Baseline Model

- 2-players on a network: “influencer” (maximize contagion) vs “mitigator” (minimize impact of influencer); **zero-sum**
- Spread of influence: variation on independent cascades model (Kempe et al., 2003)
 - Both players simultaneously select up to R nodes
 - Activate each edge (i,j) with probability p_{ij}
 - Any uninfluenced node is “activated” if it has an edge to an influenced node
 - If node i has influenced neighbors of both kinds, flip unbiased coin

Computing Equilibria in Baseline Model

- Formulate as a (very large) Linear Program (actions = subsets of nodes to influence)
- Double-oracle algorithm
 - In each iteration, for each player, compute best response to the other, add this strategy to the LP, repeat
- Greedy best response
- LSMI heuristic to approximate influence

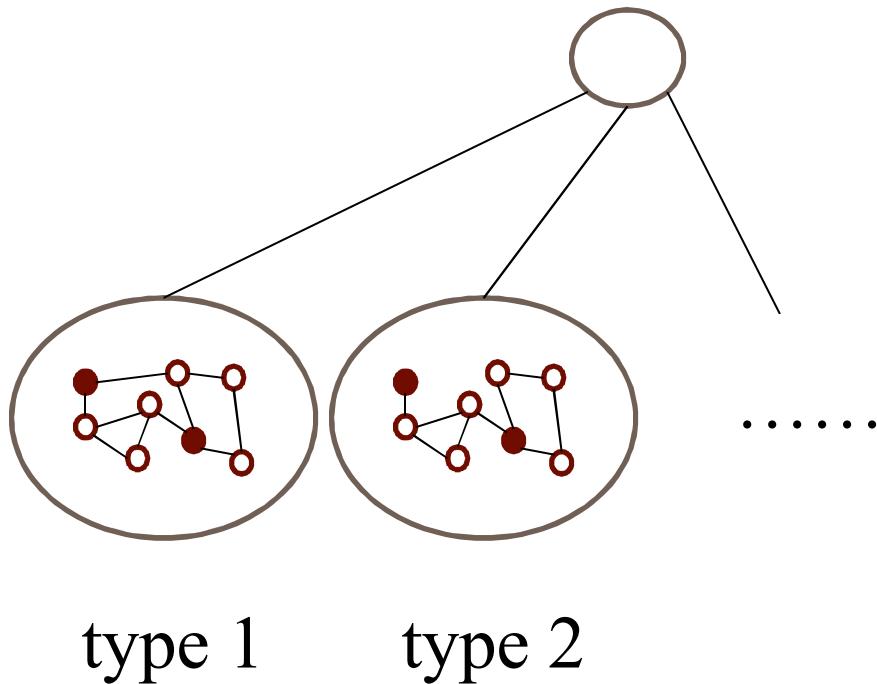


Uncertainty about the Network

- Previous art assumes complete certainty (or symmetric information) about the network structure
- In reality, the influencer typically knows the network much better than the mitigator
 - **Influencer** is “local”, better attuned to culture, reputation, better understands how the influence actually flows
 - **Mitigator** is a foreigner, needs considerable effort to obtain accurate information, and may still be wrong due to misunderstanding local culture and linguistic nuances
- Model this information asymmetry as uncertainty of the mitigator about network structure
 - Existence of nodes/edges, influence probabilities, relative value of nodes (who are the real opinion leaders)

Model: Bayesian Zero-Sum Game

types are possible graphs



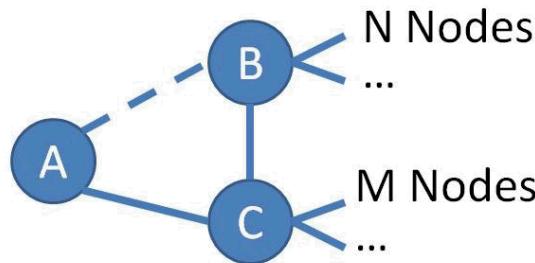
Assume that the influencer knows the true graph

Past Literature:

Graph Structure has Significant Impact

- Watts and Strogatz (1998) – a few short cuts make a dramatic difference
- Costenbader and Valente (2003) – numerous centrality measures shown to be highly sensitive to data errors
- Kossinets (2008) – Missing data can dramatically alter network-level statistics
- Lahiri et al. (2008) – network changes can have dramatic and unpredictable impacts on influence spread

Observation: Loss from Ignoring Uncertainty Can be Unbounded!



TODO: argument

#P-Hard even to estimate the impact of uncertainty about *a single edge*!

Scalability Challenges

- Can formulate our problem as an LP (since zero-sum), but there are 3 major scalability challenges
- Challenge 1: Estimating influence is #P-Hard
 - Solution: use LSMI heuristic (Tsai et al, 2012)
- Challenge 2: Strategy spaces of both players are exponential
 - Solution: Bayesian double-oracle (Halvorson et al., 2009)
- Challenge 3: Exponential number of influencer types (each type = graph)
 - Solution: ?

