

The Moral Character of Cryptographic Work^{*}

Phillip Rogaway

Department of Computer Science
University of California, Davis, USA

December 3, 2015

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 plea 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 inbred and divorced from real-world concerns. Is this what cryptography *should* be like? Is it how we *should* expend the bulk of our intellectual capital?

For me, these questions came to a head with the Snowden disclosures of 2013. If cryptography’s most basic aim is to enable secure communications, how could it *not* be a colossal failure of our field when ordinary people lack even a modicum of communication privacy when interacting electronically? Yet I soon realized that most cryptographers didn’t see it this way. Most seemed to feel that the disclosures didn’t even implicate us cryptographers.

I think that they do. So I want to talk today about the moral obligations of cryptographers, and my community as a whole. This is not a topic cryptographers routinely discuss. In this post-Snowden era, I think it needs to be.

^{*} This is an essay written to accompany an invited talk (the 2015 IACR Distinguished Lecture) given at Asiacrypt 2015 on December 2, 2015, in Auckland, New Zealand. The essay and talk are addressed to the cryptographic community—my community—and the words “we” and “our” should be so interpreted. I apologize in advance if I offend anyone with any of my comments; nothing of the sort is my intent.

Part 1: Social responsibility of scientists and engineers

A famous manifesto. I'd like to begin with a story—a true story.¹ To set the stage, it is London, the summer of 1955. A roomful of reporters have assembled for a press conference in Caxton Hall, a red brick building in Westminster. The media have summoned in a plan hatched by Bertrand Russell, with some help from the editor of *The Observer* newspaper. The reporters don't know just *why* they are here, having only been told that a team of the world's leading scientists were ready to release something of world-wide significance. The press knows that Sir Bertrand Russell is involved. With Einstein's recent death, Russell has become the world's most famous living intellectual.

Russell has been in his home, hiding, all week. All day long the phone rings, the doorbell rings. Reporters are trying to find out what is this big announcement. Russell's wife and his housekeeper make excuses and shoo the reporters away.

As the press conference begins, the reporters learn from Russell and accompanying physicist Joseph Rotblat that they have *not* been assembled to hear of some new scientific discovery, but to receive a prepared, political statement. It's a fairly brief statement, but it's been signed by eleven² of the world's leading scientists—nine of them Nobel laureates. Albert Einstein is among the signatories, signing just days before he became ill and died.

The document would become known as the Russell-Einstein manifesto.³ I hope that its contents are known to you. It speaks of the existential threat to mankind posed by nuclear weapons. Its final passage sounds desperately plaintive as Russell writes:

We appeal, as human beings, to human beings: Remember your humanity, and forget the rest. If you can do so, the way lies open to a new Paradise; if you cannot, there lies before you the risk of universal death.⁴

The reporters ask questions and soon warm to the manifesto's importance. The next day, the manifesto is carried as front-page news of most the world's major newspapers. For the next several days, at least, it is the talk of the world.

The Russell-Einstein manifesto galvanized the peace movement. It launched the Pugwash Movement, for which Joseph Rotblat and the Pugwash Conferences themselves would eventually share the Nobel Peace Prize (1995). Rotblat credits the manifesto for helping to create the conditions that gave rise to the Nuclear Non-Proliferation Treaty (NPT, 1970).⁵ In his Nobel Peace Prize acceptance speech, Rotblat explains:

From my earliest days I had a passion for science. But science, the exercise of the supreme power of the human intellect, was always linked in my mind with benefit to people. I saw science as being in harmony with humanity. I did not imagine that the second half of my life would be spent on efforts to avert a mortal danger to humanity created by science.⁶

Two modes of behaving politically. I begin today's talk with the Russell-Einstein manifesto to remind you of two things: first, that technical work *itself*

can implicate politics; and second, that some scientists, in response, do take on overtly political roles. These two ways to behave politically are different (even if, to people like Rotblat, they go hand-in-hand). Let's look at each.

Implicit politics. A scientist engages in what I'll call *implicit* politics by influencing power relations as a byproduct of technical work. Politics is about power—who has how much of it, and what sort. The nuclear bomb is the ultimate expression of coercive power; it is politics incarnate. Had Rotblat shunned every ostensibly political role in his life, his life's work would still have been political. Immensely, if implicitly, so.

But we don't need the specter of mushroom clouds to be dealing with politically relevant technology: scientific and technical work *routinely* implicates politics. This is an overarching insight from decades of work at the crossroads of science, technology, and society.⁷ Technological ideas and technological things are not politically neutral: routinely, they have strong, built-in tendencies. Technological advances are usefully considered not only from the lens of *how* they work, but also *why* they came to be as they did, *whom* they help, and *whom* they harm. Emphasizing the breadth of man's agency and technological options, and borrowing a beautiful phrase of Borges, it has been said that *innovation is a garden of forking paths*.⁸

Still, *cryptographic* ideas can be quite mathematical; mightn't this make them relatively apolitical? Absolutely not. That cryptographic work is deeply tied to politics is a claim so obvious that only a cryptographer could fail to see it. Thus I'll devote considerable time to this claim. But let me first speak of the second way for the scientist to behave politically.

Overt politics. A scientist can engage in *overt* politics through the mechanisms of activism and participatory democracy. In writing the Russell-Einstein manifesto and in rolling it out the way he did, Russell was working masterfully in this domain. Russell was not *only* a mathematician: he had broad contributions across philosophy, had won the Nobel prize in literature, and was a well-known social critic and anti-war activist.

The ethic of responsibility. Bertrand Russell's breadth was extraordinary. But the mere *existence* of the politically engaged intellectual doesn't suggest that this pairing is at all representative. To what extent *are* scientists and engineers socially engaged? And to what extent do societal norms demand that they be?⁹

Nowadays, an *ethic of responsibility* is preached in university courses and advocated by professional organizations. It is the doctrinal view. The putative norm contends that scientists and engineers have an obligation to select work that promotes the social good (a *positive right*), or, at the very least, to refrain from work that damages mankind or the environment (a *negative right*).¹⁰ The obligation stems from three basic truths: that the work of scientists and engineers transforms society; that this transformation can be for the better or for the worse; and that what we do is arcane enough that we bring an essential perspective to public discourse. The socially engaged scientist is expected to bring a normative

vision for how his field's work should impact society. He or she aims to try to steer things in that direction.

To be sure, decision making under the ethic of responsibility is not easy. It can be impossible to anticipate if a line of work is going to be used for good or for ill. Additionally, the for-good-or-for-ill dichotomy can be simplistic and subjective to the point of meaninglessness. Still, despite such difficulties, the socially engaged scientist is supposed to investigate, think, and decide what work he will or will not do, and what organizations he will or will not work for. The judgment should be made without over-valuing his own self-interest.

Historical events shaping the ethic of responsibility. The ascendancy of the ethic of responsibility was shaped by three historical events of World War 2 and its aftermath.

1. The first, already touched on, was the experience of the atomic scientists. After the war, with science left in a position both revered and feared, prominent physicists became public figures. Some became outspoken in their advocacy for peace, or their opposition to further weapons development. Recall the widespread concerns from physicists to Reagan's Strategic Defense Initiative (SDI)¹¹ or Hans Bethe's famous letter to Bill Clinton where he argued against another round of U.S. nuclear-weapons development.¹² A willingness to *speak truth to power*¹³ became a tradition among physicists—one that, I think, continues to shape physicists' identity.¹⁴

As an example, recall the pepper-spray incident of 2011 at my own campus, the University of California, Davis.¹⁵ Carrying out the Chancellor's instructions to clear "Occupy" protesters, police officer John Pike pepper-sprayed students who sat, arms linked, on the university's central quad. Videos of the event went viral,¹⁶ while memes of Officer Pike casually pepper-spraying *anything* became a second Internet sensation. But the observation I'd like to make is that, in the aftermath of the incident, the *only* UCD department outside the humanities to condemn the Chancellor or call for her resignation was Physics.¹⁷ The Chancellor was mystified. She understood the strong reaction from our (underfunded and politically liberal) English department, but she didn't anticipate complaints from a (well-funded and generally conservative) Physics department.¹⁸ What the Chancellor might not have internalized is that physicists retain a post-war legacy not only of snuggling up close to power, but also of nipping at its ankles.

2. A second historical event that helped shape the post-war view of moral responsibility was the Nuremberg trials (1945–1946). While the defense repeatedly proffered that the accused were simply following orders, this view was almost universally *rejected*: following orders did not efface legal or moral culpability. The Nuremberg trials began with the Medical Case, the prosecution of 23 scientists, physicians, and other senior officials for gruesome and routinely fatal medical experiments on prisoners.¹⁹

Years later, as though in sequel, the world would watch in nervous fascination the trial of Adolf Eichmann (1961). Hannah Arendt's controversial portrayal of Eichmann would come to be formative in shaping our understanding of

what, ethically, had transpired during the Holocaust. She wrote of the utter *ordinariness* of the man.²⁰ Arendt's book on the trial, memorably subtitled *The Banality of Evil*, would be published the same year (1963) as Stanley Milgram's classical experiments on obedience, where Milgram produced the stunning (and widely repeated) finding that a large fraction of volunteers would follow a white-coated scientist's gentle urging to deliver apparently life-threatening shocks to someone they thought was a fellow test subject.²¹

3. Finally, I would mention the rise of the environmental movement as contributing to the rise of an ethic of responsibility. While environmentalism dates to the mid-nineteenth century and before, as a significant social movement, the 1962 publication of Rachel Carson's *Silent Spring* is a milestone. Her book painted a picture of the end of life not by the drama of nuclear warfare, but the disappearance of songbirds, silenced by the routine if oversized activities of chemical manufacturing and non-specific pesticides.

The good scientist. The three experiences I have just described implied a *democratization of responsibility*. Scientist *had* to assume responsibility for what they did, for technology would take us to a very dark place if he did not. Stripped of ethical restraint, science would bring a world of nightmare bombs, gas chambers, and macabre human experiments. It would bring a dying, poisoned world.

And so, in the decades following the war, the ethic of responsibility became—at least rhetorically—the doctrinal norm. Increasing numbers of scientists and engineers, as well as their professional organization, began to engage on issues of social responsibility. The Pugwash Conferences began in 1955. The National Society of Professional Engineers adopted a code of ethics in 1964 that gave primacy to social responsibility. As its first imperative, the code says that “Engineers, in the fulfillment of their professional duties, shall hold paramount the safety, health, and welfare of the public.” Similar language would spread across other codes of ethics, including those of the ACM and IEEE.²² The Union of Concerned Scientists was formed at MIT in 1969—the same year a work stoppage at MIT, coordinated with 30 other universities, enjoyed substantial student, faculty, and administrative support. It called for a realignment of research directions away from military pursuits and towards human needs. Computer Professionals for Social Responsibility (CPSR) began its work opposing the SDI in 1983.²³ That same year, the IACR was founded, its self-described mission not only to advance the theory and practice of cryptology but also, lest we forget, to serve the public welfare.²⁴ The Electronic Frontier Foundation (EFF) and Privacy International (PI) were both formed in 1990, and became effective advocates in such matters as the defeat of the Clipper Chip. All of this is but a sampling of the overt politics from scientists and engineers.

Against this backdrop, the figure of the brilliant but humane scientist become a cultural motif. Jonas Salk had wiped out polio. Einstein became a cultural icon, one unfazed by the inconvenience of his death. The image of him sticking out his tongue may be the most widely recognizable photograph of any scientist, ever.

Richard Feynman would be painted in equally colorful ways, the no-nonsense genius pounding on bongo drums and shoving black rubbery stuff into ice water. Gene Roddenberry's *Star Trek* imagined a future that featured the scientist-humanist-hero as one team, if not one individual. Carl Sagan, speaking gently to the camera in episodes of *Cosmos* (1980), seemed the real-life embodiment of this aspirational package.

The ethic of responsibility in decline. And yet, for all I have said, the scientist or engineer seriously concerned about the social impact of his work is, I think, so rare as to be nearly a matter of myth. Never during the cold war, nor in any of the subsequent US wars, did US companies have difficulty recruiting or retaining the hundreds of thousands of scientists and engineers engaged in building weapons systems.²⁵ Universities like my own were happy to add their support; the University of California would, for decades, run the USA's nuclear weapons design laboratories.²⁶ In nearly 20 years advising students at my university, I have observed that a wish for *right livelihood*²⁷ almost never figures into the employment decisions of undergraduate computer science students. And this isn't unique to computer scientists: of the five most highly ranked websites I found on a Google search of *deciding among job offers*, not one suggests considering the institutional goals of the employer or the social worth of what they do.²⁸

Nowadays I ask computer-science faculty candidates to explain their view on the ethical responsibilities of computer scientists. Some respond like a deer in headlights, unsure what such a question could even mean. One recent faculty candidate, a data-mining researcher whose work seemed a compendium of DoD-funded projects for socially reprehensible aims, admitted that she felt no social responsibility. "I am a body without a soul," she earnestly explained. It was sincere—and creepy.

Stanley Fish, a well-known literary theorist, professor, and dean, admonishes faculty *not* to pursue research programs rooted in values. (His 2012 book is titled *Save the World on Your Own Time*.) Fish advises professors to

do your job; don't try to do someone else's job . . . ; and don't let anyone else do your job. In other words, don't confuse your academic obligations with the obligation to save the world; that's not your job as an academic . . .

Marx famously said that our job is not to interpret the world, but to change it. In the academy, however, it is exactly the reverse: our job is not to change the world, but to interpret it.²⁹

Perhaps such amorality, however revolting, is harmless in Fish's intellectual realm: one doesn't particularly expect literary theory to change the world. But scientists and engineers do just that. A refusal to direct the change we do is both morally bankrupt and ingracious. Our work as academics, we should never forget, is subsidized by society.³⁰

So far I have not said *why* the post-war ethic-of-responsibility didn't catch on. I could give multiple answers, starting with the rise of radical individualism.³¹ But I prefer to focus on something else: extreme technological optimism.

Technological optimism. Technological optimists believe that technology makes life better. According to this view, we live longer, have more freedom, enjoy more leisure. Technology enriches us with artifacts, knowledge, and potential. Coupled with capitalism, technology has become this extraordinary tool for human development. At this point, it is central to mankind's mission. While technology does bring some unintended consequences, innovation itself will see us through.

Technological optimism animates everyone from school children to Turing Award winners. Accepting his 2012 Turing Award, Silvio Micali, said that

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 of human knowledge . . . are benefiting from our conceptual and technical contributions. . . . Long live computer science!³²

If you're a technological optimist, a rosy future flows from the wellspring of your work. This implies a limitation on ethical responsibility. The important thing is to do the work, and do it well. This even becomes a *moral* imperative, as the work *itself* is your social contribution.

But what if computer science is *not* benefiting man? Technological pessimists like Jacques Ellul, Herbert Marcuse, and Lewis Mumford certainly didn't think that it was. They saw modern technology as an interlocking, out-of-control system that, instead of fulfilling human needs, engendered pointless wants and deadlier weapons. Man is becoming little more than the sex organs of the machine world.³³

Taking a less "extreme" view,³⁴ technological contextualists³⁵ acknowledge the concerns of the pessimists, but emphasize man's essential agency and the malleability of technology. Contextualism dominates the dialectic of technology studies.

The ethic of responsibility is always paired with the contextualist view of sociotechnology. At some level, this must be so: a normative need vanishes if, in the garden of forking paths, all paths lead to good (or, for that matter, to bad). But it is technological optimism that most people buy into, especially scientists and engineers. Unbridled technological optimism undermines the basic *need* for social responsibility.

Conclusion to part 1. Ultimately, I think the post-war turn towards social responsibility in science and engineering was less a turn than a sideways glance. While the rhetoric of responsibility would provide cover from technology's critics, few scientists or engineers would ever come to internalize that *their* work embodied socially relevant values. If researchers like us were actually supposed to know or care about this stuff in any operationally significant way, well, I think we didn't get the memo.

So let me retransmit it. It says that your moral duties extend beyond the imperative that you personally do no harm: you have to try to promote the social good, too. Also, it says that your moral duties stem not just from your stature

as a moral individual, but, also, from the professional communities to which you belong: cryptographer, computer scientist, scientist, technologist.

With few exceptions, the atomic scientists who worked on disarmament were not the same individuals as those who built the bomb. Their colleagues—fellow physicists—did that. Cryptographers didn't turn the Internet into an instrument of total surveillance, but our colleagues—fellow computer scientists and engineers—did that.

