Market Design & Analysis of a Hidden P2P Backup Market

October 26th, 2011
ECON 1465 - Market Design, Brown University

Sven Seuken
University of Zurich
Department of Informatics

Joint work with:
- David C. Parkes (Harvard)
- Denis Charles (MS)
- Sidd Puri (MS)
- Max Chickering (MS)
- Kamal Jain (MSR)
- Mary Czerwinski (MSR)
- Desney Tan (MSR)

New Technologies Enable new Markets

- Answering/asking questions on the Internet
- Driving a car

- Consuming online content
- Storing files online

Amazon S3
Simple Storage Service

- Consuming electricity in the smart grid

Properties of these Future Markets

- These markets can be complex:
  - Income, expenditure
  - Budgets
  - Prices
  - Trading
  - Complicated constraints

- Users of these new markets:
  - Non-experts (it is costly to become an expert)
  - Millions of users
  - Might not expect markets
  - Might find money unnatural

Designing Simple but Efficient Markets

- To summarize:
  - New, complex markets are emerging
  - + Unsophisticated users

- **Research Question**:
  How can we design *efficient* but *simple* electronic markets?

→ **New Market Design Approach Necessary**

- Hide market complexities
- Make the interaction seamless

Hidden Market Design

The 3 Components of Hidden Market Design

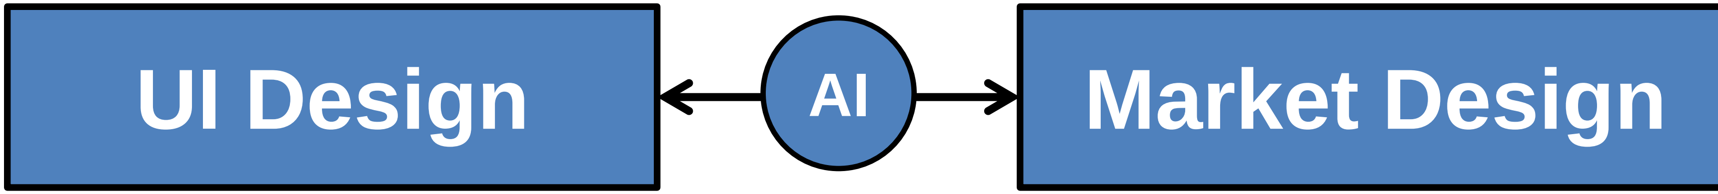

1. UI Design:
   - Constrains market design & defines how users express preferences
   - Defines the amount of cognitive load

2. Market Design:
   - Rules of the market → incentives (need for strategic behavior?)
   - Determines economic efficiency, revenue, and fairness

3. Artificial Intelligence
   - Automate/Simplify decision-making (e.g., learning)
   - Translate between users and market (e.g., proxy agents)

The Hidden Market Design Paradigm

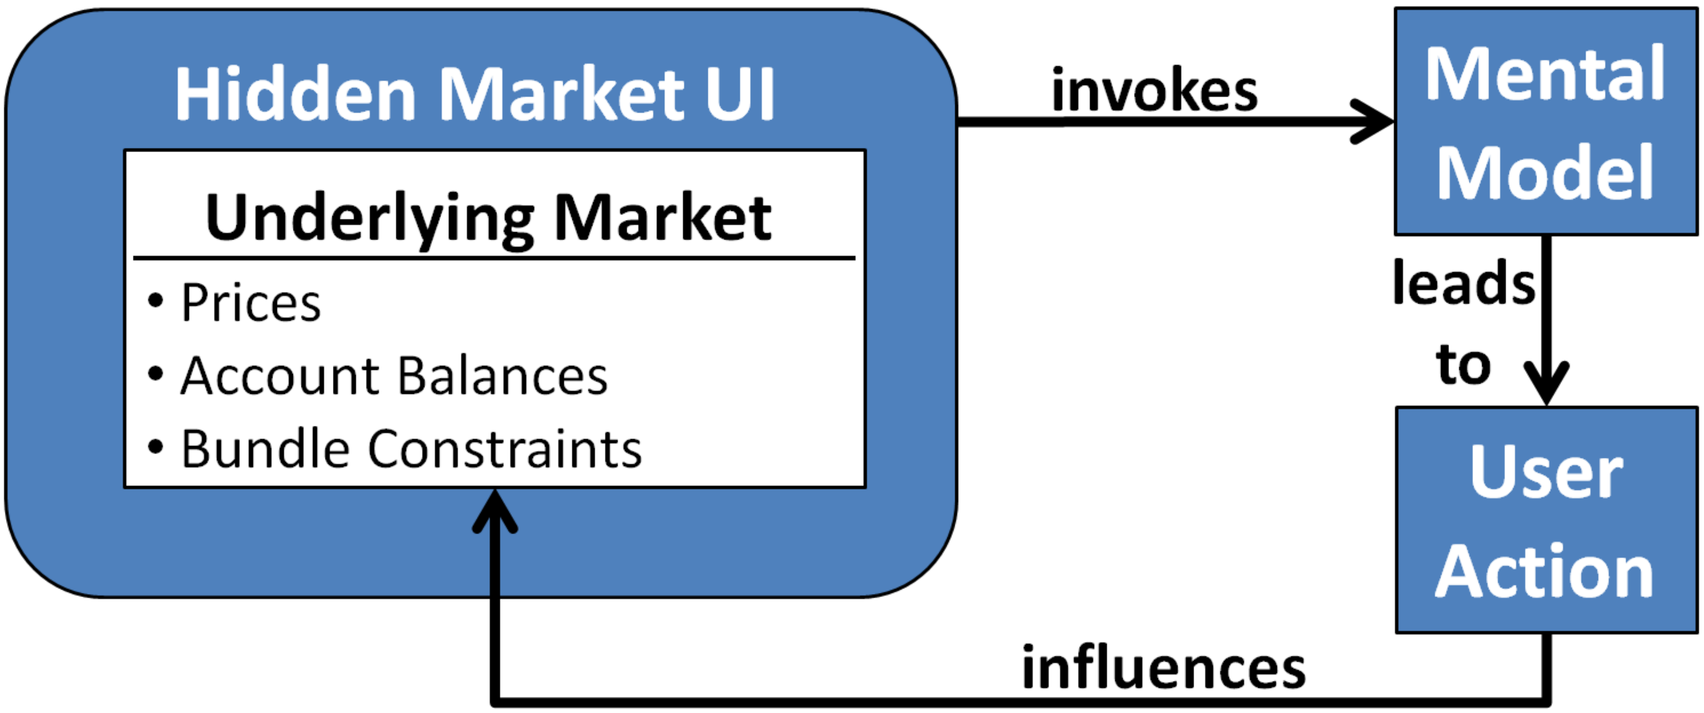

Related Work

- Paul Milgrom. **Simplified Mechanisms with an Application to Sponsored-Search Auctions**. In Games and Economic Behavior 70:1, pp. 62-70, 2010.