“Naïve” Idea

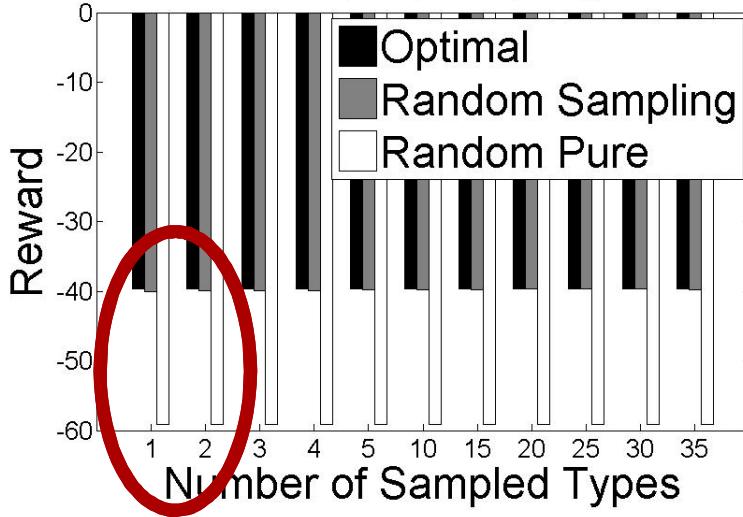
- Naïve idea:
 - take K sample types from the distribution over graphs
 - Use Bayesian double-oracle with only these types
- How well does it work? How large does K need to be (relative to the total number of possible types) to get good solutions?
- These are **empirical questions**

Experiments

- Consider variations of:
 - Graph topologies (generative models and real graphs)
 - Uncertainty/distribution over graphs
- Graph topologies:
 - “Scale-free”: preferential attachment model
 - “Small-world”: Watts-Strogatz model
 - BTER: generate K dense Erdos-Renyi subgraphs (communities), add random inter-community edges; can specify arbitrary degree distribution and clustering coefficient distribution
 - “Indian villages”: social network for several Indian villages
- Models of uncertainty:
 - “Random Edge Uncertainty”: uncertain about K randomly chosen edges
 - “Influential Node Uncertainty”: uncertain about the identity of a highly influential node (add 4 edges to 3 randomly chosen nodes)
 - “Intercommunity Edge Uncertainty” (BTER only): uncertain about K edges that connect communities (dense subgraphs)

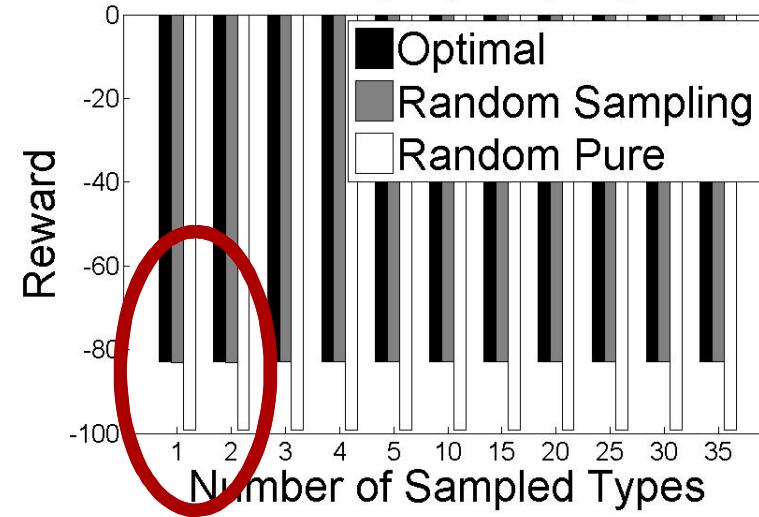
Scale-free, Small-world graphs (random edge uncertainty)

Scale-Free Graph (Sampling Scale-up)



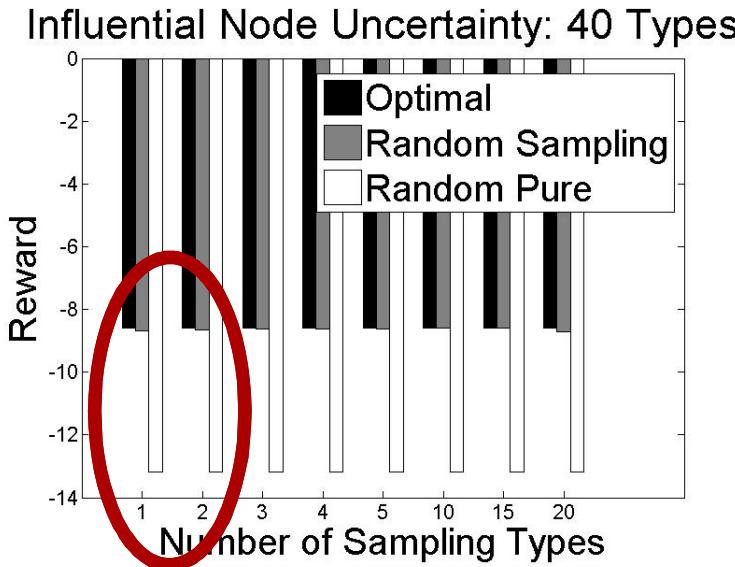
Scale-free Graphs

Small-World Graph (Sampling Scale-up)

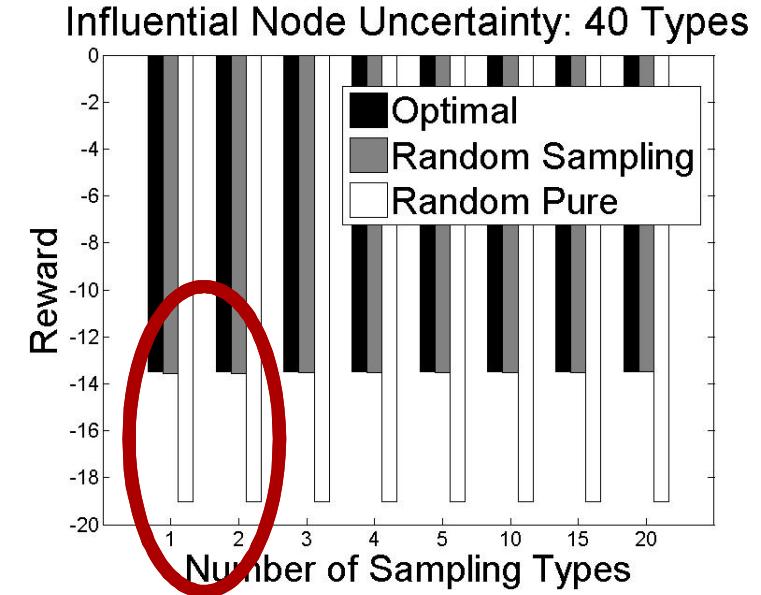


Small-world Graphs

Scale-free, Small-world graphs (influential node uncertainty)