But you will only believe that claim if you recognize that cryptography *can* influence power relations. I suspect that many of you see no real connection between social, political, and ethical values and what you work on. You don't build bombs, experiment on people, or destroy the environment. You don't spy on populations. You hack math and write papers. This doesn't sound ethically laden. I want to show you that it is.

Part 2: The political character of cryptographic work

Scientist or spy? There's an irony in discussing the claim that cryptographic work is political, and it is this: to someone unconnected to the field, and also to the crypto-hobbyist, the claim may seem obviously true. But the young researcher who spends his life writing papers in cryptography, the claim may seem just as obviously false. What gives?

The outsider's view of cryptography might be based on cinematic portrayals. Films like *Sneakers* (1992), *Pi* (1998), *A Beautiful Mind* (2001), *Enigma* (2001), *Traveling Salesman* (2012), *Citizenfour* (2014), and *The Imitation Game* (2014) depict cryptography as a field intertwined with politics. Cryptographers are the brilliant and handsome mathematicians that power needs to have working on its side. We are, I am happy to report, heroic geniuses. A little crazy, to be sure, but that just adds to the luster.

Similarly, the *crypto hobbyist* may have read historical accounts dealing with cryptography, like the books of James Bamford or David Kahn.³⁶ Such accounts demonstrate that, historically, cryptography *is* about power. It's a realm in which governments spend enormous sums of money,³⁷ and maybe not unwisely: the work shapes the outcome of wars, and undergirds diplomatic and economic maneuvering.³⁸

Yet no academic cryptographer would confuse historical or fictional accounts of cryptography with what we actually do. Our discipline investigates academic problems that fall within our disciplinary boundaries. Pick up a Springer proceedings or browse ePrint papers and our field looks utterly *non*-political. If power is anywhere in the picture, it is in the abstract capacities of notional adversaries³⁹ or, in a different branch of our field, the power expenditure, measured in watts, for some hardware. We work on problems that strike us as interesting or scientifically important. We're not aiming to advance the interests of anything but science itself (or, perhaps, one's own career).

So distinct claims about cryptography's connectedness to power stem, at least in part, from radically different archetypes of what the cryptographer *is*: scientist

or spy. The NSA/GCHQ employee who hacks Gemalto to implant malware and steal SIM keys⁴⁰ is every bit as deserving of being called a “cryptographer” as the MIT-trained theorist who devises a new approach for functional encryption. Both are dealing in questions of privacy, communications, adversaries, and clever techniques, and we would do well to emphasize these commonalities if we want to see our disciplinary universe in context or nudge it towards greater relevance.

Academic cryptography used to be more political. The ascendance of a new cryptographer archetype—the academic cryptographer—fails to really explain our politically detached posture. For one thing, academic cryptographers were once more concerned with our field’s sociopolitical dimensions. Some even came to cryptography for such reasons. Consider, for example, this fragment of Whit Diffie’s testimony at the Newegg trial. Speaking of his wife, Diffie says:

I told her that we were headed into a world where people would have important, intimate, long-term relationships with people they had never met face to face. I was worried about privacy in that world, and that’s why I was working on cryptography.⁴¹

Diffie and his advisor, Martin Hellman, have long evinced a concern for sociopolitical problems touching technology. You see it in their criticism of DES’s key length,⁴² in Hellman’s activism on nuclear disarmament,⁴³ in Diffie’s book on the politics of wiretapping with Susan Landau,⁴⁴ and in his co-invention of forward secrecy.⁴⁵ You see it in the *New Directions* paper:⁴⁶ when the authors boldly begin “We stand today on the brink of a revolution in cryptography,” the anticipated revolution was not, at least primarily, the theory community bringing forth mind-altering notions of provable security, simulatability, or multiparty computation.⁴⁷ The authors were interested in technological changes that were transpiring, and concomitant social opportunities and needs.⁴⁸

Still more ostensibly political is David Chaum’s body of scientific work, which thoroughly embeds concerns for democracy and individual autonomy. Chaum’s 1981 paper⁴⁹ *Untraceable Electronic Mail, Return Addresses, and Digital Pseudonyms*, [Chaum81] suggests that a crucial privacy goal when sending an email is to hide who is communicating with whom. The *metadata*, in modern political parlance. The author offered *mix nets* for a solution.⁵⁰

Chaum would go on to provide the founding ideas for anonymous electronic cash and electronic voting. His papers would routinely draw on overtly political motivations.⁵¹ In a recent conversation, Chaum expressed surprise at the extent to which academics gravitated to a field—cryptography—so connected to issues of power.⁵²

Stripping out the politics. But as academics gravitated to cryptography, they tended to sanitize it, stripping it of ostensible connectedness to power. Applied and privacy-related work drifted outside of the field’s core venues, the IACR conferences. It is as though a chemical synthesis would take place, transforming this powerful powder into harmless dust.

Consider that there is now a conference named “Real World Cryptography” (RWC).⁵³ There is humor—but maybe gallows humor—that a field with a genesis

and capability as real-world as ours should find reason to create a venue so named.⁵⁴ Ask a colleagues in Graphics or Cloud Computing how it would fly in *their* community if someone started a conference called Real World Computer Graphics (RWCG 2015) or Real World Cloud Computing (RWCC 2016). They will laugh.

An especially problematic excision of the political is the marginalization within the cryptographic community of the *secure-messaging problem*,⁵⁵ an instance of which was the problem addressed by [Chaum81]. Secure-messaging is the most fundamental privacy problem in cryptography: how can parties communicate in such a way that nobody knows who said what. More than a decade after the problem was introduced, Rackoff and Simon would comment on the near-absence of attention being paid to the problem.⁵⁶ Another 20-plus years later, the situation is this: there is now a mountain of work on secure-messaging, but it's unclear what most of it actually *does*. A recent systemization-of-knowledge article⁵⁷ paints a picture of a cryptographic task enjoying a flourishing of *ad hoc* solutions, but little of it arising from the cryptographic community, as narrowly construed, or tied to much theory.⁵⁸ While one could certainly claim that this is true for almost all practical security goals that employ cryptography, I think the case is different for secure-messaging: here the work feels almost intentionally pushed aside.

Children of [Chaum81] and [GM82]. Why would I make such a claim? An illuminating case study is provided by comparing the venues of the most cited papers citing [Chaum81] and [GM82], Goldwasser and Micali's *Probabilistic Encryption*.⁵⁹ The two papers appeared around the same time and have comparable citation counts.⁶⁰

The [GM82] paper put forward the definitionally centered, reduction-based approach to dealing with cryptographic problems. It became a seminal work of the cryptographic community. The most cited papers that cite it appear in Crypto and Eurocrypt, FOCS, STOC, and ACM CCS.⁶¹ The [Chaum81] paper put forward the secure email problem and suggested a solution. This paper would be just as seminal—but in spawning work mostly outside the core cryptographic community. The ten most cited papers that cite Chaum's paper appear in venues that, mostly, I had never heard of. Venues not crypto-focused, like MobiSys and SIGOPS. In fact, the venues for the ten most cited papers citing [GM82] and the venues for the ten most cited papers citing [Chaum81] have void intersection. I find this is fairly remarkable. It reflects a research community split into fragments that include a GM-derived one and a Chaum-derived one, the second fragment not really being a part of the cryptographic community at all.⁶²

Why did this fragmentation occur? The most obvious explanation has to do with rigor: [GM82] offered a mathematically precise approach to its subject, while [Chaum81] did not. So a partitioning might seem to make sense: cryptographic work that can be mathematically formal goes to the right; ad hoc stuff, over to the left.

The problem with this explanation is that it's wrong. The [Chaum81] paper supports rigor just fine. Indeed provable security would eventually catch up to

mix nets, although the first definition would take more than 20 years to appear (2003), in a paper by Abe and Imai.⁶³ That the [Chaum81] paper itself didn't provide a formal treatment says nothing about the formalizability of the problem or what communities would later embrace it; after all, Diffie and Hellman's paper⁶⁴ only informally described trapdoor permutations, public-key encryption, and digital signatures, but all would be absorbed into the cryptographic fold. Now one might well counter that the problem addressed by [Chaum81] is more difficult to formalize than any of the examples just named. That's true. But it's simpler to formalize than MPC,⁶⁵ say, which would quickly gain *entrée* and stature in the cryptographic community—even *without* definitions or proofs. So, ultimately, neither formalizability nor complexity goes far to explain *why* secure-messaging has been marginalized.

A better answer (but by no means the only answer) is made obvious by comparing the introductions to the most-cited papers citing [GM82] and the most-cited papers citing [Chaum81]. Papers citing [GM82] frame problems scientifically. Authors claim to solve important technical questions. The tone is assertive, with hints of technological optimism. In marked contrast, papers citing [Chaum81] frame problems socio-politically. Authors speak about some social problem or need. The tone is reserved, and explicitly contextualist views are routine. One observes the exact same distinction in tone and stated motivations when comparing survey articles.⁶⁶

In 2015, I attended PETS (Privacy Enhancing Technologies Symposium) for the first time. Listening to people in this community interact is a bit like watching the cryptographic community through a lens that magically inverts most things. The PETS community attends closely to the values embedded in work. They care about artifacts that support human values. They aim to serve the users of those artifacts. They're deeply concerned with the politics and technology of surveillance. Where, after Chaum, did the moral soul of academic cryptography go? Maybe it moved to PETS.

There is a lesson in all this. Some might think that a community's focus is mostly determined by the technical character of the topic it aims to study. It is not. It is extra-scientific considerations that shape what gets treated where.

The cypherpunks. Now there is a group that has long worked at the nexus of cryptography and politics: the cypherpunks.⁶⁷ The cypherpunks emerged in the late 1980's, unified by a mailing list and some overlapping values. The core belief is that cryptography can be a key tool for protecting individual autonomy threatened by power.⁶⁸

The cypherpunks believed that a key question of our age was whether state and corporate interests would eviscerate liberty through electronic surveillance and its consequences, or if, instead, people would protect themselves through the artful use of cryptography. The cypherpunks did not seek a world of universal privacy: many wanted privacy for the individual, and transparency for the government and corporate elites. The cypherpunks envisioned that one could hack power relations by writing the right code. Cypherpunk-styled creations—think of Bitcoin, PGP, Tor, and WikiLeaks—were to be transformative because

they challenge authority and address basic freedoms: freedom of speech, movement, and economic engagement.⁶⁹

Exactly *how* such tools are to shape society is not always obvious. Consider WikiLeaks. The hope is not just that a better informed public will demand accountability and change. Rather, Assange sees governmental and corporate abuse as forms of conspiracy that could be throttled by the mere threat of leaks. Conspiracies are like graphs, the conspirators nodes, the pairwise relations among them, the edges. Instead of removing nodes or disrupting links, you can weaken *any* conspiracy by suffusing it in an ever-present threat of leaks. The more unjust the conspiracy, the more likely leaks will occur, and the more the damage they will do. As elites become fearful to conspire, they do so reservedly. The conspiratorial creature's blood thickens and it dies.⁷⁰ It is a fascinating vision.

It is cypherpunks, not cryptographers, who are normally the strongest advocates for cryptography. Julian Assange writes:

But we discovered something. Our one hope against total domination. A hope that with courage, insight and solidarity we could use to resist. A strange property of the physical universe that we live in.

The universe believes in encryption.

It is easier to encrypt information than it is to decrypt it.

We saw we could use this strange property to create the laws of a new world.⁷¹

Similarly, Edward Snowden writes:⁷²

In words from history, let us speak no more of faith in man, but bind him down from mischief by the chains of cryptography.⁷³

When I first encountered such discourse, I smugly thought the authors were way over-promising: they needed to tone down this rhetoric to be accurate. I no longer think this way. More engaged in implementing systems than I'll ever be, top cypherpunks understand more than I about insecure operating systems, malware, programming bugs, subversion, side channels, poor usability, small anonymity sets, and so on. Cypherpunks believe that *despite* such obstacles, cryptography can *still* be transformative.

Cryptography favors whom? Cypherpunk discourse seems sometimes to assume that cryptography will benefit ordinary people. But one has to be careful here. Cryptography can be developed in directions that tend to benefit the weak *or* the powerful. It can also be pursued in ways likely to benefit nobody but the cryptographer. Let's look at some examples.

Encryption. One reason people might assume cryptography to benefit the weak is that they're thinking of cryptography as encryption. Individuals with minimal resources can encrypt plaintexts in a manner that even a state-level adversary, lacking the key, won't be able to decrypt.

But does it necessarily come out that way? To work, cryptographic primitives must be embedded into systems, and those systems can realize arrangements of

power that don't trivially flow from the nature of the tool. In his typically pithy way, Schneier reminds people that "encryption is just a bunch of math, and math has no agency."⁷⁴ If a content-provider streams an encrypted film to a customer who holds the decryption key locked within a hardware or software boundary she has no realistic ability to penetrate,⁷⁵ we've empowered content providers, not users. If we couple public-key cryptography with a key-escrow system that the FBI and NSA can exploit, we empower governments, not people.⁷⁶

All that said, I do believe it accurate to say that conventional encryption *does* embed a *tendency* to empower ordinary people. Encryption directly supports freedom of speech. It doesn't require expensive or difficult-to-obtain resources. It's enabled by a thing that's easily shared. An individual can refrain from using backdoored systems.⁷⁷ Even the customary language for talking about encryption suggests a worldview in which ordinary people—the world's Alices and Bobs—are to be afforded the opportunity of private discourse. And coming at it from the other direction, one has to *work* to embed encryption within an architecture that props up power, and one may encounter major obstacles to success. The Clipper Chip completely failed. Trusted Computing mostly did.⁷⁸

IBE.⁷⁹ What about identity-based encryption, IBE? The setting was proposed by Shamir, with Boneh and Franklin, years later, providing a satisfying, provably secure realization.⁸⁰ The aim is to allow a party's email address, for example, to serve as his public key. So if Alice already knows the email address for Bob, she won't need to obtain his public key to send him an encrypted message: she just encrypts under Bob's email address.

But this convenience is enabled by a radical change in the trust model: Bob's secret key is no longer self-selected. It is issued by a trusted authority. That authority knows everyone's secret key in the system. IBE embeds key escrow—indeed a form of key escrow where a single entity implicitly holds all secret keys—even ones that haven't yet been issued. And even if you *do* trust the key-generating authority, a state-level adversary now has an extremely attractive locus to subpoena or subvert. In the end, from a personal-privacy point of view, IBE might seem like an enormous leap backwards.

Descriptions of IBE don't usually emphasize the change in trust model.⁸¹ And the key-issuing authority seems never to be named anything like that: it's just the PKG, for Private Key Generator. This sounds more innocuous than it is, and more like an algorithm than an entity. In papers, the PKG further recedes from view because it is the tuple of algorithms, not the entities that one imagines to run them, that grounds formal definitions and proofs.

To be clear, I am not condemning IBE as some sort of fascist technology. That sounds silly. Nor am I suggesting that IBE can't be refined in ways to make the trust model less authoritarian.⁸² Yet one can easily see the authoritarian tendency built into IBE. And technologies, while adaptable, are not infinitely so. As they evolve, they tend to retain their implicit orientations.

Differential privacy. Let's consider differential privacy.⁸³ Dwork says that ϵ -differential privacy "addresses concerns that any participant might have about

the leakage of her personal information: even if the participant removed her data from the data set, no outputs . . . would become significantly more or less likely.”⁸⁴

At some level, this sounds great: don’t we *want* to protect individuals from privacy-compromising disclosures from corporate or governmental datasets? But a more critical and less institutionally friendly perspective makes this definitional line seem off.⁸⁵ Most basically, the model implicitly paints the database owner (the curator) as the good guy, and the users querying it, the adversary. If power would just agree to fudge with the answers in the right way, it would be fine for it to hold massive amounts of personal data about each of us. But the history of data-privacy breaches suggests that the principle threat to us is from the database owner itself, and those that gain wholesale access to the data (for example, by theft or secret government programs). Second, the harm differential-privacy seeks to avoid is conceived of in entirely individualistic terms. But privacy violations harm entire communities. The individualistic focus presupposes a narrow conception of privacy’s value. Finally,⁸⁶ differential privacy implicitly presupposes that the data collection serves some public *good*. But, routinely, this is a highly contestable claim. The alternative of less data collection, or no data collection at all, is rarely even mentioned. In the end, one must compare the reduction in harm actually afforded by using differential privacy with the increase in harm afforded by corporations having another means of whitewash and policy-makers believing, quite wrongly, that there is some sort of cryptomagic to protect people from data misuse.