- Michael J. Freedman, Christina Aperjis, and Ramesh Johari. **Prices are right: Managing resources and incentives in peer-assisted content distribution.** In *Proceedings of the 7th International Workshop on Peer-to-Peer Systems*, 2008.

- Sharad Goel, David Pennock, Daniel M. Reeves, and Cong Yu. **Yoopick: A Combinatorial Sports Prediction Market**. In Proc. of the 23rd Conference on Artificial Intelligence (AAAI), 2008.

Outline

1. Introduction: Hidden Markets

2. **Case Study: A Hidden P2P Backup Market**
   a) Market Design Preliminaries
   b) The User Interface
   c) Theoretical Analysis
   d) User Study

3. Ongoing and Future Work

4. Conclusion

Peer-to-peer Backup

Data Loss Costs in the US (2003): $18.2 Billion

Traditional Backup (CDs) → too cumbersome

Online Backup → easy, but expensive data centers

P2P Backup → each user must also provides resource: backup space, upload + download bandwidth, online time

→ Using idle resources and decentralization saves costs: 1. storage, 2. energy, 3. maintenance, 4. peak bandwidth

Problem: **Inefficient,** if every user must provide the same!

Max Efficiency: Market-Based P2P Backup

- **Efficiency: use those resource that users value *least*!**
  - User A: large hard disk, slow Internet connection
  - User B: small hard disk, fast Internet connection

Max Efficiency!

No Market	Market
• Users A and B have to provide the same resources	• User A can provide lots of space and little bandwidth • User B can provide lots of bandwidth and little space
• Real-world system deployed: Wuala.com	• Microsoft research project • Internal alpha implemented & tested

- Efficient market:
  - important for competiveness
  - maximizes the total value to users from the P2P system

P2P Backup: System + Market View

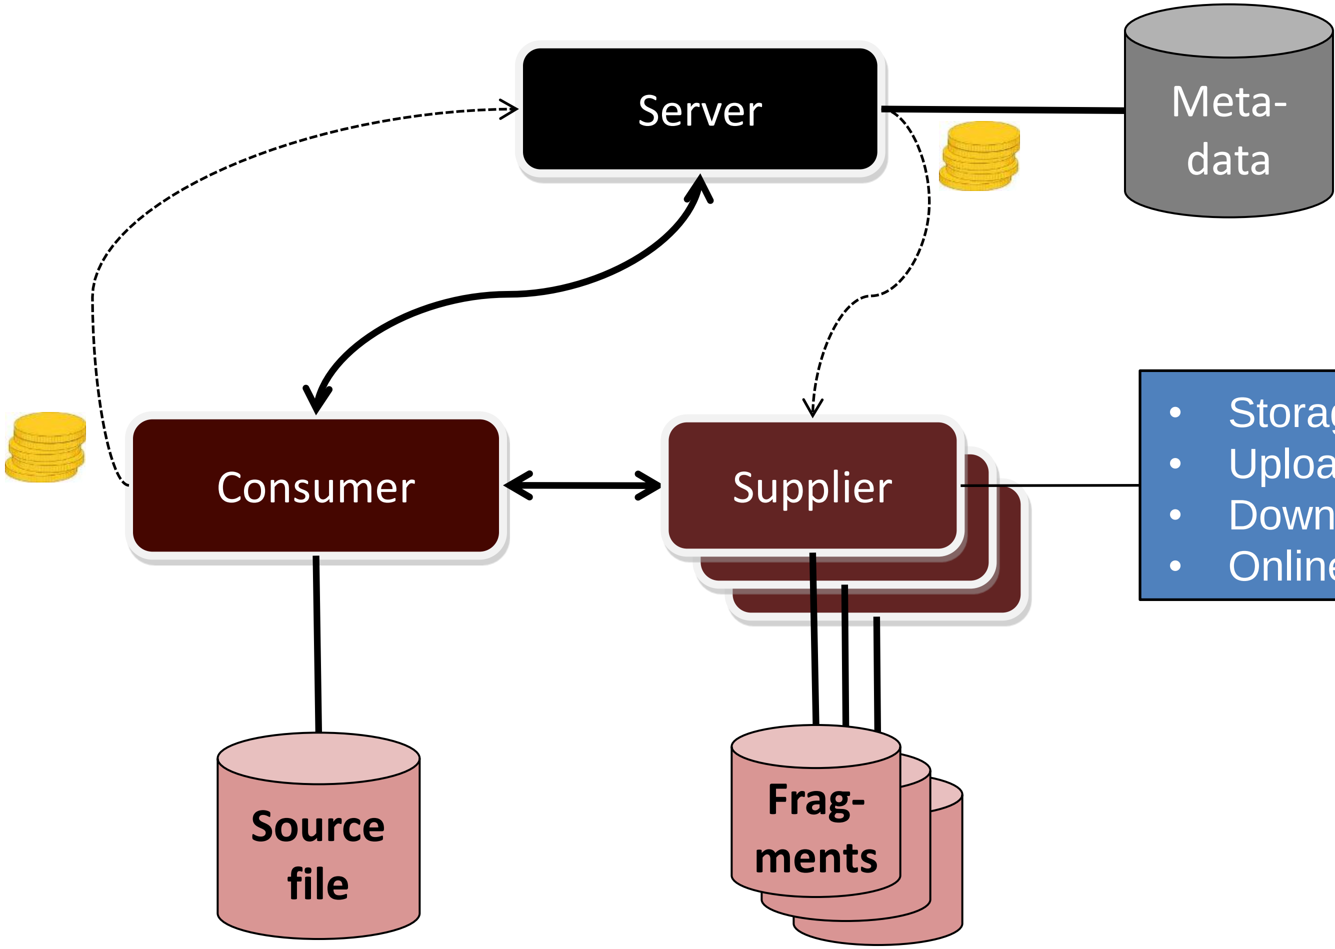

- Storage Space
- Upload Bandwidth
- Download Bandwidth
- Online Time

Reliable Backup in P2P Systems

- Goal: same or higher quality of service compared to centralized data centers → 99.999% instantaneous availability

- Method: Encryption, and Erasure Coding
  - Each file is encrypted..
  - …then encoded…
  - …and then small file pieces are distributed over 100s of peers

- Overhead Comparison (with average online time 12h/day)
  - Simple Replication: 17x
  - Erasure Coding: 3x