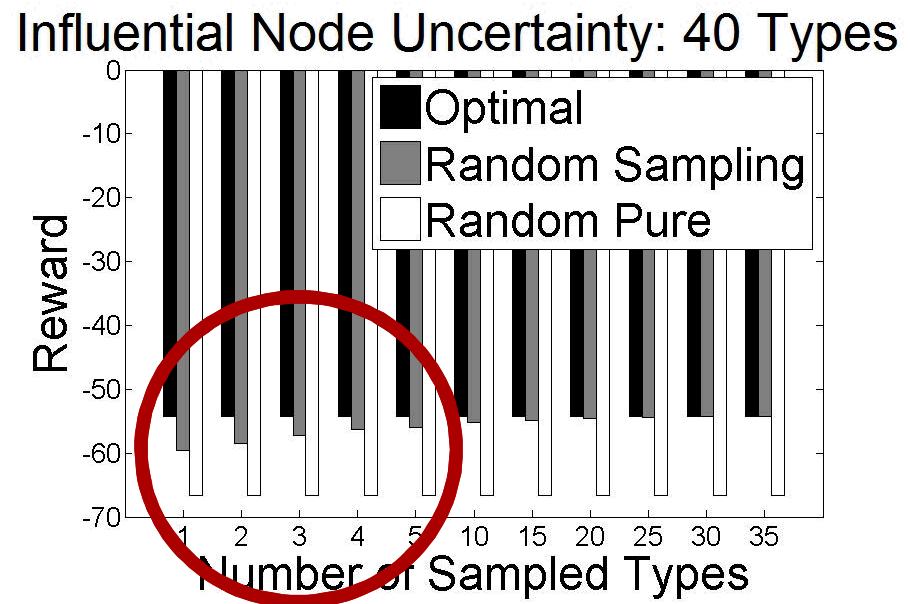
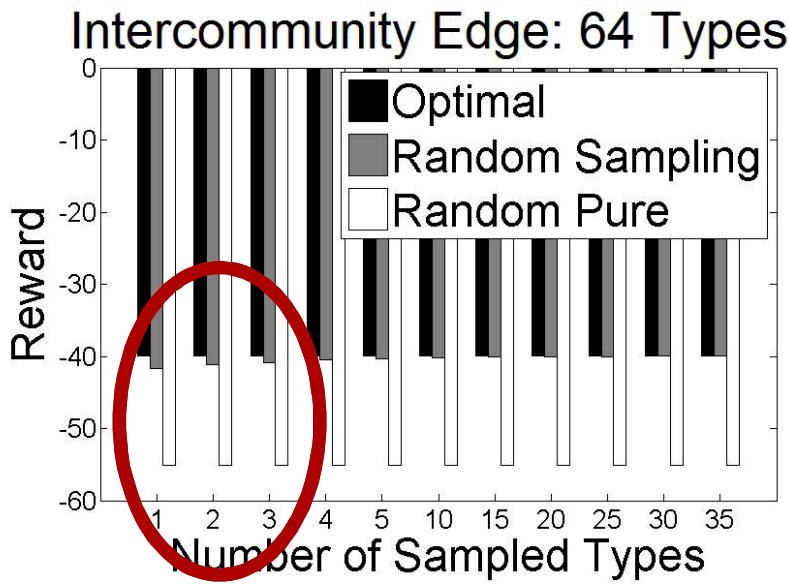


Scale-free Graphs



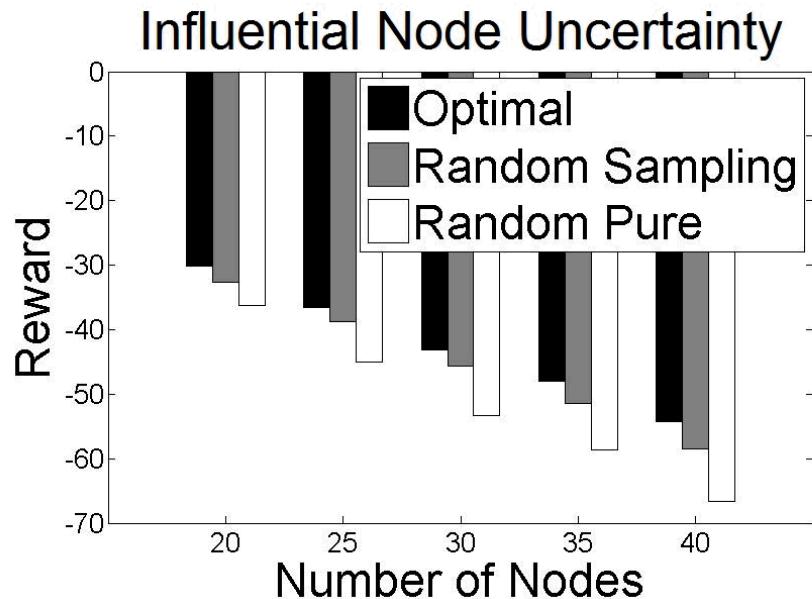
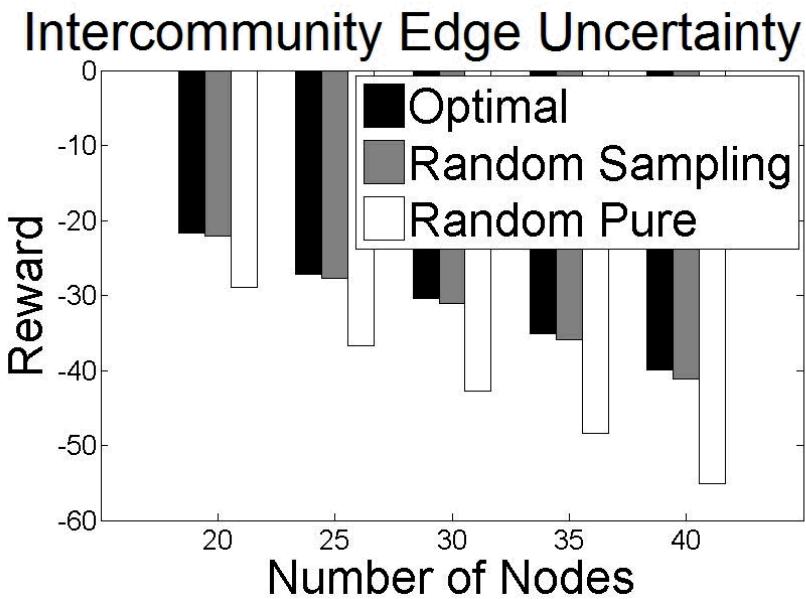
Small-world Graphs