I recently asked an expert in differential privacy, Ilya Mironov, for his reaction to my harsh critique. He explained that the abstract roles in differential-privacy settings need not correspond to business relationships in the obvious way. For example, a privacy-conscious organization might choose to make its own analysts access sensitive data through a API that provides some differential-privacy guarantee, effectively treating its own employees as “adversaries.” Mironov also explained that there are variant notions of differential privacy that do *not* implicitly regard the database owner as good, and those querying it as bad. He described differential privacy in the *local model*,⁸⁷ where everyone keeps his data to himself. They can distributively compute the responses to queries. Fundamentally, Mironov explained, the definition of differential privacy is agnostic to the data model.

While everything explained makes good sense, I don’t think it changes the landscape. No actual mechanism can be agnostic to what data resides where. And at the point when a data-mining architecture and mechanism is laid down, considerations of efficiency, familiarity, and economics—not to mention authorities’ fundamental desire to have and to hold the data—make it easy to predict what will happen: almost always, a centralized design will emerge. To me, differential privacy may be as authoritarian in its conceptual underpinnings as IBE.

FHE and iO. Ever since Craig Gentry’s groundbreaking work,⁸⁸ fully homomorphic encryption (FHE) has been a target of enormous intellectual capital.

In brief, FHE allows you to outsource your data, encrypted under a public key of your choosing, to a service provider. Later, you can ask that party whatever you'd like about your plaintext. The service provider computes the encrypted answer, not knowing what it means. This is returned to you for decryption.

From a political perspective, FHE sounds empowering—even utopian. The powerful party, say a cloud service provider, is denied access to your data. You sidestep the Faustian bargain that routinely underlies cloud computing.⁸⁹

But the analysis above is specious. It is quite speculative if FHE will ever evolve into something practically useful. If you want to assess the political leanings of something of speculative utility, you shouldn't just assume that it will give rise to the touted applications, and then try to see who would win and who would lose. It's too conjectural. It is better to focus on how the pursuit changes us in the here and now.

And on that, I would say that FHE, along with iO,⁹⁰ have engendered a new wave of exuberance. In grant proposals, media interviews, and talks, leading theorists speak of FHE and iO as game-changing indications of where we have come.⁹¹ Nobody seems to emphasize just how speculative it is that any of this will ever have any impact on practice. Nor do people emphasize our vanishing privacy, our lousy computer security, or how little modern cryptography has really done to change this landscape. And this has consequences. (a) It misleads the public about where we stand. (b) It shifts financial resources away from areas more likely to have social utility. (c) It encourages bright young researchers to work on highly impractical directions. (d) And it provides useful cover to the strongest opponents of privacy: defense and intelligence agencies.

Let me expand on the last claim. Here is what DARPA Program Director Dan Kaufman had to say about FHE in a 2014 interview:

Imagine a future that says: OK, I have to collect everything for big data to work because if I knew what wasn't relevant it wouldn't be big data. But I don't want the government to just willy-nilly look through my emails: that feels creepy. . . .

So this guy, Craig Gentry, . . . showed that you could . . . take a piece of data, encrypt it, send it down the wire, never decrypt it, [but] still perform [computation] . . . on it. It sounds crazy, except he showed you can do it

You could imagine the following: . . . [Organizations] collect . . . data but only in . . . encrypted form Now let's say you believe there is a bad guy hiding somewhere in this encrypted data. So, I come up with a bunch of search terms I could then go to a court . . . [and they] could say "yeah, that looks reasonable." I put the search into the engine but . . . all that comes out is a number: how many people meet that criteria . . . You go back to the FISA court, and say O.K. guys, we have 12. . . . I picture FISA putting in a key, and then the Agency putting in a key, and they both turn it. And [at] that point, for the first time, . . . are those 12 now revealed.⁹²

Of course, it's utter nonsense. To begin with, there's no way to make sense of who holds what key and what data for FHE to even apply. We're also told: that we need to collect everything because, if we didn't, we wouldn't have enough data to have lots of data; that the government will be careful, as it would be "creepy"

if they weren't; that they'll get court orders—even, apparently, to discover the number of people in datasets who satisfy some specified search criteria; and that to get personally identifiable information, they'll need to have the cooperation of the NSA *and* the FISA court.

Kaufman's inchoate interview is but a tiny patch of discourse from an ocean of misdirection on privacy. It doesn't impugn FHE, but it does suggest how power aims to use such work: to let them mumble words that sound privacy-friendly. Providing strong funding for FHE and iO provides risk-free political cover. It supports a storyline that cloud storage and computing is safe. It helps entrench favored values within the cryptographic community: speculative, theory-centric directions. And it helps keep harmless academics who could, if they got feisty, start to innovate in more sensitive directions.

Unthreateningly engaged. Of course it hasn't escaped the notice of intelligence agencies that the vast majority of the academic cryptographic community is unthreateningly engaged. In a declassified trip-report about Eurocrypt 1992, the NSA author opines, for example:⁹³

There were no proposals of cryptosystems, no novel cryptanalysis of old designs, even very little on hardware design. I really don't see how things could have been better for our purposes.

The NSA's newsletter in which this report appears would never again mention that academic cryptographic community.⁹⁴ Nor did any released Snowden-derived document discuss anything of our community.⁹⁵ It's as though we progressed from a band of philosophers⁹⁶ worth a few pages of snarky commentary⁹⁷ to an assemblage too insignificant even for that.

Conclusion to part 2. A 2013 essay by Arvind Narayanan suggests a simple taxonomy for cryptographic work:⁹⁸ there's *crypto-for-security* and *crypto-for-privacy*. Crypto-for-security is crypto for commercial purposes. It's the crypto in TLS, payment cards, and cell phones. Crypto-for-privacy has social or political aims. Here the author distinguishes between *pragmatic crypto*—which is about trying to use cryptography to retain our predigital privacy—and *cypherpunk crypto*—the grander hope of using cryptography to precipitate sweeping social or political reforms. The author suggests that crypto-for-security has done well, but crypto-for-privacy has fared badly.

I think Narayanan's division is illuminating, but he fails to mention that most academic cryptography isn't really crypto-for-security *or* crypto-for-privacy: it is, one could say, *crypto-for-crypto*—meaning that it doesn't ostensibly benefit commerce *or* privacy, and it's quite speculative if it will ever evolve to do either. Perhaps every field eventually becomes primarily self-referential. Maybe this is even necessary, to some extent. But for cryptography, much is lost when we become so inward-looking that almost nobody is working on problems we *could* help with that address some basic human need. Crypto-for-crypto starves crypto-for-privacy, leaving a hole, both technical and ethical, in what we collectively do.

Part 3: The dystopian world of pervasive surveillance

Mass surveillance has motivated the contents of this essay, but is it so serious a thing? Before the Snowden revelations,⁹⁹ I myself didn't really think so. Environmental problems seemed more threatening to man's future, and my country's endless wars seemed more deserving of moral consternation. It wasn't until Snowden that I finally internalized that the surveillance issue was grave, was closely tied to our values and our profession, and was being quite misleadingly framed.

Law-enforcement framing. The *framing* of mass surveillance determines what one thinks it is about.¹⁰⁰ And mass surveillance has been brilliantly framed by power so as to slant discourse in a particular and predictable direction. Let me describe what I'll call the *law-enforcement framing*, as regularly communicated by people like (U.S.) FBI Director James Comey:¹⁰¹

1. Privacy is *personal* good. It's about your desire to control personal information about you.
2. Security, on the other hand, is a *collective* good. It's about living in a safe and secure world.
3. Privacy and security are inherently in conflict. As you strengthen one, you weaken the other. We need to find the right *balance*.
4. Modern communications technology has destroyed the former balance. It's been a boon to privacy, and a blow to security. Encryption is especially threatening. Our laws just haven't kept up.¹⁰²
5. Because of this, *bad guys* may win. They are terrorists, murderers, child pornographers, drug traffickers, and money launderers.¹⁰³ The technology that we good guys use—the bad guys use it too, to escape detection.
6. At this point, we run the risk of Going Dark.¹⁰⁴ Warrants will be issued, but, due to encryption, they'll be meaningless. We're become a country of unopenable closets. Default encryption may make a good marketing pitch, but it's reckless design. It will lead us to a very dark place.

The narrative is inconsistent with the history of intelligence gathering, and with the NSA's own mission statement.¹⁰⁵ Yet the narrative's uneasy coexistence with reality hasn't mattered. It is, in fact, beautifully crafted to frame matters in a way guaranteed to lead discourse to where power wants it to go. It is a brilliant discourse of fear: fear of crime; fear of losing our parents' protection; even fear of the dark. The narrative's well-honed deceptiveness is *itself* a form of tradecraft.¹⁰⁶

Surveillance-studies framing. Of course there are radically different ways to frame mass surveillance. Consider the following way to do so, which follows often-heard thoughts from cypherpunks and surveillance studies.¹⁰⁷

1. Surveillance is an instrument of *power*.¹⁰⁸ It is part of an apparatus of control. Power need not be in-your-face to be effective: subtle, psychological, nearly invisible methods can actually be more effective.

2. While surveillance is nothing new, technological changes have given governments and corporations an unprecedented capacity to monitor everyone's communication and movement. Surveilling everyone has become cheaper than figuring out whom to surveil, and the marginal cost has dropped.¹⁰⁹ The Internet, once seen by many as a tool for emancipation, has been transformed into the most dangerous facilitator for totalitarianism ever seen.¹¹⁰
3. Governmental surveillance is strongly linked to cyberwar. Security vulnerabilities that enable one enable the other. And, at least in the USA, the same individuals and agencies handle both jobs. Surveillance is also strongly linked to conventional warfare. As Gen. Michael Hayden has explained, "we kill people based on metadata."¹¹¹ Surveillance and assassination by drones are one technological ecosystem.
4. The law-enforcement narrative is wrong to position privacy as an individual good when it is, just as much, a social good. It is equally wrong to regard privacy and security as conflicting values, as privacy *enhances* security as often as it rubs against it.
5. Mass surveillance will tend to produce uniform, compliant, and shallow people.¹¹² It will thwart or reverse social progress. In a world of ubiquitous monitoring, there is no space for personal exploration, and no space to challenging social norms, either. Living in fear, there is no genuine freedom.
6. But creeping surveillance is hard to stop, because of interlocking corporate and governmental interests.¹¹³ Cryptography offers at least some hope. With it, one might carve out a space free of power's reach.

History teaches that extensive governmental surveillance becomes political in character. As civil-rights attorney Frank Donner and the Church Commission reports thoroughly document, domestic surveillance under U.S. FBI director J. Edgar Hoover served as a mechanism to protect the status quo and neutralize change movements.¹¹⁴ Very little of the FBI's surveillance-related efforts were directed at law-enforcement: as the activities surveilled were rarely illegal, unwelcome behavior would result in sabotage, threats, blackmail, and inappropriate prosecutions, instead. For example, leveraging audio surveillance tapes, the FBI attempted to get Dr. Martin Luther King, Jr., to kill himself.¹¹⁵ U.S. universities were thoroughly infiltrated with informants: selected students, faculty, staff, and administrators would report to an extensive network of FBI handlers on anything political going on on campus. The surveillance of dissent became an institutional pillar for maintaining political order. The U.S. COINTELPRO program would run for more than 15 years, permanently reshaping the U.S. political landscape.¹¹⁶

Our dystopian future.¹¹⁷ Where mass surveillance leads has been brilliantly explored in fictional accounts, starting with Yevgeny Zamyatin's 1921 novel *We* (which inspired Orwell's *1984*). Set in a future of total surveillance, the denizens of the "One State" have internalized lessons such as: "we" is from God, and "I" is from the devil; that imagination is illness; and that the key to ridding man of crime is ridding him of freedom.

You don't have to reach to fictional or historical accounts to anticipate where we are headed. In a 2012 newsletter column, NSA's "SIGINT Philosopher," Jacob Weber, tells us his vision. After failing an NSA lie-detector test, he says:

I found myself wishing that my life would be constantly and completely monitored. It might seem odd that a self-professed libertarian would wish an Orwellian dystopia on himself, but here was my rationale: If people knew a few things about me, I might seem suspicious. But if people knew everything about me, they'd see they had nothing to fear. This is the attitude I have brought to SIGINT work since then.

We tend to mistrust what we do not understand well. A target that has no ill will to the U.S., but which¹¹⁸ is being monitored, needs better and more monitoring, not less.

So if we're in for a penny, we need to be in for a pound.¹¹⁹

Shrouded in enormous secrecy and complexity, the basic contours of the surveillance state are fundamentally unknowable. What is the individual to do? Millions of observations are made of his life. He is analyzed by techniques he cannot remotely understand. He knows that today's data, and yesterday's, will be scrutinized by tomorrow's algorithms. He knows that these algorithms will employ natural-language processing but that it probably won't *actually* understand natural language or human discourse. With all this, the rational individual has no choice but to desperately *try* to act like everyone else.

At its best, the film "Citizenfour" managed to sketch the shape of this world as it told the story of Edward Snowden. One reviewer writes of the film

evoking the modern state as an unseen, ubiquitous presence, an abstraction with enormous coercive resources at its disposal. . . .

It is everywhere and nowhere, the leviathan whose belly is our native atmosphere. Mr. Snowden, unplugging the telephone in his room, hiding under a blanket when typing on his laptop, looking mildly panicked when a fire alarm is tested on his floor, can seem paranoid. He can also seem to be practicing a kind of avant-garde common sense. It's hard to tell the difference, and [this] . . . can induce a kind of epistemological vertigo. What do we know about what is known about us? Who knows it? Can we trust them?¹²⁰

To be more prosaic: I pick up the phone and call my colleague, Mihir Bellare, or I tap out an email to him. How many copies of this communication will be stored, and by whom? What algorithms will analyze it—now and in the future? What other data will it be combined with in an attempt to form a picture of me? What would trigger a human analyst to get involved? Might my call or email contribute to a tax audit, a negative grant-funding decision, some Hoover-style dirty tricks, or even an assassination? There is not a single person who knows the answer to these questions, and those who know most aren't about to tell.

Conclusion to part 3. Ultimately, I'm not much interested in individual grievances over privacy; I am far more concerned with what surveillance does to society. Totalized surveillance vastly diminishes the possibility of effective political dissent. And without dissent, social progress is unlikely.

Consider an event like the 1971 burglary of the FBI branch office in Media, Pennsylvania.¹²¹ With the degree of surveillance we now live under, the whistleblowers—beginning with that feisty physics professor who led the effort¹²²—would be promptly arrested, and even charged with espionage. They would have spent years in prison, or even faced execution. Facing such outcomes and odds, the activists would not have attempted their daring burglary.

Less hypothetically, online and telephone surveillance already results in the imprisonment of political dissidents around the world,¹²³ and it undergirds my own country’s drone-assassination program.¹²⁴ In the U.S., Miami-model policing¹²⁵ has made attending political protests (or being near one in your car, or with your phone) an intimidating proposition. Is social progress even possible in such an environment?

But, despite all these arguments, I am skeptical about rationalist accounts of ethical affronts, be it mass surveillance or anything else. If we behave morally, it is not because of rational analyses, but an instinctual preference for liberty, empathy, or companionship.¹²⁶ As Schneier points out, animals don’t like to be surveilled because it makes them feel like prey, while it makes the surveillor feel like—and act like—a predator.¹²⁷ I think people know at an instinctual level that a life in which our thoughts, discourse, and interactions are subjected to constant algorithmic or human monitoring is no life at all. We are sprinting towards a world that we know, even without rational thought, is not a place where man belongs.

Part 4: Creating a more just and useful field

What can we cryptographers realistically do to collectively up our contribution to crypto-for-privacy? I claim no easy answers. I can offer only modest ideas.

Secure messaging in the untrusted-server model. Problem selection is the most obvious aspect in determining our community’s impact, and secure messaging, in all its forms, remains the most outstanding problem in crypto-for-privacy. While mix nets, onion routing, and DC nets have all proven to be highly useful,¹²⁸ it is not too late to be thinking on new architectures for secure communications.

