Carnegie Mellon University

Measuring Perceived Privacy Risk in Cybersecurity Information Sharing

Travis Breaux with Jaspreet Bhatia, Liora Friedberg (U. Penn),
Hanan Hibshi, and Daniel Smullen

June 27, 2016

institute for SOFTWARE RESEARCH

Cybersecurity Information Sharing

- During 1995-2005,[1] IT was 30% of US GDP and contributed 50% of economic growth: IT = Critical Infrastructure
- Last year, $1B in losses [FBI], over 431 million malware variants [Symantec]

What are we doing?

- 1998: WH introduces PDD-63 to establish ISACs
- 2015: WH introduces EO13691 to coordinate NCCIC, ISACs
- 2016: Congress passes Cybersecurity Sharing Act

How can we share incident indicators with others, including government, while preserving privacy?

Jorgenson, Ho, Stiroh. *Productivity, Volume 3: Information Technology and the American Growth Resurgence,* vol. 3 Postwar U.S. Economic Growth. MIT Press, 2005.

Presentation Overview

Problem:

- How can we share incident indicators with others, including government, while preserving privacy?

Research Question:

- What is the trade-off between data use (to develop indicators) and data privacy?

Talking Points:

- What is risk, and how do we measure it in privacy?
- What kind of information is used, and for which purposes?
- What is the trade-off between data use and privacy risk?

What is Risk?

- Bauer (1960) defined risk as comprised of two dimensions: uncertainty and adverse consequences
  - Assumptions differ about how much decision makers know about each dimension
- NIST (2015) defines risk "as a function of the likelihood that an adverse outcome multiplied by the magnitude of the adverse outcome should it occur"
  - *Risk = Likelihood* x *Impact (e.g., CVSS)*
- Empirical risk research in judgement and decision sciences (1970-) distinguishes *revealed preferences*, and *expressed preferences,* or so-called *perceived risk*
- Knightian uncertainty distinguishes risk, uncertainty and ambiguity aversion, which is a preference for known outcomes (1921)

R. A. Bauer "Consumer Behavior as Risk Taking," *American Marketing Association*, R. S. Hancock (ed.), Chicago, IL: American Marketing Association, 389-398, 1960.

S. Brooks & E. Nadeau. "Privacy Risk Management for Federal Information Systems," NISTIR 8062, May 2015.

Revealed Preferences (Starr, 1972)