BTER Graphs



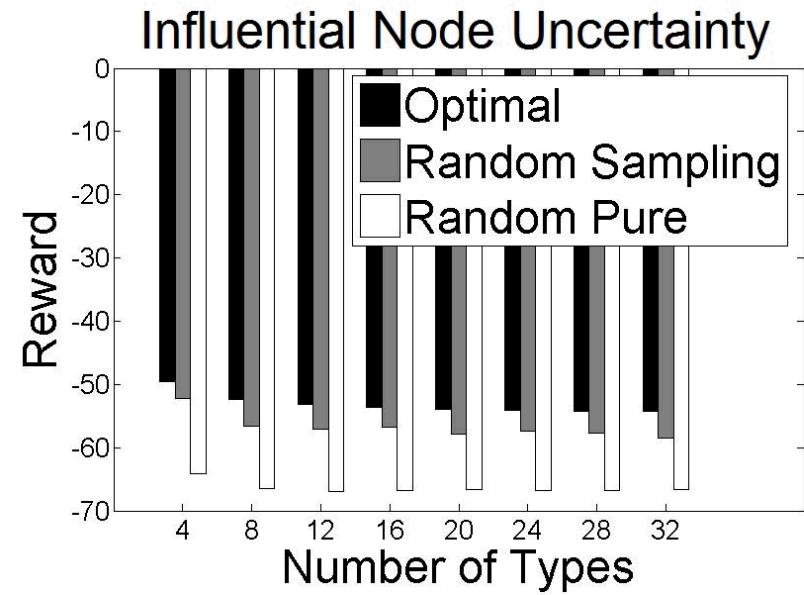
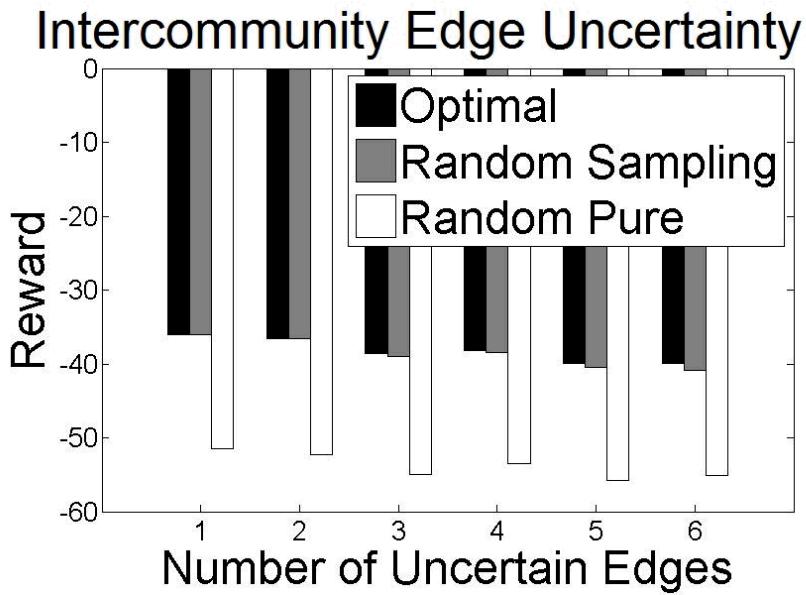
Suffices to sample a small number of types (2-5)

