

Radical CS

Phillip Rogaway

University of California, Davis, USA

Thanks to Morris Dworkin, and the entire program committee, for the kind invitation to come give a talk.

The Third NIST Workshop on Block Cipher Modes of Operation


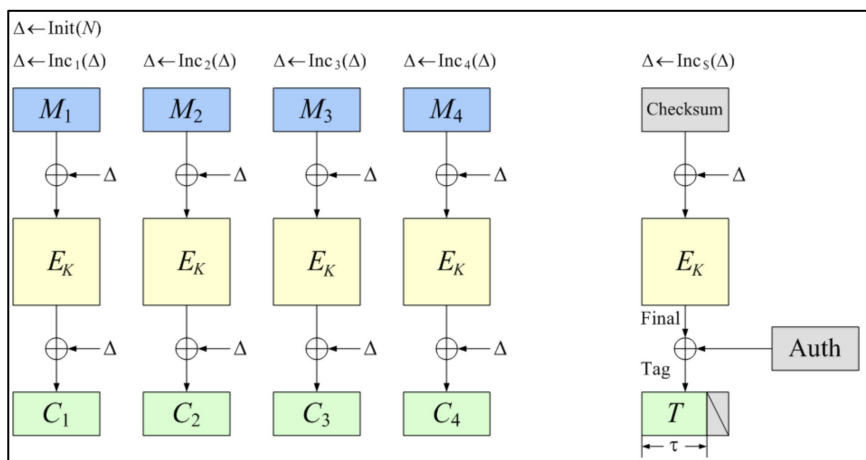
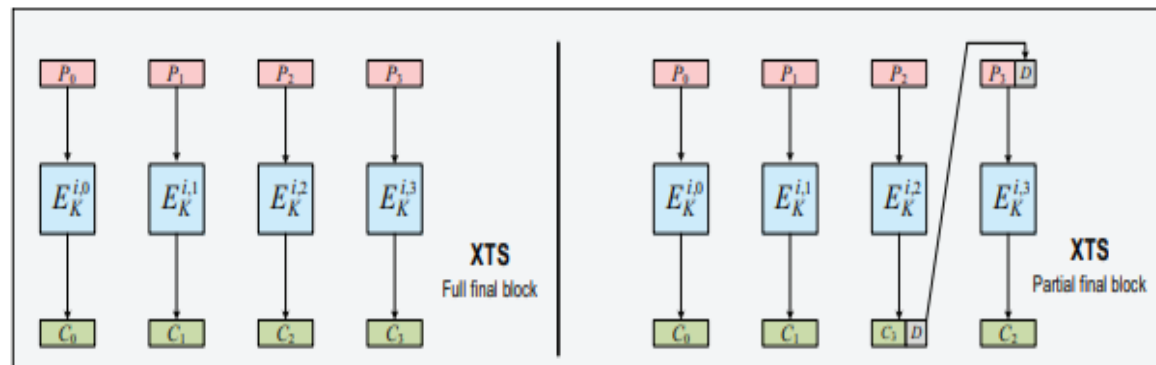
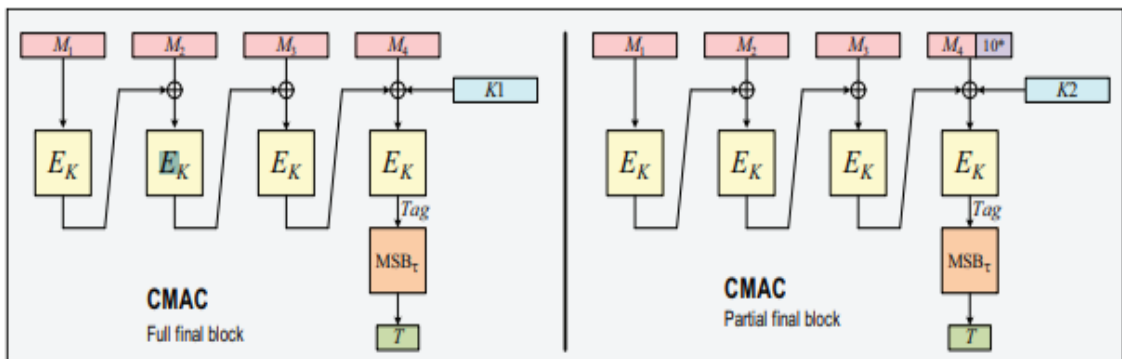
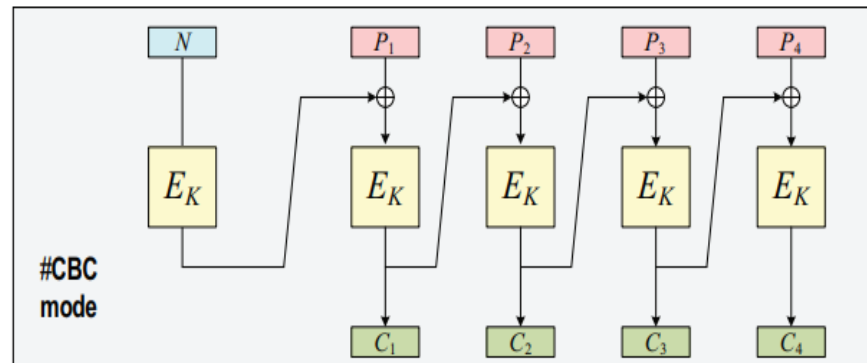

Rockville, Maryland, USA

3 October 2023

Today:

1. Rejecting the STN — embracing radical CS
2. Some attempts at radical CS
 - a) The Moral Character of Cryptographic Work (2015)
 - b) Reimagining Secret Sharing (2020)
 - c) A class: Ethics in an Age of Technology (2004–23)

I have special fondness for blockcipher modes



Modes meant intellectual independence — serious cryptographers didn't look at such things

A problem can be

- conspicuous outside the disciplinary community to which it seemingly belongs, yet
- invisible or ignored within that community.

The
“Invisibility
Phenomenon”

Random oracles are practical [BR93]

Entity authentication and key distribution [BR93]

The security of the CBC MAC [BKR94]

Optimal asymmetric encryption [BR94]

A concrete security treatment for symmetric encryption [BDJR97]

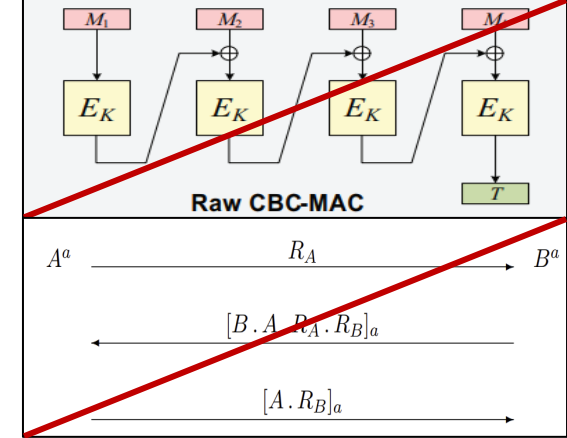
Cryptographic hash-function basics [RS04]

Formalizing human ignorance [R06]

A framework for code-based game-playing proofs [BR06]

Foundations of garbled circuits [BHR12]

The moral character of cryptographic work [R15]



Brandon Ogbunu, The Liberation of RNA, interview, June 2020: *Sometimes the most revolutionary thing ... you can do is just focus on the right things in life.*

Sometimes the most radical thing you can do in life — and sometimes the most worthwhile — is simply to *pay attention* to that which others fail to see. (Or that they do see, but choose to ignore.)