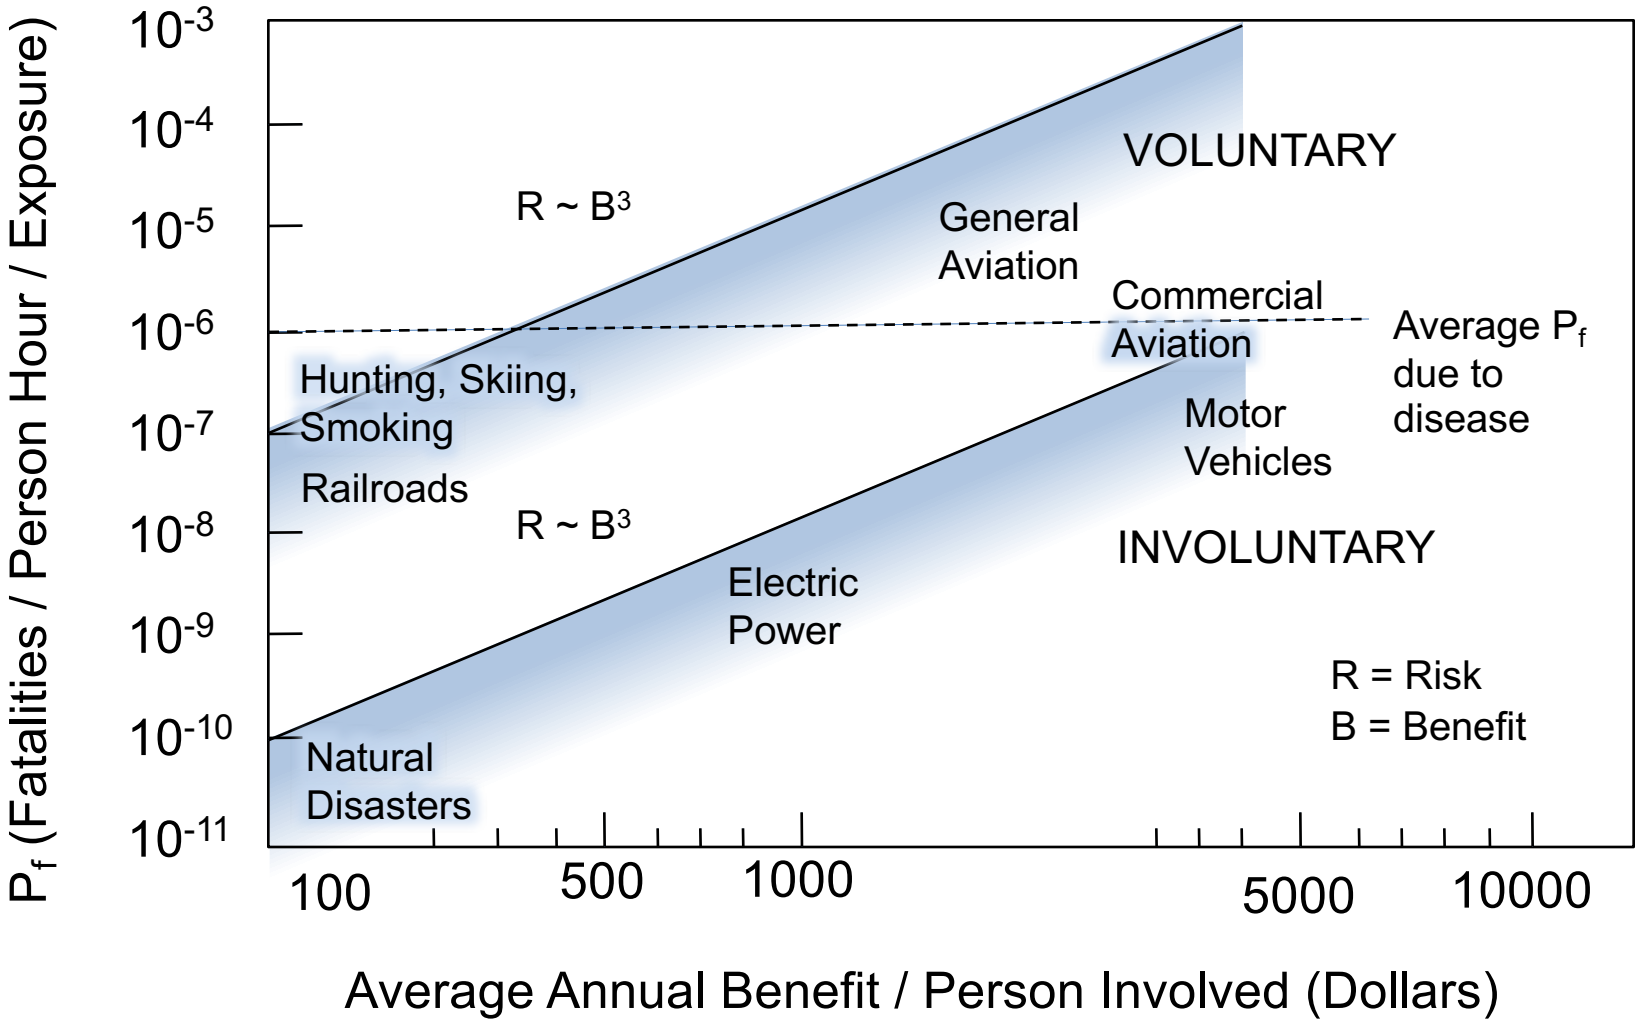

Critique of Starr's Revealed Preferences

- Assumes that past behavior is a valid indicator of present preferences (Fischhoff et al. 1978)
  - As technology changes or new technologies emerge, preferences are more likely subject to bias and heuristics than to indicators of past behavior
- Revealed preferences assume that individuals have complete information, and can use that information
  - Unless design space is known by the public, and vendor provides best design choices, the public cannot choose in the market
- Starr's results are sensitive to the way that risk measures are computed from historical data (Otway and Cohen, 1975)