BTER Graphs



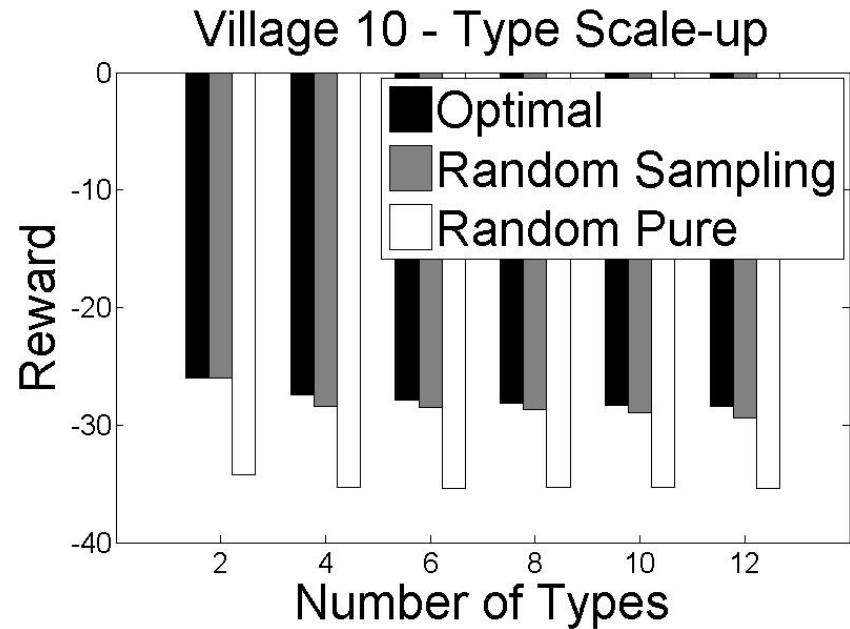
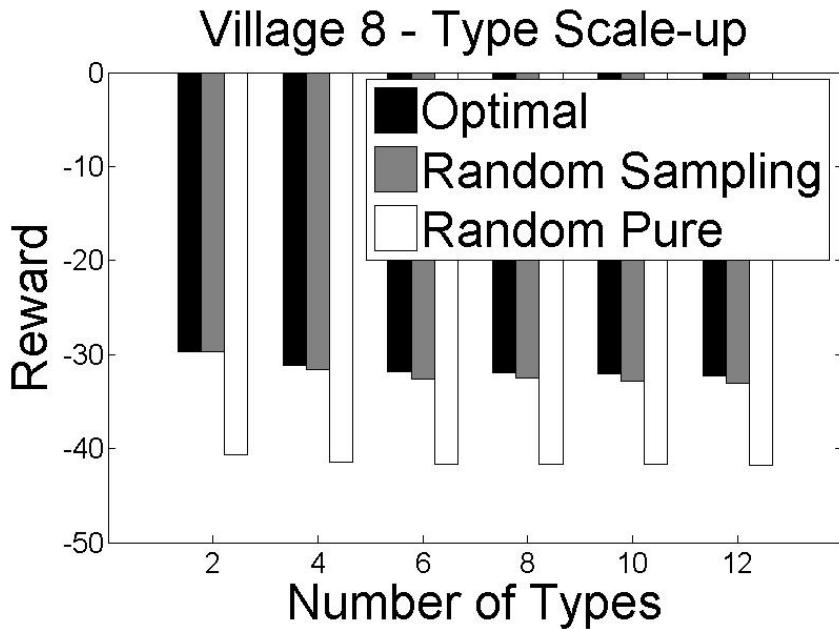
Suffices to sample a small number of types (2-5)

BTER Graphs



Not much affected as we increase the number of uncertain edges or total number of types

Indian Village Graphs



Same story if we look at Indian village graphs

Summary of Results

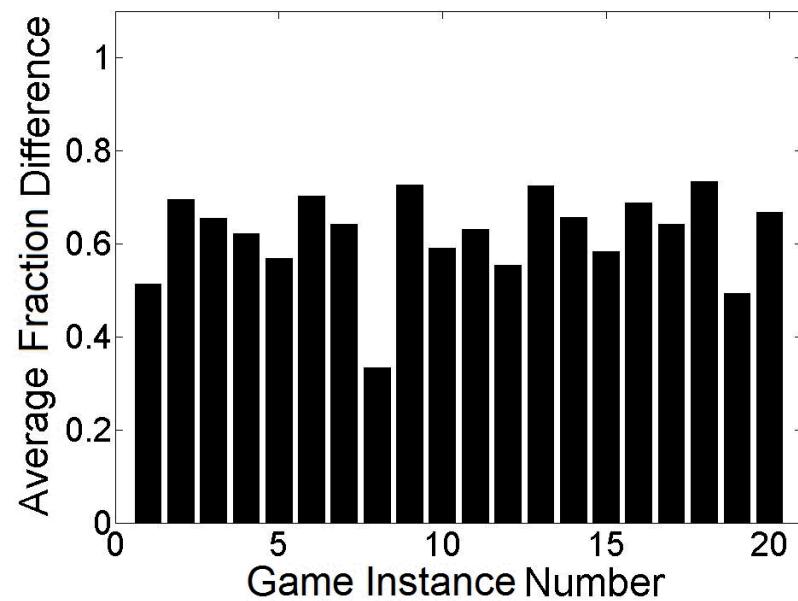
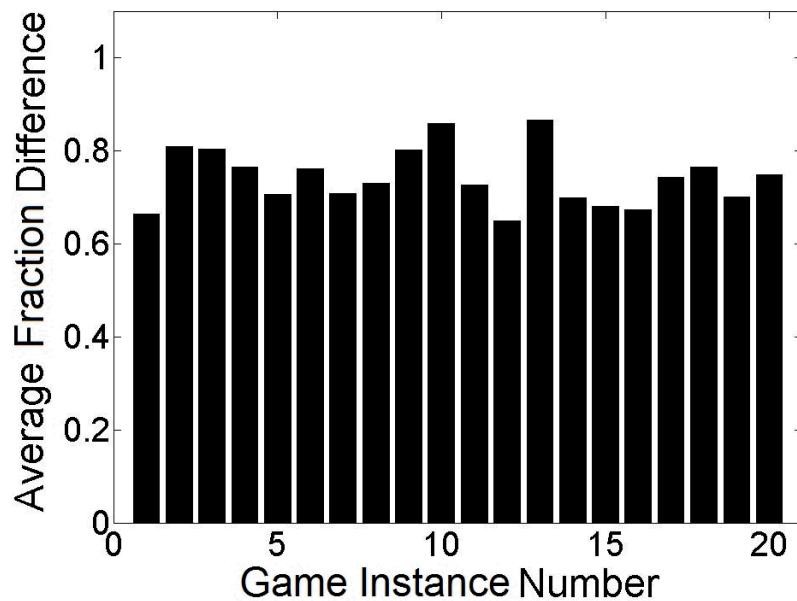
Parameter	Variations	Parameter	Variations
Graph type	Scale-free, Small-world,	Number of types	4-40
Graph size			4, 10-50
Uncertainty	Intercommunity edge, Intercommunity edge set, Inter/Intracommunity edge, Influential node	<i>Bottom-line: naïve strategy works!</i>	5-40
Sampled Types	1-40	Avg contagion probability	0.1, 0.4, 0.7
		Community density	0.6, 0.9
		Indian village graph weight determination	Uniform, Weighted