A growing ennui



I spent most my career

- Writing technical papers,
- giving technical talks,
- teaching technical subjects.

It was fun, and I am grateful.

But doing these things these days has come to feel increasingly misguided. Even self-indulgent.



The climate crisis is here. The biodiversity crises. 6th mass-extinction. Pandemic disease. Huge wildfires. Tipping points. And with these things: social, political, and economic turmoil; civilizational collapse. For young people: the future is bleak.

**Our world already feels
radically diminished**



My university
Spring 2020


Outside my
apartment
Portland
9/2020



What's left to hike
Desolation Wilderness, 7/2023



TOP 20 LARGEST CALIFORNIA WILDFIRES



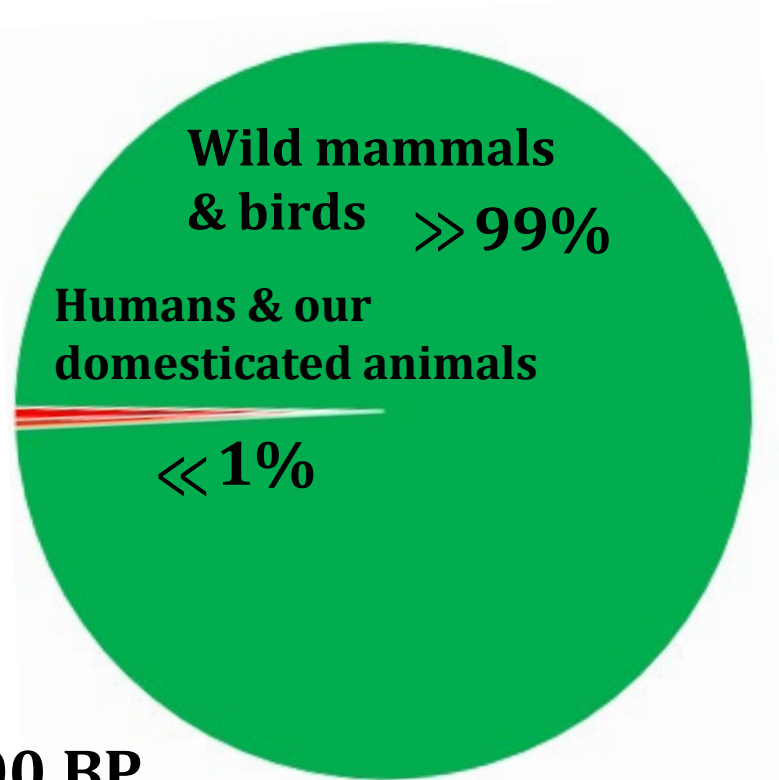
Not just
anecdotal
sensibilities
CalFire graphic
2021

Our assault on animal life

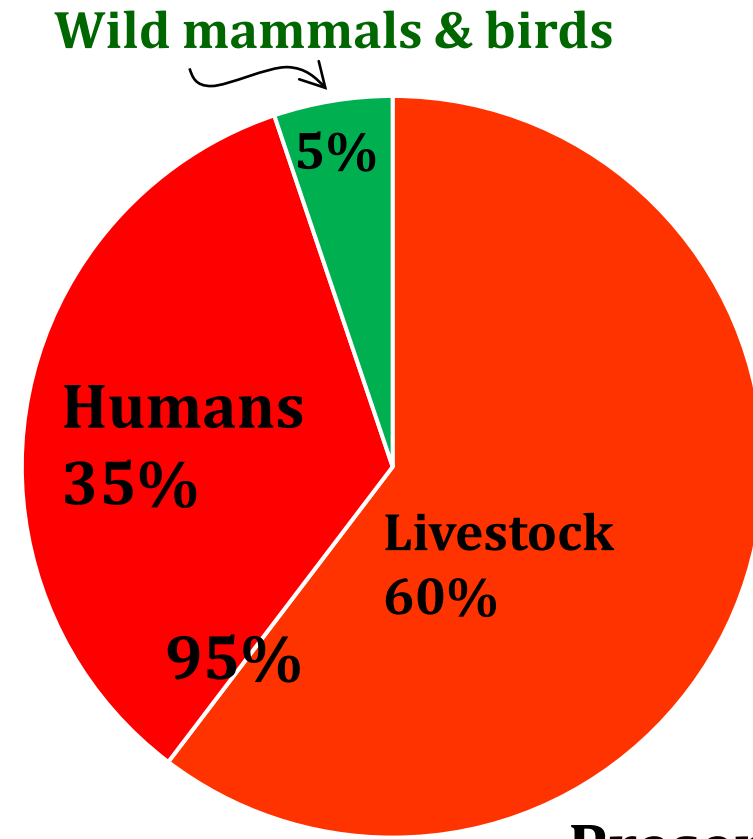
Biomass of land chordates

All for ourselves, and nothing for other people, seems, in every age of the world, to have been the vile maxim of the masters of mankind.

Adam Smith (1776)



10,000 BP



Present

Are we worth saving?




UC Davis
Spring 2019

And the role of CS?

Bringing enormous harms and risks — that mostly get ignored from within

The distraction economy




Face recognition

Killer robots



Imperiling democracy

Surveillance capitalism



Governmental mass-surveillance



Unaccountable AI



Wait! How about some optimism, instead?

The dominant narrative — techno-optimism — says that modern technology is not the problem — it's the solution

“I really do believe when ingenuity gets involved, when invention gets involved, when people get determined and when passion comes out, when they make strong goals — you can invent your way out of any box. That’s what we humans need to do right now. I believe we’re going to do it. I’m sure we’re going to do it.” J. Bezos, 2019



“Computer science is marking an epic change in human history. We are conquering a new and vast scientific continent. ... Virtually all areas of human activity ... [and] virtually all areas all areas of human knowledge ... are benefitting from our conceptual and technical contributions. ... Long live computer science!” S. Micali, 2013

Why does techno-optimism dominate?

Bezos, Gates, Jobs, Musk

Cognitive biases: optimism bias, the bandwagon-effect

Quick rewards; slow, nearly invisible harms, especially to the environment

Vaccines

The economy

Did you ever try to *read* Jacques Ellul of Lewis Mumford?

Antibiotics

Don't worry, Be happy

It's the culture, stupid.

Anesthesia

Cool gadgets. Washing machines, cars, smartphones, washing machines, ...

Just read Steven Pinker, man

Still here almost 80 years after nuclear weapons – way to go!

If we can send a man to the moon, we can send a man to the sun!