Expressed Preferences (Fischhoff, 1978)

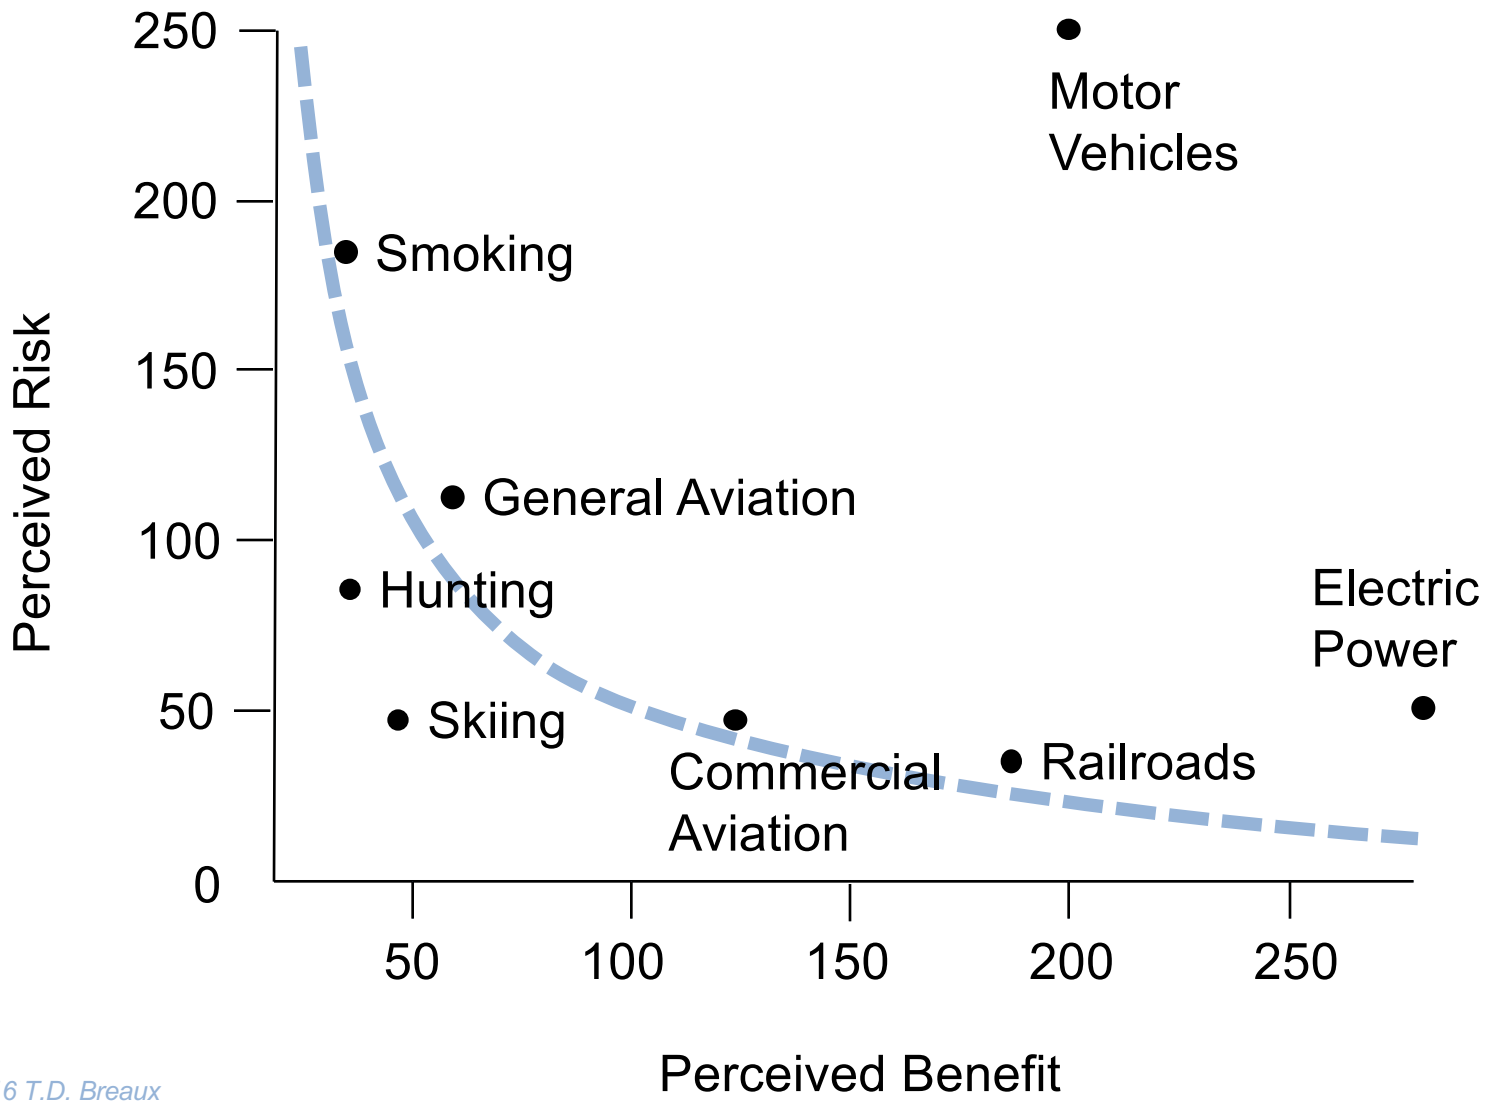

Revealed vs. Expressed Preferences

- Expressed preferences account for benefits not measured in dollars, such as “greater flexibility in patterns of living” that are more likely the benefits received from increased privacy
- People view risk levels as more acceptable after they have ordered benefits in depth: as benefits increase, perceived risk decreases
- Affect Heuristic (Alhakami and Slovic, 1994)
  - If an activity is liked, then people judge the benefits as high, and risks as low; inverse relationship, if an activity is disliked

Survey design should explicitly order technological benefits, and the dependent variable should be the acceptance of risk

Fischhoff et al. “How safe is safe enough? A psychometric study of attitudes toward technological risks and benefits.” *The Perception of Risk*, Slovic (ed.), 2000.

Why and how to measure likelihood?

- When estimating risks, people struggle to map probabilities to portions of the population, but they do better with ratios of people harmed (Fischhoff, 1978)
  - 0.01% of the U.S. population (32,000 people)
  - 1 / 10,000 people in the U.S.
- Our pilot study shows no significant effects due to ratios of people harmed for 1/4, 1/10, 1/100* and 1/1000 people

An alternative to proportions and ratios…

- Construal level theory (CLT) proposes that we form abstract mental construals of distal objects
  - Wakslak and Trope (2008) showed that unlikely events are more distal relative to time, space and social distance
  - Example: a rare cat blood type was more expected to occur in cats in spatially remote areas, than a common cat blood type that was more expected to occur in a near location
- Adopted spatial and social distances to estimate likelihood ($RL):
  - only one person in your family
  - only one person in your workplace
  - only one person in your city
  - only one person in your state
  - only one person in your country