Over 200 parameter variations tested

Why does it work?

consider overlap in nodes used by type-optimal strategies against:

a random type *optimal strategy*



appears to be a core set of nodes used by optimal strategies for individual types and optimal Bayesian strategy

Illustration: “Equal-Community” Graphs

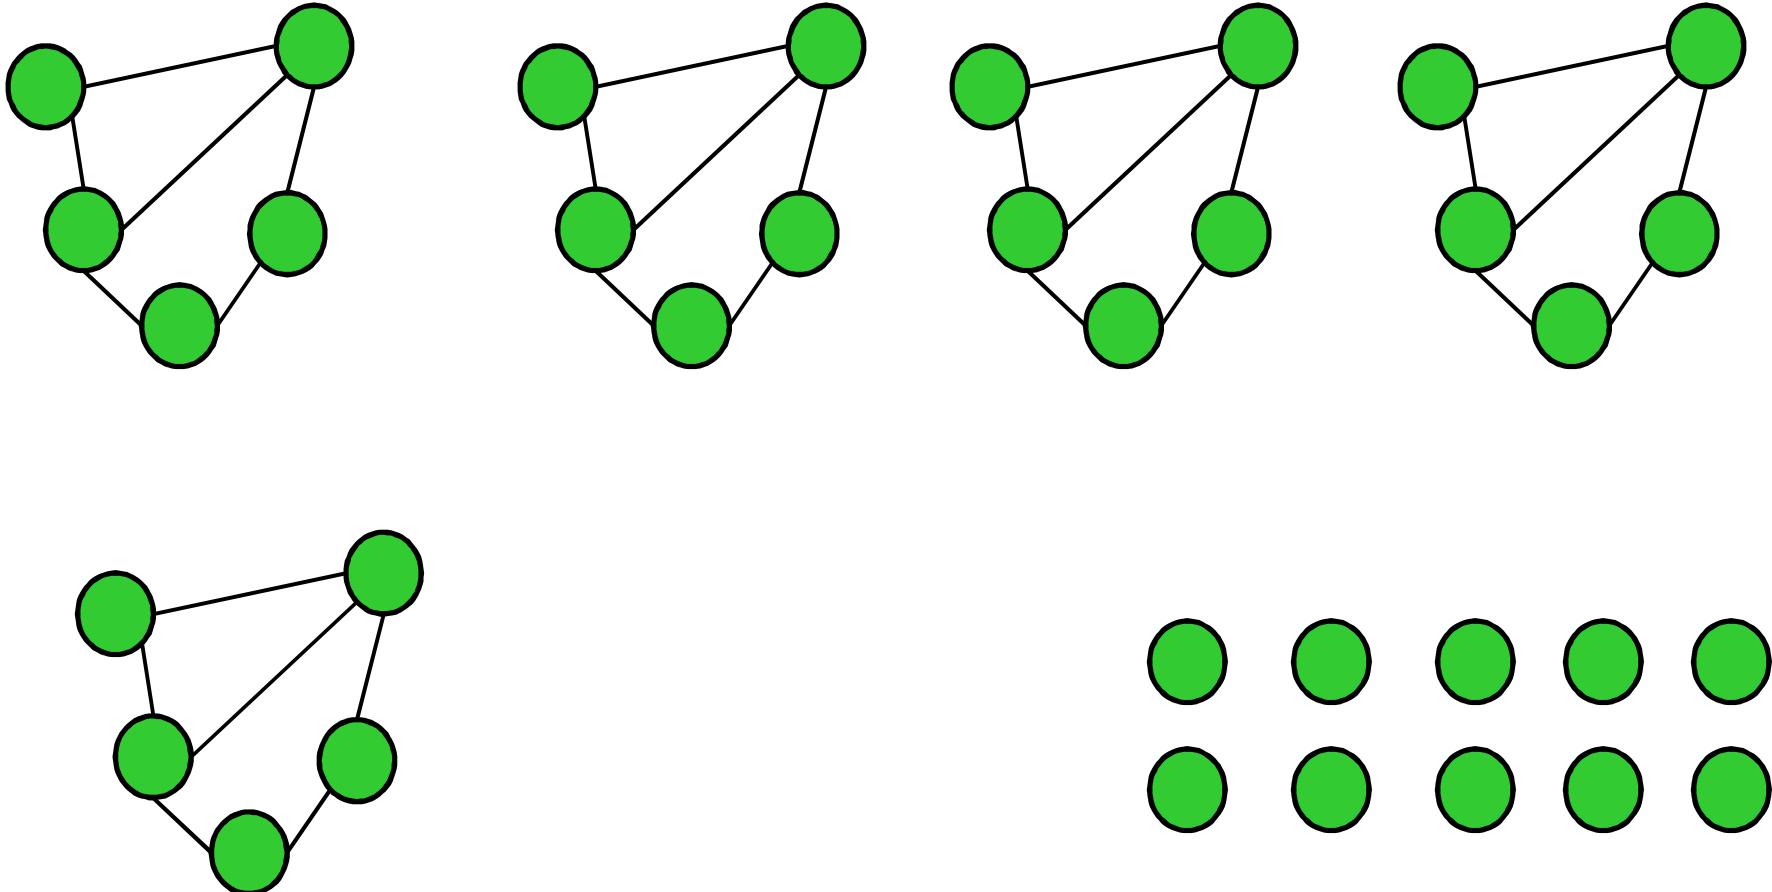
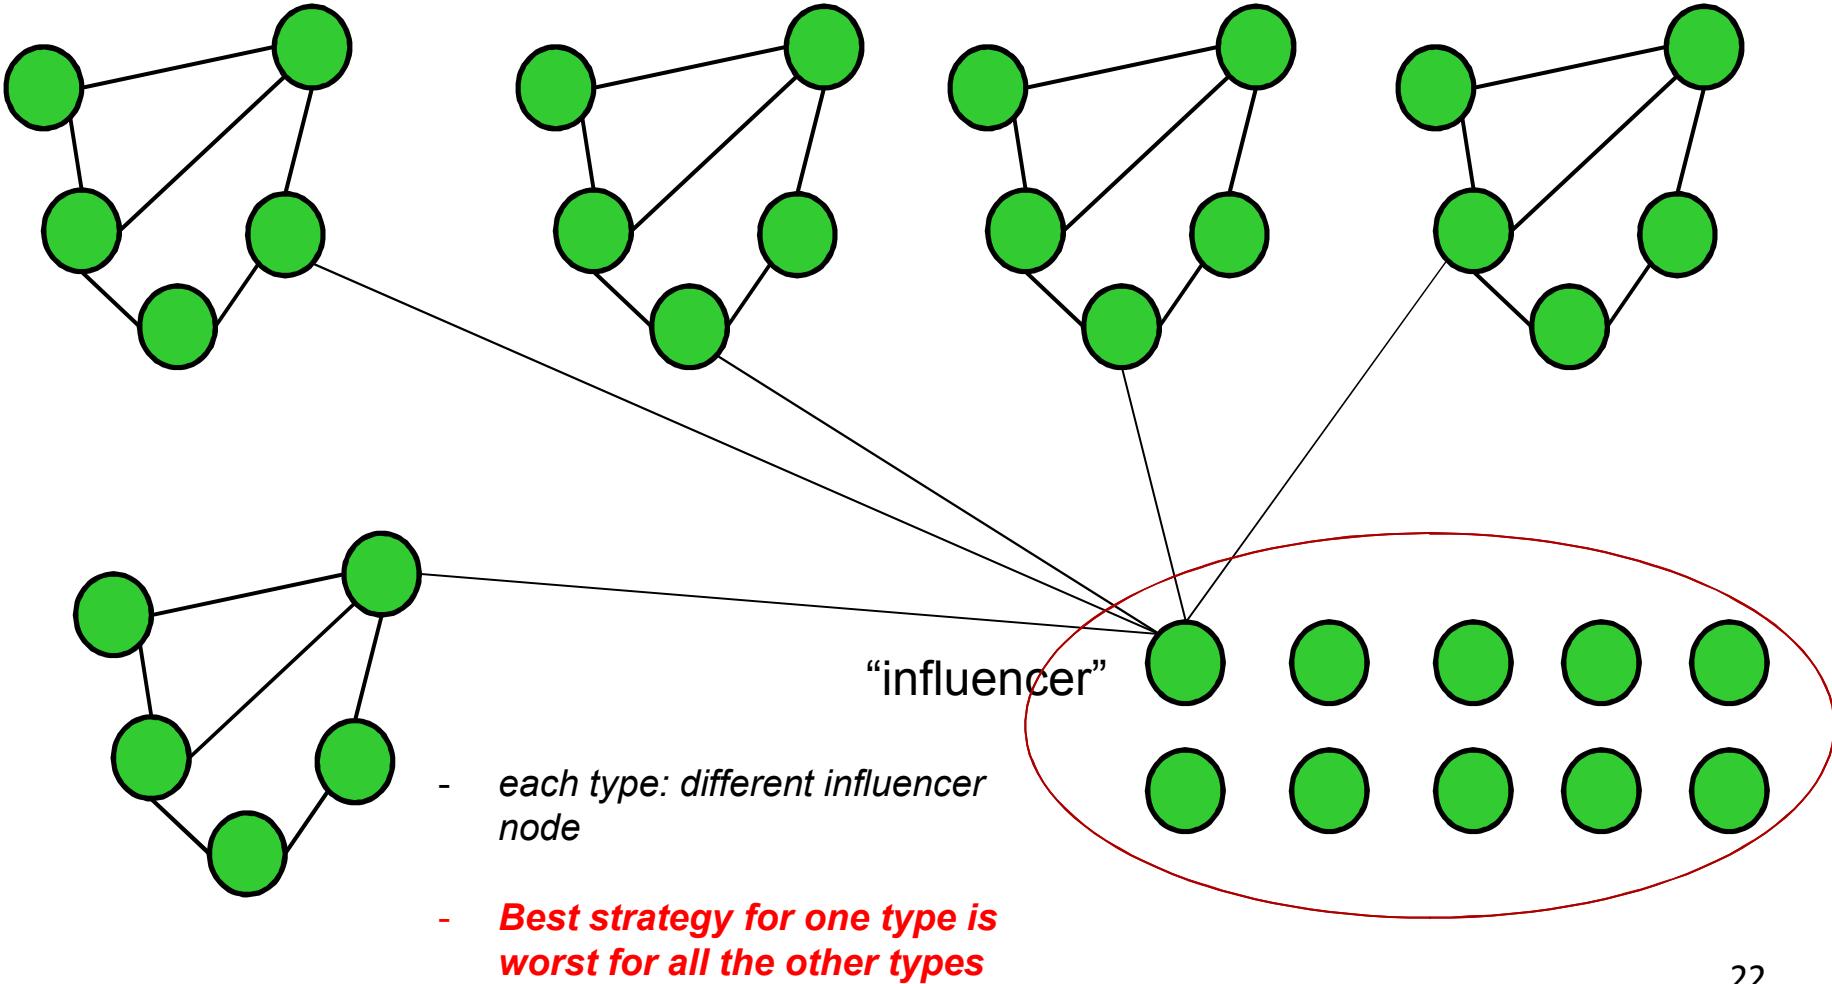
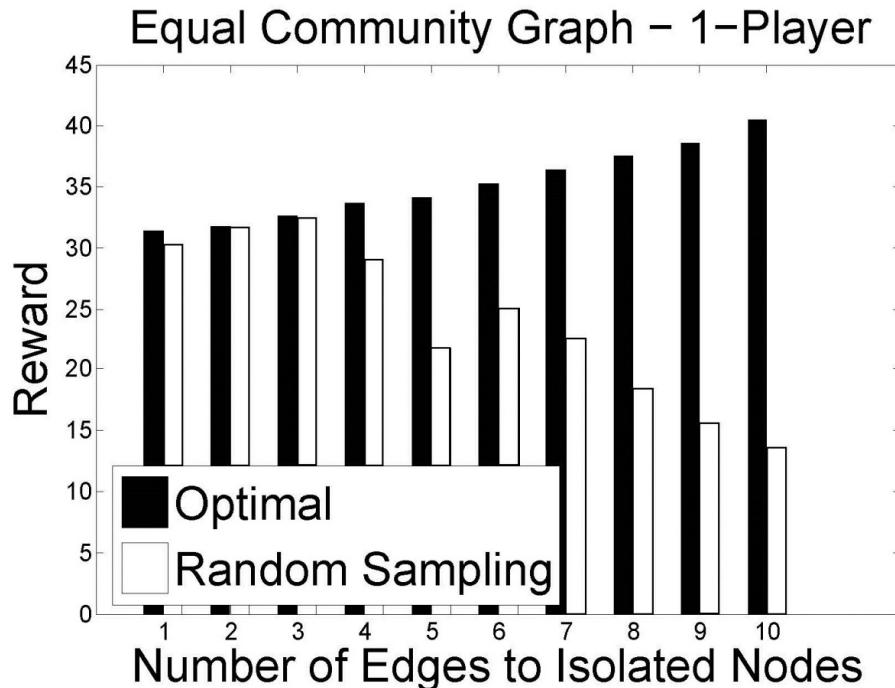