System Demo

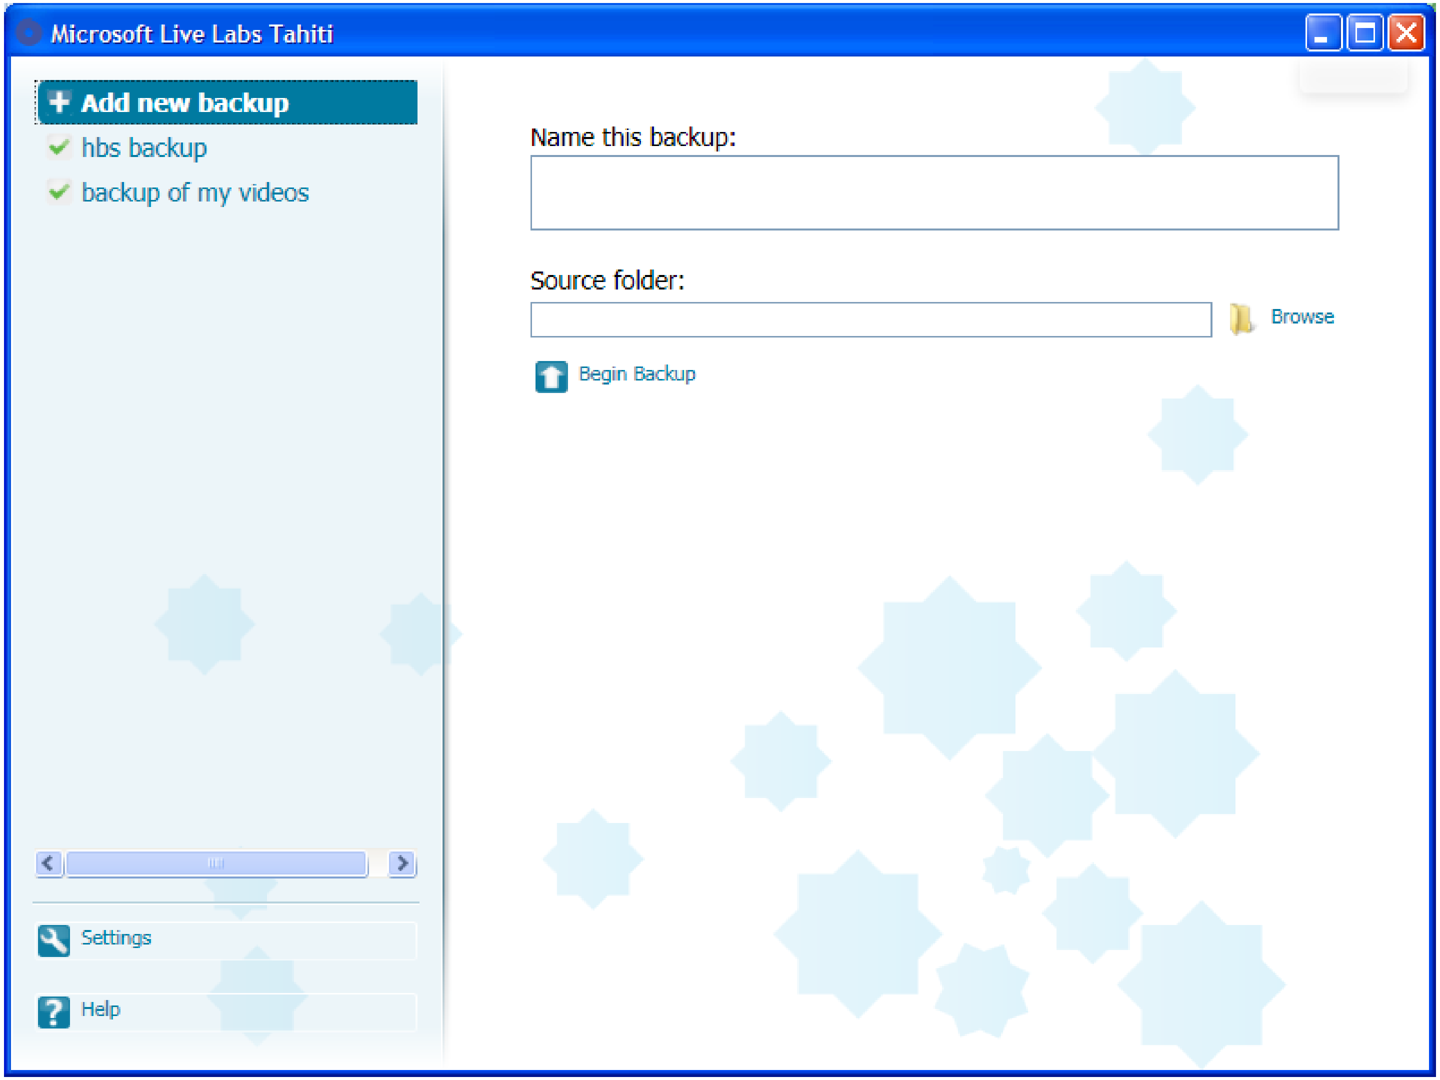

Microsoft Live Labs Tahiti
Add new backup
hbs backup
backup of my videos
Settings
Help
Name this backup:
Source folder:
Browse
Begin Backup

Challenges for the P2P Backup Market

- For UI Design:
  - Money/market unnatural in this domain
  - UI must be simple (millions of users)
  - **Complementarities between resources**

- For Market Design:
  - Elicitation of user preferences (relative value of resources)
  - Users may be unresponsive
  - **Safety: always enough supply to satisfy future demand**

Contributions: Hidden P2P Backup Market

- Designed market + UI in concert
  - The market is *efficient*
  - The UI is *simple* (market complexities are hidden)

- Safety of the System: Theoretical Analysis
  - In equilibrium, the system's safety is maximal
  - How and when can we achieve the equilibrium?
  - How robust is the equilibrium to irrational behavior?

- Usability of the System: User Study
  - Simple to use?
  - Achieves efficient allocations?

Macroeconomic design

- Account Balances:
  - Each user starts with 0 balance
  - Each user can take on a bounded deficit

- Design Implication:
  - Server assigns work to drive users' account balance to 0
  - Users do not have to monitor account balances
  - Expenditure for consumption = Income from supplying resources

Microeconomic Design

- Server is only producer → model as exchange economy
- Each resource has a price, updated once per day

- Allow each user to supply different ratio of resources
  - “cold files”
  - “hot files”

- **Value Optimization:**
  users must react to prices, i.e., adjust their supply and demand, obeying the complementarity constraints

Sounds complicated? It is!!

