Slides by SlidesCarnival

ETHICAL TRADE-OFFS IN OPERATING SYSTEMS DESIGN

ELIZA WELLS, EMBEDDED ETHI**CS** APRIL 21, 2025

SMALL GROUP DISCUSSIONS TODAY!

Take a moment to make sure you're sitting near 2 or 3 people to talk to.

Share your **name** and **favorite TV show** with each other.



WHO AM I?

- Eliza Wells, <u>elizawells@fas.harvard.edu</u>
- **Research**: moral and social philosophy
- **Postdoctoral Fellow**, Edmond & Lily Safra Center for Ethics, Embedded EthiCS @ Harvard

DEVELOPING ETHICAL SKILLS

- What kinds of skills do you need to be a good engineer?
 - Programming, problem-solving, technical knowledge, creativity, precision, teamwork...
- Like engineering, **ethics involves skills**.
- Today, we're going to work with a **particular tool** for ethical problem-solving: **cost-benefit analysis.**

GOALS FOR TODAY

- . **Evaluate** cases involving trade-offs in OS design
- 2. **Apply** cost-benefit analysis (CBA) to these cases
- 3. **Identify** considerations needed to use CBA effectively
- 4. **Assess** some strengths and limitations of CBA as a method for resolving trade-offs

TRADE-OFFS IN OS DESIGN

Good operating systems

are easy to use, efficient, reliable, safe, fast, robust, have useful features, easy to maintain, flexible...

support privacy, security, accessibility...

 But often, we can't realize all these positive features simultaneously.

TRADE-OFFS IN OS DESIGN

How should we make trade-offs?

One intuitive, useful tool: **cost-benefit analysis (CBA)**

CBA provides a **structured process** to guide decision-making.

Roughly, CBA says you should resolve trade-offs by selecting the option that has the highest cost-to-benefit ratio.

CASE: GOOSE

Your company develops operating systems. One OS, Goose, is aging and vulnerable to new security threats, in part because it lacks kernel features used by anti-malware providers. However, many of your customers still run Goose and have been resistant to updates.

What should your company do?

https://www.emsisoft.com/ en/blog/20416/why-we-be lieve-its-not-ethical-to-sell -antivirus-software-for-win dows-xp-any-longer/

- 1. Specify your options
- 2. Identify stakeholders
- 3. Determine costs and benefits
- 4. Account for uncertainty
- 5. Identify a common measure
- 6. Add up and make a recommendation

Note: for simplicity, we're skipping *discounting* here.

1. Specify your options

- What choices are open to you?
- 2. Identify stakeholders
- 3. Determine costs and benefits
- 4. Account for uncertainty
- 5. Identify a common measure
- 6. Add up and make a recommendation

1. Specify your options

2. Identify stakeholders

- Which individuals or groups could be affected (directly or indirectly) by your choices?
- 3. Determine costs and benefits
- 4. Account for uncertainty
- 5. Identify a common measure
- 6. Add up and make a recommendation

- 1. Specify your options
- 2. Identify stakeholders

3. Determine costs and benefits

- For each option and each stakeholder, what are the possible short-term and long-term outcomes? How important are they?
- 4. Account for uncertainty
- 5. Identify a common measure
- 6. Add up and make a recommendation

- 1. Specify your options
- 2. Identify stakeholders
- 3. Determine costs and benefits

4. Account for uncertainty

- How likely are different outcomes? What don't you know?
- 5. Identify a common measure
- 6. Add up and make a recommendation

- 1. Specify your options
- 2. Identify stakeholders
- 3. Determine costs and benefits
- 4. Account for uncertainty

5. Identify a common measure

- In order to compute and compare net benefits, we need a single measure of value (often: monetary)
- 6. Add up and make a recommendation

- 1. Specify your options
- 2. Identify stakeholders
- 3. Determine costs and benefits
- 4. Account for uncertainty
- 5. Identify a common measure
- 6. Add up and make a recommendation
 - The option with the highest net value is the option you should choose!

CASE: SPARROW

You work at a startup that makes an OS called Sparrow, designed for global K-12 educators. 60% of your funding currently comes from venture capitalists.

Sparrow targets desktop and laptop machines and uses virtualization technology to run applications made for both Windows and MacOS. As you prepare for Sparrow's initial release, you discover a bug in the way it implements virtualization.

CASE: SPARROW

When Sparrow runs a VM for a MacOS app and a Windows app simultaneously, there is a 3% chance that the Windows app will not be properly isolated; if this happens, then the Windows app (if malicious) can tamper with the state of the MacOS app.

It would take two weeks to fix the bug before release, and two months to fix afterwards. If Sparrow is not released by its target date, the venture capitalists may refuse to provide subsequent funding.



CASE: SPARROW

When Sparrow runs a VM for a MacOS app and a Windows app simultaneously, there is a 3% chance that the Windows app will not be properly isolated; if this happens, then the Windows app (if malicious) can tamper with the state of the MacOS app.

It would take two weeks to fix the bug before release, and two months to fix afterwards. If Sparrow is not released by its target date, the venture capitalists may refuse to provide subsequent funding.

Should you delay and fix or release on time?

Specify your options

- 2. Identify stakeholders
- 3. Determine costs and benefits
- 4. Account for uncertainty
- 5. Identify a common measure
- 6. Add up and make a recommendation

If there's information you wish you knew, note it, then proceed with your best guess.

8

WORRY 1: UNCERTAINTY

CBA is sensitive to two types of uncertainty: **risk** and **outcomes**.

In July 2024, cybersecurity company CrowdStrike distributed a faulty update that caused ~8.5 million Windows systems to crash.

- 5,000+ flights canceled
- Government and health services impacted
- Estimated cost: \$10 billion

Does this change your thinking about Sparrow?



CASE: SPARROW II

You discover another bug. It's triggered when a user configures Sparrow to use a non-English keyboard setup (e.g., Japanese or Arabic). Such users have a 3% risk of their file data being corrupted by Sparrow's file system (which has subtle errors in the way that it handles filenames that don't use English characters). You expect 10% of your users to use non-English keyboards.

Should you delay and fix or release on time?



CASE: SPARROW II

You discover another bug. It's triggered when a user configures Sparrow to use a non-English keyboard setup (e.g., Japanese or Arabic). Such users have a 3% risk of their file data being corrupted by Sparrow's file system (which has subtle errors in the way that it handles filenames that don't use English characters). You expect 10% of your users to use non-English keyboards.

Should you delay and fix or release on time?

Specify your options

- 2. Identify stakeholders
- 3. Determine costs and benefits
- 4. Identify a common measure
- 5. Account for uncertainty
- 6. Add up and make a recommendation

If there's information you wish you knew, note it, then proceed with your best guess.

WORRY 2: INTERPERSONAL AGGREGATION

CBA combines all costs and benefits into one balance, even if they're problematically **distributed** across individuals or groups.

Uncritically using CBA can lead us to neglect **justice**.



WORRY 3: INCOMMENSURABILITY

CBA relies on **common measures of value**.

Often, we use monetary value. A standard way to monetize costs and benefits is to identify the corresponding *willingness to pay*—the maximum price someone would pay to experience that benefit or avoid that cost.

Can you think of a situation where monetizing all values is unsatisfying? How would you approach that case?

TAKEAWAYS

We can face technical and ethical **trade-offs** in OS design.

Today, we've developed one **skill** for dealing with these trade-offs: the ability to use **cost-benefit analysis.**

CBA is a helpful problem-solving tool, even though it's not perfect.

We need to be particularly attentive when using CBA to look at problems that might involve **risk**, **justice**, or **different values**.



elizawells@fas.harvard.edu