Green Revolution

Plastics

Make stuff, make money

Benefits are concrete and immediate; risks are abstract and long-term

I want to say one word to you. Just one word. ... Are you listening? ... *Bitcoin*. ... There's a great future in Bitcoin. Think about it.

Moore's law (see, it's even be legislated)

8.1 billion people

In Quinn's telling, "mother culture" envelops us in a suspect story

"ONCE WHEN I WAS IN college ... I wrote a paper for a philosophy class. ... Here's What I said ... Guess what? The Nazis didn't lose the war after all. They won it and flourished. They took over the world and wiped out every last Jew, every last Gypsy, black, East Indian, and American Indian. Then ... they out the Russians and the Poles and the Bohemians [and so on]. ... [W]hen it was all over, everyone in the world was one hundred percent Aryan, and they were all very, very happy.



"Naturally the textbooks used in the schools no longer mentioned any race but the Aryan or any language but German or any religion but Hitlerism or any political system but National Socialism. ... After a few generations of that, no one could have put anything different into the textbooks even if they'd wanted to, because they didn't *know* anything different.

"But one day two young students were conversing at the University of New Heidelberg in Tokyo. Both were handsome in the usual Aryan way, but one of them looked vaguely worried and unhappy. ... His friend said, 'What's wrong, Kurt? Why are you always moping around like this?' Kurt said 'I'll tell you, Hans. There *is* something that's troubling me—and troubling me deeply. ... It's this,' Kurt said. '**I can't shake this crazy feeling that there is some small thing that we're being lied to about.**' ..."

Daniel Quinn, *Ishmael*, 1992

What *is* the small thing we've been lied to about?

The “Standard Technological Narrative” (STN)


1. Technology is a **tool**. It is **apolitical** and **ethically neutral**.
2. Due to technology, things are **great** and **getting better**.
3. Better technology will **fix** what inferior technology broke.
4. We will overcome the **climate/environmental** challenges.
5. Tech is driven by brilliant **individuals**, advanced by the **marketplace**.
6. We have risen far above **animals**, are creating a **technological utopia**.



My problem with the STN:

It's a fantasy

The thing about the STN



1. Even if you don't believe it, you might behave as if you do
2. Rejecting the STN will have a profound impact on your views.
Eg: What work is worthwhile? What faculty should we hire? What should we do in the classroom?
3. The STN is fundamentally a *religious* point of view
4. It paints the technologist as the savior / hero
5. It serves corporations and the elite
6. It de-politicizes and de-moralizes our current crises
7. In its most extreme form, it devolves into the TESCREAL bundle of beliefs
(= Transhumanism, Extropianism, Singularitarianism, Cosmism, Rationalism, Effective Altruism, and Longtermism) [Timnit Gebru, Émile Torres 2023]

Drawing by a student in my ethics class. 2019

Viewed negatively: ~~-STN-~~

Viewed positively: **Radical CS**

[1] **Radical CS** recognizes that CS — and technology more broadly — embeds **values**. It is never **neutral**. It **rearranges power**. It has tended to disproportionately empower **big corporations, tech workers, and the elite**. Doing so, it creates significant **peril** for people and the planet.

[2] **Radical CS** aims to confront this. We want to **reinvent** CS in ways that empower **ordinary people** and **disempower** the already powerful. We want to **reverse** the environmental, social, and political peril we have helped create. We want to **stop** creating **new** risks.

[3] **Radical CS** accepts that it may be better to **dismantle** a system than to tweak it. It recognizes that some projects ought **not** to be pursued at all — at least not now.

Suggestions for a **radical CS**

1. Stop pretending that things are not seriously messed up. It's disempowering and dishonest

Isn't it *better* to be optimistic?

No — at least not for society.

Excessive optimism — not pessimism — undermines social progress. It obviates

- the need for broad thinking
- the recognition of emergency
- the basis for social-change movements



Regardless: “better” isn’t the point — there’s that annoying honesty-thing


Also: existential threats motivate giving primacy to predictions of doom over prophecies of bliss even if one is skeptical of the former.

Suggestions for a **radical CS**

1. Stop pretending that things are not seriously messed up. It's disempowering and dishonest
2. See the STN for what it is. A story. A culturally-fabricated narrative.
3. Identify the embedded values. They're often explicit. Or easily coaxed out.

The Values Encoded in Machine Learning Research

[Birhane, Kalluri, Card, Agnew, Dotan, Bao 2021/22]



Suggestions for a **radical CS**

1. **Stop pretending that things are not seriously messed up.** It's disempowering and dishonest
2. **See the STN for what it is.** A story. A culturally-fabricated narrative.
3. **Identify the embedded values.** They're often explicit. Or easily coaxed out.
4. **Stop pretending that CS holds answers it does not.** AI is going to fix the climate crisis, food insecurity, lousy schools, ... Blockchain is going to be democratizing, stabilizing, ... *Give me a break.*
5. **Don't try to instill improved characteristics into rotten enterprises.** "21st century liberalism is ensuring a panel at a defense industry conference called *Building a Deadlier Drone* has adequate gender diversity." *Fredrik DeBoer*
6. **The first question to ask: should you build the thing at all.** When we emphasize properties like fairness, accountability, and transparency we skip this question and get to lower-level ones. This is unthreatening to power and careers.
7. **Attend to the *primary* reason for the thing; follow the money.** Sure, AI might read x-rays better than radiologists. But that's not from where the push comes.
8. **Move slow and fix things.** Flip the FB motto. Caveat: don't move slowly on things that imperil us, like environmental collapse.
9. **Foreground your employer's social impact.** Your own positive social impact *outside* the workplace won't compensate for negative social impact *in* the workplace.
10. **Stop the Orwellian double-speak.** A whole slide for that!

CS doublespeak

← Language designed to deceive or distort its actual meaning, normally for the benefit of those in power

Could we invent more deceptive language were this the explicit goal?



Algorithm (a) A program to compute some unknown function. (b) An opinion rendered in code.

Cloud computing Putting your data on somebody else's servers so that it can be stored in an unknown jurisdiction and mined by unknown parties for unknown ends. But at least it sounds fluffy and cool.

Crypto Used to mean *cryptology* — the art and science of secure communication. Now it refers to a massive Ponzi scheme wrapped in technobabble. (P. Klugerman, 5/21/2020)

Deep learning Learning devoid of depth due to an absence of foundations *and* domain expertise *and* sociopolitical thinking .

Smartphone A phone that is not smart and that pushes its users to be just as stupid. Also, the device should barely function as an actual phone.

Social media Systems designed to sunder social interactions.