→ **Goal: Hide these complications from the users!**

Outline

1. Introduction: Hidden Markets

2. Case Study: A Hidden P2P Backup Market
   a) Market Design Preliminaries
   b) The User Interface
   c) Theoretical Analysis
   d) User Study

3. Ongoing and Future Work

4. Conclusion

Market User Interface Demo

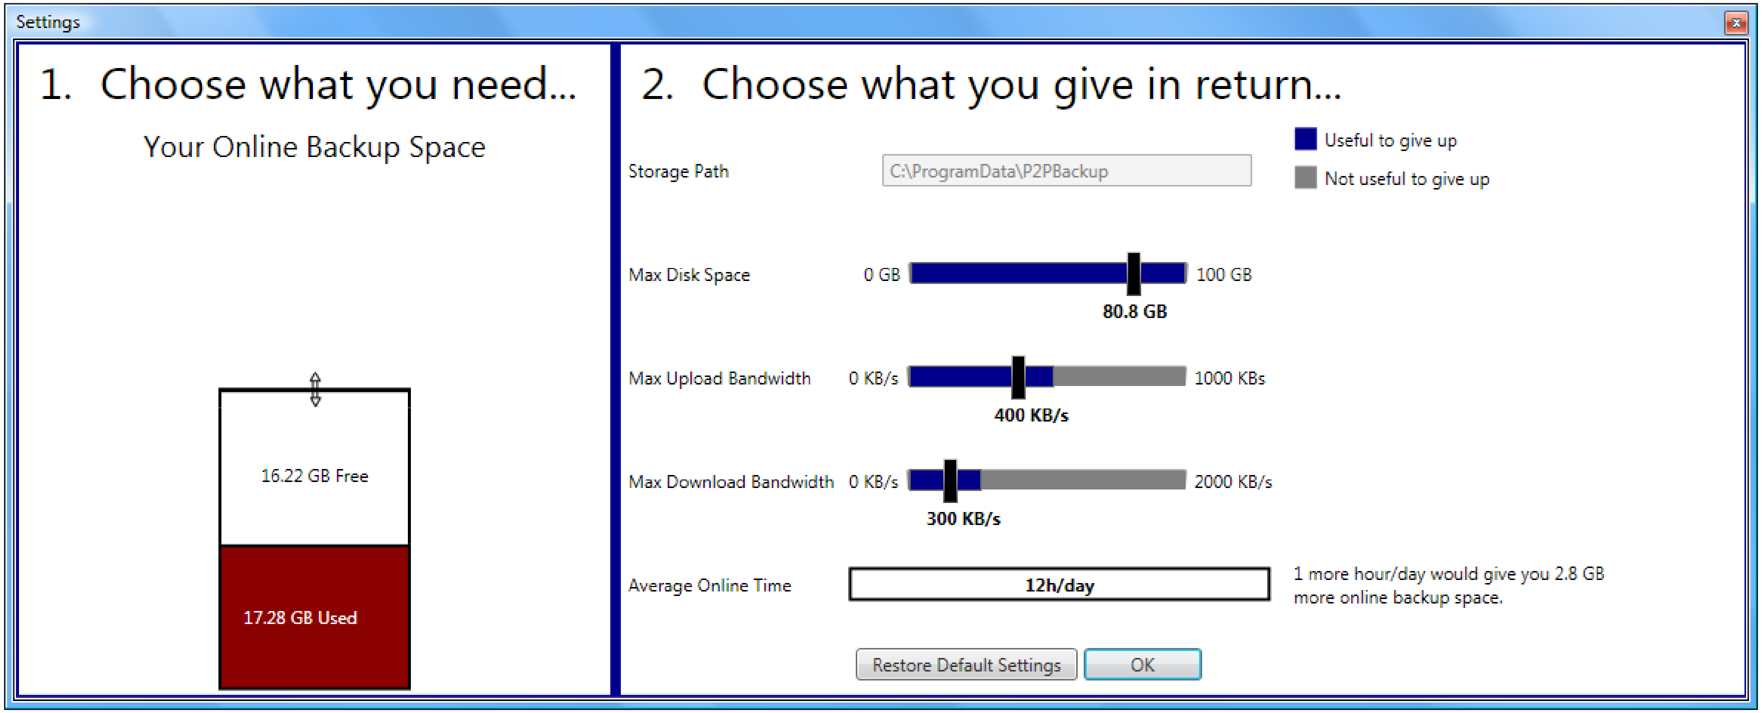

Outline

1. Introduction: Hidden Markets

2. Case Study: A Hidden P2P Backup Market
   a) Market Design Preliminaries
   b) The User Interface
   c) **Theoretical Analysis**
   d) User Study

3. Ongoing and Future Work

4. Conclusion

System Safety

- **Remember**: efficiency requires/means that users can supply *different ratios* of their resources

- Resulting problem/concern: **Safety buffer** (we always need enough supply to satisfy new demand)

→ Update prices p, such that users change their supply settings, and the *buffer* between supply and demand increases

- **Definition (Buffer Size):** *The size of the buffer between maximum supply and current demand for resource k is:* $B_k(p) = \frac{Supply_k(p)}{demand_k(p)}$*; the overall buffer is* $B(p) = \min_{k \in \{S,U,D\}} B_k(p)$.