Consider the following problem, which is inspired by Pond and the PANDA protocol that it can use.¹²⁹ The aim is similar to Adam Langley’s Pond protocol: to create an alternative to email or instant messaging but where “big brother” is unable to figure out who is communicating with whom. Unlike Pond, I don’t want to rely on Tor, for we seek security in the face of a global, active adversary (as well as a clean, provable-security treatment). Tor can always be layered on top, as a heuristic measure, to hide system participants.

The intent is this. Pairs of people who want to communicate are assumed to initially share a password. They won’t directly talk with one another; rather, all communications will go through an *untrusted* server. First, parties upgrade their shared password to a strong key with an *anonymous rendezvous* protocol.

Thereafter, the sender can deposit a (constant-length) encrypted message at the server. When a party wants to retrieve his i th message, he'll interact with the same server, which gives him a string computed from the database contents. The value permits the receiver to recover the intended message—or, alternatively, an indication that there is no such i th message for him. But, throughout, all the server ever sees are parties depositing random-looking strings to the server, and parties collecting random-looking strings from the server, these computed by applying some non-secret function to the server's non-secret database. Neither the server nor an active, global adversary can figure out who has communicated with whom, or even whether a communications has taken place. The goal is to do all this as efficiently as possible—in particular, much more efficiently than the server just handing each recipient its entire database of encrypted messages.

In ongoing work, colleagues and I are working out a provable-security treatment for the approach above. It uses conventional, game-based definition, not the fuzzy concepts or vocabulary from much of the anonymity literature.¹³⁰ We hope that the anonymous messaging in this untrusted-server model will eventually prove practical for the high-latency setting. We will see.

Bigkey cryptography Let me next describe some recent work by Mihir Bellare, Daniel Kane, and me that we call *bigkey cryptography*.¹³¹

The intent of bigkey cryptography is to allow cryptographic operations to depend on enormous keys—megabytes to terabytes long. We want our keys so long that it becomes infeasible for an adversary to exfiltrate them. Yet using such a bigkey mustn't make things slow. This implies that, with each use, only a small fraction of the bigkey's bits will be inspected.

The basic idea is not new: the concept is usually referred to as security in the *bounded retrieval model*.¹³² But our emphasis *is* new: practical and general tools, with sharp, concrete bounds. We have no objection to using the random-oracle model to achieve these ends.

Suppose you have a bigkey \mathbf{K} . You want to use it for some protocol P that has been designed to use a *conventional-length* key K . So choose a random value R (maybe 256 bits) and hash it to get some number p of *probes* into the bigkey:

$$i_1 = H(R, 1) \quad i_2 = H(R, 2) \quad \dots \quad i_p = H(R, p).$$

Each probe i_j points into \mathbf{K} : it's a number between 1 and $|\mathbf{K}|$. So you grab the p bits at those locations and hash them, along with R , to get a derived key K :

$$K = H'(R, \mathbf{K}[i_1], \dots, \mathbf{K}[i_p]) = \text{XKEY}(\mathbf{K}, R).$$

Where you would otherwise have used the protocol P with a shared key K , you will now use P with a shared bigkey \mathbf{K} , a freshly chosen R , this determining the conventional key $K = \text{XKEY}(\mathbf{K}, R)$.

We show that derived-key K is indistinguishable from a uniformly random key K' even if the adversary gets R and can learn lots of information about the bigkey \mathbf{K} . The result is quantitative, measuring how good the derived key is as

a function of the length of the bigkey, the number of bits leaked from it, the number of probes p , the length of R , and the number of random-oracle calls.

At the heart of this result is an information-theoretic question we call the *subkey-prediction problem*. Imagine a random key \mathbf{K} that an adversary can export $\ell < |\mathbf{K}|$ bits of information about. After that leakage, we select p random locations into \mathbf{K} , give those locations to the adversary, and ask the adversary to predict those p bits. How well can it do?

It turns out that the adversary *can* do better than just recording ℓ bits of the key \mathbf{K} and hoping that lots of probes fall there. But it can't do *much* better. Had nothing been leaked to the adversary, $\ell = 0$, then each probe would contribute about one bit of entropy to the random variable the adversary must guess. But if, say, half the key is leaked, $\ell \leq |\mathbf{K}|/2$, each probe will now contribute about 0.156 bits of entropy.¹³³ The adversary's chance of winning the subkey-prediction game will be bounded by something that's around $2^{-0.156p}$. One needs about $p = 820$ probes for 128-bit security, or twice that for 256-bit security.

I think that the subkey prediction problem, and the key-encapsulation algorithm based on it, will give rise to nice means for exfiltration-resistant authenticated-encryption and pseudorandom generators.¹³⁴ In general, I see bigkey cryptography as one tool that cryptographers can contribute to make mass surveillance harder.

More examples. Here are a few more examples of crypto-for-privacy work.

Consider the beautiful paper on **Riposte**, by Corrigan-Gibbs, Boneh, and Mazières.¹³⁵ A user, speaking with others on the Internet, wants to broadcast a message, such as a leaked document, without revealing his identity. The network is subject to pervasive monitoring. The authors develop definitions, protocols, and proofs for the problem, attending closely to efficiency.¹³⁶ They implement their schemes. Combining all these elements is rare—and very much needed.¹³⁷

Or consider the insightful work of Colin Percival in which he introduced **scrypt**.¹³⁸ Percival explained that, when applying an intentionally slow-to-compute hash function to a password and salt so as to up the cost of dictionary attacks,¹³⁹ it is better if the hash function can't be sped up all that much with custom hardware. It shouldn't just take lots of time, but lots of memory, too. Quite recently, a Password Hashing Competition (PHC) concluded having chosen a scheme, **Argon2**,¹⁴⁰ that follows this lead. Meanwhile, the theory for this sort of hash function has nicely progressed.¹⁴¹ While we don't yet have good bounds on schemes like scrypt and Argon2, I think we're getting there.¹⁴²

Or consider the paper on the susceptibility of symmetric encryption to mass surveillance by colleagues and me.¹⁴³ We discussed **algorithm-substitution attacks**, wherein big brother replaces a *real* symmetric encryption algorithm by a *subverted* one. Big brother's aim is to surreptitiously decrypt all encrypted traffic. The idea goes back to Young and Yung;¹⁴⁴ all we did was to rigorously explore the idea in the context of symmetric encryption. Yet what we found was disturbing: that almost all symmetric encryption schemes can be easily subverted. Still, we showed that it is easy to make schemes where this isn't true.

And then there's the **Logjam** paper, showing, for the umpteenth time, that we must watch out for high cryptanalytic value of precomputation.¹⁴⁵ Attacks should routinely be regarded as a two-step process: an expensive one that depends on widely shared parameters, then a cheaper, individualized attack.¹⁴⁶ Such thinking goes back to early time-memory tradeoffs,¹⁴⁷ and to many cryptographer's preference for nonuniform adversaries. It occurs in practical work, as in attacks on A5/1 in GSM phones.¹⁴⁸ And it is also the model that intelligence agencies seem to gravitate to, as suggested by the NSA's attack on FPE scheme FF2¹⁴⁹ and the fact that they regarded this attack as serious.

Choose well. As I hope the examples I have given illustrate, there are important crypto-for-privacy problems out there, and they are quite diverse. Choose your problems well. Let values inform your choice. Many times I have spoken to people who seem to have no real idea *why* they are studying what they are. The real answer is often that they can do it, it gets published, and that people did this stuff before. These are lousy reasons for doing something.

Introspection can't be rushed. In the rush to publish paper after paper, who has the time? I think we should breathe, write fewer papers, and have them matter more.

▷ *Attend to problems' social value. Do anti-surveillance research.*

▷ *Be introspective about why you are working on the problems you are.*

In enumerating example directions for anti-surveillance research, I didn't include the kind of work, rather common in the PET (privacy-enhancing technology) literature, that assumes that there *will* be pervasive collection, and then tries to do what one can to minimize misuse.¹⁵⁰ Since the immorality occurs at the point of data collection, the aim here is to try to blunt the impact of the wrong already done. But it is hard to know how this plays out. I am concerned that the work can play into the hands of those who seek technical support for a position that says, in effect, "the collect-it-all approach is inevitable and only temporarily problematic, for, once we figure this all out, privacy will be handled downstream, when the data is used." But pervasive collection *itself* chills free-speech and threatens liberal democracy, regardless of what one claims will happen downstream.¹⁵¹

Practice-oriented provable security. It's not just the topics we work on, but how we execute on them that shapes our field's direction. For nearly 25 years Mihir Bellare and I have developed that we call *practice-oriented provable security*. In a 2009 essay and talk,¹⁵² I discussed how various inessential choices engendered a theory of cryptography that was less useful than necessary. Today, I might number among the important historical choices **(1)** a preference for asymptotic analyses and theorems, and the correspondingly coarse conceptualizations of security with which this is paired; **(2)** a preference towards minimalism, aesthetically construed, as a starting point for reductions; **(3)** the dismissal of symmetric primitives and finite functions as targets of rigorous inquiry; **(4)** a tradition of using nonconstructive language for stating

results; **(5)** the marginalization of secure messaging; and **(6)** a condemnatory attitude towards the random-oracle model, the random-permutation model, the ideal-cipher model, Dolev-Yao models,¹⁵³ and any other model deemed “non-standard”.

Practice-oriented provable security inverts such choices. It retains provable-security’s focus on definitions and proofs, but these are understood as tools that earn their value mostly by their utility to security or privacy. The approach is equally at home in those two realms, but it has been underused for privacy problems like secure messaging. Better treating mix-nets and onion routing is an obvious place to start, which students and I are doing.

▷ *Apply practice-oriented provable security to anti-surveillance problems.*

Funding.¹⁵⁴ In the United States, it would seem that the majority of extramural cryptographic funding may now come from the military.¹⁵⁵ From 2000 to 2010, fewer than 15% of the papers at CRYPTO that acknowledged U.S. extramural funding acknowledged DoD funding.¹⁵⁶ In 2011, this rose to 25%. From 2012 to 2015, it rose to 65%.¹⁵⁷ Nowadays, many cryptographers put together a large patchwork of grants, the largest of which are usually DoD. The following funding acknowledgment isn’t so very atypical:

This work was supported by NSF, the DARPA PROCEED program, an AFOSR MURI award, a grant from ONR, an IARPA project provided via DoI/NBC, and by Samsung.¹⁵⁸

The military funding of science invariably redirects it¹⁵⁹ and creates moral hazards.¹⁶⁰ Yet suggesting to someone that they might want to reconsider their taking DoD funding may anger even a placid colleague, for it will be perceived as an assault both on one’s character and his ability to succeed.

No matter what people say, our scientific work *does* change in response to sponsor’s institutional aims. These aims may not be one’s own. For example, the mission of DARPA is “to invest in the breakthrough technologies that can create the next generation of [U.S.] national security capabilities.” Having begun in the wake of Sputnik, the agency speaks of avoiding *technological surprise*—and creating it for America’s enemies.¹⁶¹ In the USA, the NSA advises other DoD agencies on crypto-related grants. At least sometimes, they advise the NSF. Back in 1996, the NSA tried to quash my own NSF CAREER award. I learned this from my former NSF program manager, Dana Latch, who not only refused the NSA request, but, annoyed by it, told me. An internal history of the NSA reports on the mistake of theirs that allowed funding the grant leading to RSA.

NSA had reviewed the Rivest [grant] application, but the wording was so general that the Agency did not spot the threat and passed it back to NSF without comment. Since the technique had been jointly funded by NSF and the Office of Naval Research, NSA’s new director, Admiral Bobby Inman, visited the director of ONR to secure a commitment that ONR would get NSA’s coordination on all such future grant proposals.¹⁶²

People are often happy to get funding, regardless of its source. But I would suggest that if a funding agency embraces values inconsistent with your own, then maybe you shouldn't take their money. Institutions *have* values, no less than men. Perhaps, in the modern era, they even have more.

Large organizations have multiple and sometimes conflicting aims. Military organizations with offensive and defensive roles in cybersecurity have COIs built into their design. Individuals are wrong to assume that their work is non-military work errantly funded by the military.

In his farewell address of 1961, President Dwight D. Eisenhower introduced the phrase, and concept, of the military-industrial complex. In an earlier version of that speech, Eisenhower tellingly called it the military-industrial-*academic* complex.¹⁶³ If scientists wish to reverse our complicity in this convergence of interests, maybe we need to step away from this trough.

None of this was clear to me when I first joined the university. A few years ago I joined in on a DoD grant proposal (fortunately, unfunded), which I would not do today. It took me a long time to realize what eventually became obvious to me: that the funding we take both impacts our beliefs and reflects on them.

In the end, a major reason that crypto-for-privacy has fared poorly is that funding agencies don't want to see progress in this direction,¹⁶⁴ and most companies don't want progress here, either. Cryptographers have internalized this. Mostly, we've been in the business of helping business and government keep things safe. Governments and companies have become our "customers," not some ragtag activists, journalists, or dissidents, and not some abstract notion of *the people*. Crypto-for-privacy will fare better when cryptographers stop taking DoD funds and, more than that, start thinking of a very different constituency for our output.

▷ *Think twice, and then again, about accepting military funding.*¹⁶⁵

▷ *Regard ordinary people as those whose needs you ultimately aim to satisfy.*

Academic freedom. Many of us here are academics at universities, where there is a long tradition of *academic freedom*. This refers to your right—and even obligation—to think about, speak about, and write about whatever you want that is connected to your work, even if it goes against the wishes of power: your university, corporations, or the state. While academic freedom seems to be in decline,¹⁶⁶ at least for now, it recognizably persists.

Normally, scientists and other academics don't actually need or use their academic freedom: all they really need is technical skills and funding.¹⁶⁷ But crypto-for-privacy may be a rare topic where academic freedom *is* useful.¹⁶⁸ I suggest that people use this gift. Unexercised, academic freedom will wither and die.

Many nonacademics also have something akin to academic freedom: sufficient autonomy to work on what they think is important, without losing their jobs, even if it's not what their employer really wants or likes.

▷ *Use the academic freedom that you have.*

Against dogma. I think that many cryptographers would do well to foster a more open-minded attitude to unfamiliar models, approaches, and goals. The disciplinary narrowing within cryptography’s tier-1 venues has been pronounced.¹⁶⁹ Many people seem to hold rather strident beliefs about what kinds of work are *good*. Sometimes it borders on silliness, as when people refuse to use the word *proof* for proofs in the random-oracle model. (Obviously a proof in the random-oracle model is no less a proof than a proof in any other model.)

As cryptographers, we must always be sensitive, and skeptical, about the relationship between our models and *actual* privacy or security. This doesn’t mean that we should not take models seriously. It means that should see them as tentative and dialectical. There’s a lovely aphorism from statistician George Box, who said that *all models are wrong, but some are useful*.¹⁷⁰

Cryptography needs *useful* models. But the assessment of a model’s utility is itself problematic. We ask of definitions: How clean? How understandable? How general? What aspects of the computing environment are covered? What does and doesn’t it imply? The definitional enterprise sits at a juncture of math, aesthetics, philosophy, technology, and culture. So situated, dogma is disease.

It has been claimed that the mission of theoretical cryptography is to define and construct provably secure cryptographic protocols and schemes.¹⁷¹ But this is an activity of theoretical cryptography, not its mission. There are many other activities. One might work on models and results that are completely rigorous but fall outside of the provable-security framework.¹⁷² Or one can take an important protocol as fixed and then analyze it, in whatever framework works best. The aim for my own work has been to develop ideas that I hope will contribute to the construction of secure computing systems. In the symbology of Amit Sahai’s lovely flower-garden,¹⁷³ theory-minded cryptographers can be gardeners, growing seeds (hardness assumptions) into flowers (cryptographic goals); but they can do many other things as well. Which is fortunate, as cryptographic practice hasn’t benefited all that much from our horticultural activities.

▷ *Be open to diverse models. Regard all models as suspect and dialectical.*

A more expansive view. I would encourage cryptographers—especially young people in our field—to try to get a systems-level view of what is going on when cryptography is used. You need a way better view of things than a technophobe like me will ever have.