Suggestions for a **radical CS**

1. **Stop pretending that things are not seriously messed up.** It's disempowering and dishonest
2. **See the STN for what it is.** A story. A culturally-fabricated narrative.
3. **Identify the embedded values.** They're often explicit. Or easily coaxed out.
4. **Stop pretending that CS holds answers it does not.** AI is going to fix the climate crisis, food insecurity, lousy schools, ... Or: blockchain is going to be democratizing, stabilizing, ... *Give me a break.*
5. **Don't try to instill improved characteristics into rotten enterprises.** "21st century liberalism is ensuring a panel at a defense industry conference called *Building a Deadlier Drone* has adequate gender diversity." *Fredrik DeBoer*
6. **The first question to ask: should you build the thing at all?** When we emphasize properties like fairness, accountability, and transparency we skip this question and get to lower-level ones. This is unthreatening to power and careers.
7. **Attend to the *primary* reason for the thing; follow the money.** Sure, a good ML-based system might read x-rays better than most radiologists. But that's not from where the push comes.
8. **Move slow and fix things.** Flip the FB motto. Caveat: don't move slowly on things that imperil us, like environmental collapse.
9. **Foreground your employer's social impact.** Your own positive social impact *outside* the workplace won't compensate for negative social impact *in* the workplace.
10. **Stop the Orwellian double-speak.** A whole slide for that!
11. **Don't sleep with the enemy.** Don't work for or accept money from those whose values you disagree with.
12. **It's the system, stupid.** Growthism; industrial-growth capitalism.

The phrase **radical CS** is adapted from the

radicalai.net

Radical AI Network

[What is radical AI](#) [Radical AI principles](#) [Radical AI work](#) [The Radical AI Network](#)

What is Radical AI?

Radical simply means 'grasping things at the root' -Angela Davis

Radical AI exposes how AI rearranges power and dreams up and builds human/AI systems that put power in the hands of the people.

radicalAI.net

Radical AI Principles

1. ...
2. ...
3. We recognize that all technologies rearrange power.
4. We are critical of how AI shifts power. In particular, we recognize AI is frequently extractive, exploitative, surveilling, controlling, prescriptive, and reductionist. We recognize AI frequently prevents consent, deliberation, investigation, intervention, resistance, and agency.
5. ...

2 a) The Moral Character of Cryptographic Work (2015)

The Moral Character of Cryptographic Work*

Phillip Rogaway

Department of Computer Science
University of California, Davis, USA
rogaway@cs.ucdavis.edu

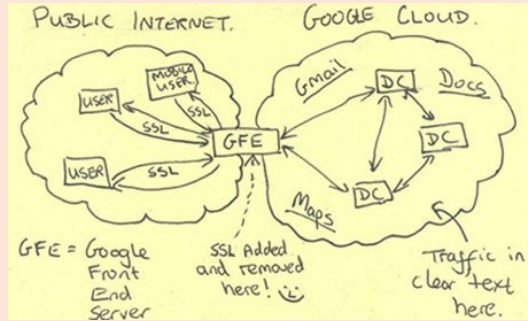

December 2015
(minor revisions March 2016)

Abstract. Cryptography rearranges power: it configures who can do what, from what. This makes cryptography an inherently *political* tool, and it confers on the field an intrinsically *moral* dimension. The Snowden revelations motivate a reassessment of the political and moral positioning of cryptography. They lead one to ask if our inability to effectively address mass surveillance constitutes a failure of our field. I believe that it does. I call for a community-wide effort to develop more effective means to resist mass surveillance. I plead for a reinvention of our disciplinary culture to attend not only to puzzles and math, but, also, to the societal implications of our work.

Keywords: cryptography · ethics · mass surveillance · privacy · Snowden · social responsibility

Preamble. Most academic cryptographers seem to think that our field is a fun, deep, and politically neutral game—a set of puzzles involving communicating parties and notional adversaries. This vision of who we are animates a field whose work is intellectually impressive and rapidly produced, but also quite

Cryptography – the science of **secure communications**.



Mass surveillance – the spectacular **failure to secure communications**.

So you might **think** that **cryptographers** would be **ashamed** and **aghast** about **mass surveillance** revelations.



You'd be wrong.

My community thinks things are going **great**, and that mass surveillance is **not our concern**.

Suggestions from the essay

1. Do more **crypto-for-privacy / anti-surveillance** research.
2. Attend to problems' **social value**.
3. Be **introspective** as to **why** you're working on what you are.
4. Look to **current security practice** and **privacy needs** as a source of probs.
5. Be open to **diverse models**. Regard **all** models as suspect and dialectical.
6. Think **twice** before accepting military funding.
7. Regard **ordinary people** as those whose needs you aim to satisfy.
8. Figure out what research would **frustrate the NSA**. Then do it.
9. Stop with the **cutesy pictures**. Take adversaries seriously.
10. Use the **academic freedom** you have.
11. Get a **systems-level** view.
12. Learn some **privacy tools**. Use them. **Improve** them.
13. Design and build a broadly useful **cryptographic commons**.

How did this work out?

Lots of friendly emails, which continue until this day.

Widely read in undergrad CS programs.

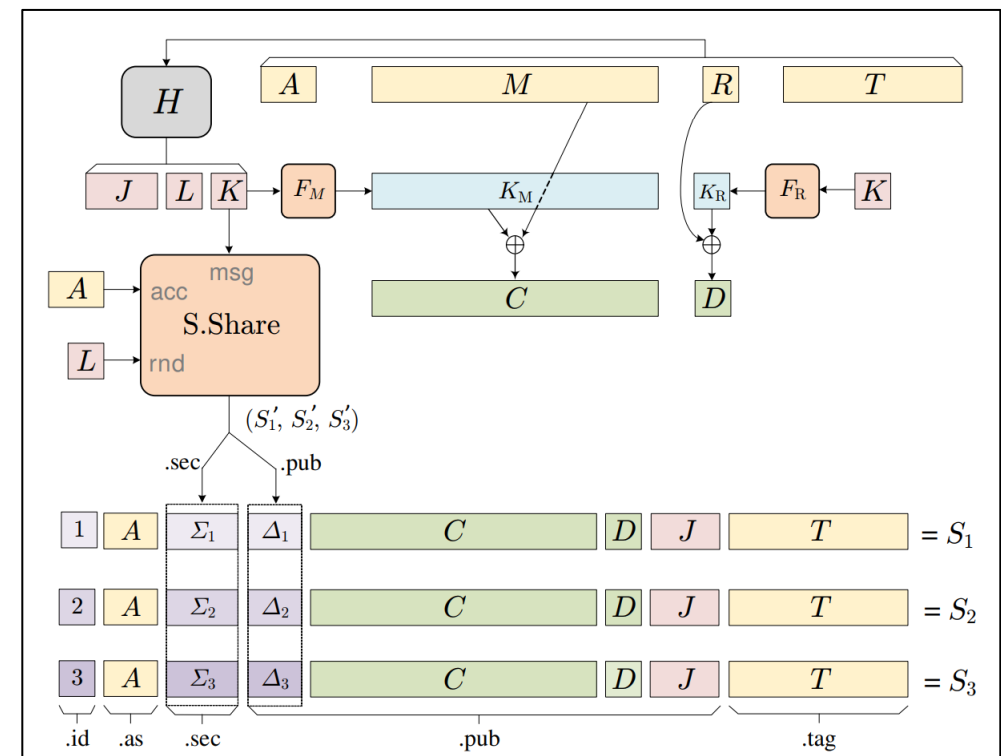
No open disagreement from within the crypto community of the essay's central claims

No recognizable change within the field.