Unsafe System = Low Supply Buffer

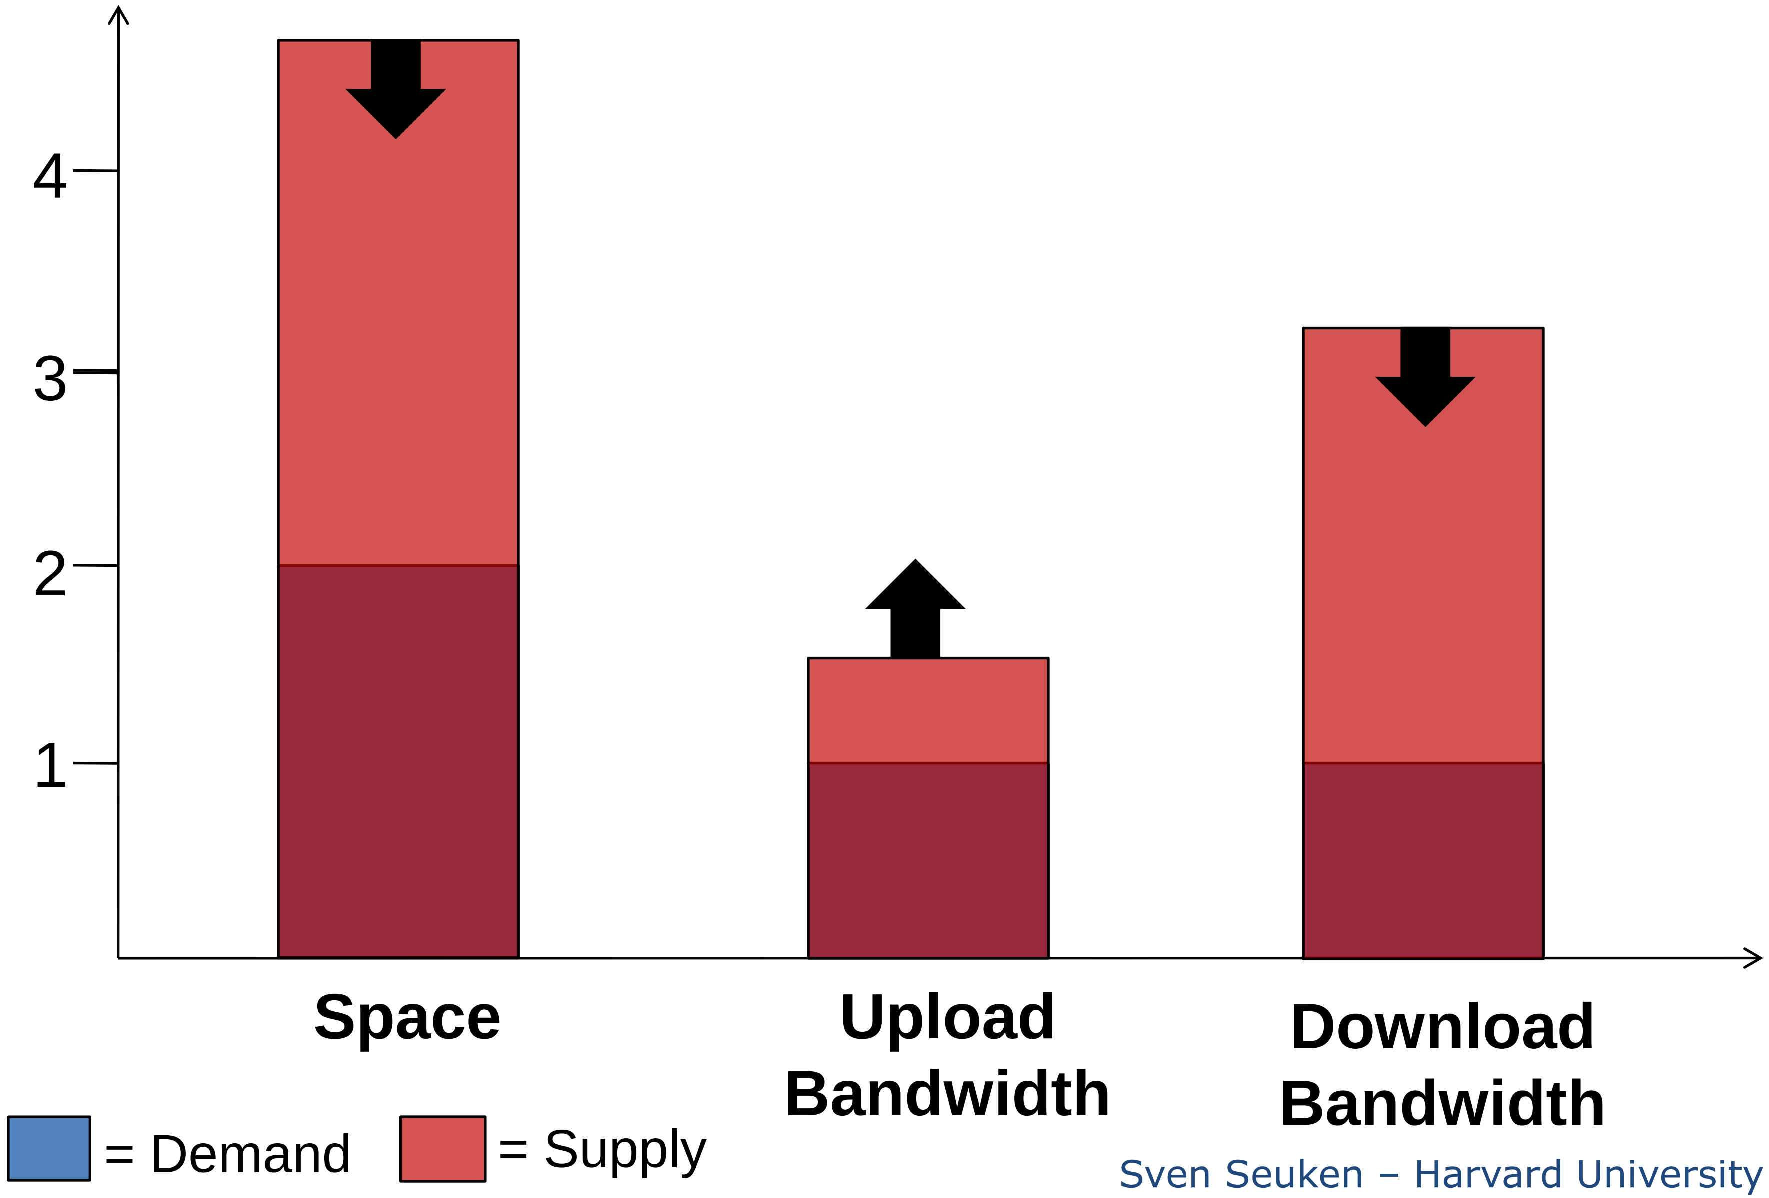

Maximally safe = Maximum Buffer

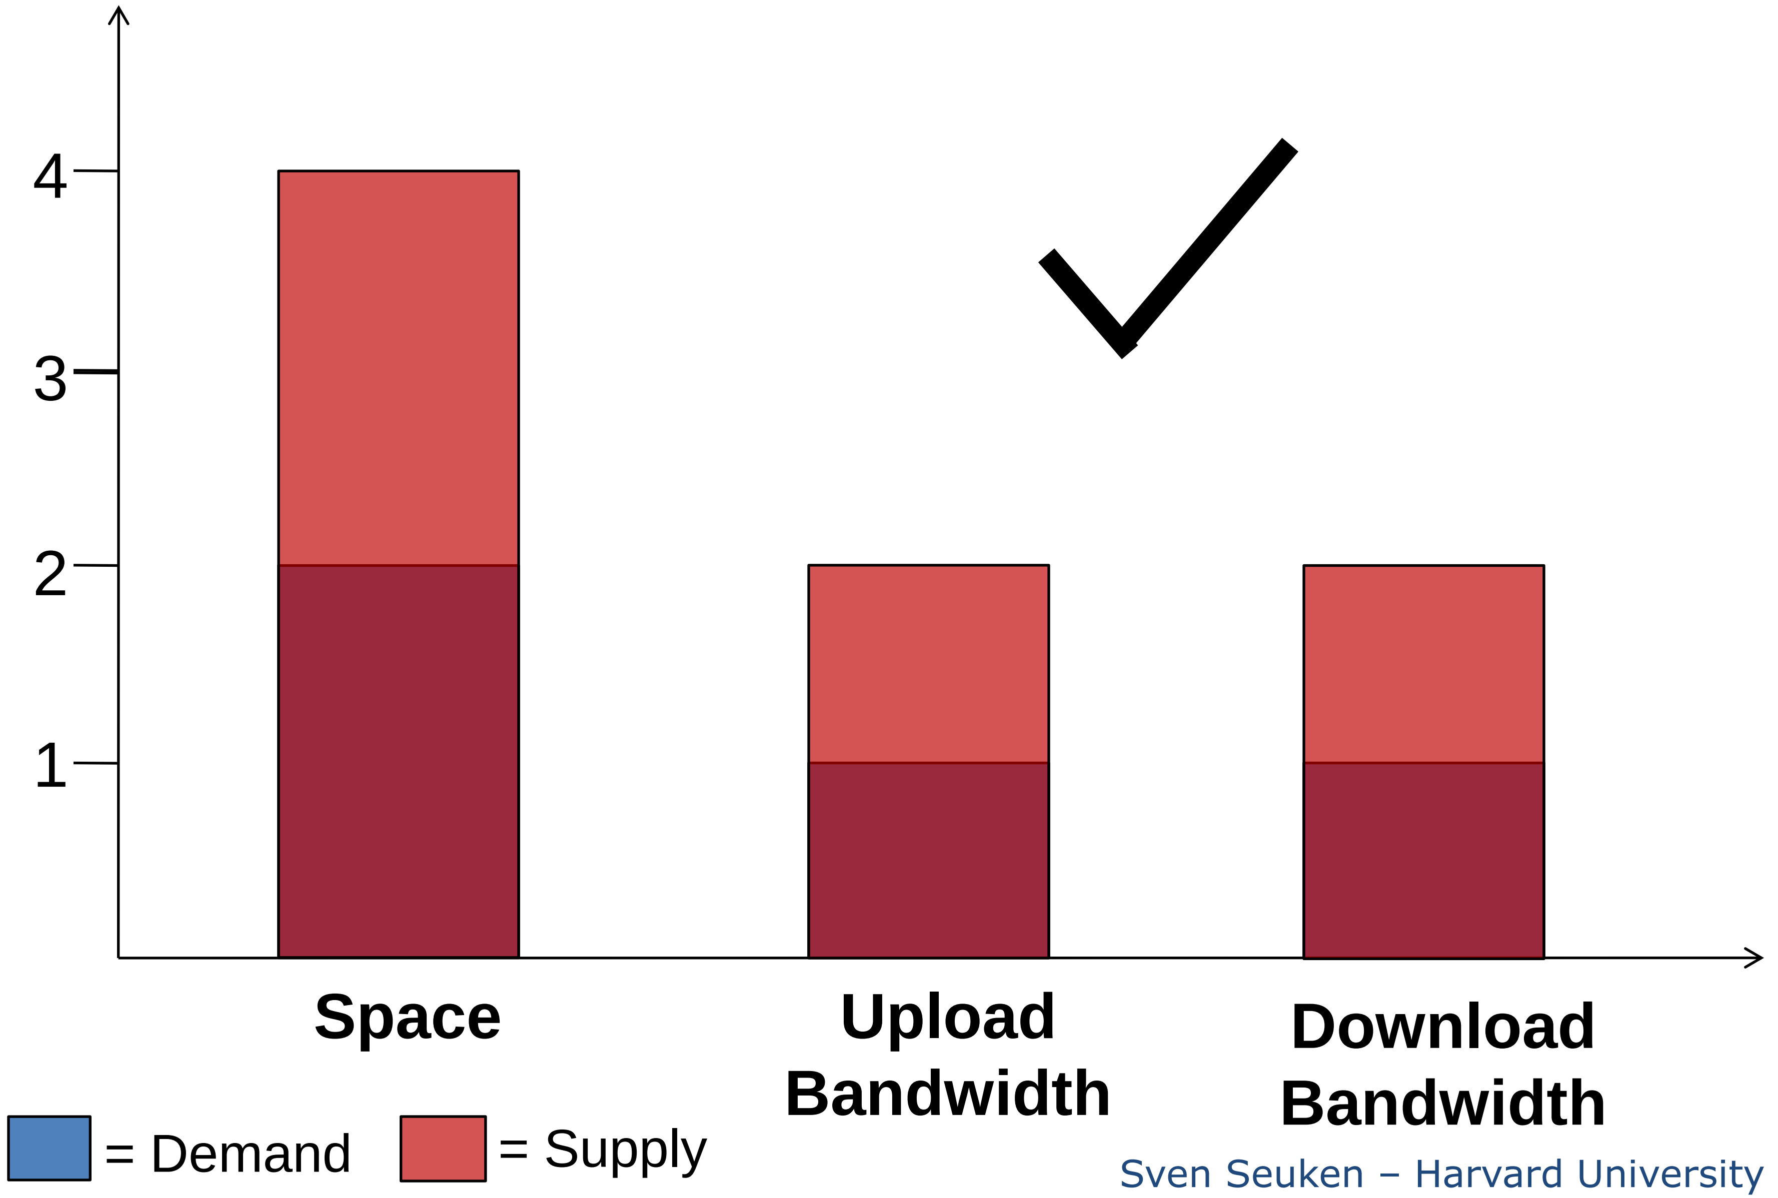

Maximum Safety: The Buffer Equilibrium

- **Definition (Buffer Equilibrium):** A buffer equilibrium is a price vector p, a maximum supply vector Supply(p) and a current demand vector *demand(p)*, such that:

$$B_S(p) = B_U(p) = B_D(p).$$

- **Theorem 1 (Buffer Maximality):** If the buffer equilibrium is unique, and given that buffers are Gross substitutes, then the size of the buffer is maximal in the buffer equilibrium (over all possible price vectors).

Price Update Algorithm

- The server adjusts prices towards buffer equilibrium

- Multi-dimensional quasi-Newton method to iteratively approximate equilibrium prices

$$p_X^{t+1} = p_X^t - \frac{p_X^t - p_X^{t-1}}{f_X(p^t) - f_X(p^{t-1})} f_X(p_t)$$

- Algorithm converges
  - if all users update their settings in every iteration
  - under some technical assumptions
  - and if initial prices are not too far from the equilibrium

Equilibrium Existence

- **Definition (Admissible User Preferences): We call preferences "admissible" if they are**
  - *continuous*
  - *strictly convex*
  - *monotone w.r.t. service products*
  - *strongly monotone w.r.t. supply resources*

- **Theorem 2 (Buffer Equilibrium Existence):** *Given that all users have admissible preferences, a buffer equilibrium is guaranteed to exist.*

Proof Sketch

- **Definition (Relative Buffer Function):** Let $Z_k(p)$ denote the relative size of buffer on resource k compared to the average on all resources.