I remember reading that 2012 paper of Dan Boneh and his coauthors, *The Most Dangerous Code in the World*,¹⁷⁴ and feeling humbled by the fact that there was this entire universe of code—this *middleware*—that I didn’t even know *existed*, but that could, and routinely did, annul the cryptography that was there. When the NSA revelations caused people to speculate as to how Internet cryptography was being defeated, it occurred to me that perhaps the NSA didn’t need any clever cryptanalysis—what they needed, most of all, was to buy exploits and hire people with a systems-level view of the computing ecosystem.

One approach that might be useful for gaining a good vantage is to take an API-centric view of things.¹⁷⁵ Not only are API-misunderstandings a common

security problem, but gaps between cryptographic formalizations and APIs can produce serious cryptographic problems.¹⁷⁶ And in the constructive direction, the notion of online-AE, for example,¹⁷⁷ effectively flows from taking an API-centric view. APIs and “serious” cryptography need stronger bonds.

Research communities have a general tendency to become inward-looking. As a community, we have fostered strong relationships to algorithms and complexity theory, but have done less well attending to privacy research, programming languages, or the law. We will play a larger social role if we up our connections to neighbors.

I recently saw a nice talk by Chris Soghoian in which he described his frustration in trying to get media to report on, or anyone else to care about, the well-known fact (that is actually *not* well known) that cell-phone conversations have essentially no privacy.¹⁷⁸ Cryptographers should be helping with such communications. But I wonder how much we have even paid attention. For most of us, if it’s not what one’s working on, one doesn’t really care. There isn’t time.

▷ *Get a systems-level view. Attend to that which surrounds our field.*

Learn some privacy tools. I would like to gently suggest that we cryptographers would do well to learn, and use, contemporary privacy tools. Very few of us use tools like OTR, PGP, Signal, Tails, and Tor. It’s kind of an embarrassment—and I suspect our collective work suffers for it. Christopher Soghoian insightfully remarks: “It’s as if the entire academic medical community smoked 20 cigarettes a day, used intravenous drugs with shared needles, and had unprotected sex with random partners on a regular basis.”¹⁷⁹

I’m a bizarre person to advocate in this direction—it’s definitely a case of the pot calling the kettle black. I am dispositionally disinterested in using technology, and am incompetent at doing so if I try. I don’t even own a smartphone. Yet I suspect that there is nothing like experience to motivate cryptographers to identify and solve the privacy problems that will help us to transform hard-to-used tools for nerds into transparently embedded mechanisms for the masses. The first problem I suggested in Section 4 is something I thought of within days of starting to use Pond.

▷ *Learn some privacy tools. Use them. Improve them.*

No cutesy adversaries. There is a long tradition of cutesiness in our field. People spin fun and fanciful stories. Protocol participants are a caricatured Alice and Bob. Adversaries are little devils, complete with horns and a pitchfork. Some crypto talks are so packed with clip-art you can hardly find the content. I have never liked this, but, after the Snowden revelations, it started to vex me like never before.

Cryptography is serious, with ideas often hard to understand. When we try to explain them with cartoons and cute narratives, I don’t think we make our contributions easier to understand. What we *actually* do is add in a layer of obfuscation that must be peeled away to understand what has actually been done. Worse, the cartoon-heavy cryptography can reshape our internal vision

of our role. The adversary as a \$53-billion-a-year military-industrial-surveillance complex and the adversary as a red-devil-with-horns induce entirely different thought processes. If we see adversaries in one of these ways, we will actually see at a different set of problems to work on than if we see things in the other. Whimsical adversaries engender a chimerical field.¹⁸⁰

As a graduate student, I wanted our field to feel fantastical. I wanted a discipline full of space aliens and communicating millionaires. Not only was it fun, but it stroked my ego, effectively embodying the sentiment: *I am a scientist too smart to have to deal with small-minded concerns.*

At this point, I think we would do well to put ourselves in the mindset of a *real* adversary, not a notional one: the well-funded-intelligence agency, the profit-obsessed multinational, the drug cartel. You have an enormous budget. You control lots of infrastructure. You have teams of attorneys more than willing to interpret the law creatively. You have a huge portfolio of zero-days.¹⁸¹ You have a mountain of self-righteous conviction. Your aim is to *Collect it All, Exploit it All, Know it All.*¹⁸² What would frustrate you? What problems do you *not* want a bunch of super-smart academics to solve?

▷ *Stop with the cutesy pictures. Take adversaries seriously.*

A cryptographic commons. Many people see the Internet as some sort of magnificent commons. This is a fantasy. There are some successful commons within the Internet: Wikipedia, the free software movement, Creative Commons, OpenSSL, Tor, and more. But most people turn almost exclusively to services mediated by a handful of corporations that provide the electronic mail, instant messaging, cloud storage, and cloud computing, for example, that people use. And they provide the hardware on which all this stuff sits.

We need to erect a much expanded commons on the Internet. We need to realize popular services in a secure, distributed, and decentralized way, powered by free software and free/open hardware. We need to build systems beyond the reach of super-sized companies and spy agencies. Such services must be based on strong cryptography. Emphasizing that prerequisite, we need to expand our *cryptographic commons*.

Dreams for such a commons go back to the cypherpunks, who built remailers, for example, as a communitarian service to enable secure communications. More recently, Feigenbaum and Koenig articulate such a vision.¹⁸³ After explaining that centralized cloud services play a central role in enabling mass surveillance, they call for a grass-roots effort to develop new, global-scale cloud services based on open-source, decentralized, configuration-management tools.

We might start small by doing our piece to improve the commons we do have: Wikipedia. It could become a routine undertaking at IACR conferences and workshops, or at Dagstuhl meeting, for folks to gather around for an afternoon or evening to write, revise, and verify selected Wikipedia pages dealing with cryptography. It's the sort of effort that will pay off in many unseen ways.

▷ *Design and build a broadly useful cryptographic commons.*

Communications. In advancing our field, well-named notions have always been important. One has only to think back to *zero-knowledge* (and the competing term *minimal-disclosure*) to recall how a beautiful phrase could help catapult a beautiful idea into prominence. Similarly, the six-letter phrase *33 bits* does a remarkably good job of embodying an important concept without going anywhere near contested vocabulary.¹⁸⁴ In both cryptography and privacy, language is both formative and fraught.

The word *privacy*, its meaning abstract and debated, its connotations often negative, is not a winning word. Privacy is for medical records, toileting, and sex — not for democracy or freedom. The word *anonymity* is even worse: modern political parlance has painted this as nearly a flavor of terrorism. *Security* is more winning a word and, in fact, I spoke of *secure messaging* instead of *private messaging* or *anonymous messaging* because I think it better captures what I want conveyed: that a communication whose endpoints are manifest is not at all *secure*. A person needs to feel *insecure* if he uses such a channel.

But even the word *security* doesn't support a good framing of our problem: we should try to speak of thwarting mass surveillance more than enhancing privacy, anonymity, or security. As discussed before, we know instinctively that ubiquitous surveillance is incompatible with freedom, democracy, and human rights.¹⁸⁵ This makes surveillance a thing against which one can fight. The surveillance camera and data center make visual our emerging dystopia, while privacy, anonymity, and security are so abstract as to nearly defy visual representation.

Concretely, research that aims to undermine objectionable surveillance might be called *anti-surveillance research*.¹⁸⁶ Tools for this end would be *anti-surveillance technologies*.¹⁸⁷ And choosing the problems one works on based on an ethical vision might be called *conscience-based research*.

▷ *Choose language well. Communication is integral to having an impact.*

Institutional values. This essay might seem to focus on the ethical weight of each scientist's personal, professional choices. But I am actually more concerned about how we, as cryptographers and computer scientists, act in aggregate. Our collective behavior embodies values—and the institutions we create do, too.

I do not intend to criticize any particular individual. People should and will work on what they think to be most valuable. The problem occurs when our community, as a whole, systematically devalues utility or social worth. Then we have a collective failure. The failure falls on no one in particular, and yet it falls on everyone.

Conclusion to it all. As computer scientists and cryptographers, we are twice culpable when it comes to mass surveillance: computer science created the technologies that underlie our communications infrastructure, and that are now turning it into an apparatus for surveillance and control; while cryptography contains within it the underused potential to redirect this tragic turn.¹⁸⁸

Authors and filmmakers, futurists and scientists, have laid out many competing visions for man’s demise. For example, Bill Joy worries about nanotechnology turning the biosphere into gray goo, or super-intelligent robots deciding that man is a nuisance, or a pet.¹⁸⁹ I don’t lose sleep over such possibilities; I don’t see this as our likely end. But a creeping surveillance that grows organically in the public and private sectors, that becomes increasingly comprehensive, entwined, and predictive, that becomes an instrument for assassination, political control, and the maintenance of power—well, this vision doesn’t merely seem possible, it seems to be happening before our eyes.

I am not optimistic. The figure of the heroic cryptographer sweeping in to save the world from totalitarian surveillance is ludicrous.¹⁹⁰ And in a world where intelligence agencies stockpile and exploit countless vulnerabilities, obtain CA secret keys, subvert software-update mechanisms, infiltrate private companies with moles, redirect online discussions in favored directions, and exert enormous influence on standards bodies, cryptography alone will be an ineffectual response. At best, cryptography might be a tool for creating possibilities within contours circumscribed by other forces.

Still, there are reasons to smile. A billion users are getting encrypted instant messaging using WhatsApp and its embedded Axolotl protocol.¹⁹¹ Two million clients connect using Tor.¹⁹² Cryptography papers inspired by the Snowden revelations are starting to come out apace. More than 50 crypto and security researchers from the U.S.A. signed an open letter I co-organized deploring society-wide surveillance.¹⁹³ The 15-author *Keys Under Doormats* report¹⁹⁴ is an explicit attempt to have cryptographic expertise inform policy.

And it’s not as though crypto-for-privacy is something new or deprecated within our community. Cryptographers like Ross Anderson, Dan Bernstein, Matt Blaze, David Chaum, Joan Feigenbaum, Matt Green, Nadia Heninger, Tanja Lange, Arjen Lenstra, Kenny Paterson, Ron Rivest, Adi Shamir, Nigel Smart, and Moti Yung, to name just a few, have been attending to practical privacy long before it started to get trendy (if this *is* happening!). The RWC (Real World Cryptography) conference is creating a new and healthy mix of participants.

Talks, workshops, and panel discussions on mass surveillance are helping cryptographers see that dealing with mass surveillance *is* a problem within our discipline. Bart Preneel and Adi Shamir have been going around given talks entitled *Post-Snowden Cryptography*, and there were panel discussions with this title at Eurocrypt 2014 and RSA-CT 2015.

Articles are emerging with titles like “Cryptographers have an ethics problem.”¹⁹⁵ When an attack on Tor by CMU researchers was allegedly used to provide bulk anonymized data to the FBI, CMU and the researchers involved were publicly shamed.¹⁹⁶ The IACR itself has been getting more vocal, both with the Copenhagen Resolution¹⁹⁷ and the statement on Australia’s Defence Trade Controls Act.¹⁹⁸

While our community has embraced crypto-for-privacy less than I would like, this has been a cultural issue—and culture can change.

I have heard it said that “if you think cryptography is your solution, you don’t understand your problem.”¹⁹⁹ If this quip is true, then our field has gone seriously astray. But we can correct it. We need to make cryptography the solution to the problem: “how do you make surveillance more expensive?”

Dan Bernstein speaks of *interesting crypto* and *boring crypto*. Interesting crypto is crypto that supports plenty of academic papers. Boring crypto is “crypto that simply works, solidly resists attacks, [and] never needs any upgrades.” Dan asks, in his typically flippant way,

What will happen if the crypto users convince some crypto researchers to actually create boring crypto?

No more real-world attacks. No more emergency upgrades. Limited audience for any minor attack improvements and for replacement crypto.

This is an existential threat against future crypto research.²⁰⁰

If this is boring crypto, I think we should go make some.

Cyberpunk cryptography has been described as *crypto with an attitude*.²⁰¹ But it is much more than that, for, more than anything else, what the cyberpunks wanted was crypto with *values*. And values, deeply felt and deeply embedded into our work, is what the cryptographic community needs most. And perhaps a dose of that cyberpunk verve.²⁰²

It has been said that “just because you don’t take an interest in politics, doesn’t mean politics won’t take an interest in you.”²⁰³ Since cryptography is a tool for shifting power, the people who know this subject well, like it or not, inherit some of that power. As a cryptographer, you can ignore this landscape of power, and all political and moral dimensions of our field. But that won’t make them go away. It will just tend to make your work less relevant or socially useful.

My hope for this talk is that you will internalize this fact, and recognize it as the starting point for developing an ethically driven vision for what you want to accomplish with your scientific work.

I began this essay speaking of the Russell-Einstein manifesto, so let me end there as well, with Joseph Rotblat’s plea from his Nobel prize acceptance speech:

At a time when science plays such a powerful role in the life of society, when the destiny of the whole of mankind may hinge on the results of scientific research, it is incumbent on all scientists to be fully conscious of that role, and conduct themselves accordingly. I appeal to my fellow scientists to remember their responsibility to humanity.²⁰⁴

Acknowledgments

My thanks go first to Mihir Bellare for countless discussions on the topic of this essay. For years, not only have we collaborated closely on technical matters, but we have also much discussed the values and sensibilities implicitly embedded within cryptographic work. Without Mihir, not only would I have done far less technically, but I would also understand far less about who cryptographers are.

Ron Rivest not only provided useful comments, but has been much on my mind as I have agonized over this essay. Many other people have given me important suggestions and ideas. I would like to thank Jake Appelbaum, Tom Berson, Dan Boneh, David Chaum, Joan Feigenbaum, Seda Gürses, Tanja Lange, Chip Martel, Chanathip Namprempre, Ilya Mironov, Chris Patton, Charles Raab, Tom Ristenpart, Amit Sahai, Rylan Schaeffer, Jessica Malekos Smith, Christopher Soghoian, Colleen Swanson, Björn Tackmann, and Yusi (James) Zhang for their comments, discussions, and corrections.

Thanks to the Schloss Dagstuhl staff and to the participants of workshop 14401, *Privacy and Security in an Age of Surveillance*, where ideas related to this talk were discussed.²⁰⁵

This work was supported by NSF Grant CNS 1228828; many thanks to the NSF for their continuing support. Some of the work on this essay was done while I was a guest professor at ENS, Paris, hosted by David Pointcheval.

My thanks to the IACR Board for the privilege of giving this year's IACR Distinguished Lecture. It is an honor that happens at most once in a cryptographer's career, and I have tried my best to use this opportunity wisely.

Notes

¹ This account is largely taken from Sandra Butcher: The origins of the Russell-Einstein manifesto. Pugwash Conference on Science and World Affairs, May 2005.

² At the time of the press conference, Russell had heard from only eight.

³ The name would seem to be an instance of the "Matthew effect," as signatories Max Born, Frédéric Joliot-Curie, and Joseph Rotblat all played roles at least as large as Einstein's.

⁴ Quoted in Joseph Rotblat, ed., *Proceedings of the First Pugwash Conference on Science and World Affairs*, Pugwash Council, 1982.

⁵ Butcher, *op. cit.*, Foreword, p. 3.

⁶ Joseph Rotblat, "Remember Your Humanity." Acceptance and Nobel lecture, 1995. Text available at Nobelprize.org

⁷ The literature in this direction is too vast to possibly survey. University programs in this space often go by the acronym STS, for *science and technology studies* or *science, technology, and society*. The work of Langdon Winner is particularly concerned with the relation between technological artifacts and their implicit political dimension.

⁸ The quote is from: R. Williams and D. Edge: The social shaping of technology. *Research Policy* (25), 865–899, Elsevier Science B.V., (1996). The implicit reference is to the eerie Borges short story "El jardín de senderos que se bifurcan" (The Garden of Forking Paths) (1941).

⁹ Some of my comments in the remainder of this section were informed by Matthew Wisnioski: *Engineers for Change: Competing Visions of Technology in 1960s America*. MIT Press, 2012.

¹⁰ While these two possibilities are very different, distinguishing between them will not be important for the discussion of this essay.

¹¹ The debate within the scientific community got significant media attention following release of a book-length study from the American Physical Society (APS) Study Group (N. Bloembergen, C. K. Patel, cochairmen), *Report to The American*

Physical Society of the Study Group on Science and Technology of Directed Energy Weapons, APS, New York (April 1987). The debate was not one-sided: many physicists supported SDI, and many felt the APS study to be too negative or too political.

¹² See <http://fas.org/bethecr.htm#letter> for Bethe's 1995 letter.