Early results on social distance and risk

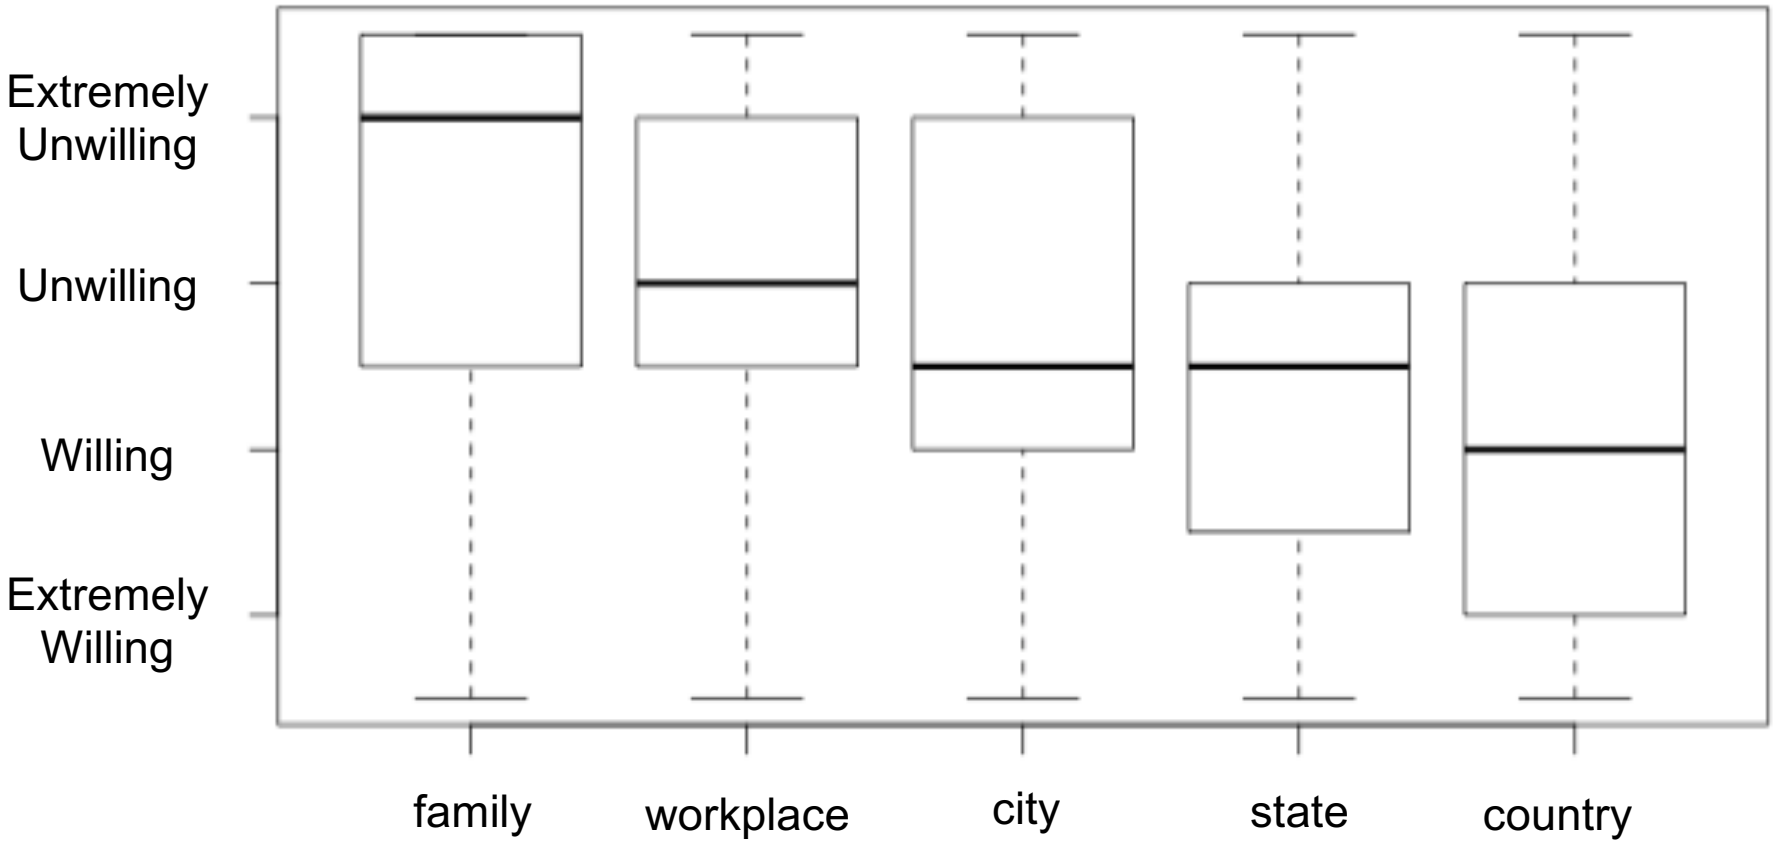

Box-chart represents $RL as a within-subjects factor.