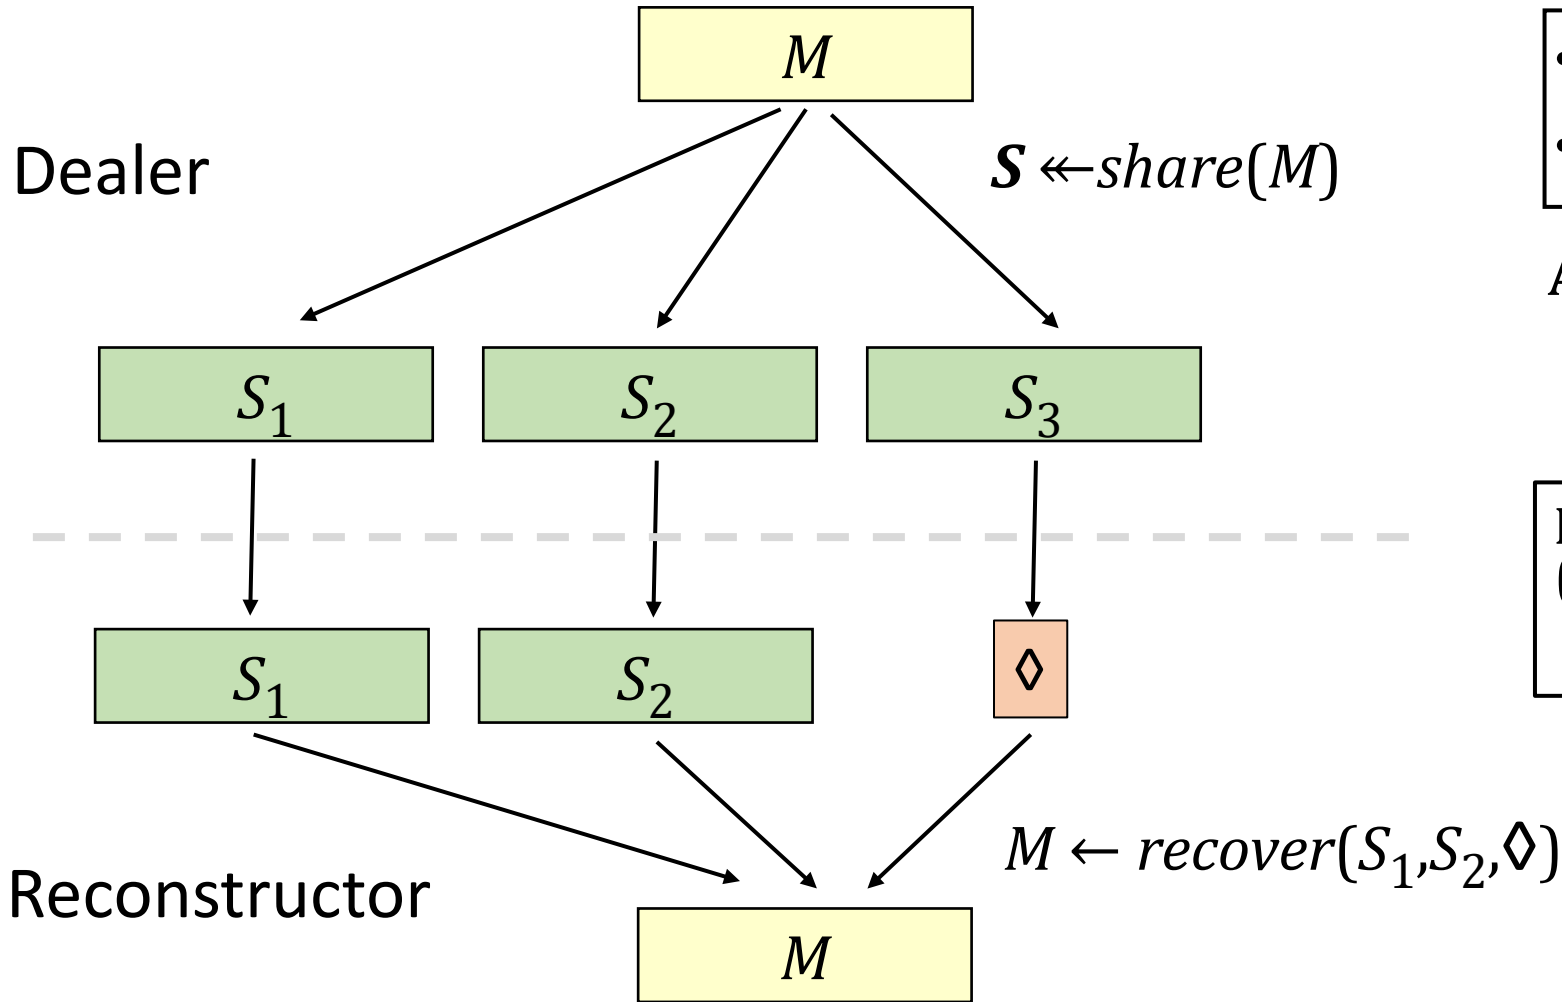
Grade: B-

2 b) Reimagining Secret Sharing (2020)

“Our initial reason for developing ADSS was to address use cases involving journalists and whistleblowers. We were motivated by a conversation with journalist Laurent Richard [36,22], by the Snowden revelations [24], and by the development of Sunder [39]. We recognized that unadorned Shamir secret-sharing [40] wouldn't do ...”



Secret Sharing



- $\text{share} : \text{Message} \rightarrow \text{Share}^n$
- $\text{recover} : (\text{Share} \cup \{ \})^n \rightarrow \text{Message}$

Access structure A

A set of subsets of $[1..n]$
for some $n = n(A)$ monotone

Privacy requirement:
 $(\forall [1..n] \ni B \notin A) (\forall M, M' \in \text{Message})$
 $(\text{share}(M))_B = (\text{share}(M'))_B$