¹³ The well-known expression comes from a pamphlet "Speak truth to power: a Quaker search for an alternative to violence," 1955. <http://www.quaker.org/sttp.html>

¹⁴ Of course not every prominent physicist was oppositional. Edward Teller famously championed the development of fusion bombs and SDI. These positions were divisive, with some of Teller's peers seeing him as irresponsible, jingoistic, or insane.

¹⁵ The definitive report on the incident is: Reynoso, C., Blando, P., Bush, T., Herbert, P., McKenna, W., Rauchway, E., Balcklock, P., Brownstein, A., Dooley, D., Kolesar, K., Penny, C., Sterline, R., Sakaki, J.: UC Davis Nov. 18, 2011 "Pepper Spray Incident" Task Force Report: The Reynoso Task Force Report. (March 2012) Available at http://www.ucsf.edu/sites/default/files/legacy_files/reynoso-report.pdf (visited 2015.08.07)

¹⁶ There are 2 million views of <https://www.youtube.com/watch?v=6AdDLhPwpp4>

¹⁷ Nathan Brown: Op-ed: Reviewing the case for Katehi's resignation. Davis Enterprise newspaper (Dec. 18, 2011). UC Davis Physics Department, untitled press release (Nov. 22, 2011), available at <http://tinyurl.com/ucd-physics-pepper-spray> (last visited 2015.08.07).

¹⁸ The Chancellor expressed this sentiment in a meeting between her and interested faculty of the College of Engineering (date unknown, probably early 2012).

¹⁹ See, for example, Israel Gutman: *Encyclopedia of the Holocaust*, Macmillan Library Reference USA, 1990, entries: "Medical Experiments" by Nava Cohen, pp. 957–966, and "Physicians, Nazi" by Robert Jay Lifton and Amy Hackett, pp. 1127–1132.

²⁰ Hannah Arendt: *Eichmann in Jerusalem: A Report on the Banality of Evil*. Viking Adult (1963)

²¹ Stanley Milgram: Behavioral study of obedience. *Journal of Abnormal and Social Psychology*, 67(4), pp. 371–378 (1963).

²² In the US, only mining engineers failed to adopt a code of ethics, according to Wikipedia article "Engineering ethics."

²³ The organization dissolved in 2013.

²⁴ Article II of the Bylaws of the International Association for Cryptologic Research: "The purposes of the IACR are to advance the theory and practice of cryptology and related fields, and to promote the interests of its members with respect thereto, and to serve the public welfare." Last revised Nov. 18, 2013. Available at www.iacr.org/docs/bylaws.pdf

²⁵ See, for example, Thomas P. Hughes: *Rescuing Prometheus*, 1998 (on the Atlas project).

²⁶ Following some security breaches, they now do this in partnership with industry, particularly with Bechtel. See <http://www.bechtel.com/projects/us-national-laboratories/>

²⁷ The term is associated to Buddhism, right livelihood being one of the virtues of the Noble Eightfold Path.

²⁸ My top five Google search results to "deciding among job offers" (no quotes) on Aug 26, 2015 were: (1) "Help! How Do I Choose Between Two Job Offers." CareerCast website, www.career-cast.com/career-news/help-how-do-i-choose-between-two-job-offers. (2) "4 questions to help you decide between job offers." Natalie Wearstler, Sep. 16, 2013. theweek.com/articles/460019/4-questions-help-decide-between-job-offers. (3) "You Got the Jobs! How to Decide Between Offers."

Forbes, June 6, 2011. www.forbes.com/sites/prettyyoungprofessional/2011/06/06/you-got-the-jobs-how-to-decide-between-offers/. (4) “6 Secrets to Choosing Between Job Offers.” July 31, 2013. www.aegistech.com/how-to-choose-between-multiple-job-offers/. (5) “How to Choose Between Multiple Job Offers.” Sam Tomarchio, Aegistech. www.aegistech.com/how-to-choose-between-multiple-job-offers/. The indicated websites do not contain the words *ethics*, *moral*, or *social*, nor any comments concerning the ethical character or social value of ones work. The author acknowledges that Google search results are not reproducible.

²⁹ Stanley Fish: Why we built the ivory tower. NY Times, Opinion Section, May 21, 2004.

³⁰ Elisabeth Pain: The social responsibility of scientists. *Career Magazine of Science*, Feb. 16, 2013. Report on the 2013 AAAS meeting and, in particular, the comments of Mark Frankel. The speaker asks listeners, especially graduate students, to keep firmly in mind that scientific research is a social institution, and one that is subsidized by society.

³¹ Radical individualism is the belief that ones personal interests are more important than those of society. It is well captured by the “Greed is Good” speech from Oliver Stone’s “Wall Street” (1987), although what the individual cares about can be something other than personal wealth.

³² Silvio Micali, ACM Turing Award Acceptance speech. ACM awards ceremony, San Francisco, June 15, 2013. <https://www.youtube.com/watch?v=W0N4WnjGHwQ>. Ending his speech with a cheer, there is, perhaps, an element of facetiousness in Silvio’s unbridled optimism. Yet over-the-top optimism seems endemic to discourse on computer science.

The awards ceremony took place a little over a week after the first news story based on a Snowden document (June 6, 2013). But the threat of computing technology was not to be acknowledged, only its promise praised.

³³ *Understanding Media*, 1964: “Man becomes, as it were, the sex organs of the machine world, as the bee of the plant world, enabling it to fecundate and evolve ever new forms.”

³⁴ To be clear, I do *not* think there’s anything extreme in the views of prominent technological pessimists; if anything, prevailing optimistic views seems to me far less reasoned and more extreme.

³⁵ This is the phrase used by Ian Barbour: *Ethics in an Age of Technology: The Gifford Lectures, Volume Two*, HarperCollins (1993), which contains a good description of the writings of the contextualists (which include the author himself). Langdon Winner speaks of an *ideology of technological politics* for roughly the same concept as contextualism; see *Autonomous Technology: Technics-out-of-Control as a Theme in Political Thought*. MIT Press, 1977.

³⁶The series of books by James Bamford on the history of the NSA are *The Puzzle Palace: A Report on America’s Most Secret Agency* (1982), *Body of Secrets: Anatomy of the Ultra-Secret National Security Agency* (2001), and *The Shadow Factory: The NSA from 9/11 to the Eavesdropping on America* (2008). The classic book of David Kahn on the history of code breaking is *The Codebreakers: The Story of Secret Writing* (1967).

³⁷ The U.S. “Consolidated Cryptologic Program” includes about 35,000 employees. NSA budget for 2013 was \$10.3 billion, with \$1 billion of this marked as “cryptanalysis and exploitation services” and \$429 million “research and technology.” intelligence operations between 2004–2013. See Barton Gellman and Greg Miller, ‘Black budget’

summary details U.S. spy network's successes, failures, and objectives; *The Washington Post*, April 29, 2013.

³⁸ For an account of the extreme utility of surveillance and cryptography in recent U.S. wars, see Shane Harris: *@War: The Rise of the Military-Internet Complex*. Eamon Dolan/Houghton Mifflin Harcourt (2014).

³⁹ For example, the phrase “an all-powerful adversary” is usually meant as an agent having no computational restrictions, but which conforms to a specified model.

⁴⁰ Jeremy Scahill and Josh Begley: The great SIM heist: how spies stole the keys to the encryption castle. *The Intercept*, Feb. 19, 2015.

⁴¹ Joe Mullin: Newegg trial: crypto legend takes the stand, goes for knockout patent punch. *Ars Technica*, Nov. 24, 2013.

⁴² Whitfield Diffie and Martin Hellman: Exhaustive cryptanalysis of the NBS Data Encryption Standard. *Computer* 10(6), pp. 74–84 (1977).

⁴³ See, for example: Anatoly Gromyko and Martin Hellman: *Breakthrough: Emerging New Thinking: Soviet and Western Scholars Issue a Challenge to Build a World Beyond War*, 1988.

⁴⁴ Whitfield Diffie and Susan Landau: *Privacy on the Line: The Politics of Wiretapping and Encryption*, MIT Press (2007).

⁴⁵ Whitfield Diffie, Paul van Oorschot, Michael Wiener: Authentication and authenticated key exchanges. *Designs, Codes and Cryptography*, 2(2), pp. 107–125 (1992).

⁴⁶ Whitfield Diffie and Martin Hellman: New directions in cryptography. *IEEE Transactions on Information Theory*, 22(6), pp. 644–654 (1976).

⁴⁷ Yet the authors did seem to anticipate such a possibility: “At the same time, theoretical developments in information theory and computer science show promise of providing provably secure cryptosystems, changing this ancient art into a science.” *Ibid.*, p. 29.

⁴⁸ They write, for example, that “The development of cheap digital hardware has freed it from the design limitations of mechanical computing and brought the cost of high grade cryptographic devices down to where they can be used in such commercial applications as remote cash dispensers and computer terminals. . . . The development of computer controlled communication networks promises effortless and inexpensive contact between people or computers on opposite sides of the world . . .”. *Ibid.*, p. 29.

⁴⁹ The paper appears in *Communications of the ACM* (CACM), 24(2), pp. 84–88 (1981).

⁵⁰ In a nutshell, a mix net works like this. Each user's communications will pass through a series of routers. (For the approach to add value, these routers, or *mixes*, should have diversity in administrative control, jurisdictional control, or the code they run.) Plaintexts will be multiply encrypted, using one key for each router in the series. Each router will peel off one layer of encryption. Before passing the result to the next router, it will wait until it has some number of outgoing messages, and then reorder them. An adversary who observes network traffic—even one who controls some subset of the routers—should be unable to identify who has sent what to whom.

⁵¹ For example: “The foundation is being laid for a dossier society, in which computers could be used to infer individuals' life-styles, habits, whereabouts, and associations from data collected in ordinary consumer transactions. Uncertainty about whether data will remain secure against abuse by those maintaining or tapping it can have a ‘chilling effect,’ causing people to alter their observable activities.” David Chaum: Security without identification: transaction systems to make big brother obsolete. *Communications of the ACM* (CACM), 28(10), pp. 1030–1044, 1985.

⁵² Speaking of cryptography, Chaum told me “I’ve been aware from the beginning of the field, that it was a powerful field, and would become more and more as information technology advanced. I was surprised that academics were as interested as they were in cryptography, because it was an area so connected to power.” Conversation with David Chaum, 17 Aug 2015. Santa Barbara, California.

⁵³ Starting in 2012, this annual event already attracts more attendees than Crypto. RWC’s steering committee is Dan Boneh, Aggelos Kiayias, Brian LaMacchia, Kenny Paterson, Tom Ristenpart, Tom Shrimpton, and Nigel Smart. Webpage <http://www.realworldcrypto.com/> explains: “This annual conference aims to bring together cryptography researchers with developers implementing cryptography in real-world systems. The conference goal is to strengthen the dialog between these two communities. Topics covered focus on uses of cryptography in real-world environments such as the Internet, the cloud, and embedded devices.”

⁵⁴ Some theorists may take offense by an implication that cryptographic practice is “real world,” suggesting those who do theory must be studying something ungrounded or imaginary (the vapors of poltergeists)? In truth, the implicit dichotomies of cryptographic theory vs. cryptographic practice, real-world and its complement (whatever that might be), would not survive a critical assessment. Yet discourse on what communities do is unavoidably full of capricious boundaries.

⁵⁵ Tellingly, the problem lacks even a widely recognized name. The version of the secure-messaging problem Chaum was interested in is a high-latency, public-key version.

⁵⁶ Charles Rackoff and Daniel Simon: Cryptographic defense against traffic analysis. STOC ’93, pp. 672–681. The quote is from pp. 672–673.

⁵⁷ Nik Unger, Sergej Dechand, Joseph Bonneau, Sascha Fahl, Henning Perl, Ian Goldberg, Matthew Smith: SoK: Secure messaging. *IEEE Symposium on Security and Privacy 2015*, pp. 232–249.

⁵⁸ Little, but not none. See, for example: Nik Unger and Ian Goldberg: Deniable key exchanges for secure messaging. *ACM CCS 2015*. N. Borisov, I. Goldberg, and E. Brewer: Off-the-Record communication, or, why not to use PGP. *Privacy in the Electronic Society*, pp. 77–84, 2004. Ron Berman, Amos Fiat, Marcin Gomulkiewicz, Marek Klonowski, Miroslaw Kutylowski, Tomer Levinboim, Amnon Ta-Shma: Provable unlinkability against traffic analysis with low message overhead. *J. of Cryptology*, 28(3), pp. 623–640, 2015. Jan Camenisch and Anna Lysyanskaya: A formal treatment of onion routing. *CRYPTO 2005*. Joan Feigenbaum, Aaron Johnson, and Paul F. Syverson: Probabilistic analysis of onion routing in a black-box model. *ACM Transactions on Information and System Security (TISSEC)*, 15(3), 2012.

⁵⁹ Chaum, *op. cit.* Shafi Goldwasser and Silvio Micali: Probabilistic encryption and how to play mental poker keeping secret all partial information. STOC 1982, pp. 365–377. Journal version as: Probabilistic encryption, *Journal of Computer and System Science (JCSS)*, 28(2), pp. 270–299, April 1984.

⁶⁰ According to Google scholar: 4481 citations to the Chaum paper, and 3818 citations for the Goldwasser-Micali paper (both versions, combined). Data as of Oct 14, 2015.

⁶¹ The one outlier is a paper by Perig, Szwedczyk, Tygar, Wen, and Culler in *MobiCom*.

⁶² In greater detail, the ten most cited [GM82] (both versions combined) appeared at CCS, Crypto 3×, Eurocrypt, FOCS (2×) *MobiCom*, STOC (2×). Only *MobiCom* is an “unexpected” venue. The ten most cited [Chaum81]-citing papers appeared at ACM Comp. Surveys, ACM J. of Wireless Networks, ACM *MobiSys*, ACM *Tran. on*

Inf. Sys., ACM SIGOPS, IEEE SAC, Proc. of the IEEE USENIX Security Symposium, and workshops named IPTPS and Designing Privacy Enhancing Technologies. None of these ten venues is crypto-focused.

⁶³ Masayuki Abe and Hideki Imai: Flaws in some robust optimistic mix-nets. *Information Security and Privacy (ACISP)*, pp. 39–50, 2003. Journal version: Masayuki Abe and Hideki Imai: Flaws in robust optimistic mix-nets and stronger security notions. *IEICE Trans. on Fundamentals of Electronics, Communications and Computer Science*, vol. E89A, no. 1, pp. 99–105, Jan 2006.

⁶⁴ Diffie and Hellman, *op. cit.*

⁶⁵ Multiparty computation. Andrew Chi-Chih Yao: Protocols for secure computations (Extended Abstract). FOCS 1982, pp. 160–164. Oded Goldreich, Silvio Micali, Avi Wigderson: How to play any mental game or a completeness theorem for protocols with honest majority. STOC 1987, pp. 218–229. Michael Ben-Or, Shafi Goldwasser, Avi Wigderson: Completeness theorems for non-cryptographic fault-tolerant distributed computation (extended abstract). STOC 1988, pp. 1–10. David Chaum, Claude Crépeau, Ivan Damgård: Multiparty unconditionally Secure protocols (extended abstract). STOC 1988, pp. 11–19.

⁶⁶ For example, compare the opening pages of: Oded Goldreich: Foundations of cryptography—a primer. *Foundations and Trends in Theoretical Computer Science*, 1(1), pp. 1–116, 2004. And: George Danezis and Seda Gürses: A critical review of 10 years of privacy technology. *Proceedings of Surveillance Cultures: A Global Surveillance Society?*, 2010.

⁶⁷ Andy Greenberg: *This Machine Kills Secrets: How WikiLeaks, Cypherpunks, and Hacktivists Aim to Free the World's Information*. Dutton, 2012. Steven Levy: *Crypto: How the Code Rebels Beat the Government Saving Privacy in the Digital Age*. Viking Adult, 2001. Julian Assange, Jacob Appelbaum, Andy Müller-Maguhn, Jérémie Zimmermann: *Cypherpunks: Freedom and the Future of the Internet*. OR Books, 2012. Robert Manne: The cypherpunk revolutionary: on Julian Assange. *The Monthly: Australian Politics, Society, & Culture*. March 2011.

⁶⁸ It is the same belief embedded in the work already discussed by David Chaum.