Factorial Vignette Design

You were informed by your workplace IT department that your **workplace computer** was compromised by a cyber attack that allowed an attacker to gain unauthorized access to your employer's internal network.

Please rate your willingness to share your information below with the Federal government for the purpose of **investigating terrorism**, given the following risk.

Risk: In the last 6 months, while using this website, **only one person in your state** experienced a privacy violation due to government surveillance.

When choosing your rating for the information types below, consider the workplace computer, purpose and the risk, above.

	Extremely Willing	Very Willing	Willing	Somewhat Willing	Somewhat Unwilling	…
$DT	○	○	○	○	○	

Vignette Factor Levels

Factors	Factor Levels
Computer Type ($CT)	Personal smart phone
	Workplace computer
Data Purpose ($DP)	Investigating intellectual property and trade secrets
	Investigating economic harm, fraud and identity theft
	Investigating imminent threat of death
	Investigating terrorism
Privacy Harm ($PH)	A privacy violation due to government surveillance
Data Types ($DT)	Group 1: Usernames & passwords, device information device ID, UDID / IMEI, sensor data, network information, IP address & domain names, packet data, MAC address
	Group 2: OS information, OS type & version, memory data, temporary files, registry information, running processes, application information, application session data
	Group 3: emails, chat history, browser history, websites visited, contact information, keyword searches, keylogging data, video & image files

Identifying Relevant Incident Data

- **Who did we survey?**
  - 76 respondents with mean 8 years experience in incident analysis
  - Job titles range from security analyst, security architect to director
- **Practices**
  - 36% use STIX for indicator sharing
  - 73% conduct network monitoring and forensic investigations
  - 69% prepare incident reports
- **Information Types**
  - >66% collect network and OS information, and usernames
  - >50% collect application data, temporary files and device IDs
  - 25-50% collect browser history, keyword searches, web sites visited, e-mails, contact information

Based on qualitative responses, there is confusion about what constitutes personally identifiable information (PII), and whether it is collected

Rank-order Benefit to Society of Data Purpose

- Participants were first asked to rank-order the benefits, before they were asked to estimate the multiplicative distances between benefits

Rank	Data Purpose	Consensus	Value
1	Investigating imminent threat of death	68.8%	68.2
2	Investigating terrorism	60.0%	39.9
3	Investigating economic harm, fraud and identity theft	68.8%	21.4
4	Investigating intellectual property and trade secrets	63.8%	10.0

Data purposes summarized from the 2016 Cybersecurity Information Sharing Act

Multilevel Privacy Risk Model

Interaction Term	Coefficient	Std. Error
Intercept (family + workplace PC + intellectual)	4.200***	0.162
Risk Level – 1 person in your workplace	-0.029	0.201
Risk Level – 1 person in your city	-0.083	0.201
Risk Level – 1 person in your state	-0.024	0.201
Risk Level – 1 person in your country	-0.126	0.201
Data Purpose – Economic Harm	0.132*	0.072
Data Purpose – Terrorism	0.310***	0.072
Data Purpose – Imminent Death	0.452***	0.072
Computer Type – Personal Smart Phone	0.015	0.127

*p ≤ 0.05 **p ≤ 0.01 ***p≤ 0.001

Data Usage Estimation

Information Type	Estimated Usage	Estimated Risk
IP addresses / domain names	0.741	3.364
OS type and version	0.673	6.811
Device identifiers	0.543	3.474
Web sites visited	0.545	5.080
MAC address	0.519	3.430
E-mails	0.524	4.549
Packet data	0.505	4.021
Contact information	0.442	5.083
Keyword searches	0.319	5.130
Passwords	0.244	5.380
Chat history	0.203	4.586
Keylogging data	0.144	4.439