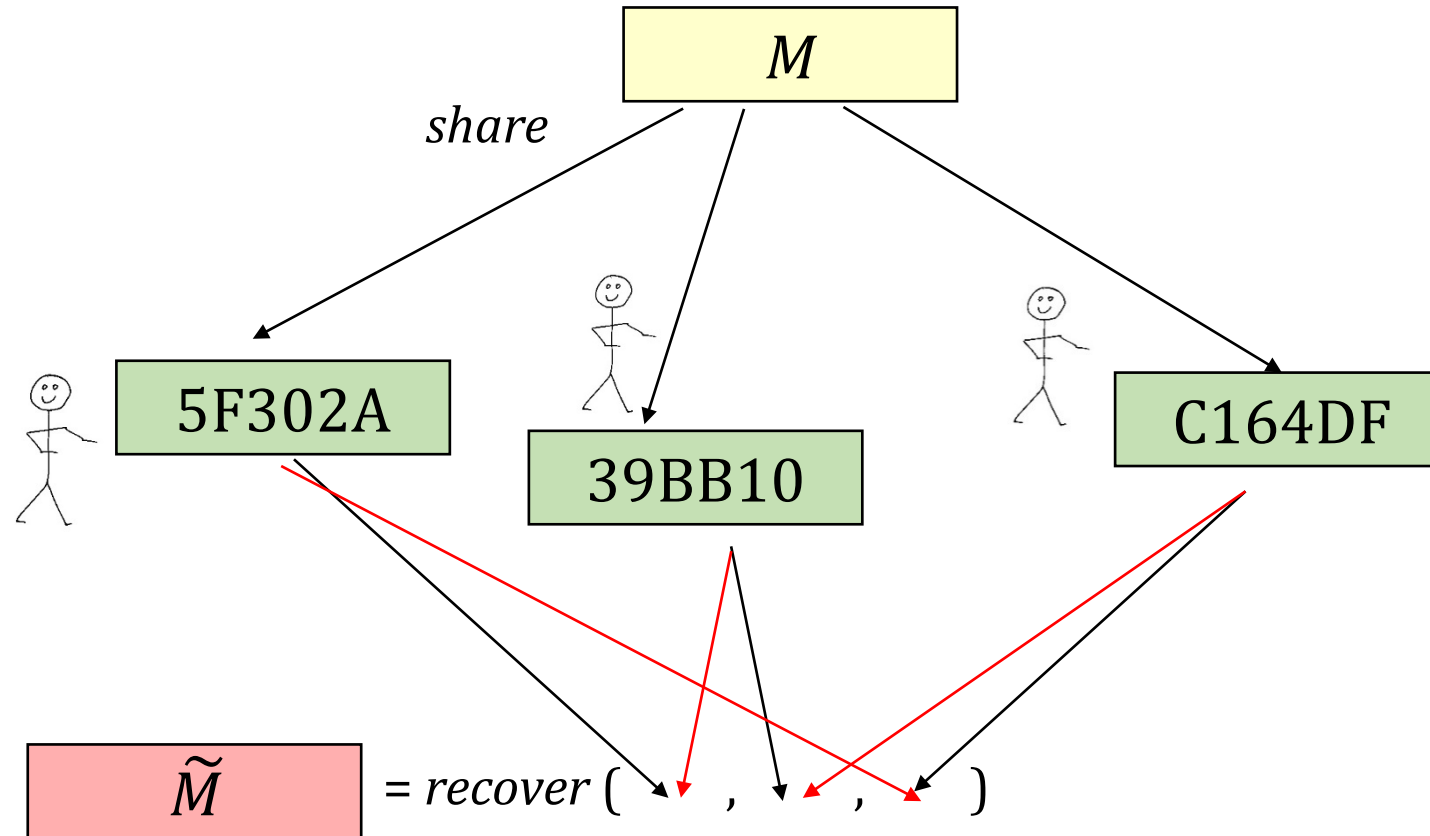
where $S_U[i] = \begin{cases} S[i] & \text{if } i \in U \\ \} & \text{if } i \notin U \end{cases}$

Problems?

Simple, elegant, 45-year-old notion and technique—
what could possibly be wrong or unsaid?

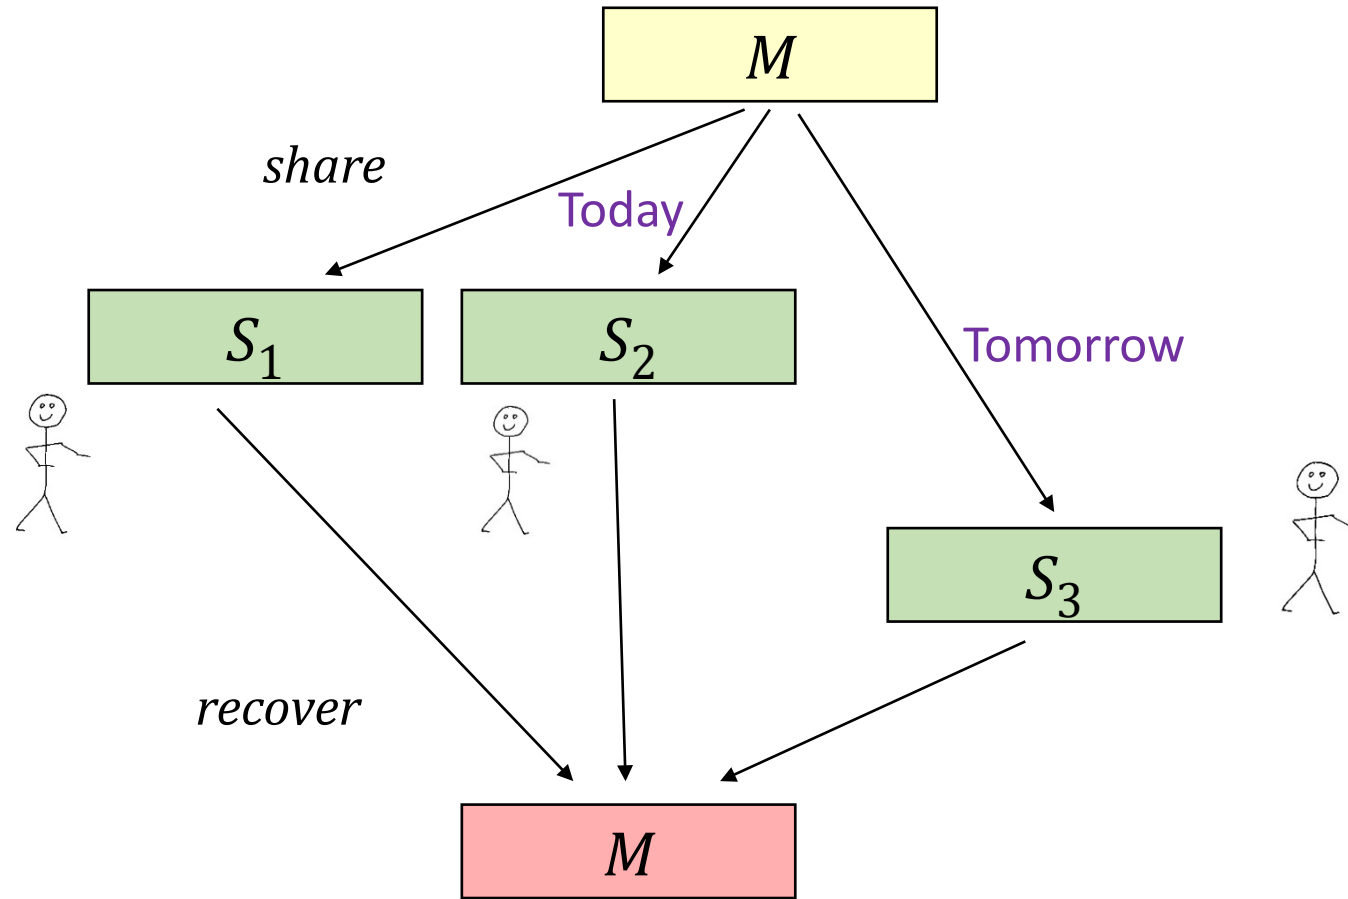
A lot. Classical-SS has a ton of unexplored problems
that wreck its utility for what it's ostensibly for.