- **Lemma**: Given users with admissible preferences, if we decrease the price of resource *k* to 0 and keep the other prices constant, then

$$\exists l\colon Z_l(p) > Z_k(p).$$

- Equipped with this Lemma:

1. Let Δ denote the space of non-zero price vectors
2. Construct a correspondence *f* from Δ to Δ
3. Prove that *f* is convex-valued
4. Prove that *f* is upper hemicontinuous
5. Apply Kakutani‘s fixed point theorem: a fixed point of *f* in Δ exists
6. Prove that each fixed point constitutes a buffer equilibrium

Robustness against Irrational Users

- Our previous results assumed rational users, in particular:
  - Strictly convex preferences
  - (Strongly) Monotone preferences

- In reality, some users will be "irrational":
  - They might update their settings with a delay
  - They might be price-insensitive
  - They might be adversarial

→ Analyze robustness of the equilibrium to **adversarial** users!

Equilibrium Existence with Adversarial Users

- **Definition (Slack Parameter):** Let γ denote the system's slack parameter**,** that is, how far away from the system-wide averages an individual user's supply ratios are allowed to be:

$$\forall k, \forall l \in \{S, U, D\}: \frac{Supply_k}{Supply_l} \leq \gamma \cdot \frac{System\text{−}wide\text{−}Demand_k}{System\text{−}wide\text{−}Demand_l}, \gamma \geq 1$$

- **Theorem 3 (Equilibrium Existence):**
If $Demand(rational) \geq (\gamma^2 + \gamma) \cdot Demand(adversarial)$*, and if the rational users' preferences are admissible, then a buffer equilibrium is guaranteed to exist.*

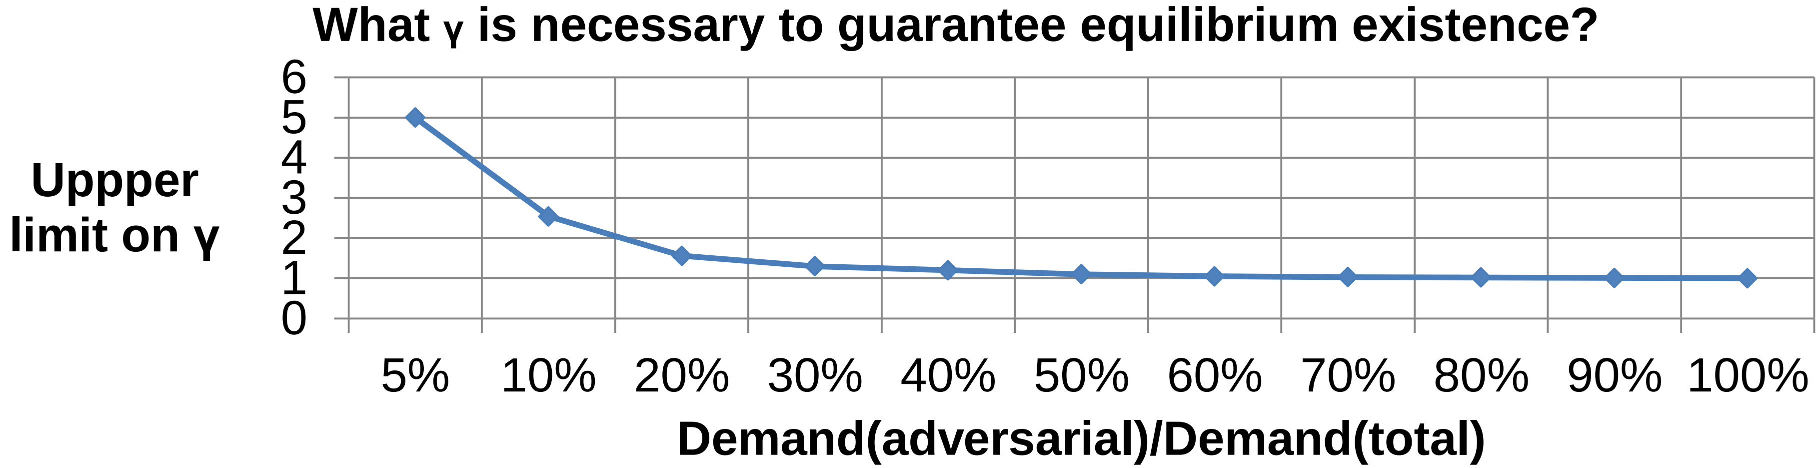

Equilibrium Uniqueness

- **Theorem 4 (Equilibrium Uniqueness):**
  Given that a buffer equilibrium exists, and
  - *the aggregate demand function satisfies the perfect complements property and*
  - *the aggregate supply function satisfies the relative gross substitutes property*

  then the buffer equilibrium is unique.

(Un-)Controllability of Buffer Size

- **Corollary:** *With price changes alone, we cannot influence the size of the buffer in equilibrium.*

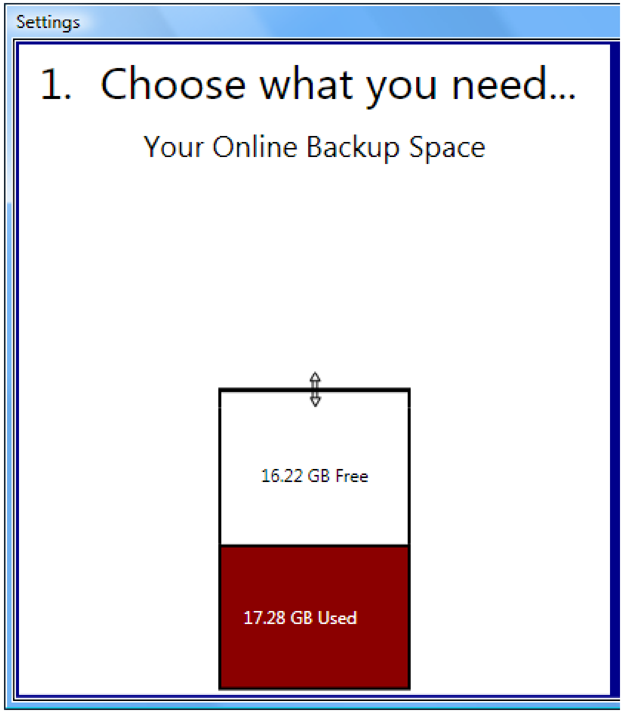

- **Definition (Minimum Free Space Requirement):** *Let δ denote a UI requirement regarding how much minimum free space each user must have.*

- **Theorem 5:** *For any desired buffer size b, there exists a δ, that induces a buffer equilibrium with buffer size b.*

Outline

1. Introduction: Hidden Markets

2. Case Study: A Hidden P2P Backup Market
   a) Market Design Preliminaries
   b) The User Interface
   c) Theoretical Analysis
   d) **User Study**

3. Ongoing and Future Work

4. Conclusion

Formative Usability Study

- Main research question:

  Is the UI a usable instantiation of the hidden markets paradigm?

