Why “No Worse Off” *is* Worse Off

John Aycock
Department of Computer Science
University of Calgary, Canada

John Sullins
Department of Philosophy
Sonoma State University, CA

Our contention

- Justifying security research by claiming that affected users are “no worse off” is a very low bar
- Reflects a naively utilitarian bias

Furthermore

- “IRB-approved” is not necessarily ethical
- “Legal” is not necessarily ethical

Case study #1: Phishing

Disclaimer: we are not stating or implying that the work in these case studies was not valid or interesting, or that the researchers exceeded the bounds of what their IRBs allowed.

Phishing experiments

- Trying to study natural user behavior
- Deception studies
- Did not seek consent beforehand (with IRB approval)
- One experiment did not debrief either (again, with IRB blessing)

Phishing experiments

‘Both the use of deception […] and a complete waiver of consent […] clearly challenge the principle of respect for persons.’

– Finn and Jakobsson

Phishing experiments

‘…the risks inherent in a phishing study – as long as the researchers can ensure complete security for any information released by the subjects – are lower than those involved in a real phishing attack, to which online users are commonly exposed.’

– Finn and Jakobsson

Phishing experiments

‘…it is likely that many users are so accustomed to being subjected to phishing attacks that they are not likely to be upset by the fake attack (not knowing that it, indeed, is fake.)’

– Finn and Jakobsson

Phishing experiments

‘…it is likely that many users are so accustomed to being subjected to phishing attacks that they are not likely to be upset by the fake attack (not knowing that it, indeed, is fake.)’

– Finn and Jakobsson

To paraphrase: users are no worse off by being subjected unknowingly to these experiments.

Case study #2: Botnets

Botnet work

- UCSB researchers took control of Torpig in 2009 and operated it for ten days, obtaining data from 180,000 machines
- IRB approval obtained *after* work had begun, once it became apparent that PII was being captured

Botnet work

'…we protected the victims according to the following: […] The sinkholed botnet should be operated so that any harm and/or damage to victims and targets of attacks would be minimized.'

– Stone-Gross *et al.* (2009)

Botnet work

'…we protected the victims according to the following: […] The sinkholed botnet should be operated so that any harm and/or damage to victims and targets of attacks would be **minimized**.'

– Stone-Gross *et al.* (2009)

Botnet work

‘We established two main ethical criteria for our underground economy research: no user should be worse off as a result of our activities, and our activities should be beneficial for society at large.’

– Stone-Gross *et al.* (2011)

Botnet work

‘We established two main ethical criteria for our underground economy research: **no user should be worse off** as a result of our activities, and our activities should be beneficial for society at large.’

– Stone-Gross *et al.* (2011)

Botnet work: society benefits?

‘information [provided] to law enforcement […] might help in eradicating this threat’

‘we believe that we improved the understanding of this type of malware infrastructure’

– Stone-Gross *et al.* (2011)

Botnet work: society benefits?

‘information [provided] to law enforcement […] **might** help in eradicating this threat’

‘we **believe** that we improved the understanding of this type of malware infrastructure’

– Stone-Gross *et al.* (2011)

Ethical issues

Utilitarian bias

- Claim: an action is ethical if and only if there is a social benefit or at least no net harm
  - This sets the bar too low
  - Some problems with utilitarian ethical analysis
    - Very difficult to predict if actions will or will not leave the subject “no worse off” in the future
    - Necessarily biased to the viewpoint of the researcher given that the subject was not even consulted
    - It is hard to judge the value of any positive utility for the users/subjects/victims that is a direct or indirect consequence of the research

Will deontology help?

- Duty based ethical justifications
  - Professional codes of ethics for IT professionals and researchers exist
    - e.g. the ACM code of ethics and professional conduct
      - Does not support deception and lack of consent
  - Treats subjects as a means to an end and violates their moral agency
      - This violates Kantian ethical duties as well
  - Is our duty to protect users from botnets truly upheld by taking over the botnet and running it for ten days?

When you stare into the abyss…

- Maintaining the ethical high ground is difficult in this type of research
  - We use the same tools and techniques as moral miscreants
  - Gaining these habits can corrupt our best intentions unless we guard against this tendancy
  - These studies need an ethical foundation that goes beyond what is simply legal
    - Or what can be slipped past an IRB

Research ethics proposal

- Inspired by computer ethics and information ethics
  - Step one—assess the potential utility of the research on all those involved including the subjects, **but we do not stop here**
  - Step two— the safety net; always respect the moral rights of the subjects regardless of imagined utility
  - Step three – fulfill our duties as computing professionals
    - In the collection, storage, and synthesis of the information collected from the study
  - If we gain monetarily or in academic status from the research, we must be particularly diligent in this analysis

Beware of the "efficiency" counterargument

- This process works best in a collaborative context with much discussion about proposed research methods leaving no unstated moral commitments
  - Example: Some of the researchers might like to protect the subjects more; others argue against it on the grounds that it would be too difficult or expensive to achieve a higher level of protection
    - We must then determine if these "efficiency" arguments are based on real technological constraints or do they actually mask certain commitments (or dismissals) of unstated moral values held by those making the counterargument

Results

- We now have an open and honest research method with multiple layers of protection for the subjects
  - The standard moral intuitions of utilitarian analysis
  - Plus strict protections granted by deontological and professional duties

Conclusion

- Security research must maintain the moral high ground
- We look back on some early medical research with revulsion, e.g., the Tuskegee syphilis experiments…

  …will today's security research withstand the moral scrutiny of future generations?

Why “No Worse Off” *is* Worse Off

John Aycock
Department of Computer Science
University of Calgary, Canada

John Sullins
Department of Philosophy
Sonoma State University, CA