Scenario #1



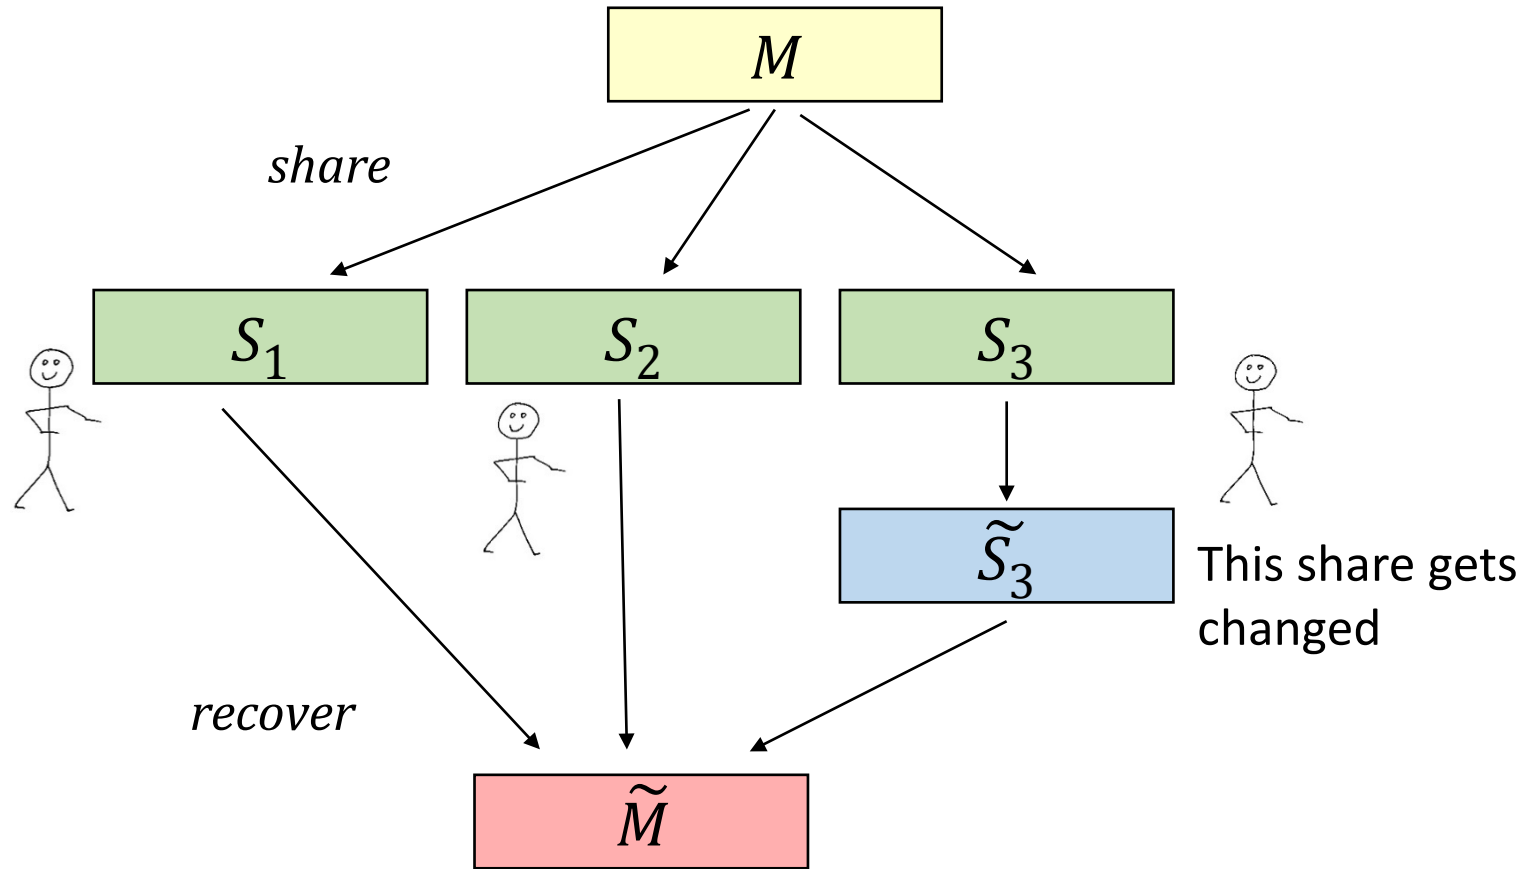
Shareholders must know their "position" — shares have implicit metadata

Scenario #2



Algorithm *share* is **randomized**, so a share can't be regenerated without retaining the coins. But the coins can't be retained without destroying security.

Scenario #3



You'll recover **something** — and get no indication anything is wrong.

Such issues can be fixed, of course, which is what **adept SS** aims to do.

But your ending point — definitions, properties, and constructions — will be quite unlike classical secret sharing.

Just from taking seriously that you are trying to craft a practical tool to actually split up a secret.

How did this work out?

The paper was technically successful; it solved everything it aimed to solve.

It was ignored. 6 citations. Even the journalists who brought the problem to our attention didn't really seem to really need a technical solution.


Grade: C-

2 c) A class: Ethics in an Age of Technology (2004-2023)



“I want you to think about **and act upon** the ethical implications of

- **your personal and professional choices, and**
- **our collective work as technologists.”**