⁶⁹ I have heard academic cryptographers saying that we deserve credit for these artifacts, for, regardless of who wrote the code, the ideas stem from cryptographic notions that come from us. While there is some truth to this, the claim still feels ungeniue. The work, in each case, arose from somewhere else, employed only old and basic tools, and was at most modestly embraced by our community.

⁷⁰ Julian Assange: Conspiracy as governance. Manuscript, December 3, 2006. <http://cryptome.org/0002/ja-conspiracies.pdf>. Finn Brunton: Keyspace: reflections on WikiLeaks and the Assange papers. *Radical Philosophy* 166, pp. 8–20, Mar/Apr 2011.

⁷¹ Julian Assange: *Cypherpunks, op. cit.*

⁷² I do not know if Snowden considers himself a cypherpunk, but the sentiment expressed in this quote nicely reflects cypherpunk discourse.

⁷³ Edward Snowden, quoted in Glenn Greenwald, *No Place to Hide: Edward Snowden, the NSA, and the U.S. Surveillance State*, 2014. Snowden is reformulating a quote of Thomas Jefferson: “In questions of power then, let no more be heard of confidence in man but bind him down from mischief by the chains of the Constitution.”

⁷⁴ Bruce Schneier: *Data and Goliath: The Hidden Battles to Collect Your Data and Control Your World*. W.W. Norton & Company, 2015. The quote is from p. 131.

⁷⁵ This was the dream behind the the Trusted Platform Module (TPM), a coprocessor that implements a standard set of cryptographic operations within a tamper-resistant boundary. See Wikipedia entry: Trusted Platform Module, and Ross Anderson’s FAQ:

‘Trusted Computing’ Frequently Asked Questions. <http://www.cl.cam.ac.uk/~rja14/tcpa-faq.html>

⁷⁶ For a nice discussion of the Clipper Chip, see Michael Froomkin: The metaphor is the key: cryptography, the clipper chip, and the constitution. *University of Pennsylvania Law Review*, 143(3), pp. 709–897, 1995.

⁷⁷ This may be theoretically possible but practically difficult: it might be hard to accomplish for technically unsophisticated people. This is why it is said: *When crypto is outlawed only outlaws will have crypto* (or *privacy*, or *security*) (original attribution unknown).

⁷⁸ On the other hand, some other forms of cryptographically enabled DRM (Digital Rights Management), including the CSS (Content Scrambling System) used to make more difficult the copying of DVDs, has been reasonably effective (from the point of view of content owners).

⁷⁹ Ideas expressed on this topic are joint with Mihir Bellare.

⁸⁰ Adi Shamir: Identity-based cryptosystems and signature schemes. *Crypto ’84*, pp. 47–53, 1984. Dan Boneh and Matthew Franklin: Identity-based encryption from the Weil pairing. *SIAM J. of Computing*, 32(3), pp. 586–615, 2003.

⁸¹ For example, the Wikipedia entry “ID-based encryption” (Nov. 2015) has a lead section that fails to even mention the presence of a key-issuing authority.

⁸² For many IBE schemes it is easy to distribute the PKG’s secret, the master key, across a set of parties. One could further arrange for distributed generation of this key, so no one party ever held it. Both of these points are made in the original Boneh and Franklin paper: *op. cit.*, Section 6, Distributed PKG.

⁸³ See, for example: Cynthia Dwork: Differential privacy: a survey of results. *Theory and Applications of Models of Computation*, pp. 1–19, Springer, 2008.

⁸⁴ *Ibid*, p. 2.

⁸⁵ For a quite different critique of differential privacy, see Jane Bambauer, Krishnamurthy Muralidhar, Rathindra Sarathy: Fool’s gold: an illustrated critique of differential privacy. *Vanderbilt Journal of Entertainment and Technology Law*, 16(4), pp. 701–755, Summer 2014.

⁸⁶ Colleen Swanson, personal communications.

⁸⁷ Cynthia Dwork, Krishnamurthy Muralidhar, Frank McSherry, Ilya Mironov, and Moni Naor: Our data, ourselves: privacy via distributed noise generation. *Eurocrypt 2006*, pp. 486–503. For an exploration of limitations of local model, also see: Shiva Prasad Kasiviswanathan, Homin K. Lee, Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith: What can we learn privately? *SIAM Journal of Computing*, 40(3), pp. 793–826, 2011. Amos Beimel, Kobbi Nissim, and Eran Omri: Distributed private data analysis: on simultaneously solving how and what. *CRYPTO 2008*, pp. 451–468. Also arXiv:1103.2626

⁸⁸ Craig Gentry: Fully homomorphic encryption using ideal lattices. *STOC 2009*.

⁸⁹ I refer, of course, to the exchange of personal information for a network service.

⁹⁰ The acronym is for *indistinguishability obfuscation*. Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters: Candidate indistinguishability obfuscation and functional encryption for all circuits. *FOCS 2013*, pp. 40–49. Earlier work introducing iO: Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil Vadhan, and Ke Yang. On the (im)possibility of obfuscating programs. *JACM*, 59(2):6, 2012. Earlier version in *CRYPTO 2001*.

⁹¹ See, for example: Erica Klarreich: Perfecting the art of sensible nonsense. *Quanta Magazine*, Jan. 30, 2014. “Researchers are hailing the new work as a watershed moment for cryptography. ... If the problem of obfuscation has been solved, what remains

for cryptographers?” Kevin Hartnett: A New Design for Cryptography’s Black Box. *Quanta Magazine*, Sep. 2, 2015. “New advances show how near-perfect computer security might be surprisingly close at hand.”

⁹² Evelyn M. Rusli: A Darpa [sic] director on fully homomorphic encryption (or one way the U.S. could collect data). *Wall Street Journal* blog: blogs.wsj.com. March 9, 2014. Quote edited for punctuation.

⁹³ Eurocrypt 1992 reviewed. Author name redacted. National Security Agency, CRYPTOLOG. First issue of 1994. Available at <http://tinyurl.com/eurocrypt1992>

⁹⁴ 136 editions of the NSA newsletter, CRYPTOLOG, are available, in redacted form. The period covered is 1974–1997. An archive can be found at <https://www.nsa.gov/public.info/declass/cryptologs.shtml>

⁹⁵ Of course Tor is extensively discussed in Snowden documents—but it would be wrong to view Tor as a contribution springing from the cryptographic community, nor as something much addressed within it.

⁹⁶ This is the term NSA author above prefers.

⁹⁷ The use of the word “snarky” here is from Bruce Schneier: Snarky 1992 NSA report on academic cryptography. Blog post, Nov. 18, 2014. https://www.schneier.com/blog/archives/2014/11/snarky-1992_nsa.html

⁹⁸ Arvind Narayanan: What happened to the crypto dream? Part 1 in *IEEE Security and Privacy Magazine*, 11(2), pp. 75–76, 2013. Part 2 in *IEEE Security and Privacy Magazine*, 11(3), pp. 68–71, 2013.

⁹⁹ For a review of key information learned from the Snowden revelations, see the EFF webpages “NSA Spying on America,” <https://www.eff.org/nsa-spying>, and the ACLU webpage, “NSA Surveillance,” <https://www.aclu.org/issues/national-security/privacy-and-surveillance/nsa-surveillance>

¹⁰⁰ Collin Bennett: *The Privacy Advocates: Resisting the Spread of Surveillance*. The MIT Press, 2008.

¹⁰¹ James Comey: “Going Dark: Are Technology, Privacy, and Public Safety on a Collision Course?” Speech at the Brookings Institute, October 16, 2014. <http://tinyurl.com/comey-going-dark>

¹⁰² Comey speaks particularly of CALEA, the Communications Assistance for Law Enforcement Act, the 1994 U.S. law mandating lawful intercept capabilities be built into telecommunications equipment—but does not mandate such abilities for encrypted email, instant messaging, and the like.

¹⁰³ The traditional cast of evil-doers, the “four horsemen of the info-pocalypse,” omits murderers from this list of five; see Assange *et al.*, *Cypherpunks*, *op. cit.*. Note that in Comey’s Brookings Institute speech, the speaker adds murderers to this list, and subtracts money launderers. The move is shrewd; money laundering is technocratic and legalistic crime compared to murder.

¹⁰⁴ The phrase (including the capitalization) is from the Comey speech just cited.

¹⁰⁵ Church Committee Reports, 1975–1976. Available at <http://www.aarclibrary.org/publib/church/reports/contents.htm>. Frank Donner: *The Age of Surveillance: The Aims and Methods of America’s Political Intelligence System*. Vintage Press, 1981. Classified NSA document: SIGINT Mission Strategic Plan FY2008–2013. Oct. 3, 2007. <http://tinyurl.com/sigint-plan>

¹⁰⁶ “Verbal deception is itself an intelligence practice,” explains Donner, *op. cit.*, p. xiv. See to his Appendix I, pp. 464–466.

¹⁰⁷ Michel Foucault: *Discipline and Punish: The Birth of the Prison*. Alan Sheridan, translator. 1975/1977. Frank Donner, *op. cit.* Assange *et al.*, *op. cit.* David Lyon:

Surveillance Studies: An Overview. 2007. David Murakami Wood and Kirstie Ball: A Report on the Surveillance Society. Sep. 2006.

¹⁰⁸ Even the etymology of *surveillance* suggests this: from the French: *sur*, meaning *over*, plus *veiller*, meaning *to watch*. Thus: to watch from above, from a position of power. An interesting turn on this conception is the notion of *sousveillance*, as coined by Steve Mann. See Wikipedia entry: *Sousveillance*.

¹⁰⁹ Adding one more phone number to a wiretap is now essentially free. Kevin S. Bankston and Ashkan Soltani: Tiny constables and the cost of surveillance: making cents out of *United States v. Jones*. The Yale Law Journal, vol. 123, Jan. 2014.

¹¹⁰ This sentence is a slightly modified from Assange *et al.*, *Cypherpunks, op. cit.*.

¹¹¹ Michael Hayden: The Johns Hopkins Foreign Affairs Symposium Presents: The Price of Privacy: Re-Evaluating the NSA. May 9, 2014. Video, available at <https://www.youtube.com/watch?v=kV2HDM86XgI>

¹¹² For the last of these claims: “A life spent entirely in public, in the presence of others, becomes, as we would say shallow.” Hannah Arendt, *The Human Condition*. Chicago University Press, 1959.

¹¹³ Of course corporations and governments also, on occasion, have competing interests, as when publicity on surveillance scares away customers. This is the force motivating the principles enunciated in <https://www.reformgovernmentsurveillance.com/>

¹¹⁴ Church, *op. cit.*; Donner, *op. cit.*; and Julian Sanchez: Wiretapping’s true danger. Los Angeles Times, March 16, 2008. <http://articles.latimes.com/2008/mar/16/opinion/op-sanchez16>

¹¹⁵ Nadia Kayyali: FBI’s “Suicide Letter” to Dr. Martin Luther King, Jr., and the Dangers of Unchecked Surveillance. Nov. 12, 2014. EFF website, <http://tinyurl.com/fbi-suicide-letter>

¹¹⁶ Church Committee Reports, *op. cit.*, and Frank Donner, *op. cit.*

¹¹⁷ Or present?

¹¹⁸ The native-speaker’s English-language error of *which* instead of *who* is revealing: the target is never a person (the person *who*) but always an object (the thing *which*).

¹¹⁹ Jacob Weber (deanonymized by Guardian readers): The SIGINT Philosopher Is Back — with a New Face! <https://www.documentcloud.org/documents/2259077-the-sigint-philosopher-is-back-with-a-new-face.html> Quoted in: Peter Maass: The Philosopher of Surveillance. The Intercept. Aug. 11, 2015. This article may be my favorite from the entire corpus of articles coming out of the Snowden revelation. It functions at multiple levels, giving the reader the uncomfortable sense of doing to another (and even enjoying to do to another) precisely what the reader would not like done to himself. “Modern life is such an unholy mix of voyeurism and exhibitionism,” says the character of Stella Gibson (Gillian Anderson) in season 2, episode 4, of Allan Cubitt’s TV series *The Fall* (2014).

¹²⁰ A. O. Scott: Intent on defying an all-seeing eye: ‘Citizenfour,’ a documentary about Edward Snowden. New York Times movie review, Oct. 23, 2014.

¹²¹ Betty Medsger: *The Burglary: The Discovery of J. Edgar Hoover’s Secret FBI*. Knopf, 2014.

¹²² William Davidon, as revealed in Medsger, *The Burglary, Ibid.*

¹²³ Reporters without Borders: The Enemies of Internet: Special Edition: Surveillance. 2013. <http://surveillance.rsf.org/en/>

¹²⁴ Jeremy Scahill and Glenn Greenwald: The NSA’s secret role in the U.S. assassination program. The Intercept, Feb 9, 2014.

¹²⁵ Greg Elmer and Andy Opel: *Preempting Dissent: The Politics of an Inevitable Future*. Arbeiter Ring Publishing, 2008. Also see the film, *Preempting Dissent* (2014), by the same team. <http://preemptingdissent.com/>

¹²⁶ Michael Gazzaniga: *The Ethical Brain: The Science of Our Moral Dilemmas*. Dana Press, 2005.

¹²⁷ Bruce Schneier, *op. cit.*, paraphrasing a sentence of p. 127.

¹²⁸ Joan Feigenbaum and Bryan Ford: Seeking anonymity in an Internet panopticon. *Communications of the ACM*, 58(10), October 2015.

¹²⁹ Adam Langley: Pond. Webpages rooted at <https://pond.imperialviolet.org/>. For PANDA, see Jacob Appelbaum and “another cypherpunk”: Going dark: phrase automated nym discovery authentication: or, finding friends and lovers using the PANDA protocol to re-number after everything is lost: or, discovering new Pond users easily. Manuscript, Feb 21, 2014. <https://github.com/agl/pond/tree/master/papers/panda>

¹³⁰ An ambitious but informal attempt to unify terminology and concepts in privacy and anonymity is provided by Andreas Pfitzmann and Marit Hansen: Anonymity, unlinkability, undetectability, unobservability, pseudonymity, and identity management — a consolidated proposal for terminology. Version v0.34. Aug. 10, 2010. Manuscript at http://dud.inf.tu-dresden.de/Anon_Terminology.shtml

¹³¹ Mihir Bellare, Daniel Kane, Phillip Rogaway: Bigkey cryptography: symmetric encryption with enormous keys and a general tool for achieving key-exfiltration resistance. Manuscript, 2015.

¹³² Stefan Dziembowski: Intrusion-resilience via the bounded-storage model. TCC 2006. Giovanni Di Crescenzo, Richard J. Lipton, Shabsi Walfish: Perfectly secure password protocols in the bounded retrieval model. TCC 2006. David Cash, Yan Zong Ding, Yevgeniy Dodis, Wenke Lee, Richard J. Lipton, Shabsi Walfish: Intrusion-resilient key exchange in the bounded retrieval model. TCC 2007. Joël Alwen, Yevgeniy Dodis, Daniel Wichs: Survey: leakage resilience and the bounded retrieval model. ICITS 2009. Work in the bounded-retrieval model has earlier roots in Maurer’s bounded-storage model: Ueli Maurer: Conditionally-perfect secrecy and a provably-secure randomized cipher. *Journal of Cryptology*, 5(1), pp. 53–66, 1992.

¹³³ The peculiar-looking constant is the (approximate) value of $-\lg(1-c)$ where $c \approx 0.1100$ is the real number in $[0, 1/2]$ satisfying $H_2(c) = 0.5$ where H_2 is the binary entropy function, $H_2(x) = -x \lg x - (1-x) \lg(1-x)$.

¹³⁴ The sort of PRG that takes input and maintains state, which will now include a bigkey. See: Boaz Barak, Shai Halevi: A model and architecture for pseudo-random generation with applications to `/dev/random`. ACM CCS 2005. Yevgeniy Dodis, David Pointcheval, Sylvain Ruhault, Damien Vergnaud, Daniel Wichs: Security analysis of pseudo-random number generators with input: `/dev/random` is not robust. ACM CCS 2013.

¹³⁵ Henry Corrigan-Gibbs, Dan Boneh, David Mazières: Riposte: An anonymous messaging system handling millions of users. *IEEE Symposium on Security and Privacy*, pp. 321–338, 2015.