- Users:
  - 16 Users (8 female, 8 male)
  - No UI design experts
  - 50/50 mix of tech experts and novices

- 11 tasks/scenarios with increasing complexity

Sample Tasks for User Study

- **Task 2:**
  "Please change the settings such that you have approximately 20GB of free online backup space available, but you are not allowed to give up more than 70GB of your own hard disk space."

- **Task 6:**
  *"Imagine you are a user who does a lot of photo and video editing on his computer and thus needs a lot of free hard disk space to do this kind of work. Assume that **you don't care about your upload and download bandwidth**. Please change your settings so that you have **20GB of free online backup space** available while **minimizing the amount of hard disk space you give up**."*

Did we Invoke the Right Mental Model?

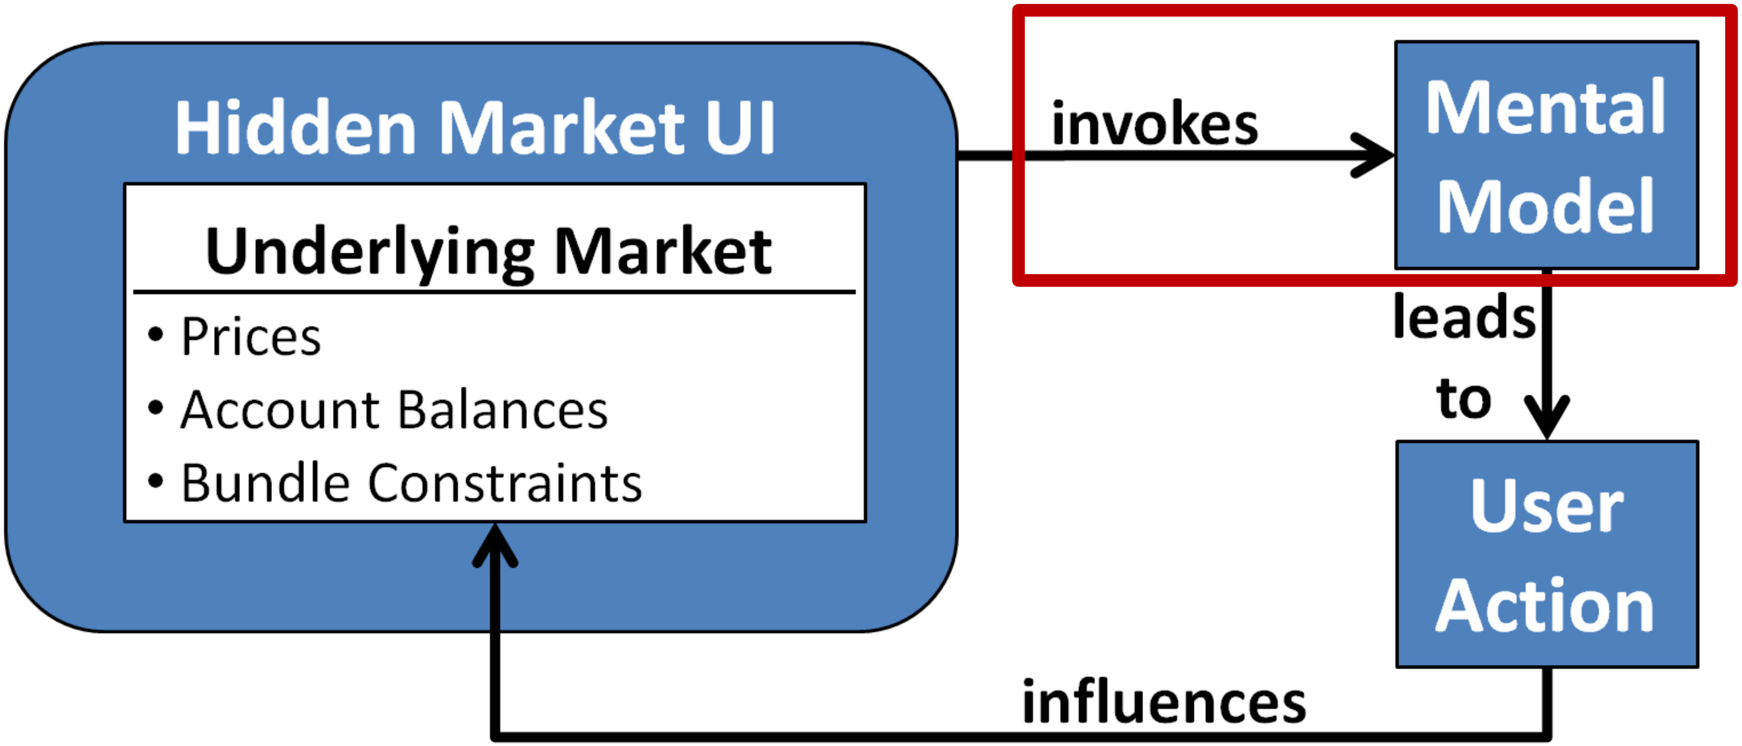

Summary of User Study Results

Category	Mental Model	Success Rate
1	Give & Take	16/16
2	Bundled Resources	9/16
3	Prices/Trade-offs	7/16

Summary: UI & Market Design Features

UI Design

1. Low frequency of interaction
2. Bundling of resources is hidden
3. Hidden account balances
4. Hidden payments
5. Prices are semi-hidden
6. "Was there a market at all"?

↔

Market Design

1. Buffer Equilibrium
2. Equilibrium robustness to adversaries depends on slack parameter
3. Simple pricing within slack regions
4. Price updates work with point-wise user feedback

Outline

1. Introduction: Hidden Markets

2. Case Study: A Hidden P2P Backup Market
   a) Market Design Preliminaries
   b) The User Interface
   c) Theoretical Analysis
   d) User Study

3. **Ongoing and Future Work**

4. Conclusion

Outline

1. Introduction: Hidden Markets

2. Case Study: A Hidden P2P Backup Market
   a) Market Design Preliminaries
   b) The User Interface
   c) Theoretical Analysis
   d) User Study

3. Ongoing and Future Work

4. **Conclusion**

Conclusion

- Hidden P2P Backup Market:
  - Paradigm: hide market complexities
  - Buffer equilibrium analysis
  - Robustness depends on slack parameter (amount of freedom)
- User Study Results:
  - Promise for hidden markets
  - Eluded to interesting questions at intersection of UI & market design
- Future Work:
  - New hidden market applications: **smart grids**, display advertising,…
  - A principled study of market user interface design
- Thank you!