Risk: 1=extremely unwilling, 8=extremely willing

Carnegie Mellon University

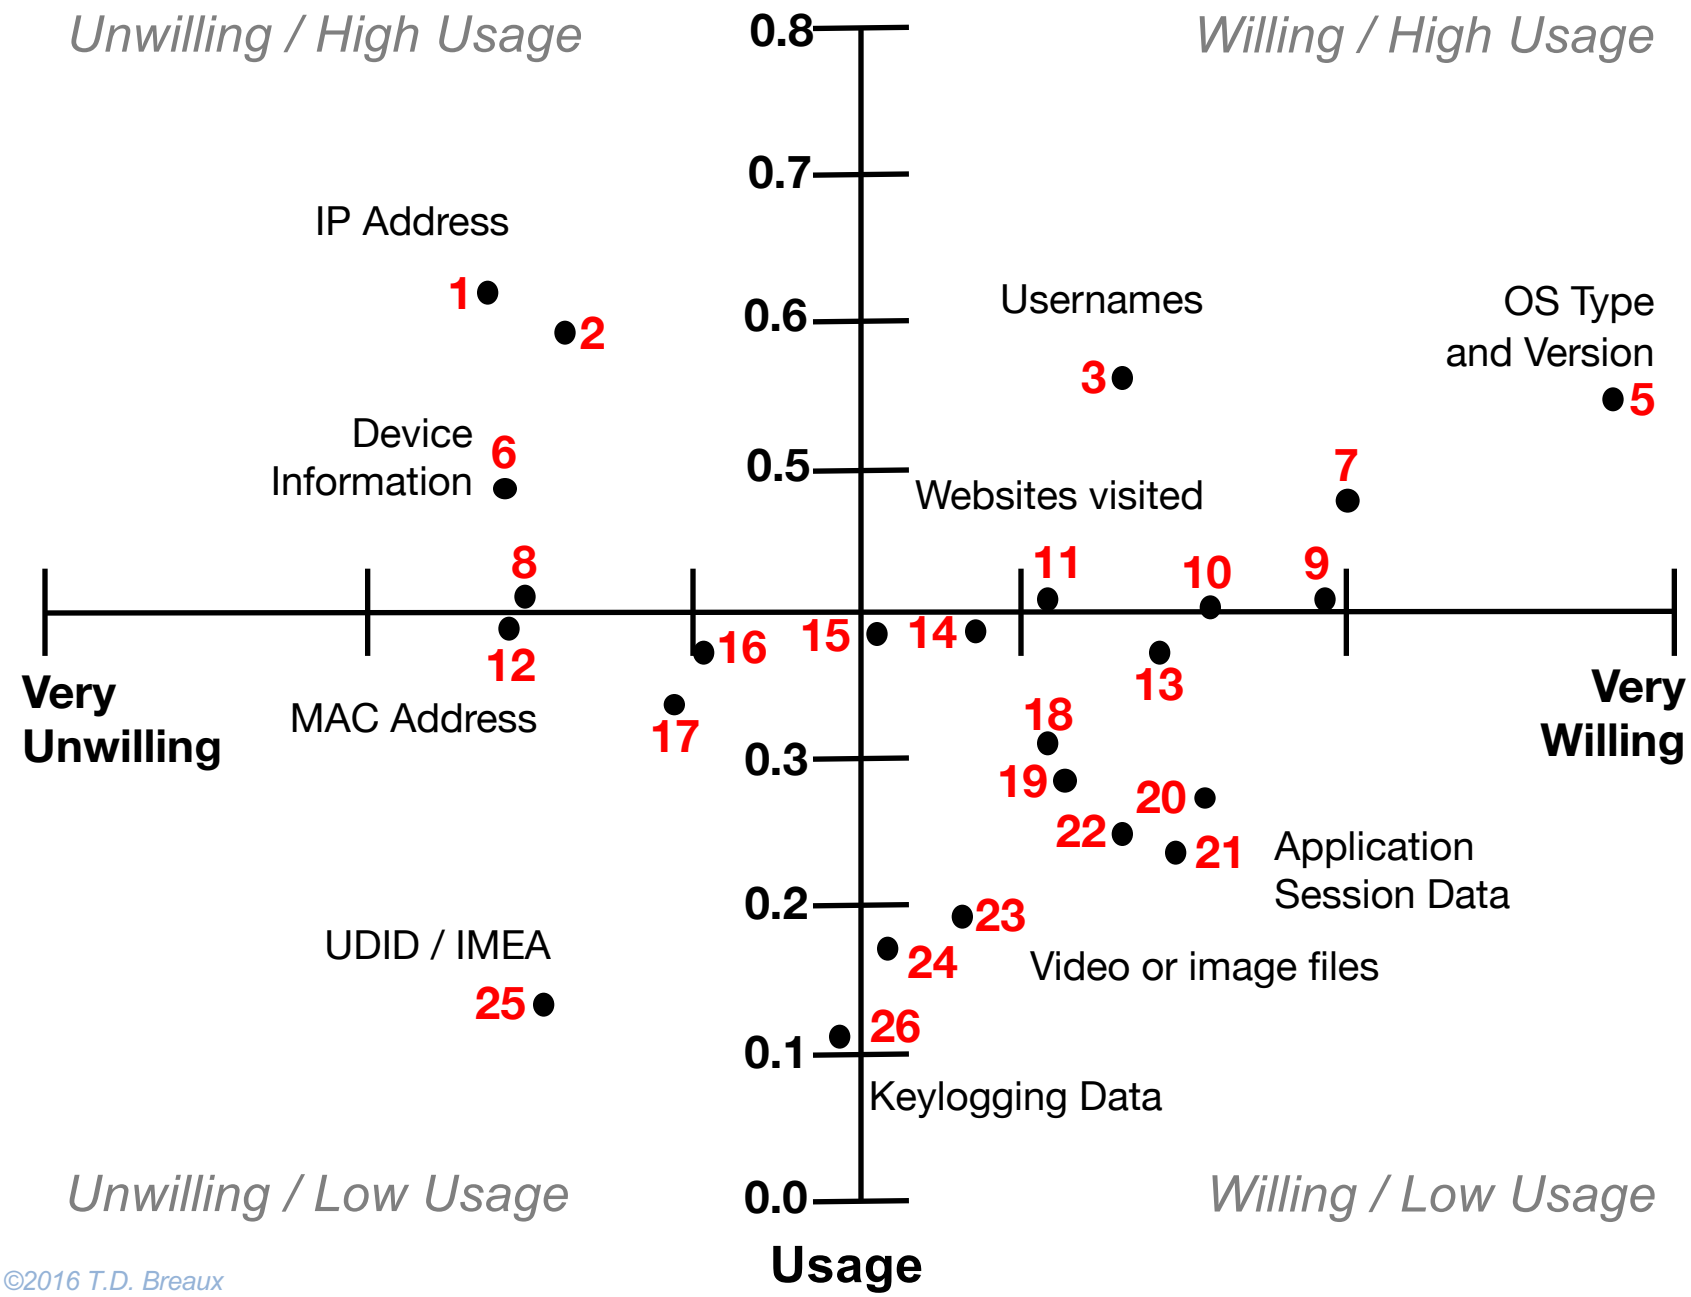

Contributions and Future Work

- Contributions:
  - Empirical method to measure perceived privacy risk by data type
  - Indirect identifiers alone may increase perceived risk (who you are, versus what you do)
  - No significant interaction between harm likelihood and perceived risk
- Future work:
  - Trade-off between privacy and data utility (benefits)
  - Interactions with mitigations to increase privacy
    - Data Redaction
    - Statistical Aggregation
    - Data Perturbation
    - Increased Restriction
  - Estimating privacy risk for data append (composition)

Acknowledgements

- We thank the CMU Requirements Engineering Lab, including João Caramujo, Morgan Evans, Mitra Bokaei Hosseini, and David Widder
- This work was supported by NSA Award #141333, NSF Frontier Award #133059, ONR Award #N00244-16-1-0006