¹³⁶ The techniques come mostly from PIRs, the body of work to make PIRs more efficient, and the recent notion of a distributed point function, from Gilboa and Ishai. Benny Chor, Eyal Kushilevitz, Oded Goldreich, Madhu Sudan: Private information retrieval. *JACM* 45(6), pp. 965–981, 1998. Earlier version from *FOCS 1995*. Niv Gilboa, Yuval Ishai: Distributed point functions and their applications. *EUROCRYPT 2014*, pp. 640–658.

¹³⁷ Another recent paper that does a beautiful job at the nexus of systems, privacy, and cryptography is: Nikita Borisov, George Danezis, Ian Goldberg: DP5: A private

presence service. *Proceedings on Privacy Enhancing Technologies* (PETS) vol. 2, pp. 4–24, 2015. That paper again depends crucially on PIRs—this time for creating a service to tell you—but not the service provider—which of your “friends” are currently online, using the same, Facebook-like service.

¹³⁸ Colin Percival: Stronger key derivation via sequential memory-hard functions. BSDCan’09, May 2009. <http://www.tarsnap.com/scrypt/scrypt.pdf>

¹³⁹ The technique goes back to the UNIX `crypt(3)` functionality.

¹⁴⁰ Alex Biryukov, Daniel Dinu, Dmitry Khovratovich: Argon2. July 8, 2015. <https://www.cryptolux.org/index.php/Argon2>

¹⁴¹ Joël Alwen and Vladimir Serbinenko: High parallel complexity graphs and memory-hard functions. *STOC 2015*, pp. 595–603.

¹⁴² Why is this topic crypto-for-privacy? First, it’s about helping individuals avoid getting their accounts compromised. Second, for electronic currency, it helps keep small-scale mining more cost-competitive with large-scale operations. That isn’t at all true for bitcoin, where mining tends to be centralized and energy-intensive The GHash.IO mining pool boasts on its website a mining rate of 6.35 Ph/s ($2^{52.5}$ hashes per second), and the overall rate is about 465 Ph/s ($2^{58.7}$ hashes per second). Data taken from <https://ghash.io/> and <https://blockchain.info/charts/> on Nov. 1, 2015.

¹⁴³ Mihir Bellare, Kenneth G. Paterson, and Phillip Rogaway: Security of symmetric encryption against mass surveillance. *CRYPTO 2014*.

¹⁴⁴ Adam Young, Moti Yung: The dark side of black-box cryptography, or: should we trust capstone? *CRYPTO 1996*. Adam Young, Moti Yung: Kleptography: using cryptography against cryptography. *EUROCRYPT 1997*.

¹⁴⁵ David Adrian, Karthikeyan Bhargavan, Zakir Durumeric, Pierrick Gaudry, Matthew Green, J. Alex Halderman, Nadia Heninger, Drew Springall, Emmanuel Thomé, Luke Valenta, Benjamin VanderSloot, Eric Wustrow, Santiago Zanella-Béuclin, Paul Zimmermann: Imperfect forward secrecy: how Diffie-Hellman fails in practice. *Computer and Communications Security (CCS ’15)*, 2015.

¹⁴⁶ Additionally, whenever possible, the active attack should be something that *detectable*—something that requires interaction.

¹⁴⁷ Martin Hellman: A cryptanalytic time-memory trade-off. *IEEE Trans. on Information Theory*, 26(4), pp. 401–406, 1980.

¹⁴⁸ Karsten Nohl: Breaking GSM phone privacy. Black Hat USA 2010. Available on youtube, URL <https://www.youtube.com/watch?v=0hjn-BP8nro>

¹⁴⁹ Morris Dworkin, Ray Perlner: Analysis of VAES3 (FF2). Cryptology ePrint Archive Report 2015/306. April 2, 2015. FPE stands for Format-Preserving Encryption.

¹⁵⁰ See, for example, Seny Kamara: Restructuring the NSA metadata program. *Financial Cryptography Workshops 2014*, pp. 235–247, 2014.

¹⁵¹ Reflecting this, the U.S. Fourth Amendment speaks not only of particularized warrants being required for search, but also for seizure.

¹⁵² Phillip Rogaway: Practice-oriented provable security and the social construction of cryptography. Manuscript, May 2009, and invited talk at *Eurocrypt 2009*.

¹⁵³ Danny Dolev, Andrew C. Yao: On the security of public key protocols. *IEEE Trans. on Information Theory*, IT-29, pp. 198–208, 1983.

¹⁵⁴ This section deals exclusively with academic funding of cryptography in the U.S. I know very little about cryptographic funding in other countries.

¹⁵⁵ I have been unable to locate statistics on this.

¹⁵⁶ DoD = Department of Defense. This includes organizations like AFOSR, IARPA, DARPA, and ONR.

¹⁵⁷ This data is based on an accounting I did myself, by hand, going through all these old proceedings.

¹⁵⁸ The acronyms are: AFOSR = Air Force Office of Scientific Research; DARPA = Defense Advanced Research Projects Agency; DoI/NBC = Department of Interior National Business Center; IARPA = Intelligence Advanced Research Projects Activity; MURI = Multidisciplinary University Research Initiative; NSF = National Science Foundation; ONR = Office of Naval Research; and PROCEED = Programming Computation on Encrypted Data. Following the statement came another 35 words of legalistic language, including a statement that the paper had been cleared “Approved for Public Release.” <http://eprint.iacr.org/2013/403.pdf>

¹⁵⁹ Daniel S. Greenberg: *Science, Money, and Politics: Political Triumph and Ethical Erosion*. University of Chicago Press, 2003.

¹⁶⁰ A *moral hazard* is a situation in which one party gets the benefits and another takes the risk. The term is common in economics.

¹⁶¹ Darati Prabhakar: Understanding DARPA’s Mission. <http://tinyurl.com/darpa-mission> YouTube version <http://tinyurl.com/darpa-mission2>

¹⁶² Tom Johnson: *Book III: Retrenchment and Reform*, 1998. Formerly classified book, available, as a result of a FOIA request, at <http://cryptome.org/0001/nsa-meyer.htm>

¹⁶³ Henry A. Giroux: *The University in Chains: Confronting the Military-Industrial-Academic Complex*, Routledge, 2007.

¹⁶⁴ People will of course point to Tor as a counterexample; it has received funding from DARPA, ONR, the State Department. I don’t think there’s much to explain. Ever large bureaucracy has within it competing and conflicting directions. Some segments of the U.S. government can think Tor is great even when others would like to defund, dismantle, or subvert it.

¹⁶⁵ In the USA, this means AFOSR, ARO, DARPA, IARPA, MURI, NSA, ONR, and more.

¹⁶⁶ Frank Donoghue: *The Last Professors: The Corporate University and the Fate of the Humanities*. Fordham University Press, 2008.

¹⁶⁷ Lorren R. Graham: Money vs freedom: the Russian contradiction. *Humanities*, 20(5), Sept/Oct 1999.

¹⁶⁸ For a description of an incident involving Matthew Green, see: Jeff Larson and Justin Elliott: Johns Hopkins and the Case of the Missing NSA Blog Post. ProPublica, Sep. 9, 2013. For a description of an incident at Purdue involving Barton Gellman, see his article: I showed leaked NSA slides at Purdue, so feds demanded the video be destroyed. *Ars Technica*, Oct. 9, 2015. <http://tinyurl.com/gellman-at-purdue>

¹⁶⁹ For the most part, hardware is gone, formalistic approaches to cryptography are gone (unless they claim to bridge to “real” crypto), cryptanalysis of real-world schemes is little to be seen, and so on.

¹⁷⁰ George E. P. Box: Robustness in the strategy of scientific model building. In: Launer, R. L.; Wilkinson, G. N., *Robustness in Statistics*, Academic Press, pp. 201–236, 1979. Box was not the first to express this sentiment. For example, Georg Rasch explained, in 1960, that “When you construct a model you leave out all the details which you, with the knowledge at your disposal, consider inessential. . . . Models should not be true, but it is important that they are applicable, and whether they are applicable for any given purpose must of course be investigated. This also means that a model is never accepted finally, only on trial.” Georg Rasch: Probabilistic models for some intelligence and attainment tests. Copenhagen: Danmarks Paedagogiske Institut, pp. 37–38, 1960. republished by University of Chicago Press, 1980.

¹⁷¹ Shafi Goldwasser, Yael Tauman Kalai: Cryptographic assumptions: a position paper. Cryptology ePrint Archive Report 2015/907, Sep. 16, 2015.

¹⁷² The entire information-theoretic tradition of cryptography is in this vein.

¹⁷³ Amit Sahai: Obfuscation II. Talk at the Simons Institute. May 19, 2015. Available at <https://simons.berkeley.edu/talks/amit-sahai-2015-05-19b>

¹⁷⁴ Martin Georgiev, Subodh Iyengar, Suman Jana, Rishita Anubhai, Dan Boneh, Vitaly Shmatikov: The most dangerous code in the world: validating SSL certificates in non-browser software. *ACM Conference on Computer and Communications Security*, pp. 38–49, 2012.

¹⁷⁵ API is *application programming interface*, the architected interfaces among component code.

¹⁷⁶ Serge Vaudenay: Security flaws induced by CBC padding: applications to SSL, IPSEC, WTLS . . . *Eurocrypt 2002*.

¹⁷⁷ Viet Tung Hoang, Reza Reyhanitabar, Phillip Rogaway, Damian Vizár: Online authenticated-encryption and its nonce-reuse misuse-resistance. *Crypto 2015*, vol. 1, pp. 493–517, 2015. Guido Bertoni, Joan Daemen, Michaël Peeters, Gilles Van Assche: Duplexing the sponge: single-pass authenticated encryption and other applications. *Selected Areas in Cryptography 2011*, pp. 320–337, 2011.

¹⁷⁸ Chris Soghoian, Workshop on Surveillance and Technology (SAT 2015), Drexel University, June 29, 2015. See also: Chris Soghoian: How to avoid surveillance. . .with your phone. TED talk. <https://www.youtube.com/watch?v=ni4FV5zL6lM>. Stephanie K. Pell and Christopher Soghoian: Your secret Stingray’s no secret anymore: the vanishing government monopoly over cell phone surveillance and its impact on national security and consumer privacy. *Harvard Journal of Law and Technology*, 28(1), Fall 2014.

¹⁷⁹ Christopher Soghoian, personal communications, Nov. 28, 2015.

¹⁸⁰ Despite all these comments, I think a graphic novel on cryptography could be great. Something like the work of Keith Aoki, James Boyle: *Bound By Law: Tales from the Public Domain*, 2006.

¹⁸¹ Exploits that nobody else knows about.

¹⁸² The phrases are from an NSA slide released by Snowden. Reprinted on p. 97 of Glenn Greenwald: *No Place to Hide: Edward Snowden, the NSA, and the U.S. Surveillance State*, 2014.

¹⁸³ Joan Feigenbaum, Jérémie Koenig: On the feasibility of a technological response to the surveillance morass. *Security Protocols Workshop 2014*, pp. 239–252, 2014.

¹⁸⁴ The value refers, of course, to the number of bit needed to identify a living human. See Arvind Narayanan’s website 33bits.org

¹⁸⁵ For a beautiful exposition of this idea, see Eben Moglen: Privacy under attack: the NSA files revealed new threats to democracy. The Guardian, May 27, 2014. The article is derived from the four-part lecture: Eben Moglen: Snowden and the future, delivered Oct. 9, Oct. 30, Nov. 13, and Nov. 4, 2013, the Columbia Law School. <http://snowdenandthefuture.info/>

¹⁸⁶ This suggestion, as well as *conscience-based research*, are from Ron Rivest.

¹⁸⁷ While the term is already in use, the more customary term of art is *privacy-enhancing technologies*.

¹⁸⁸ Of course we cryptographers are not the only ones in the thick of this. People who work on “data science” and “big data” are especially involved.

¹⁸⁹ Bill Joy: Why the future doesn’t need us. *Wired*, 8.04. April 2000.

¹⁹⁰ That said, there is considerable drama in the experiences of people like Julian Assange, William Binney, William Davidon, Tom Drake, Daniel Ellsberg, Mark Klein,

Annie Machon, Chelsea Manning, Laura Poitras, Jesselyn Radack, Diane Roark, Aaron Swartz, Edward Snowden, and J. Kirk Wiebe.

¹⁹¹ Moxie Marlinspike and Trevor Perrin: The TextSecure Ratchet (webpage). Nov. 26, 2013. <https://whispersystems.org/blog/advanced-ratcheting/>. User numbers available at URL <http://www.statista.com/statistics/260819/number-of-monthly-active-whatsapp-users/>

¹⁹² Data from <https://metrics.torproject.org/userstats-relay-country.html>, 2014-11-01 to 2015-11-29.

¹⁹³ An open letter from US researchers in cryptography and information security. Jan. 24, 2014. <http://masssurveillance.info/>

¹⁹⁴ H. Abelson, R. Anderson, S. M. Bellovin, J. Benaloh, M. Blaze, W. Diffie, J. Gilmore, M. Green, S. Landau, P. G. Neumann, R. L. Rivest, J. I. Schiller, B. Schneier, M. Specter, D. J. Weitzner: Keys under doormats: mandating insecurity by requiring government access to all data and communications (2015). Available at http://www.crypto.com/papers/Keys_Under_Doormats_FINAL.pdf. 2015. Earlier, related report: H. Abelson, R. Anderson, S. M. Bellovin, J. Benaloh, M. Blaze, W. Diffie, J. Gilmore, P. G. Neumann, R. L. Rivest, J. I. Schiller, B. Schneier: The risks of key recovery, key escrow, and trusted third-party encryption (1997). Available at <http://academiccommons.columbia.edu/catalog/ac:127127>

¹⁹⁵ Antonio Regalado: Cryptographers have an ethics problem. MIT Technology Review. Sep. 13, 2013. John Bohannon: Breach of trust. *Science Magazine*, 347(6221), Jan. 30, 2015.

¹⁹⁶ Tor security advisory: “relay early” traffic confirmation attack. July 30, 2014. <http://tinyurl.com/tor-attack1>. Tor project blog: Did the FBI pay a university to attack Tor users? Nov. 11, 2015. <http://tinyurl.com/tor-attack2>. This article included: “Whatever academic security research should be in the 21st century, it certainly does not include ‘experiments’ for pay that indiscriminately endanger strangers without their knowledge or consent.” Also see reaction like: Joe Papp: The attempt by CMU experts to unmask Tor project software was appalling. *Pittsburgh Post-Gazette*. Aug. 5, 2014. <http://tinyurl.com/papp-letter>

¹⁹⁷ The statement, adopted May 14, 2014 at the business meeting in Copenhagen, Denmark, says: *The membership of the IACR repudiates mass surveillance and the undermining of cryptographic solutions and standards. Population-wide surveillance threatens democracy and human dignity. We call for expediting research and deployment of effective techniques to protect personal privacy against governmental and corporate overreach.* <https://www.iacr.org/misc/statement-May2014.html>

¹⁹⁸ See <https://www.iacr.org/petitions/australia-dtca/>

¹⁹⁹ I do not know the correct attribution for this quote. I heard it from Tom Ristenpart, who didn’t know the quote’s source.

²⁰⁰ Dan Bernstein: Boring crypto. Talk at SPACE 2015. Malaviya National Institute of Technology, Jaipur. Slides and audio at <http://cr.yt.to/talks.html>

²⁰¹ Steven Levy: Crypto rebels. *Wired*, Feb. 01, 1993.

²⁰² John Perry Barlow: A declaration of the independence of cyberspace. Feb. 8, 1996. Eric Hughes: A cypherpunk’s manifesto. March 9, 1993. Timothy May: The cyphernomicon. Sept. 10, 1994. Aaron Swartz: Guerilla open access manifesto. July 2008.

²⁰³ Barry Popik indicates that quote “has been attributed to Greek leader Pericles (495 BC — 429 BC), but only since the late 1990s. The Greek source is never identified in the frequent citations. The quotation appears to be of modern origin.” Blog entry, June 22, 2011. <http://tinyurl.com/not-pericles>

²⁰⁴ Joseph Rotblat: Remember Your Humanity. Acceptance and Nobel lecture, 1995. Text available at Nobelprize.org

²⁰⁵ Bart Preneel, Phillip Rogaway, Mark D. Ryan, and Peter Y. A. Ryan: Privacy and security in an age of surveillance (Dagstuhl perspectives workshop 14401). *Dagstuhl Manifestos*, 5(1), pp. 25-37, 2015.