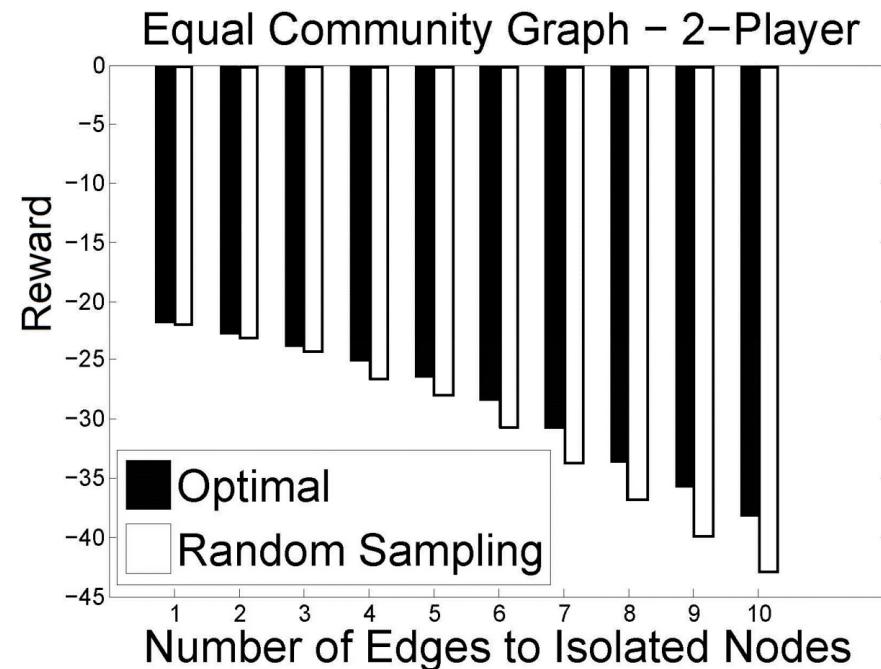
Illustration: “Equal-Community” Graphs



Real Answer: Games are Robust



Sampling a single type becomes arbitrarily worse than optimal



Sampling a single type remains close to optimal (remains robust)

Conclusions

- Uncertainty is pervasive in practice, and when uncertainty is about a graph structure, it presents severe scalability limitations in the worst case
- **While a single-player influence maximization can be very sensitive to uncertainty, influence games appear to be very robust to it**
- **As long as we have decent information about network structure, sampling a small number of types yields near-optimal mitigation strategies**