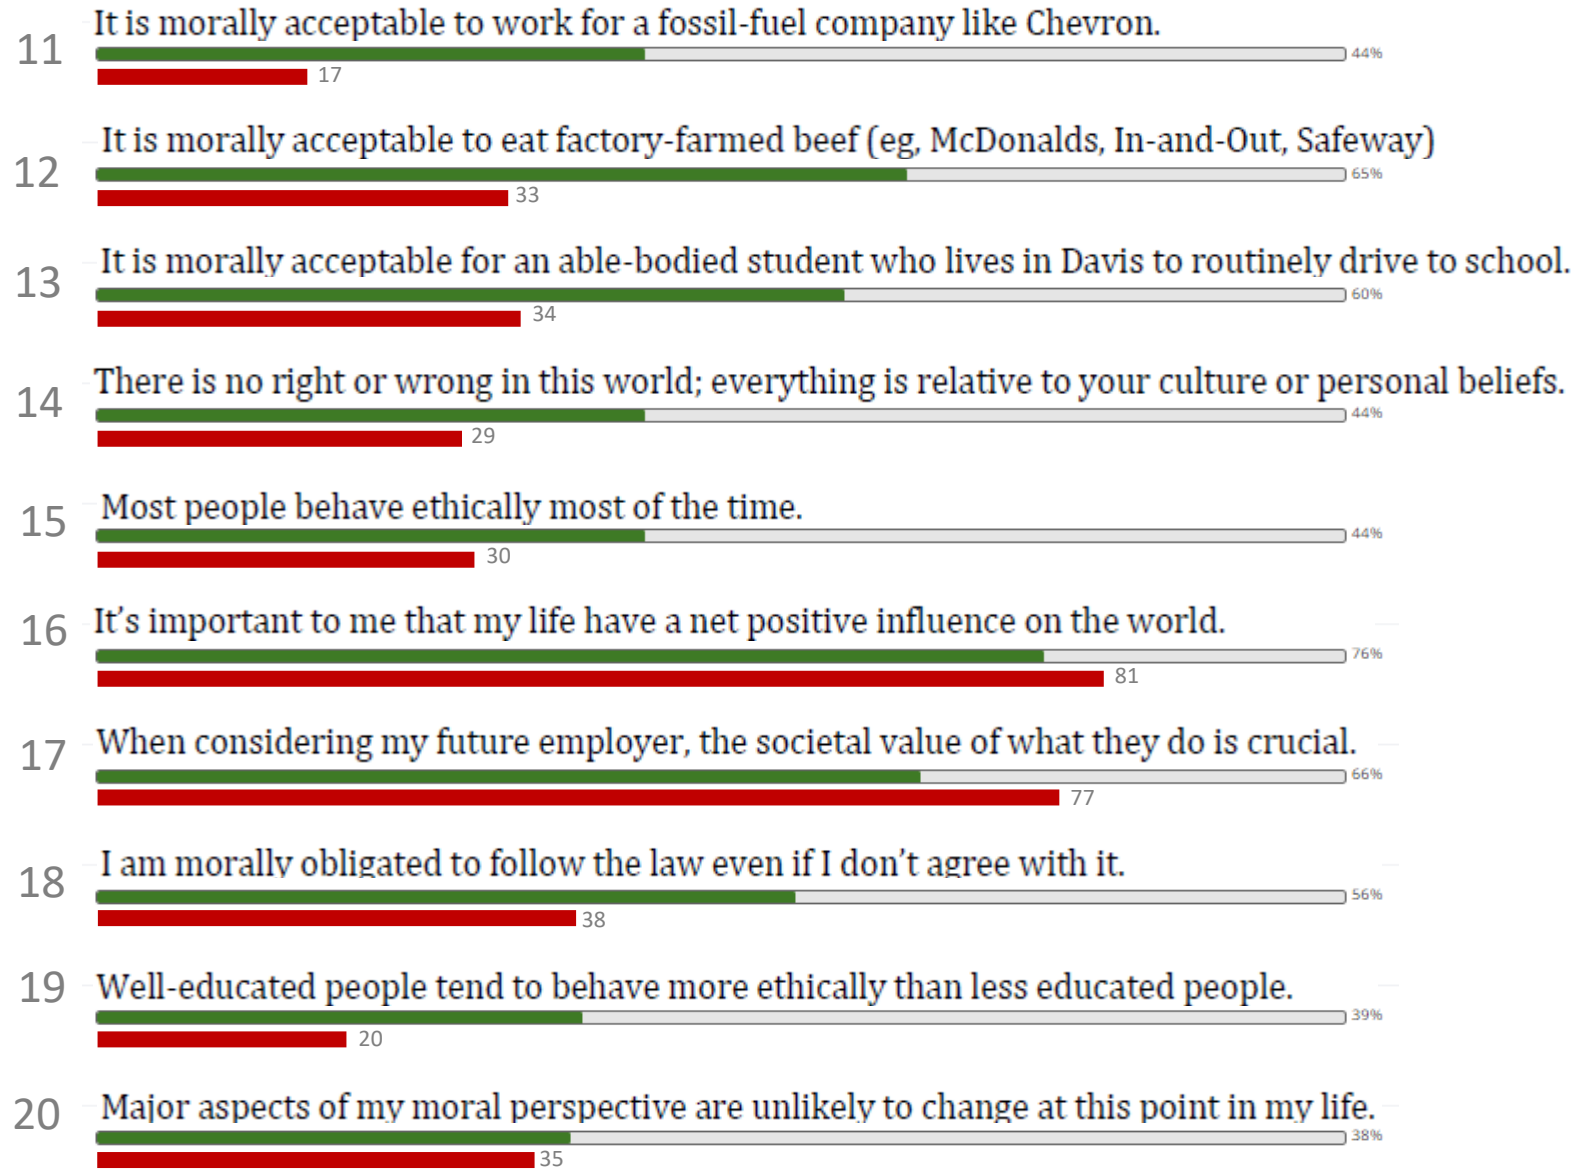
20 years of teaching ethics

1. No lectures, just facilitate.
2. Genuinely listen
3. Allow no phones, no laptops
4. Forget moral philosophy
5. Steer far away from methodological approaches to ethical analyses
6. Encourage students to *feel*, not think
7. Select disturbing films and articles; have disturbing discussions
8. Don't worry about the students feeling bad
9. Chatham House Rule
10. Urge students to be judgmental
11. Dismantle hyper-individualism, ethical relativism
12. Keep it personal: what we eat, where we work, how we die, ...

Do basic attitudes shift? **Beginning** — **End** (SQ23)



Do basic attitudes shift? **Beginning** — **End** (SQ23)



How did this work out?

Many students **do** change.

Course seems to have a profound impact on values of many, perhaps most..

But ... ~24 per class. Not remotely commensurate with the problem.

And I have never known how to scale this up ... if that is possible at all.

Grade: B+

Concluding remarks

CS, and technology more broadly, is full of smart people that are ethical morons. Don't be one.

My efforts at radical CS haven't been very successful. But you can do better. There *is* a community of people who care about these things. And a rich history of waxing and wanning efforts to make technology more responsive to human needs.

“In dark times, it does no good to pretend that you are not living in dark times”



Ira Glass, *This American Life*, 781: Watching the Watchers, 7 Oct 2022

References

- Ludwik Fleck, *Genesis and Development of a Scientific Fact* (1935)
Dahr Jamail, *The End of Ice* (2019)
Hans Jonas, *The Imperative of Responsibility* (1979/1984)
Paul Kingsnorth & Dougald Hine, *Uncivilization: The Dark Mountain Manifesto* (2009)
Daniel Quinn, *Ishmael* (1992) (+ two successor books)
Rupert Read and Samuel Alexander, *This Civilisation is Finished* (2019)
Shoshana Zuboff, *The Age of Surveillance Capitalism* (2019)

Title: Radical CS

Abstract: An unhappy reality has plagued my career: that I disagree with most everything that goes on within my field. That is true whether I am thinking of my field narrowly, as cryptography, or broadly, as computer science (CS). In this talk I own up to my grumpy discontent. I describe what I understand to be its principle cause: a rejection of the “Standard Technological Narrative” (STN). I call the negation of this view, as it applies to computer science, *radical CS*. I try to imagine what a program of radical CS might look like. I provide a post-mortem on three pieces of my prior work that were, in retrospect, attempts at radical CS: writing about technopolitics in *The Moral Character of Cryptographic Work* (2015); redefining secret sharing in *Reimagining Secret Sharing* (2020); and replacing much of my teaching with a disturbing course on ethics-and-technology (2004–2023). While I find none of these efforts to have been particularly successful, I express hope that others might do better.